

Standardized Data-Parallel Rendering Using ANARI

Ingo Wald*
NVIDIA

Stefan Zellmann†
University of Cologne

Jefferson Amstutz‡
NVIDIA

Qi Wu§
University of California, Davis

Kevin Griffin¶
NVIDIA

Milan Jaros||
IT4Innovations, VSB – Technical University of Ostrava

Stefan Wesner**
University of Cologne




Figure 1: Several examples of large sci-vis data being rendered using the data-parallel ANARI paradigm proposed in this paper. From left to right: a) Roughly one billion color-mapped spheres, rendered using HayStack and BANARI. b) The roughly 500GB DNS data set, with volume path tracing on 128 GPUs, also using HayStack and BANARI. c) An iso-surface rendered during an in-situ Ascent session, while attached to an S3D simulation. d) ParaView performing data-parallel rendering on the airplane data set, using our data-parallel ANARI integration in pvservr.

ABSTRACT

We propose and discuss a paradigm that allows for expressing *data-parallel* rendering with the classically non-parallel ANARI API. We propose this as a new standard for data-parallel rendering, describe two different implementations of this paradigm, and use multiple sample integrations into existing applications to show how easy it is to adopt, and what can be gained from doing so.

1 INTRODUCTION

Visualization is about more than rendering, but rendering nevertheless plays a large role in many tools. Rendering is difficult: it was already a hard problem when all such tools could rely on a single common API (e.g. OpenGL); today, it is further complicated through the emergence of a whole host of different vendor specific APIs—which for many good reasons visualization tools are often loath to adopt.

To make it easier for developers to adopt newer rendering technologies—and for vendors, to get their tools adopted—the Khronos organization has proposed the ANARI API for portable cross-platform 3D rendering [33]. ANARI aims to provide a single API that both 3D application developers and different rendering engine vendors can all agree on, providing important benefits to both sides. For application developers, it means they can target a single API without having to adopt vendor-specific APIs, and without compromising the availability of state-of-the-art rendering features. For those developing ANARI implementations, it means higher adoption rates and faster adaptation to technological changes as they become available in back-ends.

To the application developer, ANARI presents itself in a pragmatic way: ANARI is an API for the user to specify scene data and render frames. Content changes are expressed by updating objects

involved in rendering, such as cameras or data arrays containing geometry, materials, colors, etc. These objects ultimately represent a generic interface to the private implementation, where the mechanics of rendering frames is left up to the implementation.

ANARI is not free from challenges, though. Even with a single agreed-upon API, different implementations can and will still differ in what features exactly they will support (and in which form). Thus, applications still need to be aware of which specific implementation they may be running on—and either adopt a least common denominator approach, or have some application features only available from specific ANARI vendors. Still, standardization is encouraging as ANARI is already seeing adoption even in VTK and VTK-m, and through that, in a variety of tools that use these [2, 6, 17, 34].

For this paper, it is paramount to observe that ANARI currently does *not* explicitly cover data-parallel rendering. This does not mean that data-parallel applications such as ParaView/Catalyst or VisIt/libsim cannot use ANARI for their own data-parallel rendering. They certainly can use ANARI to render on each node, and then composite the resulting per-rank images in exactly the same way they have always done, using image-compositing libraries such as IceT. This approach is already used in practice, but is intrinsically limited to whatever image-compositing can or cannot achieve.

Rendering intermediate images independently on each node, sorting them in order and compositing them *at the end* only results in correct images for the most simple shading models. If we are interested in any global effects—even such rudimentary effects as shadows or ambient occlusion—an individual surface cannot be shaded without knowledge of neighboring surfaces; and such neighboring surfaces can, and in general do reside on other compute nodes. This fundamental problem can not be solved merely by using ghost and halo regions, as lighting effects are a global phenomena.

Distributed renderers relying on ray tracing solve this problem effectively by queuing and exchanging rays not at the end of rendering a frame, but instead constantly while the image is generated. GPU renderers based on ray wave-fronts are demonstrably efficient at doing this, making the approach considered state-of-the-art, allowing scientific visualization tools to benefit from higher-quality rendering techniques at very little extra cost compared to standard data-parallel sci-vis renderers.

What we also observe today is that tools like Paraview or VisIt *do* adopt those higher-quality rendering effects, e.g., by integrating OSPRay or VisRTX in their standard rendering pipelines. However, when visualizing larger data sets requiring distributed rendering,

*e-mail: iwald@nvidia.com

†e-mail: zellmann@uni-koeln.de

‡e-mail: jamstutz@nvidia.com

§e-mail: qadwu@ucdavis.edu

¶e-mail: kgriffin@nvidia.com

||e-mail: milan.jaros@vsb.cz

**e-mail: wesner@uni-koeln.de

even when using those rendering engines, have to resort to local shading only because when combined with ordinary image compositing cannot create artifact-free images otherwise. This, in fact, is a puzzling conundrum at the heart of scientific visualization today: sci-vis is already adopting ANARI, with the goal of standardizing around a single rendering API for ray tracing—yet current ANARI cannot properly express the kind of data-parallel rendering that high-end sci-vis requires.

In this paper, we explore and argue for the concept of *data-parallel ANARI*. We do so through the lens of using the existing ANARI API to define—for certain types of data parallel ANARI devices—a *distributed world* whose individual consistent parts are located on different *collaborating* ranks. We show that this can be done by merely defining the *semantics* of a data-parallel ANARI device, and how the different ranks’ individual API calls will *jointly* define such a common distributed world. We then describe two different sample implementations that each implement this paradigm—with very different capabilities and limitations—and show the potential of this approach using a set of diverse applications that make use of these implementations. In particular, we show several examples of what this explicitly distributed data-aware paradigm can realize that a purely application-side compositing-based approach can not.

2 BACKGROUND AND RELATED WORK

2.1 Data Parallel Rendering

Scientific visualization applications can be classified as either post-hoc or in-situ/in-transfer, where in the latter case the visualization and analysis pipeline is executed as the data is generated. Post-hoc is the alternative processing paradigm to in-situ/in-transfer, where data is saved to permanent storage and then loaded by the application to perform visualization and analysis.

At scale, the dominating operation of a traditional renderer that uses rasterization, simple shading, etc., is sort-last image compositing; the literature has focused on optimizing this operation through efficient algorithms such as parallel direct send [8, 14] or radix- k [20]. These developments culminated in the IceT image compositing library [24] which has become the de-facto standard for real-time distributed renderers. The distributed rendering back-ends of visualization packages like VisIt [6] or ParaView [2], as well as in-situ visualization frameworks like libsim [39], Catalyst2 [21], or Ascent [18, 22] all internally build on IceT.

The implication of using sort-last is that compositing happens at the end, once each rank has reduced its rendering operation to a single fragment (color, opacity, and depth) per pixel. In reality, that restricts the application in what kind of content it can render in this way; in particular, compositing with a single fragment only works correctly if data is partitioned in a way where each rank’s data is convex. To remove some of these restrictions researchers have looked at various forms of “deep” frame buffers that can store more than one fragment per pixel, which in turn allows for better handling transparency when data is partitioned in non-convex ways. A very early example of this was proposed by Ma [23], more recent ones by Binyahib et al. [4] and Sahistan et al. [30]. The same concept is also used by OSPRay [37], whose data-parallel rendering mode relies on using a *distributed frame buffer* [35] where ranks can also produce more than one fragment per pixel. One interesting contradiction becomes obvious when looking at OSPRay: while in non-data parallel OSPRay supports advanced rendering techniques realized with Monte Carlo ray tracing, when rendering in parallel within ParaView or VisIt, OSPRay has to switch to simple ray casting without secondary rays, resulting in similar quality as the typical rasterization renderers used by other sci-vis applications.

2.2 Data Parallel Ray Tracing

For ray tracing, data parallel rendering means that any ray on any rank may at any time need intersection with geometry stored by any

other rank. This can be done by using either one of two alternative techniques: fetching the remote data to the rank that traces the ray (typically involving some caching of said data, e.g. [7, 19]), or sending the ray to the rank that has the data (see, e.g., [1, 11, 28, 31, 36, 38], or some combination thereof (e.g., [26, 27, 29])).

Though some of these approaches have been around for decades, this kind of data parallel rendering has only recently seen interest from the sci-vis community, probably because only recently hardware has become capable enough to actually do this. In particular, SpRay [28], Galaxy [1], BriX [38], and RQS [36] are all been shown to achieve interactive performance for non-trivial data.

Data parallel CPU and GPU ray tracers aimed at sci-vis achieve illumination effects, including shadows or ambient occlusion, but also global illumination with diffuse reflection, by queuing and exchanging rays across ranks [1, 36, 38]. Shading effects are computed by light rays “bouncing”, i.e., they change direction upon interaction with surfaces, volumes, or lights. Each time a ray bounces it consequently has to visit scene content in completely different spots, resulting in highly incoherent patterns.

To avoid latency, when exchanging rays across ranks, data parallel ray tracers do that in batches rather than exchanging small groups or even individual rays. A core concept for that is the ray wave-front: instead of computing the recursive ray tracing function per each ray individually, compute kernels are executed per bounce, on all the rays the rank is currently responsible for. After a bounce finished, the ray wave-fronts are synchronized, can optionally be sorted and compacted; this is also the point where a *data parallel* wave-front ray tracer will synchronize its wave-fronts across ranks using MPI broad-casts or uni-casts. The ranks exchange the rays they are responsible for until a maximum number of bounces was reached, until all wave-fronts contain zero rays; etc. The main difference of state-of-the-art methods these days is how the renderers compute, cull, and assign ray wave-fronts to ranks in-between bounces.

Traditionally, the reason for distributing rays was to reduce the overall bandwidth. Optimizing ray tracers account for that by building complex culling acceleration structures [1, 38] to minimize the overall number of rays exchanged. Those data structures can be extremely unbalanced, especially when scene data is instanced, requiring additional constraints to avoid overflowing memory [40].

An alternative approach, also adopted in this work, by the data parallel ANARI back-end described in Section 5.2, is to distribute the scene data in a very simple way—e.g., round-robin, bin-packing until all GPU memory is used, etc.—and then just generate full wave-fronts on all ranks. The ranks can initially discard some of the rays in their wave-fronts by testing them for visibility against their geometry, but still generally exchange more rays than with an optimized partitioning; wave-fronts are exchanged in a ring buffer or similar pattern. The reason that this works and is not overly costly is that wave-fronts are sent in lock-step, i.e., all the communication happens at the same time and is then limited by the pair of ranks with the highest ray count to exchange. When adding more ranks, the overall ray count stays the same, resulting in more ranks exchanging fewer rays individually. This concept was evaluated by Wald et al. [36] on multi-GPU systems with NV-Link interconnect, but also extends to massively parallel architectures like GPU clusters. One compelling reason for preferring this method over approaches that optimize for bandwidth is its negligible pre-processing time; whereas building an optimized acceleration structure can often take hours and is then impractical for visualization in a production context.

2.3 ANARI

ANARI is a cross-vendor 3D rendering API maintained by the Khronos Group. ANARI connects applications from diverse domains to any 3D rendering engine implementing the ANARI API while still giving implementations a vast degree of freedom of how exactly rendering is done.

ANARI has already been adopted by major scientific visualization packages, namely VTK [32], VTK-m [25], ParaView [2], VisIt [6], VMD [17], and OVITO [34]. ANARI is not limited to scientific visualization, and has also been integrated, for example, in Blender [12], and OpenUSD’s Hydra subsystem [9].

A comprehensive introduction to ANARI is out of the scope of this work; instead we briefly summarize what we consider the main components of the API that are relevant for data parallel rendering. For a complete overview we refer the reader to [33].

ANARI’s core design centers around opaque handles to objects representing the various bespoke actors commonly found in rendering an image—surfaces, materials, volumes, cameras, lights, frames etc. These objects are parameterized through generic parameters represented by name/value pairs and are transitioned between states using parameter commit semantics—specifically that committing the object’s parameters indicates that those changes should be visible in the next rendered frame. Alas, frames are rendered asynchronously, where the application triggers a render operation to start, and then is free to synchronize with it to access resulting output buffers.

The foundational object for most API calls is the *device*, which represents the instance of the rendering engine handling ANARI API calls. After creating a device, applications make all ANARI API calls through this special handle, which provides implementations a point to reconcile any common implementation-wide state and gives applications a clear set of rules for how implementations can be used concurrently. Thus for the rest of this paper, the phrases “ANARI implementation” and “ANARI device” will be used synonymously.

Once the ANARI device is initialized, a valid `ANARIFrame` object is all that is required for rendering. A frame object has color, depth, and other optional auxiliary memory buffers allocated that it can retire rendered pixels to. A valid frame has a `ANARIRenderer` object, a `ANARICamera`, and `ANARIWorld` assigned to it. The ANARI world defines a declarative scene graph. It is a special group node serving as a collection of surfaces, volumes, and light objects. In addition it may also contain `ANARIInstance` objects that themselves contain groups of surfaces, volumes, and lights. `ANARIWorld` is a special group node that always represents the top of the (up to 2-level) scene hierarchy. Buffer objects for input and output are realized as `ANARIArray`s. Arrays deviate from the otherwise strict immediate mode model in that they can be mapped and their content altered by the user.

3 DATA PARALLEL ANARI (DP-ANARI)

The goal of this paper is to propose a paradigm for data parallel rendering using ANARI that fits the object model of a data parallel ray tracer without the limitations of image compositing. One option would be to propose a completely new API, but this would be asking the scientific visualization community to completely discard all existing ANARI progress, and start anew. Instead, we decided to look into what it takes to extend the *existing* ANARI API to also do data parallel rendering. As it turns out, this can actually be done without *any* additional new API calls, by simply proposing a new set of semantics of what the calling of different API calls on different ranks actually means.

The core of our work is defining the semantics of ANARI API usage in the context of a distributed application environment. Specifically, all ANARI API calls are usable as-is, but have additional semantics and constraints applied to them. The following subsections will outline these semantics and constraints.

3.1 Object Locality and Consistency

A concept that occurs in distributed rendering is how objects relate to one another between nodes. For some objects, they are only defined and interacted with on the node in which they are created, while for others there must be a global definition to them on every rank. The

following describes the different application of object definitions with respect to their global or local definitions.

3.1.1 Globally Consistent Objects

Some objects are considered `global` in the sense that they represent a single, cooperative entity in the ANARI object hierarchy on all ranks. We define all `ANARIFrame` and `ANARIWorld` objects to be considered global objects where their global identity is established by the order of their construction. Thus all ranks which use an `ANARIFrame` or `ANARIWorld` handle must have those objects constructed as their respective *N*’th object of that type, effectively requiring all ranks to create the same number of these objects.

Some objects must have their parameters match on every node in order to have a well-defined image—`ANARIFrame`, `ANARIRenderer`, and `ANARICamera`. Applications must use the same sub-type and parameterize these objects identically, otherwise the output of the resulting image will be undefined.

3.1.2 Locally Defined Objects

All objects under `ANARIWorld` are locally defined within the rank on which they are created and are globally visible (i.e. secondary illumination when applicable) in the final rendered image. This includes anything which can be contained with the `ANARIWorld`—instances, groups, surfaces, volumes, geometries, materials, spatial fields, samplers, and even arrays themselves. There is no application requirement that any object within the world has any knowledge or connection to an object on any other rank.

3.1.3 Locally Mapped Frame Buffers

One question is how and where the data parallel application can actually access the pixels that a distributed `anariRenderFrame` call has produced. While this sounds like a trivial problem, it is not: it is easy to imagine some applications wanting to map the entire frame buffers on all ranks, or for others to have different ranks map different regions of a frame, or ranks just having some call-back mechanism for image tiles that a given rank has produced (see, e.g., [15, 35]). On the other hand, trying to capture all these potential use cases would not only require significant extensions to the API, but also raise the cost for device developers to implement all these options.

For this trade-off, we intentionally opt for simplicity and specify that data parallel ANARI devices are only responsible for providing the final frame buffer on rank 0¹: Though all ranks participate in specifying and rendering the scene only a single rank—rank 0—contains the final image buffers (designated by a parameter on the participating `ANARIFrame`); and on this rank the application can map this using the existing `anariMapFrame` in exactly the same way a non-parallel application would have.

While it is not an error to map the frame on another rank, its dimensions, pixel type, and buffer contents are undefined. Given that this operation is inherently local, mapping frame buffer outputs has no synchronization requirements across nodes. ANARI’s asynchronous frame rendering semantics however still apply as they do in traditional single-node rendering setups. An added benefit of choosing this route is that there is no difference whatsoever in how a data-parallel application maps a frame buffer on a data parallel device, vs. how a traditional single-process application does on a non-parallel device (also see Section 7.1). Of course, any use cases not captured through this paradigm could still get added later on through ANARI’s extension mechanism.

¹This refers to rank 0 of the MPI communicator used to initialize the data-parallel ANARI device; if the application wants this to be on another rank than *its* rank 0 it can of course use MPI’s `split` operation to create a new communicator for this.

3.2 Collaborative Operations

Most ANARI API calls can be done independently on each rank, but a few will behave as a rank synchronizing operation—in other words, some ANARI API calls will necessarily require all ranks to participate in and will implicitly barrier at that call.

The first and most obvious synchronizing API call is `anariRenderFrame`, as it is the central place that all object parameter transactions must be completed and where the vast majority of the implementation’s work is done. While the mechanics of rendering a frame is intentionally left as implementation defined, applications must follow this semantic to express that every node is ready for its local `ANARIWorld` is ready to render.

Similarly, `anariGetProperty` is a synchronizing API for objects that have a global identity (`ANARIFrame` and `ANARIWorld`). This is to ensure that implementations can guarantee global consistency to those objects when properties are queried by the application. However, this constraint is only required when applications pass `ANARI_WAIT` to `anariGetProperty`, as `ANARI_NO_WAIT` would indicate it is permissible for the device to return a previously held value (or none at all) in order to prevent blocking.

Finally, `anariRelease` is a synchronizing API when called on the `ANARIDevice` which has no remaining application references. This permits implementations to rely on the device being destroyed in lock-step and guarantee no additional ANARI API calls will be made using the device.

3.3 How this works out in practice

We observe that this new data-parallel paradigm in practice actually matches how existing data parallel applications work. An application using IceT for compositing would also use MPI for synchronization. The individual ranks would operate largely independent of each other and render different content using calls to whatever library the application uses for per-rank rendering. Similarly, that application must also ensure consistency between different ranks’ global rendering information like camera, background color, etc., so it would likely already have some mechanism to synchronize such information before it calls rendering. For our paradigm, this is exactly the same, except that we apply it to ANARI, and formalize the process.

In Fig. 2 we illustrate exactly the workflow: Upon startup all ranks would (collaboratively) create their DP-ANARI device, then do whatever the application wants to do for data loading, iso-surface extraction, etc. At some point each rank creates its local ANARI world, and populate it with different ANARI objects—these calls are not collaborative, so different ranks can do as many of those—and whichever ones they want—without any other ranks even being aware. Once all ranks are ready for rendering they all call `anariRenderFrame()`, at which point they implicitly synchronize the operation until the frame is done. Rank 0 can then map the generated frame buffer, save or display it, and wait for UI events or user input, at which point it will instruct its worker ranks to do whatever scene updates are required, etc.

4 EVALUATION CHALLENGES AND METHODOLOGY

What we have discussed so far is not a specific *method*, nor a *system*. It cannot be evaluated via any one implementation, nor via any one application using it, nor via any one or more use cases thereof. Ultimately, the success of this proposed paradigm will depend on whether—and to what degree—it will actually get adopted for day-to-day data-parallel rendering, in tools such as ParaView or VisIt.

Any such adoption is challenging because there is an inherent chicken-and-egg problem that must be solved: applications will not adopt any API or paradigm that there are no compelling device implementations for, and for device implementers it makes little sense to create such compelling implementations if there is no plausible path for them to be adopted—nor is it easy to develop them if there are no use cases to exercise them.

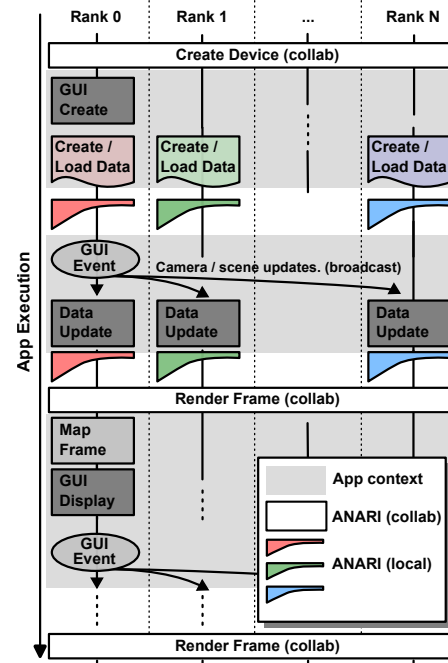


Figure 2: Illustration of how a typical data-parallel application would use our paradigm (also see Section 3.3).

Eventually—and this is the purpose of the work described in this paper—breaking this log-jam requires three *simultaneous* efforts: a) some example device implementation(s) that implement this paradigm, and that end-user applications can actually target; b) some reasonably complex applications that actually use this paradigm, that use it in a way that is reasonably representative of how the eventual applications-to-be will use it, and that device implementers can use to develop, debug, and tune their implementations; and c) some reasonably compelling proof-of-concept use cases that show that this paradigm is actually worth adopting, and thus can serve as an incentive for the other two parties to actually work towards this goal.

5 EXAMPLE REALIZATIONS

To serve as sample *implementations* of our paradigm we created two different ANARI devices that implement it. Both devices implement the same paradigm, but are by no means identical: they do not offer the same feature set, nor will they produce the same images for the same inputs. This is OK—different applications have different needs, and different devices will always offer different features. For our purposes, we intentionally chose two opposite extremes of the spectrum: one relying on compositing, and one performing true data-parallel path tracing.

5.1 ANARI-Composite: Application-transparent compositing using an ANARI Pass-through Device

The bespoke non-distributed ANARI API—i.e., ANARI without the paradigms introduced in this paper—has no concept of a “data parallel world”. However, that does not mean that applications *using* ANARI could not do data-parallel rendering of their own: such applications could use ANARI to render data locally on each rank, and then rely on depth- and/or alpha compositing to somehow composite the resulting images. This approach has some obvious limitations in terms of what effects can or cannot be rendered, but these limitations are not specific to ANARI: they are not different from those when using any other rendering back-end for the per-rank rendering.

Where these limitations are acceptable, one could use exactly that same approach also *within* an ANARI device to implement the

data-parallel ANARI paradigm we have introduced above. In such a device, almost all functions would behave in exactly the same way as in any other local device; namely, they would set up the scene to be rendered, and perform rendering of a local frame buffer. The only thing this device would need to modify is `anariRenderFrame()`, in which it would first perform its local rendering, and then composite the results. We decided to implement this approach, for which we need two main ingredients: a means of performing local-node rendering, and a means of compositing.

5.1.1 Compositing

The standard way of implementing alpha and depth compositing is through IceT [24], but for our general use case this library is too restrictive. Alpha blending with IceT requires that the application provides a fixed compositing order. With simple pass-through that compositing order is quite difficult to determine as the input data would need to be classified into opaque or transparent, the data convexly partitioned, etc., which would require data specific knowledge the pass-through device cannot have.

To circumvent this we require that any fragment generated by the local device for a pixel not only has an RGBA but also a depth component, which allows us to instead rely on another compositing library we had readily available: the *deep compositing* (`deepComp`) library originally developed for another data-parallel rendering project [30]. Using this library an application (in this case, our ANARI compositing device) can render possibly multiple RGBA-z fragments for each pixel; the compositor will then follow a *parallel direct send* [8, 14] paradigm to have all ranks exchange their fragments such that each rank gets all the fragments for some portion of all pixels. Each rank receives its pixels' fragments, and then, in a CUDA kernel, sorts and composites each pixel's fragments in proper front-to-back alpha-composited order. The resulting composited pixels are gathered at rank 0. For our purposes we only need a single RGBA-z fragment per pixel (on each rank), but then no longer need to worry about which order the different ranks generate these fragments.

5.1.2 Rendering

For rendering itself, we leave most of the heavy lifting to ANARI itself, by using another *existing* ANARI implementation for what we call a "pass-through" device. We pass all other calls through to this device and let it do local rendering, as long as it is capable of computing both color and depth buffers.

To implement this compositing device, we can simply take every ANARI API call and pass it on to the pass-through device, except for the following:

`anariLoadLibrary()`: we first load our own data parallel device, then on each rank also load the device requested by the application—and save that as a pass-through device.

`anariNewFrame()`: we first pass this through to the pass-through device, and intercept the returned `ANARIFrame` handle for that rank. We then create our own—collaborative—frame object that first creates a new deep compositing context for that frame, and then also stores that rank's intercepted pass-through frame handle (through which we can later access that rank's local frame).

`anariCommitParameters()`: we pass this through to the pass-through device, but also check if this commit affected a frame, and if so, whether it resized that frame (and if so, resize that frame's compositing context).

`anariRenderFrame()`: we pass this to the pass-through device to perform local rendering, then wait for that to finish, map the local rank's frame using the stored pass-through device frame handle, and perform compositing. Compositing in `deepComp` requires collective MPI calls, but in our data-parallel ANARI paradigm `anariRenderFrame` is collective, so this poses a problem.

`anariMapFrame()`: this is the one call we do not pass through at all, since local frames have already been read and composited in `anariRenderFrame`. We simply retrieve the composited image from the compositing context, and return this.

One advantage of this approach is that it is straightforward to implement: given an existing compositing library, we implemented a working proof of concept with very little effort.

What is particularly useful is that because our compositing device itself works by issuing ANARI calls to the pass-through device it can actually use *any* other existing ANARI device for the actual rendering, without having to know which. This makes this approach useful as an easy "fall-back" mechanism of using any other—not yet natively data parallel—ANARI device in a data parallel context.

The downside to this approach is that it is intrinsically limited to what compositing can or cannot do. Using the `deepComp` library means we can avoid some of the specific limitations of IceT—in particular, we do not need to specify a fixed compositing order—but it still relies on compositing, and thus will never be able to produce guaranteed-correct shadows or path tracing for data parallel content.

5.2 Barney and (B)ANARI

Barney is a new—and still under development—rendering engine for data-parallel path tracing on multi-node and multi-GPU hardware. For rendering, Barney relies on ray forwarding similar to what is done by the recently published Brix [38] and RQS [36] papers, where rays are sent to the node(s) that may have geometry that may intersect a given ray—and where each ray will always find its respectively closest intersection no matter which rank the ray was spawned on, or which rank holds that respective geometry.

Barney was built with parallel rendering—and in particular, data parallel rendering—in mind from its very inception. In addition to the relatively simple mode we described above—where each rank has exactly one part of the data to be rendered—Barney also offers various additional modes such as, for example, data-replicated rendering, islands-parallel rendering, non-MPI multi-GPU rendering, additional multi-GPU data-parallel rendering within a given rank, etc. Despite this bigger set of functionality, Barney follows the same general paradigm described above: it has the concept of a data-parallel world, where different ranks can independently specify different pieces of this world, and render operations are MPI synchronous.

5.2.1 BANARI

Barney is not *exclusively* built for ANARI, but targets the same end-user applications, and thus supports similar functionality: it supports both surface and volume types, and in particular also supports the more scientific visualization oriented data types of cylinders and spheres for surface data, or unstructured mesh and AMR data for volume data.

With all these pieces in place, implementing a data-parallel ANARI device was relatively straightforward: this mostly required implementing the various ANARI API functions to properly read scene data passed through this API, and passing it on to its matching data types, where applicable. As with many other ANARI implementations, the Barney ANARI device—or BANARI, for short—implements only a subset of the full set of ANARI's different data types, and simply ignores all others.

5.2.2 Local vs. Global Rendering

Barney is, by nature, designed for data-parallel rendering. However, to also be accessible to non-data parallel applications it can also be built without MPI support, in which case it simply performs local, data-replicated multi-GPU rendering.

If used in this way, the BANARI device still implements the ANARI API (just without out data-parallel semantics), and we can

thus also use this as a pass-through device for the ANARI compositing device described in Section 5.1.

6 EXAMPLE INTEGRATIONS

Whereas the previous section showed that it is possible (and in fact, not all too difficult) to write devices that implement our paradigm, in this section we look at the reverse problem of integrating such an API in an application that wants to use data-parallel (DP)-ANARI.

6.1 Minimal, Proof-of-Concept Applications

We are ultimately most interested in how challenging it is to integrate our data-parallel rendering with ANARI paradigm into actual end-user applications like VisIt or ParaView, or widely used frameworks such as VTK. However, we could not evaluate that until some implementation(s) existed, but neither could we have developed the aforementioned back-ends without applications to exercise it. To break this chicken-and-egg problem we decided to defer integration into actual end-user tools to the end, and in the meanwhile, relied on developing both back-ends and several different front-ends in parallel.

6.1.1 OSPRay and TSD Mini-Applications

As one proof-of-concept we started with several existing mini-applications from the OSPRay project. The semantics of our data parallel ANARI and those of data parallel OSPRay are (intentionally) quite similar; the API calls and function names differ, and so do some low-level concepts, but the application flow is very similar. To prove the generality of data parallel ANARI we took the data parallel sample applications coming with OSPRay and ported them to our semantics—where possible line-by-line. This allowed us to exercise our semantics on some very simple data parallel applications outside our own software ecosystem.

This provided an early proof of concept, but does not come with much interesting data to test with. To get more realistic inputs, our next step was to take an existing single-rank ANARI viewer (TSD, from visRTX [3]), and prototypically extend that to use MPI parallelism: all ranks load different parts of the input, then rank 0 runs the existing viewer, and broadcasts UI updates and render requests to the worker ranks; workers wait for such broadcasts, then perform ANARI scene updates and call `anariRenderFrame`. This all naturally follows our paradigm, meaning all the effort in this proof-of-concept was in changing this viewer to be MPI-parallel in itself—with no extra work for our paradigm at all.

6.1.2 HayStack, and HANARI

In order to have a more challenging use case we also took Barney's original HayStack viewer, and prototypically ported (parts of) that over to ANARI. HayStack was originally developed for Barney even before Barney's ANARI device was developed.

HayStack is what we consider a developer-centric minimal GUI application with the minimum at what it takes to manipulate cameras, edit transfer functions, etc. While HayStack is intentionally minimalist when it comes to the user interface, it was from the beginning designed to be able to stress-test Barney to the end of its abilities, and to be as realistic a mock-up of what a real application like ParaView would be as it possibly could. In particular, HayStack supports importers for many different data types including structured and unstructured volume data; triangle meshes, spheres, cylinders, and even production-style data with instances and textures; it contains facilities for data-parallel loading, offline and on-the-fly partitioning (including both object-space and spatial partitions), where desired, different data load balancing schemes (to simulate different ways of how an application might assign geometry to different nodes).

While originally *not* under ANARI, this wealth of different data configurations made it an attractive candidate to also add an ANARI-based render pass. As with the mini-applications, virtually all the

work required in doing so is the traditional ANARI calls to create and provision each rank's geometry, with no special effort to meet our DP-ANARI requirements whatsoever. We observe, however, that this ANARI path currently only supports some of HayStack's data types, and even then does not necessarily produce exactly the same images. Usually, this is because some of Barney's data formats and material types are slightly different from ANARI's. Nevertheless, the key insight from this exercise matched exactly what we saw for the significantly simpler mini-applications: almost all the effort required for using a DP-ANARI device lies in how a given rank specifies its ANARI—which is exactly the same as it would have been for per-rank ANARI rendering—with virtually no extra effort required for the data-parallel portion.

6.1.3 Data Parallel VTK Mini-Application

Being one of the most popular scientific visualization frameworks, interoperability with the Visualization Toolkit (VTK) [32] was an obvious choice. In preparation for further integration with ParaView [2] we started with a simple proof of concept application using VTK internal classes only.

To add a new rendering subsystem to VTK, a custom "pass" must be registered with the `vtkRenderer`. VTK already has builtin support for rendering with ANARI via the `Rendering/ANARI` module, which co-exists with other rendering modules (e.g., for core OpenGL, or OSPRay), and which we build on.

VTK itself has no concept of any data parallel rendering; any data parallel loading and rendering has to happen on the application side. To do this we started by implementing a simple MPI-parallel application in which each ranks loads its respective geometry, assuming pre-partitioned data on disk. For the user interface we created a simple VTK GUI application using the `vtkRenderWindow` class, which implements a platform-specific event loop utilizing GLX, WGL, or similar depending on the target platform, and calls the `Render()` function of the `vtkRenderer` object attached when redraw events occur. Simply attaching a `vtkAnariPass` object to the latter we achieve serial rendering.

The actual GUI window only runs on rank 0, but to make our paradigm work, the workers also need to run the same VTK render passes, in synchronous mode. We solve that by having each worker create an off-screen window of the same window class that the display rank uses: this creates the same pipeline as on the display rank, but obviously does not get any of that rank's UI events. We solve this by having the display rank perform MPI broadcasts for events that affect global data, such as resize, camera changes, or render. Workers then implement a custom event loop in which they listen for such requests from the display rank, upon which they then issue the corresponding events.

One caveat with the standard `vtkRenderWindow`, which was not implemented with data parallelism in mind, is that events that require lockstep processing can be triggered unexpectedly; on X11 for example, a resize event *plus* redraw event only gets triggered when the window size increases. To guarantee lockstep execution, we implemented a pass-through extension via a C++ class inheriting from `vtkAnariPass` that intercepts these events and communicates with the worker's event loop using a simple custom communication protocol on top of MPI.

In retrospect, most of the challenges of this exercise turned out to be related to VTK—and in particular its event system—not having any concept of other clients it might have to synchronize with, and wanting to issue supposedly-synchronous ANARI calls at seemingly-random UI events. This made the integration into nominally more complicated applications (see upcoming sections) actually *easier* than what was originally intended as a warm-up exercise for that.




Figure 3: Screenshots from our prototypical integrations of our DP-ANARI paradigm into VisIt/Libsim and ParaView. a) VisIt/Libsim data set generated by a mini-simulation used to test VisIt’s in-situ capabilities, running on 8 ranks and using our data-parallel ANARI integration (see Section 6.2.3). b) ParaView rendering the $1,225 \times 580 \times 2,000 \times \text{float}$ *aircraft* model with Q-criterion vortex field (see Section 6.2.1). This case was rendered with BANARI on 4 ranks of 1 GPU each. c) ParaView rendering an iso-surface extracted from a $2 \times$ reduced version of the DNS data set. d) ParaView rendering the $2,501 \times 3,001 \times 305 \times \text{float}$ *thunderstorm* data set, in this case rendered with BANARI on 4 ranks of 1 GPU each (using Barney’s volumetric path tracing mode).

6.2 Prototypical Production Integrations

While mini-applications and *HayStack* were invaluable from a developer’s point of view, the eventual goal for any data-parallel ANARI effort must necessarily lie with actual “real” end-user visualization applications such as ParaView, VisIt, or various in-situ frameworks.

6.2.1 ParaView

To implement our DP-ANARI semantics into ParaView [2] we start with the proof of concept application described in Section 6.1.3. In contrast to VTK, ParaView supports full data parallel rendering via a client/server architecture. Synchronization issues such as the ones encountered with `vtkRenderWindow` in our test application hence are not an issue. To run ParaView in data-parallel mode, we start the `pvserver` application with as many MPI ranks as we have workers, and connect to that using the ParaView client. In that setup, the ParaView client does not participate in the MPI-parallel rendering, but will receive remote-rendered, compressed images from the dedicated MPI worker rank that eventually holds all the pixels.

Although exposing ANARI-based rendering in ParaView is straightforward (using similar routines as for *OSPRay*), this extension has not yet found its way into the upstream repository; our work builds off a custom implementation comprising a handful of ParaView classes to expose VTK’s `Rendering/ANARI` in the GUI.

ParaView realizes compositing via *IceT*, but for our paradigm it does not make sense that the application should do that. One option would have been to take this component out of ParaView, but this would have required non-trivial changes to ParaView. Instead, we simply have rank 0 pass the already-finalized image to *IceT*, while all other ranks report an empty frame. When *IceT* “composites” these frames it will simply end up getting what we reported on rank 0; it will have performed some un-necessary computations, but since compositing is not a bottleneck, this is acceptable to us. Using this approach all our paradigm’s demands are already met: rendering and resize are synchronous (*IceT* had this requirement, too), and all per-rank rendering is done through ANARI already.

The only snag we encountered with this approach is that ParaView currently uses an optimization by communicating the local scene bounds to *IceT* on the worker ranks; this allows *IceT* to determine that certain pixels will not have any (directly) visible geometry, and thus can be excluded from compositing. This optimization obviously does not work for our approach where rank 0 always reports all pixels. We currently disable this optimization, in which case our approach works as expected. We also observe that this problem only occurs because we decided to leave *IceT* enabled in ParaView; the moment ParaView were to fully adopt data parallel ANARI it would be cleaner to disable compositing on the application layer, anyway, at which point this problem would no longer exist.

6.2.2 Ascent

To also experiment with an in-situ setup, we also integrated our data parallel ANARI approach into *Ascent*—an increasingly popular, lightweight in-situ visualization and analysis infrastructure [18].

Ascent uses VTK-m, not VTK, so the VTK work from Section 6.1.3 did not apply. We therefore decided to base our integration directly on the ANARI level, which we then did by implementing new “extracts” in *Ascent* (i.e., to follow using a customized renderer type in *Ascent*’s terminology).

In this framework, meeting our DP-ANARI requirements was trivial. Consequently, virtually all the work in this integration was in issuing the ANARI calls to create each rank’s local scene data, which is identical to what would have been needed for a local per-rank ANARI back-end, too. In fact, using DP-ANARI meant that we *only* had to do the ANARI geometry calls, and not worry about any compositing of the results, which in DP-ANARI is not required.

Unlike the previously mentioned mini-applications, this integration was not just done for the sake of evaluation—but as an actual means of enabling high-fidelity data parallel rendering for actual state of the art simulations. In Fig. 4 we show two examples of this: *NekRS*, a GPU-accelerated spectral element Navier-Stokes Solver for incompressible turbulent flows employing an unstructured hexahedral mesh [10]; and *S3D*, a scalable direct numerical solver for reactive and compressible flows, based on a rectilinear mesh [5, 16].


6.2.3 VisIt/LibSim

As a final proof of concept we are also working on an integration into *Libsim* [39]. *Libsim* is an infrastructure for in-situ visualization using *VisIt* [6], so some sense this effort is closer to the *Ascent* effort than it is to our ParaView integration. However, *libsim* connects to *VisIt* for visualization which uses VTK for rendering, so the actual steps to make it work are essentially identical to what we have described above for the VTK mini-application, and works exactly the same.

For the initial testing we used the *globalids* mini-simulation [13] and are thus limited to the relatively simple test geometry produced by this *libsim* mini-simulation. This is obviously not representative for the kind of volume or geometry data that a real *libsim* session would generate. However, what geometry and volumes *libsim* generates on a given rank should only affect what local per-rank operations the underlying ANARI VTK renderer would need to perform. While this clearly needs more testing and profiling, we believe that the key concepts are already in place. A screenshot of this effort is shown in Figure 3 (in this case, rendering through BANARI)

7 A CASE FOR ADOPTING THIS PARADIGM

In the last two sections, we have shown that it is reasonably easy to *realize* our DP-ANARI paradigm in back-ends, and that it is similarly easy to *integrate* it into new or existing data-parallel applications. Together, these two arguments show that our proposed paradigm is *realistic* in the sense that if one is seriously interested in standardizing towards an API for data-parallel ray traced rendering, then this paradigm could indeed work. In this section, we are making an argument as to *why* any given application should care.



NekRS: pb146 pressure

S3D: ammonia pressure (iso-surface)

S3D: ammonia, velocity magnitude (volume rendered)

Figure 4: Some images from our ASCENT integration, used on two different HPC simulation codes and data sets (rendered using BANARI).

7.1 Impact on Existing, Classic-ANARI Applications

Our first observation is that our paradigm is purely additive in that it does not take *anything* away from existing state of the art. A single-process application already using ANARI (e.g., TSD or Blender (cf. Fig. 5)) will use ANARI in the exact same way, which will have the exact same performance implications as it did before. Similarly, an existing data-parallel application that however already uses ANARI for local, per-rank rendering (e.g., mainline ParaView without our modifications) would also see no negative effects whatsoever when using the DP-ANARI paradigm and whatever means it uses to render images in parallel.

Let us now consider this same application to follow the (simple) steps outlined above (Section 6.2.1) to actually enable that device’s data-parallel capabilities, and follow our paradigm. Assuming it simply used our compositing device (passing through to whatever it used before) it still would not see any difference whatsoever. At this point the application could already get rid of its own compositing layer, and still not “lose” anything it could have done before. Though the application would still not have seen any benefit, it also would not have lost anything by adopting our paradigm.




Figure 5: Examples of our paradigm on traditional—i.e., *not* data parallel—applications: a) a stand alone ANARI-based viewer (in this case, using BANARI). b) blender, using ptc/visionaray. Though our paradigm does enable data-parallel path tracing, it does not change anything for traditional, non-parallel applications, which can therefore use our DP-enabled devices without any changes.

7.2 Benefits of Adopting a DP-ANARI Paradigm

Our motivation for a DP-ANARI paradigm is to enable global effects usually realized with ray tracing in data-parallel renderers. Ordinary data-parallel renderers cannot produce these effects without artifacts, as can be seen in Fig. 6. To illustrate the effect of missing global effects for some more realistic data, in Fig. 7 we show two large data sets (one volumetric, one surface based) rendered with Barney, and in exactly the same data-parallel configuration—but once with only local shading, and once with path tracing turned on. Clearly the images with global effects are not only “nicer”, but also convey more information (which is what visualization is about)—but with classical ANARI (i.e., without exposing the concept of a data parallel world) this could not be reproduced. An application could of course still decide to not use ANARI for such use cases at all, and instead integrate into libraries like Barney in the first place—but this would lose all the benefits of why applications are integrating in the first place.

Whereas Fig. 7 explicitly illustrated data-parallel rendering with and without advanced ray tracing effects, in Fig. 1 and Fig. 8 we




Figure 6: Capabilities and limitations of compositing: A test case of 4^3 boxes pseudo-randomly interleaved across 4 ranks. Left, rendered with the pass-through compositing (PTC) device, compositing images from another ANARI device that does not know that other ranks (or their data) even exist. Our ANARI PTC device can properly z-composite these images, but cannot create (global) shadows. Middle, the same data-parallel application using a natively data-parallel ANARI device (Barney, in this case), with proper global shading effects. Right: difference images generated with FLIP. Bottom row: similar experiment, conducted with the DNS data set in ParaView.

provide several more examples of what an application could expect when adopting a data-parallel path tracing paradigm.

8 DISCUSSION

In this paper we have proposed a paradigm for using ANARI for data-parallel rendering. We have not proposed any new method, nor a new API, nor even a specific system. Instead, this paper should be seen as an attempt to rally both developers and users of ANARI to *agree* on a specific way of using ANARI. Doing this is what we believe to be the key to breaking the chicken-and-egg problem in which—at least for data parallel rendering—applications are stuck with compositing, while ANARI device developers cannot develop data-parallel renderers because existing ANARI is purely single-rank.

We have formalized a workflow that is simple and flexible, yet sufficiently expressive to work for both compositing and true data-parallel path tracing. We have shown that it is quite easy to implement this paradigm (for both of these categories), and have used a variety of different (prototypical) integrations to show that this is also easy to adopt. We have also shown that there is a very easy on-ramp for applications to ease into this paradigm, by simply using our (or any similar) compositing device—if this is used with whatever ANARI device the application is currently using, the application would get exactly the same outcome as before, while being able to also run any true data-parallel path tracer when desired.

8.1 Limitations

The most obvious limitation of our approach is that it is not as easily tangleable as any new API extension would be. At its heart our paradigm only specifies a *convention*, and even that would eventually need some sort of formalization in the ANARI specification. This is, however, not all that different from how ANARI works in general: for example, ANARI does the API call for creating a material with a given name, but it does *not* make any guarantees how (or even whether) a given device will implement a given material. In practice,

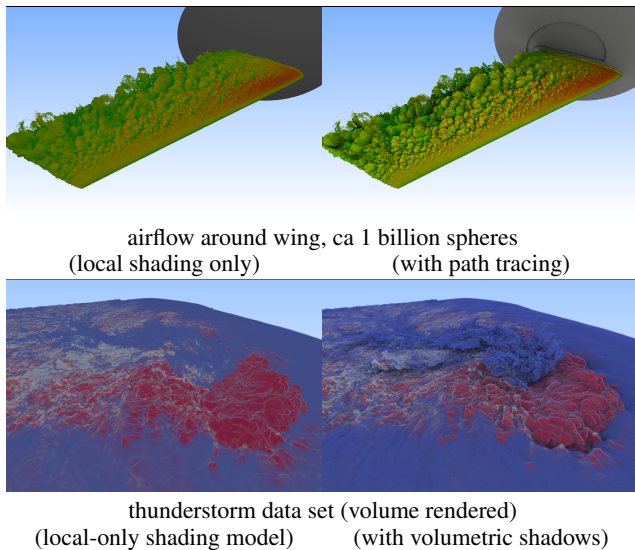


Figure 7: Examples of why path traced rendering is important in the first place, even for data-parallel visualization. Top: a roughly 1 billion spheres data set from NASA AMES (showing airflow around a wing). Bottom: the $2,501 \times 3,001 \times 305 \times \text{float}$ thunderstorm data set. Left: with local shading only; right: with path tracing turned on. All images rendered with Barney, HayStack, and in data-parallel.

this still works, due to what we call a *normative pressure*: once enough applications start to expect a given material to operate in a given way, device developers come under significant pressure to implement it the way that those applications expect. We fully expect adoption of our paradigm to work exactly the same way.

Another limitation is that our paradigm specifies how a given scene is to be created, but does not make any guarantees about how a given device will then render it. This is yet another example of ANARI being intentionally vague, and relying on said normative pressure; however, we fully expect compositing-based devices to remain in use for a considerable while, and applications will have to decide how to deal with that.

8.2 Scene/Data Partitioning

One key issue for data parallel rendering—which we have completely disregarded thus far—is how the scene is partitioned across the different ranks. This is important because many data-parallel renderers will only work for certain types of data partitioning. For example, IceT’s alpha blending mode [38] requires a spatial partitioning of the scene as well as an a-priori known compositing order; and similar limitations would apply to other renderers. In scientific visualization, this problem gets even more interesting because it is not the renderer that does the scene partitioning, but the application—such as `pvsolver` for ParaView, or `libsimg` for VisIt.

Clearly, if there is such a strong dependence on how the scene is partitioned, any API or paradigm for data parallel rendering *must* have a means of communicating what the back-end can consume, and/or what the front-end has generated. One way of solving this would be to specify a certain partitioning requirement in the API, but this would unduly restrict what kind of renderers could or could not be implemented.

Instead, we suggest to handle that by having the applications pass such meta information by setting a set of parameters on the underlying device. For example, applications that do want to use devices that use IceT could set some `int compositingOrder` and `box3 boundingBox` parameter on the device. Of course, this only works if the application can actually provide such data, but if it couldn’t

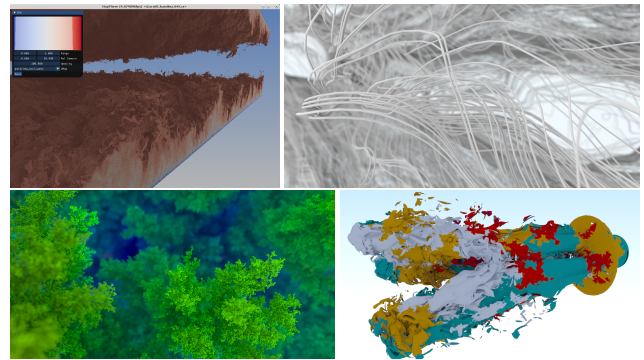


Figure 8: Several examples of data-parallel path traced data. Left to right, top to bottom: a) The $10,240 \times 7,680 \times 1,536$ DNS data set, volume path traced on either RTX 8000 cards. b) streamlines from a jet flow simulation rendered with depth of field c) 250M particle DLA simulation also rendered with depth of field d) A total of order 1 billion, extracted on-the-fly from the order 1 billion element NASA Mars Lander dataset (see [30], and color-coded by which rank owned the underlying mesh part.

it wouldn’t be able to use IceT, anyway. Here, however, we would expect the aforementioned normative pressure to eventually assert itself, too: if some devices have more constraints than others then clearly these will see some pressure to relieve these constraints.

8.3 Remaining Issues

The key remaining issue is to get the developers of actual tools like ParaView and VisIt to adopt this paradigm. Though we believe this paper to have made a strong argument that they should, this will likely not happen immediately.

Ultimately this will also require more work on the device side. OSPRay, for example, already more or less follows the same paradigm, and already has a (single-rank) ANARI interface—but would yet have to merge these two. For our own devices, much is also left to be done. For Barney, there are still rather large gaps between what Barney supports and what ANARI would expect. Adding these missing features—and changing existing ones to be more ANARI-like—will require significant effort. However, applications will not adopt it until it supports enough of the ANARI features that said application requires.

9 CONCLUSION

In this paper, we have proposed a paradigm—or convention—for how data-parallel applications and data-parallel renderers can use ANARI to jointly reason about a global scene. This by itself clearly does not completely solve the problem of data-parallel rendering in its entirety. However, we believe this paper to have made three major contributions towards that goal: First, to have proposed what is essentially a road-map towards true data-parallel path tracing in scientific visualization, which both application and device developers can follow. Second, a set of arguments *why* application developers should join in this effort, and that there is no longer a reason not to. And third, a set of devices, prototypes, and proof-of-concepts that others can build on (all of which we have made publicly available), and which we believe will be a foundation for reaching a virtuous cycle where application developers and device developers can now jointly work towards a common goal.

ACKNOWLEDGMENTS

This work was in parts supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant no. 456842964. This work was also supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254).

REFERENCES

- [1] G. Abram, P. Navrátil, P. Grosset, D. Rogers, and J. Ahrens. Galaxy: Asynchronous Ray Tracing for Large High-Fidelity Visualization. In *IEEE 8th Symposium on Large Data Analysis and Visualization*, 2018.
- [2] J. Ahrens, B. Geveci, and C. Law. *ParaView: An End-User Tool for Large Data Visualization*. Visualization Handbook. Elsevier, 2005.
- [3] J. Amstutz. VisRTX: A NVIDIA OptiX based implementation of ANARI. <https://github.com/NVIDIA/VisRTX>, 2024.
- [4] R. Binyahib, T. Peterka, M. Larsen, K.-L. Ma, and H. Childs. A Scalable Hybrid Scheme for Ray-Casting of Unstructured Volume Data. *IEEE Transactions on Visualization and Computer Graphics*, 25(7), 2019.
- [5] J. H. Chen, A. Choudhary, B. De Supinski, M. DeVries, E. R. Hawkes, S. Klasky, W.-K. Liao, K.-L. Ma, J. Mellor-Crummey, N. Podhorszki, et al. Terascale direct numerical simulations of turbulent combustion using s3d. *Computational Science & Discovery*, 2(1), 2009.
- [6] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, K. Bonnell, M. Miller, G. H. Weber, C. Harrison, D. Pugmire, T. Fogal, C. Garth, A. Sanderson, E. W. Bethel, M. Durant, D. Camp, J. M. Favre, O. Rübel, P. Navrátil, M. Wheeler, P. Selby, and F. Vivodtzev. VisIt: An End-User Tool For Visualizing and Analyzing Very Large Data. In *Proceedings of SciDAC 2011*. Denver, CO, July 2011.
- [7] D. E. DeMarle, C. Gribble, and S. G. Parker. Memory-Savvy Distributed Interactive Ray Tracing. In *5th Eurographics / ACM SIGGRAPH Symposium on Parallel Graphics and Visualization*, 2004.
- [8] S. Eilemann and R. Pajarola. Direct Send Compositing for Parallel Sort-Last Rendering. In *Proceedings of the Eurographics Symposium on Parallel Graphics and Visualization*, 2007.
- [9] G. Elkoura, S. Grassia, S. Boonyatara, P. Jeremias-Vila, M. Kuruc, and A. Mohr. A Deep Dive Into Universal Scene Description and Hydra, 2019. SIGGRAPH '19 Course Notes.
- [10] P. Fischer, S. Kerkemeier, M. Min, Y.-H. Lan, M. Phillips, T. Rathnayake, E. Merzari, A. Tomboulides, A. Karakus, N. Chalmers, et al. Nekrs, a gpu-accelerated spectral element navier–stokes solver. *Parallel Computing*, 114, 2022.
- [11] S. Fouladi, B. Shaklett, F. Poms, A. Arora, A. Ozdemir, D. Raghavan, P. Hanrahan, K. Fatahalian, and K. Winstein. R2E2: Low-Latency Path Tracing of Terabyte-Scale Scenes using Thousands of Cloud CPUs. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 2022.
- [12] B. Foundation. Blender—Free and Open 3D Creation Software, 2024. <https://www.blender.org>.
- [13] VisIt/Libsim Data/Simulation Examples. <https://github.com/visit-dav/visit/tree/develop/src/tools/data/DataManualExamples/Simulations>, 2006.
- [14] P. Grosset, M. Prasad, C. Christensen, A. Knoll, and C. Hansen. TOD-Tree: Task-Overlapped Direct Send Tree ImageCompositing for Hybrid MPI Parallelism and GPUs. *IEEE Transactions on Visualization and Computer Graphics*, 23(6), 2017.
- [15] M. Han, I. Wald, W. Usher, N. Morrical, A. Knoll, V. Pascucci, and C. R. Johnson. A Virtual Frame Buffer Abstraction for Parallel Rendering of Large Tiled Display Walls. In *IEEE VIS 2020 - Short Papers*, 2020. doi: 10.1109/VIS47514.2020.00009
- [16] E. R. Hawkes, R. Sankaran, J. C. Sutherland, and J. H. Chen. Scalar mixing in direct numerical simulations of temporally evolving plane jet flames with skeletal CO/H₂ kinetics. *Proc. Combust. Inst.*, 31(1), 2007.
- [17] W. Humphrey, A. Dalke, and K. Schulten. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics*, 14, 1996.
- [18] S. Ibrahim, T. Stitt, M. Larsen, and C. Harrison. Interactive in situ visualization and analysis using Ascent and Jupyter. ISAV '19, 2020. doi: 10.1145/3364228.3364232
- [19] T. Ize, C. Brownlee, and C. D. Hansen. Real-Time Ray Tracer for Visualizing Massive Models on a Cluster. In *Eurographics Symposium on Parallel Graphics and Visualization*, 2011.
- [20] W. Kendall, T. Peterka, J. Huang, H.-W. Shen, and R. Ross. Accelerating and benchmarking radix-k image compositing at large scale. EG PGV'10, 2010.
- [21] Kitware. Catalyst2: GPU resident workflows, 2024. <https://www.kitware.com/catalyst2-gpu-resident-workflows/>.
- [22] M. Larsen, J. Ahrens, U. Ayachit, E. Brugger, H. Childs, B. Geveci, and C. Harrison. The ALPINE In Situ Infrastructure: Ascending from the Ashes of Strawman. In *Proceedings of the In Situ Infrastructures on Enabling Extreme-Scale Analysis and Visualization*, ISAV'17, 2017. doi: 10.1145/3144769.3144778
- [23] K. Ma. Parallel volume ray-casting for unstructured-grid data on distributed-memory architectures. In *Proceedings of the IEEE Symposium on Parallel Rendering*, 1995. doi: 10.1145/218327.218333
- [24] K. Moreland, W. Kendall, T. Peterka, and J. Huang. An image compositing solution at scale. In *SC '11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011. doi: 10.1145/2063384.2063417
- [25] K. Moreland, C. Sewell, W. Usher, L.-t. Lo, J. Meredith, D. Pugmire, J. Kress, H. Schroots, K.-L. Ma, H. Childs, M. Larsen, C.-M. Chen, R. Maynard, and B. Geveci. VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures. *IEEE Computer Graphics and Applications*, 36(3), 2016. doi: 10.1109/MCG.2016.48
- [26] P. A. Navrátil. *Memory-Efficient, Scalable Ray Tracing*. PhD thesis, University of Texas, Austin, 2010.
- [27] P. A. Navrátil, H. Childs, D. S. Fussell, and C. Lin. Exploring the Spectrum of Dynamic Scheduling Algorithms for Scalable Distributed-Memory Ray Tracing. *IEEE Transactions on Visualization and Computer Graphics*, 20(6), 2014.
- [28] H. Park, D. Fussell, and P. Navrátil. SpRay: Speculative Ray Scheduling for Large Data Visualization. In *IEEE Symposium on Large Data Analysis and Visualization*, 2018.
- [29] E. Reinhard. *Scheduling and Data Management for Parallel Ray Tracing*. PhD thesis, University of East Anglia, 1995.
- [30] A. Sahitan, S. Demirci, I. Wald, S. Zellmann, J. Barbosa, N. Morrical, and U. Gündükbay. GPU-based Data-parallel Rendering of Large, Unstructured, and Non-convexly Partitioned Data, 2022. doi: 10.48550/ARXIV.2209.14537
- [31] J. Salmon and J. Goldsmith. A Hypercube Ray-Tracer. In *C3P: Proceedings of the third conference on Hypercube concurrent computers and applications - Volume 2*, 1989.
- [32] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit (4th ed.)*. Kitware, 2006.
- [33] J. E. Stone, K. Griffin, J. Amstutz, D. E. DeMarle, W. R. Sherman, and J. Günther. ANARI: A 3-D Rendering API Standard. *Computing in Science & Engineering*, 24(02), 2022. doi: 10.1109/MCSE.2022.3163151
- [34] A. Stukowski. Visualization and analysis of atomistic simulation data with OVITO—the Open Visualization Tool. *Modelling and Simulation in Materials Science and Engineering*, 18(1), 2010. doi: 10.1088/0965-0393/18/1/015012
- [35] W. Usher, I. Wald, J. Amstutz, J. Günther, C. Brownlee, and V. Pascucci. Scalable Ray Tracing Using the Distributed FrameBuffer. *Computer Graphics Forum*, 38, 2019. doi: 10.1111/cgf.13702
- [36] I. Wald, M. Jaroš, and S. Zellmann. Data Parallel Multi-GPU Path Tracing using Ray Queue Cycling. *Computer Graphics Forum*, 42(8), 2023. doi: 10.1111/cgf.14873
- [37] I. Wald, G. P. Johnson, J. Amstutz, C. Brownlee, A. Knoll, J. Jeffers, J. Günther, and P. Navrátil. OSPRay – A CPU Ray Tracing Framework for Scientific Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2017.
- [38] I. Wald and S. G. Parker. Data Parallel Path Tracing with Object Hierarchies. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 5(3), 2022. doi: 10.1145/3543861
- [39] B. Whitlock, J. M. Favre, and J. S. Meredith. Parallel In Situ Coupling of Simulation with a Fully Featured Visualization System. In *Eurographics Symposium on Parallel Graphics and Visualization*, 2011. doi: 10.2312/EGPGV/EGPGV11/101-109
- [40] S. Zellmann, N. Morrical, I. Wald, and V. Pascucci. Finding Efficient Spatial Distributions for Massively Instanced 3-d Models. In S. Frey, J. Huang, and F. Sadlo, eds., *Eurographics Symposium on Parallel Graphics and Visualization*. The Eurographics Association, 2020. doi: 10.2312/pgv.20201070