A Scalable WebGL-based Approach for Visualizing Massive 3D Point Clouds using Semantics-Dependent Rendering Techniques

Sören Discher
Hasso Plattner Institute
University of Potsdam, Germany
soeren.discher@hpi.de

Rico Richter
Hasso Plattner Institute
University of Potsdam, Germany
rico.richter@hpi.de

Jürgen Döllner
Hasso Plattner Institute
University of Potsdam, Germany
juergen.doellner@hpi.de

ABSTRACT

3D point cloud technology facilitates the automated and highly detailed digital acquisition of real-world environments such as assets, sites, cities, and countries; the acquired 3D point clouds represent an essential category of geodata used in a variety of geoinformation applications and systems. In this paper, we present a web-based system for the interactive and collaborative exploration and inspection of arbitrary large 3D point clouds. Our approach is based on standard WebGL on the client side and is able to render 3D point clouds with billions of points. It uses spatial data structures and level-of-detail representations to manage the 3D point cloud data and to deploy out-of-core and web-based rendering concepts. By providing functionality for both, thin-client and thick-client applications, the system scales for client devices that are vastly different in computing capabilities. Different 3D point-based rendering techniques and post-processing effects are provided to enable task-specific and data-specific filtering and highlighting, e.g., based on per-point surface categories or temporal information. A set of interaction techniques allows users to collaboratively work with the data, e.g., by measuring distances and areas, by annotating, or by selecting and extracting data subsets. Additional value is provided by the system's ability to display additional, context-providing geodata alongside 3D point clouds and to integrate task-specific processing and analysis operations. We have evaluated the presented techniques and the prototype system with different data sets from aerial, mobile, and terrestrial acquisition campaigns with up to 120 billion points to show their practicality and feasibility.

CCS CONCEPTS

• **Human-centered computing** → *Geographic visualization*; • **Computing methodologies** → *Computer graphics*; *Point-based models*;

KEYWORDS

3D Point Clouds, web-based rendering, point-based rendering

Web3D '18, June 20–22, 2018, Poznan, Poland

ACM ISBN 978-1-4503-5800-2/18/06…$15.00
https://doi.org/10.1145/3208806.3208816

ACM Reference Format:
Sören Discher, Rico Richter, and Jürgen Döllner. 2018. A Scalable WebGL-based Approach for Visualizing Massive 3D Point Clouds using Semantics-Dependent Rendering Techniques. In *Web3D '18: Web3D '18: The 23rd International Conference on Web3D Technology, June 20–22, 2018, Poznan, Poland.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3208806.3208816

1 MOTIVATION

3D point clouds allow for a discrete representation of real-world objects and environments. They can be time-efficiently and cost-efficiently generated by a large number of acquisition techniques using active or passive sensing technology such as LiDAR, radar, or aerial and digital cameras [Eitel et al. 2016; Ostrowski et al. 2014]. Integrated into a variety of carrier platforms such as airplanes, helicopters, UAVs, cars, trains, and robots, the sensing technology can capture data at different scales, ranging from small assets over buildings and infrastructure networks up to entire cities and countries [Kersten et al. 2016; Langner et al. 2016; Remondino et al. 2013]. The resulting data sets are essential for a growing number of applications in domains such as land surveying, urban planning, landscape architecture, environmental monitoring, disaster management, construction as well as spatial analysis and simulation [Eitel et al. 2016; Nebiker et al. 2010; Pătrăucean et al. 2015].

By their very nature, 3D point clouds are unstructured and do not contain or imply any order or connectivity between individual points. As a consequence, traditional analysis algorithms for geodata often struggle with 3D point clouds as they commonly rely on explicitly defined connectivity information. Visualization algorithms often apply a uniform pixel size and render style to each point and, therefore, are prone to visual artifacts such as holes or visual clutter which severely limits perception, interaction, and navigation [Richter et al. 2015]. As a remedy, GIS applications frequently use 3D point clouds only as input data to derive mesh based 3D models (e.g., 3D city models, terrain models) [Berger et al. 2014], leading to a loss of precision, density and data quality. In addition, deriving 3D meshes constitutes a time consuming and only semi-automatic process that does not scale for massive data sets. Moreover, improved scanning hardware and novel carrier systems, which get cheaper and easier-to-use, result in more dense 3D point clouds. Thus, there is a strong demand to store, manage, process and explore massive, arbitrarily dense 3D point clouds in order to take advantage of their full potential and to provide an unfiltered, detailed representation of captured sites.

In this paper, we present a web-based visualization system for massive 3D point clouds (Fig. 1) based on spatial data structures and level-of-detail representations that provide efficient access to arbitrary subsets of the 3D point cloud stored on a central server

Figure 1: Example of a massive 3D point cloud rendered with our web-based system. Context-providing geodata such as 2D maps and 3D terrain models can be integrated into the visualization.

component. By combining out-of-core rendering concepts with web-based rendering concepts massive data sets can be simultaneously distributed to and interactively visualized on an arbitrary number of client devices with different computation capabilities. To facilitate a collaborative inspection and to highlight task-relevant aspects of the data, different point-based rendering and interaction techniques are implemented that can be combined and configured by the user. In addition, the system can take advantage of per-point attributes generated by additional point-cloud analysis services. We evaluate our system with real-world data sets containing up to 120 billion points. Results show that the system is capable to provide a powerful component in production workflows to manage, distribute and share 3D point clouds.

2 RELATED WORK

3D point clouds represent a universal data category for a large number of geospatial applications [Eitel et al. 2016; Rüther et al. 2012]; many approaches exist to enhance information implicitly contained in 3D point clouds by deriving information about surface categories for each point, typically based on local topological analysis [Chen et al. 2017] or by deep learning concepts [Boulch et al. 2017; Huang and You 2016]. [Richter et al. 2013] show how to efficiently identify changes in multi-temporal data sets, which contain data acquired at different points in time. [Awrangjeb et al. 2015] combine surface categories and change detection results to filter detected changes based on semantics. Our approach provides efficient means to integrate such analyses as separate web processing services [Müller and Pross 2015]. Furthermore, analysis results can be shared, explored and inspected.

A general overview of point-based rendering techniques is provided by [Gross and Pfister 2011]. Photorealistic rendering techniques [Preiner et al. 2012; Schütz and Wimmer 2015a] focus on minimizing artifacts such as visual clutter or visible holes between neighboring points by applying appropriate size, orientation and color schemes to each point. Non-photorealistic rendering techniques [Simons et al. 2014; Zhang et al. 2014], on the other hand, deal with the fuzziness of a 3D point cloud and highlight edges and structures, commonly without requiring any additional attributes apart from a point's spatial position [Boucheny 2009; Pintus et al. 2011]. In our approach, we implement both rendering styles, which can be switched at runtime. State-of-the-art out-of-core rendering concepts for massive 3D point clouds, initially introduced by [Rusinkiewicz and Levoy 2000], typically use spatial data structures such as quadtrees [Gao et al. 2014], octrees [Elseberg et al. 2013], or KD-trees [Goswami et al. 2013] to subdivide the data into smaller subsets that can be selected dynamically for rendering tasks.

Recent approaches combine out-of-core and web-based rendering concepts to enable an ubiquitous visualization of 3D point clouds [Rodriguez et al. 2012; Wand et al. 2008]. With *Potree*, [Schütz and Wimmer 2015b] propose a thick-client approach for arbitrary large data sets, which is adapted by [Martinez-Rubi et al. 2015] to interactively present a massive data set of the Netherlands. An alternative thick-client renderer for 3D point clouds named *Plasio* was introduced by [Butler et al. 2014]. Using open-source libraries such as Entwine and Greyhound massive data sets can be streamed interactively. While both approaches provide effective interaction and inspection techniques for 3D point clouds, they offer only minimal support to integrate additional, context-providing geodata (e.g., shapes, 2D maps). While thin-client approaches that delegate the rendering to the server side have been successfully implemented in the past [Christen and Nebiker 2015; Döllner et al. 2012; Gutbell et al.

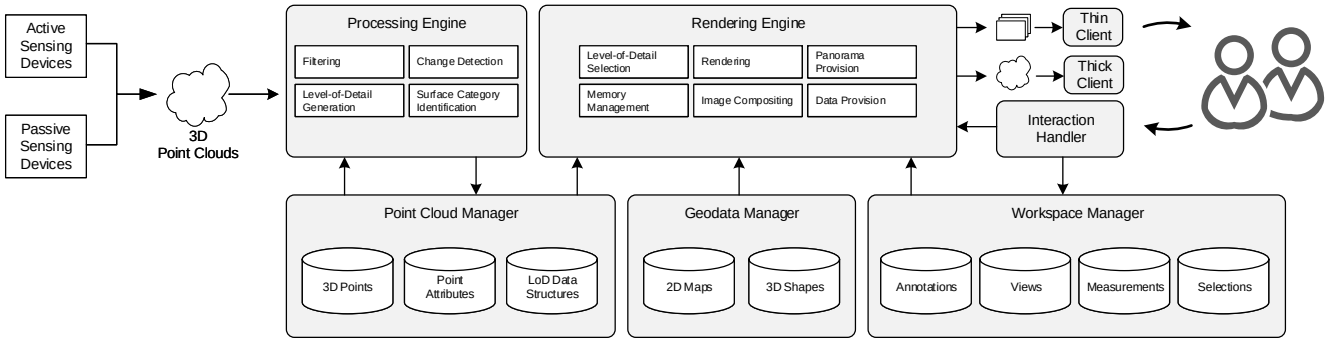

Figure 2: System architecture showing data flow between integration, processing, visualization, and interaction components.

2016], those contributions typically focus on mesh-based geometry instead of 3D point clouds. To generate stereoscopic panoramas we implemented the theoretical concepts described by [Peleg et al. 2001] by means of modern 3D computer graphics.

Systems for the efficient management of massive 3D point clouds have been recently presented and evaluated by [Cura et al. 2017], [van Oosterom et al. 2017], and [Poux et al. 2016]. However, those contributions focus on the efficient storage, retrieval and processing of the stored data sets. Less emphasis is put on the collaborative exploration, inspection and manipulation of the stored data sets.

3 REQUIREMENTS

We have identified the following requirements that need to be addressed by a system for the web-based visualization and collaborative exploration of massive 3D point clouds:

- **R1** Use of 3D point clouds as a fundamental geometry type instead of generalized mesh-based representations to enable a direct and unfiltered provision of the data.
- **R2** No limitations regarding used acquisition methods as well as density, resolution, and scale of the data (e.g., hundreds of billions of points, complete countries).
- **R3** Support for varying hardware platforms and computation capabilities, ranging from high-end desktop computers to low-end mobile devices.
- **R4** Distributed data storage to enable load balancing and to adjust for data specific requirements (e.g., certain 3D point clouds might have to be stored on a specific server).
- **R5** Capabilities to prepare and clean up 3D point clouds for the visualization (e.g., noise and outlier removal).
- **R6** Capabilities to conduct task and data specific analyses on 3D point clouds (e.g., surface category extraction) to provide adaptive and task specific content.
- **R7** Visualization of analysis results (e.g., surface categories) to enable task specific highlighting and filtering.
- **R8** Capabilities to compare and show differences between 3D point clouds from different points in time of the same site (i.e., change detection).
- **R9** Integration of supplementary, context-providing geodata such as 2D maps.
- **R10** Provision of interaction techniques to inspect (e.g., measuring of distances, areas, volumes) and annotate 3D point clouds.
- **R11** Basic user management to customize data access.
- **R12** Capabilities to share specific rendering configurations, annotations and measurements with others (e.g., via link).

4 CONCEPTS

We have addressed the aforementioned requirements in the design and implementation of our web-based system that seamlessly combines functionality to integrate, process, and collaboratively explore massive, heterogeneous 3D point clouds as well as supplementary, context-providing geodata. The proposed system (Fig. 2) consists of the following major components:

4.1 Point Cloud Manager

In our approach, 3D point clouds are organized in a single, homogeneous spatial data model. Access to that model is handled by the *point cloud manager* storing spatial information and additionally provided or computed per-point attributes (e.g., temporal information or surface categories) **(R1)**. Level-of-detail representations [Elseberg et al. 2013; Goswami et al. 2013] are required to efficiently access arbitrary data subsets of any size based on spatial, temporal or any other attributes. These representations as well as additional per-point attributes can be generated by the *processing engine* (Section 4.4) **(R2)**. While the point cloud manager logically acts as a singular component, the data itself may be stored in a distributed infrastructure, e.g., to maximize data throughput and network transfer rates or to account for data specific requirements regarding server location and data security **(R4)**.

4.2 Workspace Manager

The *workspace manager* handles information specific to a workspace, i.e., each user's private view of a specific data subset containing custom selections, measurements, annotations, view positions, and angles. Per default, each user operates in its own private workspace rather than sharing one globally with everyone else to avoid conflicting modifications **(R11)**. However, a given workspace may be shared via links **(R12)**. Each user may also own multiple workspaces.

(a) Measuring of distances between points.

(b) Measuring of areas defined by multiple points.

(c) Annotation of selected points or areas.

(d) Selecting areas of interest.

Figure 3: Overview of implemented interaction techniques for 3D point clouds.

4.3 Geodata Manager

By application-specific geodata, we refer to additional geodata that should be used and rendered in combination with a 3D point cloud to provide application-specific information layers **(R9)**. Examples are digital terrain models, aerial images, BIM models, or 3D city models. Similar to 3D point clouds, these data types also require supplemental level-of-detail representations to allow for an interactive visualization. Application-specific geodata can be stored and provided by independent geospatial databases or geodata services, access to which is handled by the *geodata manager*.

4.4 Processing Engine

The processing engine conducts task and data specific operations on a given data subset. These operations range from (a) essential preprocessing steps (e.g., converting input data sets into a homogeneous georeference system or generating level-of-detail representations), over (b) simple point cloud filtering (e.g., noise and outlier removal **(R5)**) to (c) more complex analyses (e.g., surface category extraction and change detection **(R6)**) deriving additional per-point attributes. The operations can be accessed via web processing services implemented as separate web services that are individually combined and scheduled by the processing engine. Thus, existing web processing services for 3D point clouds can be easily integrated into the system. The results of each operation are automatically stored by the point cloud manager and can be seamlessly integrated by the *rendering engine* (Section 4.5) into depictions of the corresponding site **(R7)**.

4.5 Rendering Engine

Providing the core functionality of our system, the rendering engine is responsible for interactively visualizing three types of data: (a) 3D point clouds featuring a varying number of per-point attributes, (b) task-specific geodata providing context (e.g., maps **(R9)**), and (c) workspace elements resulting from user interactions (e.g., annotations or selection and measurement indicators **(R10)**). For each of those data types the corresponding manager is queried, retrieving only data subsets that are relevant for the current view and task. To highlight certain aspects of the data (e.g., temporal changes or surface categories in an area), different point-based rendering techniques and post processing effects can be combined **(R8)**. Changes to the currently applied render configuration can be made dynamically via the *interaction handler*) (Fig. 4). In general, retrieved data subsets will be transferred to and rendered on client side, which minimizes the workload on the server (i.e., thick clients). As an alternative, server-side rendering can be applied to reduce the performance impact for clients (i.e., thin clients). Thus, the

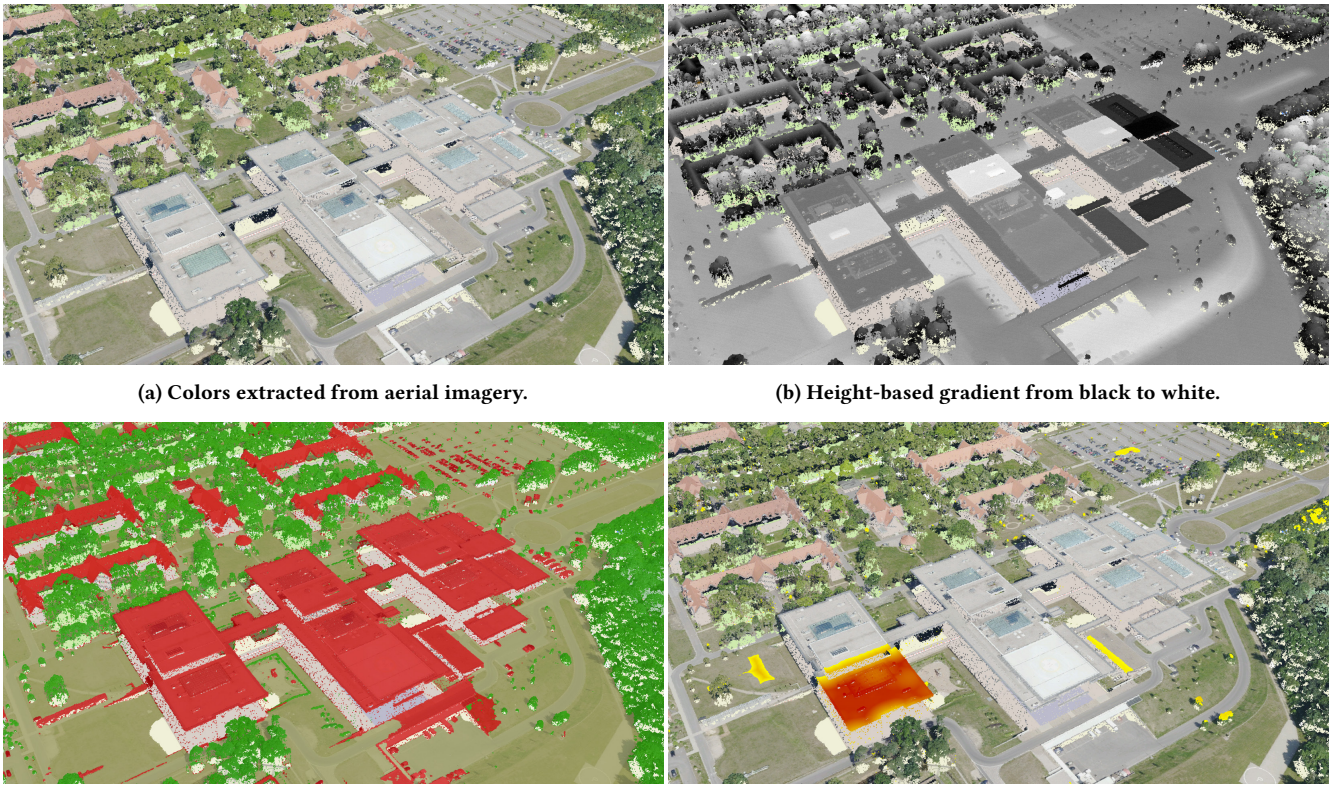

(a) Colors extracted from aerial imagery.

(b) Height-based gradient from black to white.

(c) Colorization based on surface categories, i.e., green for vegetation, red for buildings, and brown for ground.

(d) Colors from aerial imagery combined with change detection results, i.e., gradient from yellow to red indicates the degree of change.

Figure 4: Different point-based rendering styles can be selected and configured at runtime.

system scales for a broad range of devices, ranging from high-end workstations to mobile devices **(R3)**.

4.6 Interaction Handler

The interaction handler is responsible for handling user interactions as well as for updating the rendered data and workspace elements accordingly **(R10)**. Users may

- define or load workspaces,
- select which data subsets to render,
- configure the presentation of the data (with regards to applied rendering techniques),
- select, query and highlight individual points or groups of points,
- measure distances and areas between selected points,
- annotate selected points or areas,
- modify annotations,
- saving and loading view positions and angles.

5 RENDERING ENGINE IMPLEMENTATION

To seamlessly combine 3D point clouds, context providing geodata and interactive workspace elements into a homogeneous visualization, a multi-pass rendering pipeline is used that consists of three distinct stages (Fig. 5):

5.1 Level-of-Detail and Data Subset Selection

While 3D point clouds may easily contain billions of points, only a fraction of that data is required to render a frame. Subsets of the 3D point cloud that are manageable by available CPU and GPU capabilities can be queried dynamically from the point cloud manager by specifying the current view frustum as well as device specific metrics (e.g., an assigned memory budget) and task specific qualifiers (e.g., value ranges for selected per-point attributes) that further filter the corresponding data sets. To enable an efficient subset retrieval, the data is hierarchically subdivided using multiple layers, i.e., for each data set, a separate spatial data structure is generated that best compliments the spatial distribution of the corresponding points (e.g., quadtrees for airborne data sets, octrees or kd-trees for terrestrial data sets). In turn, those spatial data structures are integrated into an overarching quadtree, allowing to efficiently answer queries stretching across multiple data sets. Compared to uniform, single-layer spatial data structures, e.g., as they are used by [Schütz and Wimmer 2015b], this avoids a time consuming rebalancing when new 3D point clouds are added while simultaneously ensuring balanced tree structures and minimal data access times. Context providing geodata and workspace elements are handled in similar fashion by their corresponding manager and are queried simultaneously when required.

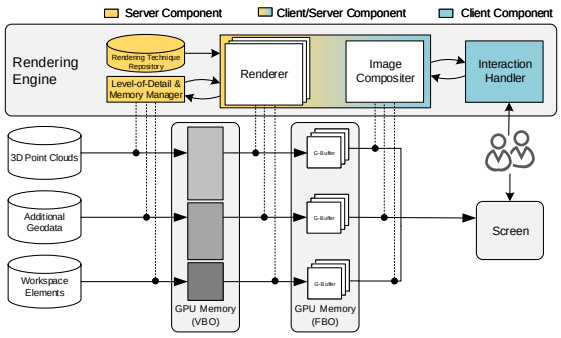

Figure 5: Overview of the rendering pipeline. Each data type is managed and rendered separately.

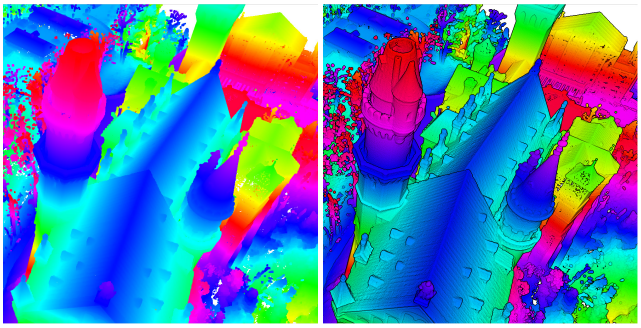

Figure 6: Post-processing effects such as Eye Dome Lighting facilitate visual filtering and highlighting.

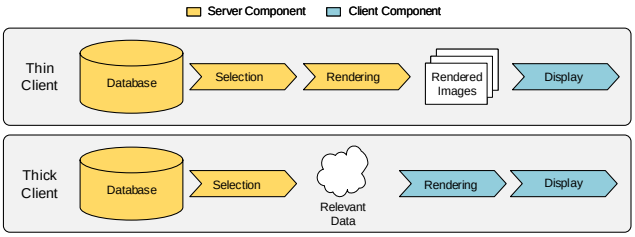

Figure 7: Comparison of web-based rendering concepts: Thin clients vs thick clients

5.2 Rendering

After being queried from the respective managers, 3D point clouds, context providing geodata and interactive workspace elements are rendered into separate g-buffers [Saito and Takahashi 1990], i.e., specialized frame buffer objects (FBO) combining multiple 2D textures for, e.g., color, depth, normal, or id values. The use of id values is important to separate point clouds from context data. Each rendered point has a unique identifier, stored in an id texture, to allow for an efficient point selection, e.g., to implement interaction features (Section 5.3). In addition, different rendering styles can be configured and applied at runtime. As an example, size and color of each point can be modified based on selected per-point attributes (e.g., surface categories, topological metrics) to enable task specific visual filtering and highlighting (Fig. 4). Similarly, several options exist to dynamically adjust the appearance of mesh-based geometry, ranging from transparency settings to changeable texture mappings.

5.3 Image Compositing

A final image compositing stage is used to merge the separate g-buffers, i.e., to combine several independently generated views of 3D point sub-clouds into a final image. For example, image-based post processing effects emphasizing edges and depth differences (e.g., Screen Space Ambient Occlusion [Mittring 2007] or Eye Dome Lighting [Boucheny 2009]) can be applied at that stage to improve the visual identification of structures within 3D point cloud depictions (Fig. 6). The id textures stored by the g-buffers provide efficient means to identify which point was rendered at a specific pixel. Thus, individual points can be selected in real-time, which is an essential requirement to support annotating points or measuring distances and areas.

5.4 Web-based Rendering

To accommodate for client devices with varying computation capabilities, different web-based rendering concepts are combined with the presented rendering pipeline (Fig. 7). We provide a thick-client application that uses a central server infrastructure to organize, process, select and distribute the data, but delegates the actual rendering of selected data subsets to the clients. This approach significantly reduces workload on server side, allowing to serve massive numbers of clients simultaneously. Transferred data subsets are cached on client side up to a device specific limit, thus, minimizing the frequency of data requests for subsequent frames. In fact, additional data subsets are only required if the view frustum changes significantly, whereas inspecting the transferred data or changing the applied rendering style triggers no such requests. Alternatively, in the sense of a thin-client approach, the data can be rendered directly on the server, supplying only the resulting images. While this comes with the drawback of increased workload on server side as any user interactions trigger a new data request, hardware requirements for clients are notably reduced. A common optimization for such thin-client applications is to render and transfer cube-maps or virtual panoramas instead of individual images [Döllner et al. 2012; Hagedorn et al. 2017]. This provides clients with efficient means to locally reconstruct the 3D scene for a specific view position. Thus, the data only has to be rendered anew whenever the view center or the rendering style are modified, which significantly reduces the frequency of data requests for subsequent frames. Our system provides a thin-client application that expands that concept, distributing not only traditional 2D panoramas but also stereoscopic panoramas. Thus, emerging virtual reality technologies allowing for an immersive exploration of 3D point clouds even on mobile devices can be easily integrated. We generate those stereoscopic panoramas by rendering several equally-sized image strips along a *viewing circle* that are stitched together in a post-processing step [Peleg et al. 2001]. The visual quality of the

Table 1: Devices used to evaluate the rendering engine. All web browsers were updated to the latest version as of 04/20/2018.

Client Device	CPU	Main Memory	GPU	Evaluated Web Browsers
Lenovo M710t	Intel Core i7-6700	32GB	GeForce GTX 1050Ti	Chrome, Firefox, Opera, Edge
Macbook Pro 13"	Intel Core i5-4278U	16GB	Intel Iris 5100	Safari, Chrome, Firefox
iPhone SE	Apple A9 @ 1.84 GHz	2GB	PowerVR GT7600	Safari Mobile, Chrome Mobile
Galaxy s7	Samsung Exynos 8890	4GB	ARM Mali-T880 MP12	Samsung Internet Browser, Chrome Mobile

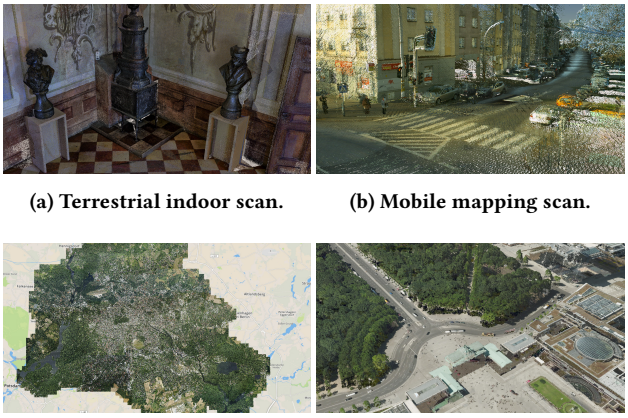

(a) Terrestrial indoor scan. **(b) Mobile mapping scan.**

(c) Airborne scan of a city (zoomed out) **(d) Airborne scan of a city (zoomed in)**

Figure 8: Scenes used during the performance evaluation.

Table 2: Average data throughput of the processing engine.

Processing Operation	Average Data Throughput
Noise & Outlier Filtering	1.26B pts/hour
Surface Category Extraction	0.10B pts/hour
Change Detection	1.42B pts/hour
Kd-Tree Generation	4.85B pts/hour

panoramas depends on the requested resolution as well as the number of image strips; both settings can be specified upon requesting a new panorama. To further reduce overall network load, both applications dynamically compress and decompress the transferred data, using common standards such as gzip (for thick clients) and png (for thin clients), respectively. We decided against using any lossy compression standards (e.g., jpeg compression) to maximize visual quality.

6 PERFORMANCE EVALUATION

We have implemented the presented concepts on the basis of several C++ and Javascript libraries. The processing engine uses *CUDA* [1] and the *Point Cloud Library* [2]. Regarding the rendering engine, we use WebGL and *Cesium* [3] for thick client applications. For thin client applications, server-side rendering is based on OpenGL, *glbinding* [4] and *GLFW* [5]. On client-side, *Three.js* [6], *WebGL*, and *WebVR Polyfill* [7] are combined to display 2D as well as stereoscopic panoramas. For data compression we use *gzip* [8] and *lodePNG* [9], respectively. Evaluated 3D point clouds are represented by separate kd-trees, that in turn are integrated into an overarching quadtree. We opted to use kd-trees to optimize the balancedness of the tree structures, speeding up the subset retrieval, albeit at the cost of a prolonged preprocessing. Those spatial data structures and corresponding data subsets are serialized into files acting as a point cloud database. Similar, file-based approaches are applied to store and access context-providing geodata and workspace elements.

All server side operations were performed on a server featuring an AMD Ryzen 7 1700 CPU, 32 GB main memory and an NVIDIA GeForce GTX 1070 with 8 GB device memory. The test data sets include a terrestrial, indoor scan of an individual site (1.33 billion points), a mobile mapping scan (2.57 billion points) and a massive, multi-temporal data set of an urban region (120 billion points) captured by airborne devices. For all data sets essential preprocessing steps (i.e., spatial data structure generation) and filtering (i.e., noise and outlier removal) were performed by the processing engine. In addition, surface categories (i.e., ground, building, vegetation) and changes in comparison to earlier scans were extracted for the airborne data set, allowing to evaluate the system's ability to dynamically combine different rendering styles. The average data throughput for the applied processing operations is listed in Table 2.

The rendering engine was evaluated based on four different scenes (Fig. 8) with client applications running on a number of different devices and web browsers (Table 1). Similar to [Schütz and Wimmer 2015b], our thick client implementation allows to render several millions of points simultaneously at interactive frame rates (i.e., >30 fps) on standard desktop computers and notebooks (Table 4). On mobile devices, frame rates are significantly lower due to the more limited computing capabilities. However, arbitrary large data sets can be visualized on all evaluated devices by assigning device-specific memory budgets, thus, limiting the density of the point cloud depiction. As an alternative, our thin client implementation provides a uniform rendering quality on all client devices since the panoramas are generated on server side, minimizing workload on client side. On all evaluated devices we measured frame rates close to the corresponding display's refresh rate (e.g., 60 fps on

[1] https://developer.nvidia.com/cuda-zone
[2] http://pointclouds.org
[3] https://cesiumjs.org
[4] https://github.com/cginternals/glbinding
[5] http://www.glfw.org
[6] https://threejs.org
[7] https://github.com/immersive-web/webvr-polyfill
[8] http://www.gzip.org
[9] http://lodev.org/lodepng/

Table 3: Average data throughput of the rendering engine based on the scenes defined in Fig. 8. For thin clients, a stereoscopic panorama was created per request. While the same, device-dependent resolution was requested for each scene, different entropies affected the compressed image size.

Scene	Thick Client		Thin Client		
	Transferred Data	Transfer Time	Transferred Data	Panorama Generation Time	Transfer Time
Terrestrial	156.2 MB	16.18s	4.68 MB	5.27s	1.36s
Mobile Mapping	140.7 MB	14.15s	4.16 MB	5.05s	1.27s
Airborne (zoomed out)	16.1 MB	3.43s	4.15 MB	4.96s	1.22s
Airborne (zoomed in)	82.4 MB	8.09s	4.54 MB	5.13s	1.32s

Table 4: Average performance rate of the thick client for different point budgets based on the airborne data set (Fig. 8d).

Number of Points	Transferred Data (uncompressed)	Transferred Data (compressed)	Lenovo M710t	Macbook Pro 13"	iPhone SE	Galaxy s7
2M pts	29.5 MB	26.2 MB	122.63fps	53.85fps	41.83fps	39.96fps
4M pts	57.4 MB	50.9 MB	84.48fps	45.63fps	36.44fps	35.29fps
6M pts	85.6 MB	76.1 MB	63.23fps	39.08fps	26.36fps	24.83fps
8M pts	113.6 MB	100.7 MB	56.87fps	35.83fps	19.65fps	18.43fps

Table 5: Panorama generation time for different configurations based on the terrestrial data set (Fig. 8a).

Resolution	Transferred Data	Panorama Generation Time		
		90 image strips	120 image strips	160 image strips
2360x1600 px	2.33 MB	1.88s	2.26s	2.29s
2360x3200 px	4.68 MB	4.20s	4.68s	5.27s
2360x6400 px	9.35 MB	6.17s	7.14s	7.88s

the Galaxy S7), making our approach applicable to state-of-the-art VR devices such as GearVR or Oculus Rift. The performance of the panorama generation is primarily influenced by the requested resolution and to a lesser degree on the number of image strips used (Table 5).

For all evaluated scenes, thick client applications require to transfer significantly more data for an individual scene than thin clients as long as no reusable data subsets have been cached from previous requests, even if gzip compression is applied (Table 3). However, they do not require all those data subsets at once, allowing to update the scene progressively. Furthermore, while exploring a 3D point cloud, the view will usually change only gradually across subsequent frames, allowing for thick clients to reuse many of the previously transferred data subsets, thus, resulting in smaller and faster scene updates over prolonged explorations. Changes to the rendering style as well interaction techniques such as picking, selecting or measuring don't trigger any additional data requests at all and can be applied even under unstable network conditions. For thin client applications on the other hand, no parts of the previously transferred data can be reused if the currently used panorama becomes invalid: Navigating -apart from merely looking around from a fixed position- as well as rendering style adjustments require the server to generate and transfer a new panorama as a replacement. Similar to thick clients however, picking, selecting or measuring can be conducted on the already transferred data and does not trigger any new data requests.

7 CONCLUSION AND FUTURE WORK

Web-based visualization and exploration of massive 3D point clouds from aerial, mobile, or terrestrial data acquisitions represent a key feature for today's and future systems and applications dealing with digital twins of our physical environment. In our web-based approach, we show a system architecture that scalably visualizes massive 3D point clouds to web-based client devices. To cope with extremely large number of points, the implementation relies on spatial data structures and level-of-detail representations, combined with different out-of-core rendering and web-based rendering concepts. Since the rendering process can be shifted from client side to server side, the system can be easily adapted to varying network conditions and to clients with different computing and graphics capabilities. Tests on data sets with up to 120 billion points show the usability of the system and the feasibility of the approach. Various rendering techniques allow us to filter and highlight subsets of the data based on any available per-point attributes (e.g., surface categories or temporal information), which is required to build task-specific or application-specific tools. Various interaction methods (e.g., for collaborative measurements and annotations), built-in support to display context-providing, mesh-based geodata, and

the possibility to conduct different processing and analysis operations provide additional features. Our system could be further extended by integrating additional analyses (e.g., for asset detection, or surface segmentation) [Jochem et al. 2012; Teo and Chiu 2015] as well as by specialized interaction techniques. For example, [Scheiblauer and Wimmer 2011] and [Wand et al. 2008] propose spatial data structures that allow for an interactive editing of 3D point clouds. In addition, sophisticated visualization techniques for multi-temporal 3D point clouds are becoming more and more important to understand captured environments.

ACKNOWLEDGMENTS

We thank Pawel Böning and Pascal Führlich for their contributions to the thin client implementation. Data sets have been provided by Illustrated Architecture, SHH sp. z o.o. and virtualcitySYSTEMS.

REFERENCES

Mohammad Awrangjeb, Clive S Fraser, and Guojun Lu. 2015. Building change detection from LiDAR point cloud data based on connected component analysis. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 2 (2015), 393–400.

Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Joshua Levine, Andrei Sharf, and Claudio Silva. 2014. State of the art in surface reconstruction from point clouds. In *EUROGRAPHICS star reports*, Vol. 1. 161–185.

Christian Boucheny. 2009. *Interactive Scientific Visualization of Large Datasets: Towards a Perceptive-Based Approach.* Ph.D. Dissertation. Université Joseph Fourier, Grenoble.

Alexandre Boulch, Bertrand Le Saux, and Nicolas Audebert. 2017. Unstructured point cloud semantic labeling using deep segmentation networks. In *Eurographics Workshop on 3D Object Retrieval*, Vol. 2. 1.

Howard Butler, David C Finnegan, Peter J Gadomski, and Uday K Verma. 2014. plas. io: Open Source, Browser-based WebGL Point Cloud Visualization. In *AGU Fall Meeting Abstracts*.

Dong Chen, Ruisheng Wang, and Jiju Peethambaran. 2017. Topologically aware building rooftop reconstruction from airborne laser scanning point clouds. *IEEE Transactions on Geoscience and Remote Sensing* 55, 12 (2017), 7032–7052.

Martin Christen and Stephan Nebiker. 2015. Visualisation of complex 3D city models on mobile webbrowsers using cloud-based image provisioning. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 2 (2015), 517–522.

Rémi Cura, Julien Perret, and Nicolas Paparoditis. 2017. A scalable and multi-purpose point cloud server (PCS) for easier and faster point cloud data management and processing. *ISPRS Journal of Photogrammetry and Remote Sensing* 127 (2017), 39–56.

Jürgen Döllner, Benjamin Hagedorn, and Jan Klimke. 2012. Server-based rendering of large 3D scenes for mobile devices using G-buffer cube maps. In *Proceedings of the 17th International Conference on 3D Web Technology*. 97–100.

Jan UH Eitel, Bernhard Höfle, Lee A Vierling, Antonio Abellán, Gregory P Asner, Jeffrey S Deems, Craig L Glennie, Philip C Joerg, Adam L LeWinter, Troy S Magney, et al. 2016. Beyond 3-D: The new spectrum of lidar applications for earth and ecological sciences. *Remote Sensing of Environment* 186 (2016), 372–392.

Jan Elseberg, Dorit Borrmann, and Andreas Nüchter. 2013. One billion points in the cloud–an octree for efficient processing of 3D laser scans. *ISPRS Journal of Photogrammetry and Remote Sensing* 76 (2013), 76–88.

Zhenzhen Gao, Luciano Nocera, Miao Wang, and Ulrich Neumann. 2014. Visualizing aerial LiDAR cities with hierarchical hybrid point-polygon structures. In *Proceedings of Graphics Interface 2014*. 137–144.

Prashant Goswami, Fatih Erol, Rahul Mukhi, Renato Pajarola, and Enrico Gobbetti. 2013. An efficient multi-resolution framework for high quality interactive rendering of massive point clouds using multi-way kd-trees. *The Visual Computer* 29, 1 (2013), 69–83.

Markus Gross and Hanspeter Pfister. 2011. *Point-based graphics.* Morgan Kaufmann.

Ralf Gutbell, Lars Pandikow, Volker Coors, and Yasmina Kammeyer. 2016. A framework for server side rendering using OGC's 3D portrayal service. In *Proceedings of the 21st International Conference on Web3D Technology*. 137–146.

Benjamin Hagedorn, Simon Thum, Thorsten Reitz, Volker Coors, and Ralf Gutbell. 2017. *OGC 3D Portrayal Service 1.0.* OGC Implementation Standard 1.0. Open Geospatial Consortium.

Jing Huang and Suya You. 2016. Point cloud labeling using 3d convolutional neural network. In *Proceedings of the 23rd International Conference on Pattern Recognition*. 2670–2675.

Andreas Jochem, Bernhard Höfle, Volker Wichmann, Martin Rutzinger, and Alexander Zipf. 2012. Area-wide roof plane segmentation in airborne LiDAR point clouds. *Computers, Environment and Urban Systems* 36, 1 (2012), 54–64.

Thomas P Kersten, Heinz-Jürgen Przybilla, Maren Lindstaedt, Felix Tschirschwitz, and Martin Misgaiski-Hass. 2016. Comparative geometrical investigations of hand-held scanning systems. *ISPRS Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* (2016).

Tobias Langner, Daniel Seifert, Bennet Fischer, Daniel Goehring, Tinosch Ganjineh, and Raúl Rojas. 2016. Traffic awareness driver assistance based on stereovision, eye-tracking, and head-up display. In *Proceedings of ICRA 2016*. 3167–3173.

Oscar Martinez-Rubi, Stefan Verhoeven, Maarten Van Meersbergen, M Schûtz, Peter Van Oosterom, Romulo Gonçalves, and Theo Tijssen. 2015. Taming the beast: Free and open-source massive point cloud web visualization. In *Proceedings of the Capturing Reality Forum 2015*.

Martin Mittring. 2007. Finding next gen: Cryengine 2. In *ACM SIGGRAPH 2007 courses*. ACM, 97–121.

Matthias Müller and Benjamin Pross. 2015. OGC WPS 2.0 interface standard. *Open Geospatial Consortium Inc.* (2015).

Stephan Nebiker, Susanne Bleisch, and Martin Christen. 2010. Rich point clouds in virtual globes–A new paradigm in city modeling? *Computers, Environment and Urban Systems* 34, 6 (2010), 508–517.

Steve Ostrowski, Grzegorz Jóźków, Charles Toth, and Benjamin Vander Jagt. 2014. Analysis of point cloud generation from UAS images. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 2, 1 (2014), 45–51.

Viorica Pătrăucean, Iro Armeni, Mohammad Nahangi, Jamie Yeung, Ioannis Brilakis, and Carl Haas. 2015. State of research in automatic as-built modelling. *Advanced Engineering Informatics* 29, 2 (2015), 162–171.

Shmuel Peleg, Moshe Ben-Ezra, and Yael Pritch. 2001. Omnistereo: Panoramic stereo imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 3 (2001), 279–290.

Ruggero Pintus, Enrico Gobbetti, and Marco Agus. 2011. Real-time Rendering of Massive Unstructured Raw Point Clouds Using Screen-space Operators. In *Proceedings of VAST 2011*. 105–112.

Florent Poux, Pierre Hallot, Romain Neuville, and Roland Billen. 2016. Smart point cloud: Definition and remaining challenges. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 4 (2016), 119–127.

Reinhold Preiner, Stefan Jeschke, and Michael Wimmer. 2012. Auto Splats: Dynamic Point Cloud Visualization on the GPU.. In *Proceedings of the EGPGV*. 139–148.

Fabio Remondino, Maria Grazia Spera, Erica Nocerino, Fabio Menna, Francesco Nex, and Sara Gonizzi-Barsanti. 2013. Dense image matching: comparisons and analyses. In *Proceedings of DigitalHeritage 2013*, Vol. 1. 47–54.

Rico Richter, Sören Discher, and Jürgen Döllner. 2015. Out-of-core visualization of classified 3d point clouds. In *3D Geoinformation Science*. Springer, 227–242.

Rico Richter, Jan E Kyprianidis, and Jürgen Döllner. 2013. Out-of-Core GPU-based Change Detection in Massive 3D Point Clouds. *Transactions in GIS* 17, 5 (2013), 724–741.

Marcos B Rodriguez, Enrico Gobbetti, Fabio Marton, Ruggero Pintus, Giovanni Pintore, and Alex Tinti. 2012. Interactive Exploration of Gigantic Point Clouds on Mobile Devices.. In *13th International Conference on Virtual Reality, Archaeology and Cultural Heritage*. 57–64.

Szymon Rusinkiewicz and Marc Levoy. 2000. QSplat: A multiresolution point rendering system for large meshes. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. 343–352.

Heinz Rüther, Christoph Held, Roshan Bhurtha, Ralph Schroeder, and Stephen Wessels. 2012. From point cloud to textured model, the zamani laser scanning pipeline in heritage documentation. *South African Journal of Geomatics* 1, 1 (2012), 44–59.

Takafumi Saito and Tokiichiro Takahashi. 1990. Comprehensible rendering of 3-D shapes. In *ACM SIGGRAPH Computer Graphics*, Vol. 24. ACM, 197–206.

Claus Scheiblauer and Michael Wimmer. 2011. Out-of-core selection and editing of huge point clouds. *Computers & Graphics* 35, 2 (2011), 342–351.

Markus Schütz and Michael Wimmer. 2015a. High-quality point-based rendering using fast single-pass interpolation. In *Proceedings of Digital Heritage 2015*. 369–372.

Markus Schütz and Michael Wimmer. 2015b. Rendering large point clouds in web browsers. *Proceedings of CESCG* (2015), 83–90.

Lance Simons, Stewart He, Peter Tittman, and Nina Amenta. 2014. Point-based rendering of forest LiDAR. In *Workshop on Visualisation in Environmental Sciences (EnvirVis), The Eurographics Association*. 19–23.

Tee-Ann Teo and Chi-Min Chiu. 2015. Pole-like road object detection from mobile lidar system using a coarse-to-fine approach. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 8, 10 (2015), 4805–4818.

Peter van Oosterom, Oscar Martinez-Rubi, Theo Tijssen, and Romulo Gonçalves. 2017. Realistic benchmarks for point cloud data management systems. In *Advances in 3D Geoinformation*. Springer, 1–30.

Michael Wand, Alexander Berner, Martin Bokeloh, Philipp Jenke, Arno Fleck, Mark Hoffmann, Benjamin Maier, Dirk Staneker, Andreas Schilling, and Hans-Peter Seidel. 2008. Processing and interactive editing of huge point clouds from 3D scanners. *Computers & Graphics* 32, 2 (2008), 204–220.

Long Zhang, Qian Sun, and Ying He. 2014. Splatting lines: an efficient method for illustrating 3D surfaces and volumes. In *Proceedings of the 18th meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 135–142.