

Scalable Multi-Platform Distribution of Spatial 3D Contents

Jan Klimke, Hasso-Plattner-Institute, University of Potsdam, Potsdam, Germany

Benjamin Hagedorn, Hasso-Plattner-Institute, University of Potsdam, Potsdam, Germany

Jürgen Döllner, Hasso-Plattner-Institute, University of Potsdam, Potsdam, Germany

ABSTRACT

Virtual 3D city models provide powerful user interfaces for communication of 2D and 3D geoinformation. Providing high quality visualization of massive 3D geoinformation in a scalable, fast, and cost efficient manner is still a challenging task. Especially for mobile and web-based system environments, software and hardware configurations of target systems differ significantly. This makes it hard to provide fast, visually appealing renderings of 3D data throughout a variety of platforms and devices. Current mobile or web-based solutions for 3D visualization usually require raw 3D scene data such as triangle meshes together with textures delivered from server to client, what makes them strongly limited in terms of size and complexity of the models they can handle. This paper introduces a new approach for provisioning of massive, virtual 3D city models on different platforms namely web browsers, smartphones or tablets, by means of an interactive map assembled from artificial oblique image tiles. The key concept is to synthesize such images of a virtual 3D city model by a 3D rendering service in a preprocessing step. This service encapsulates model handling and 3D rendering techniques for high quality visualization of massive 3D models. By generating image tiles using this service, the 3D rendering process is shifted from the client side, which provides major advantages: (a) The complexity of the 3D city model data is decoupled from data transfer complexity (b) the implementation of client applications is simplified significantly as 3D rendering is encapsulated on server side (c) 3D city models can be easily deployed for and used by a large number of concurrent users, leading to a high degree of scalability of the overall approach. All core 3D rendering techniques are performed on a dedicated 3D rendering server; and thin-client applications can be compactly implemented for various devices and platforms.

Keywords: *Image-Based Rendering, Mobile 3D Visualization, Tile-Based Provisioning, Virtual 3D City Model, Web-Based 3D Portrayal*

1. INTRODUCTION

Virtual 3D city models provide powerful user interfaces for communication of 2D and 3D geoinformation. Through the continuing desire for more detail and realism, the model complexity

of common scenes has not reached its peak by far” (Jeschke, Wimmer, & Purgathofer, 2005). The ongoing development in remote sensing and data processing technologies produces more and more 3D model data in increasing amounts and in continuously improving quality, which

DOI: 10.4018/ij3dim.2014070103

makes providing high quality visualization of massive 3D geoinformation in a scalable, fast, and cost efficient manner still a challenging task.

Today's systems for mobile or web-based visualization of virtual 3D city models, e.g., Google Earth, Apple Maps, or here.com, mostly rely on streaming 3D geometry and corresponding textures to client devices. In this way, the applications running on those devices need to implement the whole rendering part of the visualization pipeline (Haber & McNapp, 1990), i.e., rasterization of images from computer graphic primitives. The rasterization process is a resource intensive task, which requires specialized rendering hardware and software components. Rendering performance and their requirements regarding CPU, GPU, main memory, and disk space, strongly depend on the complexity of the model to be rendered. Further, high-quality rendering techniques, e.g., for realistic illumination, shadowing, or water rendering, increase the resource consumption that is necessary to provide users with interactive frame rates (more than 10 frames per second). This makes it very hard to develop applications that adapt to different hardware and software platforms while still providing a high-quality rendering and an acceptable frame rate also for large 3D datasets. Building a fast, stable application that renders large 3D models in high quality on a variety of heterogeneous platforms and devices, incorporates a huge effort in software development and maintenance as well as data processing, which usually raises with the number of different platforms to be supported.


Approaches for image-based 3D portrayal introduced recently (Doellner, Hagedorn, & Klimke, 2012) tackle these problems by shifting the more complex and resource intensive task of image synthesis to the server side, which allows for interactive thin client applications on various end user devices. Such clients can *reconstruct* lightweight representations of the server side 3D model from server-generated G-Buffer images (i.e., multi layer raster images that not only encode color values, but also other information like depth, etc.). Image-based portrayal provides two major advantages com-

pared to geometry-based approaches: a) They decouple the rendering complexity on client side from the model complexity on server-side and b) they allow to deliver a homogeneously high rendering quality to all end user platforms and devices, regardless of the 3D capabilities of these devices. Nevertheless, the operation of such visualization applications needs a 3D rendering service to generate these image-based representations of the underlying 3D model. This makes scaling the applications for many simultaneously used clients a complex and relatively expensive task, since each service instance can only effectively serve a limited number of clients.

In this paper we introduce a novel approach for provisioning of massive, virtual 3D city models on different platforms (web browsers, smartphones, tablets) by means of an interactive map showing synthetic, tiled images of the 3D city model (*oblique map*). The key concept is to synthesize these oblique views in a preprocessing step by a 3D rendering service, to store the corresponding tiles, e.g., on a web server, to be easily accessible and usable by client applications (Figure 1). Different stylizations, combinations of thematic layers, map layers, and viewing directions can be specified and applied for the tile generation process, leading to multiple tile sets, each storing not only a RGB images but also additional information such as world coordinates, depth information, surface normals, object identities or thematic information (G-Buffer). Client applications can use the additional information, e.g., for object-based interaction or additional client-side rendering effects such as object highlighting. In particular object identity information allows client applications to set up to application specific functionality, e.g., to retrieve object-based information from a remote server.

Compared to existing mobile 3D geovisualization solutions, the presented approach scales far better for large numbers of simultaneously active users as the computationally expensive task of 3D rendering of large-scale datasets is shifted from application runtime to a preprocessing step, generating the oblique views of the 3D

Figure 1. Screenshot of the iOS-App smartMap Berlin from our case study running on an iPhone. The App's user interface is kept simple: Users can use zooming and moving gestures known from 2D map interfaces. Further an interaction item on the upper left provides stepwise rotation. The mobile client can also be used to view large numbers of high detail 3D models since the client side computational complexity does not depend on the complexity of the 3D model itself.



scene as image tile sets. These image tiles can be deployed, stored, and accessed easily using, e.g., conventional web servers or cloud storage services, which are able to easily scale with the number of concurrent users. This approach, however, restricts the way users can interact with the 3D scene; It is generally known that map-based access is the most efficient way to explore and analyze spatial information. Our oblique map approach supports map-based interactions, but is restricted in terms of 3D camera control.

2. FOUNDATIONS AND RELATED WORK

A set of web standards and formats (e.g., WebGL (Jackson, 2013), X3D ("X3D Standard," 2005), X3DOM (Behr et al., 2010)) tackle the demand for plugin-free 3D rendering in web browsers and form the basis of several approaches and solutions in this field (Chris-

ten, Nebiker, & Loesch, 2012; Gesquière & Manin, 2012; Prieto & Izgara, 2012). These approaches utilize transferred 3D geometry and textures for rendering. This makes rendering large models a resource intensive task for client devices. While the camera interaction is relatively unrestricted (free exploration of 3D model), these approaches are limited in terms of the implementation of high-end rendering techniques. In order to maintain compatibility to a large range of client hardware (different desktop graphics, different mobile devices) and software (WebGL capabilities of different web browsers) the ability to stylize the underlying model is limited. Also the size and complexity of 3D data that can be rendered in interactive frame rates is restricted by the client's hardware. In contrast, our approach of a tile-based oblique 3D map, performs styling in a preprocessing step and by a dedicated high-performance 3D rendering service. As the computational and memory capabilities required on client

side do neither depend on the model size nor on the complexity of the rendering technique applied, this approach provides a solution for the distribution of massive 3D models in the same graphical quality and expressiveness on all devices.

In recent years, several approaches for visualization of massive virtual 3D city models on mobile devices or web browsers have been published (Bao & Gourlay, 2006; Boukerche & Pazzi, 2006; Lamberti & Sanna, 2007). Döllner et al. (Doellner et al., 2012) provide an approach using an image-based portrayal service to generate several image layers (e.g., color, depth information, and object identity) of a complex virtual 3D city model, that are transferred to a client application running on mobile devices. A corresponding client uses the images generated for a specific point of camera to create a lightweight reconstruction of the server-side model using either a cube map or geometry derived from the depth image delivered by the server. The client application in connection with the 3D rendering service allows users to configure the visualization of the server-side 3D model data to match their needs in terms of data selection and styling. Generally, this approach does not apply very well for a large number of concurrent users because service responses can hardly be cached due to the manifold combinations of styling options and due to the fact, that the virtual camera is not restricted but can be positioned at anytime and anywhere in the 3D environment. The approach presented in this paper combines the possibility to create preconfigured visualizations (in terms of data selection and styling) with a near optimal cacheability for single image tiles on server and client side.

Examples of tile-based 3D oblique maps for web browsers are developed, e.g., by the Chinese companies edushi.com¹ and Dushiquan². They provide access to very detailed virtual 3D city models of Chinese cities. Besides high-quality (and comic-like) visualization, these maps provide client-side highlighting of objects (e.g., buildings or areas) as well as information retrieval for selected objects. Nevertheless (as

reported in Internet articles and forums) these virtual cities are mainly crafted manually based on existing maps and satellite imagery. This means a rather low degree of automation in the creation and updating process of these maps and implies a rather high effort for keeping the model up-to-date. Also, this provides only limited possibilities for creating custom-styled maps, e.g., for specific use cases or users, and for the creation of different variants of the maps, e.g., maps from a different viewing direction and viewing angle or maps with different building variants. Possibly due to this, users cannot adjust the map orientation of these maps, which also limits the 3D impression of the overall application. In contrast, the project osm2world³ uses a much more automated approach for generating tiled image views from publicly available OpenStreetMap data. From this (still rather 2D data), 3D building geometry is automatically derived; then, oblique map tiles are rendered based on OpenGL or a POVray ray tracing backend. The corresponding viewer client allows users to rotate the view in four steps, but currently does not provide object-based interaction and data retrieval. Both of the solutions presented do not offer a solution for a native user experience on mobile devices, while we target those devices explicitly to consider the increasing importance of such devices. Compared to the systems presented above, our approach focuses on automation of processes as well as providing a high graphical quality for image tiles through implementing state of the art rendering techniques in the rendering service. Also, we focus on supporting complex and "real-world" 3D geodata by handling huge 2D and real 3D geodata (including geometries, textures, thematic information, object attributes) and combining the virtual 3D city model with additional information sources.

Service-oriented architectures (SOA) (Papazoglou, Traverso, Dustdar, & Leymann, 2007), as a paradigm for design and development of distributed information systems, represent a common approach to address these challenges. 3D geovisualization systems, based on the SOA paradigm, encapsulate resource intensive tasks,

such as management, processing, transmission, and rendering of massive 2D and 3D geodata as services that can be reused by various client applications (Basanow, Neis, Neubauer, Schilling, & Zipf, 2008). Thus, client applications do not need to implement such functionality, but can reuse this functionality implemented through services. While 2D geovisualization systems can rely on standardized and robust services such as the Web Map Service (WMS), specified by the Open Geospatial Consortium (OGC), only first approaches and implementations for service-based 3D geovisualization have been suggested. The Web 3D Service (W3DS) (Schilling & Kolbe, 2010) supports streaming of 3D scenes (i.e., 3D geometry and corresponding textures) to client applications, which implement the rendering processes on client-side. The Web View Service (WVS) (Hagedorn, 2010) instance implements the complete 3D rendering process on server side, encapsulating the complexity of 3D geodata management, processing, and rendering. The two alternative approaches for 3D portrayal services are currently undergoing an effort within the OGC to harmonize them and to provide a common base for service interfaces. A WVS compliant rendering service is the core part of our preprocessing pipeline for generating orthographic views out of 3D city model data. Due to the standards-based (WVS is currently a discussion paper within the OGC) service interface, we are able to exchange service implementations and utilize third party remote services for tile generation.

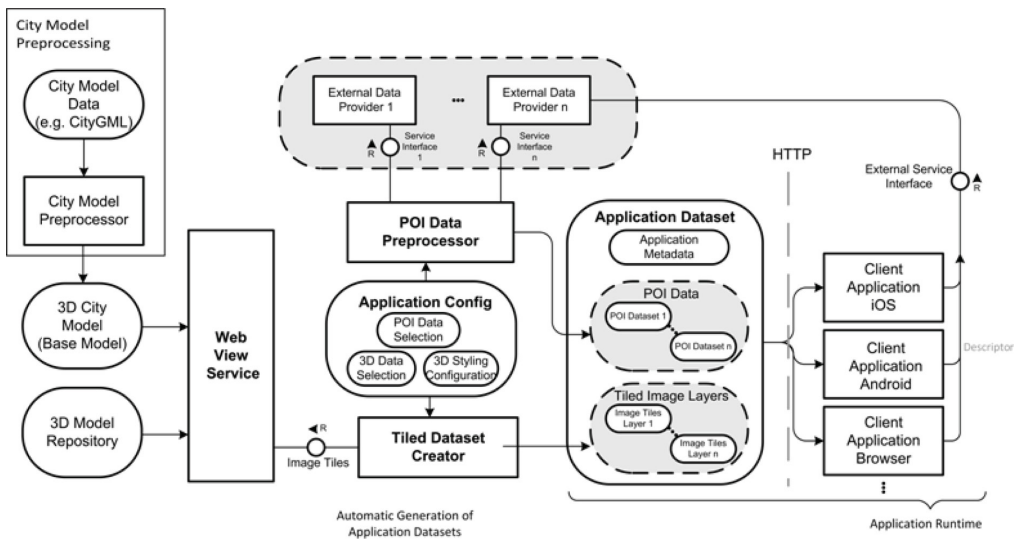
3. CONCEPT

The key concept of our approach is to synthesize tileable, oblique views on a given virtual 3D city model by a 3D rendering service in a preprocessing step. This 3D rendering service encapsulates all aspects related to 3D rendering of massive, complex structured, heterogeneous 2D (e.g., terrain map layers) and 3D (terrain data, textured building geometry) geodata. This way, we decouple all the complexity related

to processing, management, and rendering of 3D geodata from client applications and additionally from the capabilities of servers that provide data for client applications at runtime. A major advantage of this is the ability to select the contents of an oblique view as well as its stylization to match the purpose of a specific application. In our approach, the geographical spaces is separated into numbered tiles in different zoom levels, in analogy to existing tiling schemes such as Time Map Service or WMS-T. Each tile of an oblique dataset (a map layer of the oblique map in a certain direction) has a unique key that consists of three components: the tile zoom level, the tile number in horizontal axis, and the tile number in vertical axis of the geographical coordinate system. Each image tile is covering a spatial extent that is defined by these components. Serving image tiles to client applications only needs a conventional web server or a cloud storage services implementing a key value store. This allows the approach to scale to a large number of concurrent users fetching tiles in parallel, since scaling the network-based delivery provisioning of files to client applications can easily be achieved. Oblique views (Figure 1) are generated for different viewing directions (currently the four principal directions) allowing a stepwise rotation of the view by clients through exchanging the underlying image dataset (map layer). Minimal implementations of client applications do only need capabilities to fetch image tiles via HTTP and display them, therefore they can be very lightweight. Such client applications can easily be implemented on mobile platforms (i.e., smartphones or tablets) or plugin free (based on JavaScript) on web browsers. They provide a map-like user interface (also called Slippy Map (Sample & Ioup, 2010)) for exploring the city model visualization allowing users to shift the map center, to change zoom levels and to rotate the overall view into one of the directions that were pregenerated by the 3D rendering backend beforehand.

The map-based approach implies a simple and intuitive interaction, focusing on navigation the 3D and avoiding the difficulties of a

Figure 2. An overview of the system for automated generation of application datasets from different data sources. 3D city models were imported in a form optimized for interactive 3D rendering into the rendering service. Application datasets, that define contents (oblique map tiles, POI data, and external service URLs to be used) and behavior (e.g., which URL to load on tap for POI items) are generated for being consumed by client applications.



completely free 3D user interaction, in which a user generally has to manage six degrees of freedom. Most tasks, such as exploring a certain area, do not require this full 3D camera control. Up to a certain degree of necessary detail, the provided interaction is regarded sufficient. This way, users are able to explore the virtual 3D city model and connected data sets using a well-established interaction paradigm.

In order to provide use-case specific thematic information, an application configuration that is consumed by more generic client applications allows defining a set of information items that are included in each generated map application. Several kinds of thematic data can be integrated into one application. Beside conventional items, such as point of interest (POI) information and 2D shapes (e.g., roads or public transport lines), applications are able to connect server-side information belonging to specific features in a model to pixels displayed in the front-end application. An overview over the whole architecture for generating the necessary application datasets (oblique map layers,

harmonized POI data, and configurations for generic client applications) is provided in Figure 2.

3.1. Data Import 3D City Model Data and Tile Generation

Image tiles are generated using a 3D rendering service that implements the complete 3D rendering process for massive, heterogeneous 3D models. In particular we focus on CityGML (Gröger, Kolbe, Nagel, & Häfele, 2012) as standard format for city model, since “CityGML is a common information model for the representation of 3D urban objects” (Mao, 2011). The high integration of 3D geometry, appearance (materials and textures), and detailed models for object semantics make CityGML a promising base data source for applications. Further, the rendering service is able to integrate a broad range of other 3D formats for 3D data, e.g., KML/Collada, 3DS, OBJ, or the ShapeFile format. Since usual 3D geodata formats are designed for effective data exchange and not

for efficient 3D rendering, the source data has to be converted into a rendering optimized form for two reasons: (1) 3D city model features are usually modeled with regard object semantics and geometric correctness. The suitability for efficient 3D rendering of the 3D model data plays a negligible role in data acquisition and modeling. (2) Objects with individual real world textures use to have several unique textures or materials assigned to their surfaces. Each activation of such a texture during the 3D rendering process requires changes of the rendering state, which is a quite expensive operation during 3D rendering. The occurrences of this operation should therefore be minimized for an efficient 3D rendering.

The import of 3D model data is implemented in a preprocessing tool that merges the geometry from the source model that usually has a fine-grained semantic model structure, into geometry batches that can be efficiently processed by modern 3D graphics hardware. The optimal size of these batches depends on the capabilities of target graphics hardware. There is always a tradeoff between consumption of memory and computing power for rendering of unnecessary geometry (i.e., geometry that is included in a batch but is not included within the camera's view frustum, such as building interior) and the amount of CPU processing that is necessary to select and upload the geometry batches to be processed for a single frame.

Besides processing and rendering of 3D geometry with large amounts of individual textures is usually a very inefficient procedure, since the rendering state needs to be changed every time a different texture needs to be used. We address this issue by merging all model textures into a single texture atlas in order to optimize the rendering performance and minimize necessary texture switches. Modern graphics hardware currently supports textures with a maximum of 16384 x 16384 pixels. The texture atlas created during preprocessing of large, fully textured virtual 3D city models, like the one generated from the Berlin 3D city model, are usually far bigger than this value. To address this problem, we implement a *Virtual Texturing* (Lefebvre,

Darbois, & Neyret, 2004) technique that enables the rendering system to manage nearly arbitrary sized texture atlases and render those models efficiently. The technique works in analogy to the virtual memory management of modern operating systems. The complete texturing of the 3D city model uses a single physical texture allocated in GPU memory. The physical texture provides a number of equally sized slots for texture tiles created from the overall texture atlas. During a prerendering pass per frame, an image that contains the number of a visible texture tile per pixel is generated using the GPU. The image is analyzed to determine the texture tiles that are currently not available in the physical texture. The missing tiles for the current frame are loaded from HDD into the GPU memory. This process is very efficient, since only those texture tiles are loaded from disk that are necessary for texturing the geometry that is currently visible. Due to this optimized mechanism for texture management, very large city models with individual textures can be handled by the rendering system since the resource requirements, except for the required disk space, do not depend on the amount of textures.

The 3D rendering service exposes a Web View Service interface (WVS) for querying images of server-side 3D model data. As all core 3D rendering techniques are implemented and performed on a dedicated 3D rendering server, there are no multi-platform issues regarding the stability and implementation of the 3D rendering techniques for specific client hardware.

3.2. Accessing 3D City Model Data

A client application integrates application specific spatial and georeferenced datasets. Connecting application data with the virtual 3D city model and its visualization is essential for building useful applications based on virtual 3D city models as it opens up the CityGML-based virtual 3D city model as integration platform for feature-based 2D and 3D geodata. The rendering backend used for generating color image tiles is also used to render object ID images (see Figure 3(c) and (d) for example tiles). Here it follows

the G-Buffer concept (Akenine-Möller, Haines, & Hoffman, 2008; Saito & Takahashi, 1990), well known in the 3D rendering domain, that allows to create several rasterized information layers per pixel. In this way, additional attributes can be encoded in images on a per pixel basis. Object ID images, as one type of such information, assign an identification color per object that is unique for this server and 3D city model. The components of the color of a pixel encode a 24bit integer number (allowing for about 16.7 million distinct objects per model) that can be used to access data connected to the feature at a pixel position. A mapping between object id and feature id (the database identifier of the underlying feature dataset) is stored on server side. The object id number is assigned by the rendering service during the import of the model elements for rendering (i.e., CityGML datasets or planning models). Using this relatively simple mapping, a connection can be made between pixels of image tiles and feature data in the database. An application server holding the feature data in a database is able to generate a response for given object IDs (values extracted from object ID images for specific positions) or feature IDs (feature identifiers within the database) that contain generically structured data, containing named attribute groups and corresponding attributes for request. The requested attribute groups can be configured via an URL parameter. By using such a generic approach as default behavior, we are able to implement the presentation code in client applications also in a generic way of providing feature attribute lists or server-generated HTML to users. More specific applications can be implemented using external reference attributes leading to connected data sources delivering more specific data in specialized formats.

3.3. Application Contents

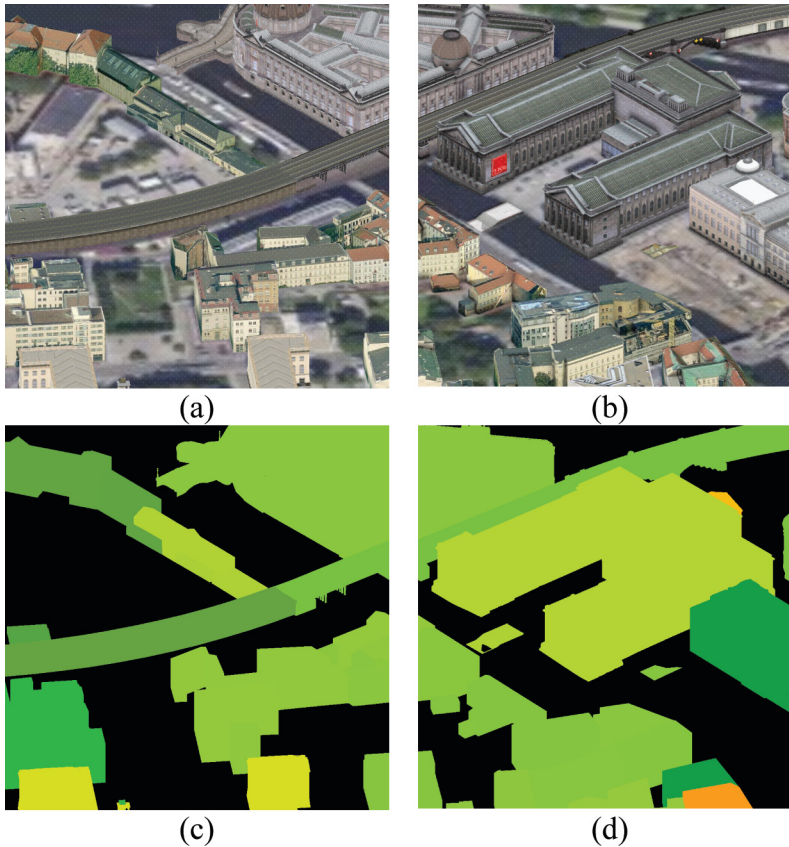
We distinguish two kinds of data to be integrated, regarding their update frequency: (a) data that changes frequently or has to be up-to-date in order to keep the application useful (e.g., vehicle positions, sensor values, etc.) and (b) data that

changes mostly infrequent or whose update date is not very relevant (e.g., infrastructural data, POIs, or landmarks). Data that changes infrequently is harmonized to a common internal data format that is optimized for data size and adjusted to the capabilities and needs of client applications. Further, the geo locations of geometries are translated into a common geographical coordinate system in order to keep client applications as simple as possible. In this way, data formats that are consumed directly by client applications can be reduced to a minimum while keeping the size of the data that has to be transferred minimal. This means, e.g., to avoid transmission of unused attributes or to aggregate data into one document which would otherwise need several client requests to a remote server per item that would imply additional overhead for data transmission and decoding on client side. The necessary service and format adaptors are implemented as configurable scripts that are executed regularly in order to keep the server-side application data as up-to-date as it is desired.

Client-side service adaptors are used for frequently changing data or other data that should be updated at runtime, such as real time traffic data, object locations, or sensor values. The services that provide this kind of data are queried directly by client applications in order to fetch the most up-to-date data that is available. Client-side caching is the only way of optimizing the amount of data that is transferred for this type of data, which allows users to control how often the services are queried. Since we cannot influence the format and complexity of the information delivered by 3rd party services, there can be a significant data overhead between the data needed for immediate visualization on the map and the data that is provided by a service. Further, the service response time for 3rd party data, which contributes significant parts to a user's subjective impression of the overall application performance, may vary significantly.

Another type of application specific data is additional 3D data that can be displayed in client applications, e.g., alternative planning

Figure 3. Examples for corresponding image tiles generated by the tileable color (subfigure a) and b)) and corresponding object map images (subfigure c) and d). The object id images allows client applications to connect each pixel position of the color image to a.



models or visualization of specific areas in the course of time. This is implemented using overlay tiles that are generated using models that are integrated into the base city model, e.g., for providing different planning variants or different combinations of them as an option to inspect. Tile sets for the affected areas can be activated or deactivated in client applications. They are rendered after the tiles of the base map (usually the 3D city model with a specific map layer). In this way, the visible image tiles of the base model are exchanged against the alternative ones covering the same area as transparent or opaque overlay. The same client-side mechanism applies for changing the styling of the map that is described in more detail in Section “Map Styling and Configuration”.

4. AUTOMATED GENERATION OF APPLICATION DATASETS

For a cost efficient provisioning of customized mobile and web-based applications based on massive, information rich virtual 3D city models, we implemented a configurable process that generates application datasets and configuration files for generic client applications. These files are encoded in JavaScript object notation (JSON⁴), which is easy to use an implemented for different target platforms (we developed a browser-based variant, as well as native applications for Android and iOS). They contain all the information that is specific for the area of interest and the application use case (e.g., POI-Data, custom map overlays, etc.). The overall

process is running automatically in large parts, including the import of 3D city model data into the rendering service, generation of oblique map tiles, and creation of according templates for application configuration files. In this way, application data can be regenerated, respectively updated efficiently.

The overall application configuration that is used for generating application specific datasets is decided into two parts: The definition of tile sets and the definition of POI datasets to prepare per application. The generation process for tile datasets can be parameterized in several ways. First, a list of tile sets is defined, each defining the 3D model data and its style.

The application metadata is consumed by client applications to configure their map contents and their interactive behavior (e.g., URL to fetch upon touch/click a map icon). Map content definition primarily contains the URL for tiles of map layers. Since map layers may either be base layers, covering the whole area of interest, or additional overlay layers that can optionally be added or removed on top of base layers.

5. MAP STYLING AND CONFIGURATION

Custom application-specific styling of a 3D view is one of the core advantages of our approach. In contrast to other solutions for mobile and web-based visualization of large-scale geodata, our solution is based on synthesized views on 3D city model data instead of providing the 3D geodata (respectively textured geometry) itself to client applications. Here, the data displayed on a map can be selected fine grained way. There is a wide range of options ranging from including or omitting single features to the replacement of complete area by planning variants in order to visualize planned scenarios, e.g., in a public participation application. Using our rendering system specific styling can be applied to complete views as well as for selected model elements (e.g. for drawing a

user's attention to certain regions of interest in specific zoom levels).

One of the most influencing parts of the visualization for oblique map tiles is the underlying terrain map. Depending on the purpose of an application, the terrain map can communicate different kind of information. Therefore, selection of rendered terrain map layers is one of the key decisions when defining a styling for an oblique map layer. The rendering service is able to integrate terrain map layers from WMS instances over the network or local maps that have been deployed to the rendering server machine previously.

Further, maps with different representations of the 3D data depending on the current zoom level could be generated. Examples for promising approaches for view dependent stylization of virtual 3D city models were provided, e.g., by Semmo et al. (Semmo, Trapp, Kyprianidis, & Döllner, 2012) or Pan et al. (Pan et al., 2013). By emphasizing features that are important for orientations on larger scales (also by partially moving from 3D to 2D representations) and providing a detailed 3D visualization on smaller scales, such techniques can help to improve a users perception and orientation in virtual 3D city models. This stylization is currently not implemented within the rendering service, but provides a promising option for providing a new type of high quality 3D visualization to broader user basis.

Modern web (e.g. using HTML5 Canvas elements) and mobile APIs (e.g., iOS or Android SDK) allow for an efficient implementation of 2D graphics effects with a minimum of system requirements. Given the fact, that object id images are transferred to clients alongside the color images, object-based rendering effects such as highlighting of single objects or groups of objects can be implemented (see Figure 4 for an example). Using such techniques, further features, i.e., visual search through highlighting objects matching a specific search query, can be applied to applications.

Figure 4. Example for client-side rendering effects implemented using a HTML5 canvas object id images. Subfigure (a) shows a debug view with object ids blended into the color image. Subfigure (b) shows the result of the highlighting effect for a specific object id.



6. CASE STUDY - REAL-ESTATE MARKETING

In our case study we have built the application “smartMap Berlin” (see Figure 1) targeting a use case in the area of business real-estate marketing. Our cooperation partner “Berlin Partner GmbH” conducts the city marketing for the city of Berlin and provides consulting services for companies interested in settling in the Berlin metropolitan region. Since finding a real-estate property that fits a company’s needs is one of the major challenges, their real-estate portal offering numerous real-estate properties in different categories is a major part of their services.

The virtual 3D city model of the city of Berlin is one of the largest and most complete city models in the world. It covers the whole urban area of Berlin that consists of about 890 square kilometers and roughly 600,000 building models, far most of them modeled in CityGML LOD2 and textured with real-world textures. About 350 building models are included as LOD3 or LOD4 with a highly detailed outer shell or even modeled interiors. All in all, the model comprises about 5 millions single tex-

tures for building models. Additionally a high resolution aerial photography and several other map layers for terrain texturing are available from standards-based Web Map Services (de La Beaujardiere, 2004). The target of the project is to apply city models with a high usability. A very strong argument for using our approach was that transmitting the original 3D data of the city model to the clients would hardly be possible due to legal restrictions and security - 3D geodata is stored and managed only by the dedicated 3D server.

We configured the process to generate three different combinations of underlying map (an aerial photo, an OpenStreetMap-based map, and an official map of Berlin using a 1:5000 scale) as base layers. Each layer was generated for all four principal directions in eight zoom levels, leading to 412885 JPEG-encoded image tiles per direction and base layer resulting in an overall data size for all base-layers and all map orientations of about 40 GB (leads to average tile size of about 8 kB). The overall area covered by the map is around 1700 km², which is far more than the area of the city alone, but necessary to cover all of the cities’ area in one rectangular map area. The tile sets

are organized in a way that the distinct map zoom levels form a hierarchical quadtree that is inspired by a conventional Tile Map Service⁵. Since we generated the tile images as oblique views with a viewing angle of 30 degrees to ground level, our tiles cover twice the area of conventional TMS tiles within an equal zoom level. To compensate that, clients compute the required tile numbers in a custom way to match with the maps georeference.

The core functionality of the system is to explore the available real-estate offers and provide information about the social, economic, and infrastructural context of the city of Berlin. For our partner, it was essential not to show out-of-date real estate offers, since Berlin Partner consultants use the application in in-situ counseling sessions together with customers. This is why we implemented a special type of client-side data adaptor to query the real-estate service, which is also the basis for the real-estate portal website. The service provides a list of available real estates as well as details per offer encoded in JSON. The overview list is fetched on start of the application. Any additional detail data for single real-estate properties, i.e., images, detailed descriptions, and pricing information, is requested on demand from the real-estate service.

The POI information for infrastructure, economy, and social information is considered to change very infrequently. Therefore, they are integrated into the application through a unified JSON format that is held on a server together with the tile data. Precaching and format homogenization assures, that the data originating from different sources can be accessed with the same simple format and the same (low) server-side latency.

7. DISCUSSION

For most purposes, such as user orientation within or exploration of urban areas, the very simple mode of presentation and user interaction used in our oblique view applications turns out to be very effective in use and efficient

in deployment. The complexity that comes with applications allowing a free 3D camera navigation in virtual 3D environments is often not needed by users to perform typical tasks using a virtual 3D city model. While a free 3D camera navigation in conventional applications enables users to perform tasks such as evaluation of lines of sights or in-detail inspection of objects in 3D, the map-like interaction and presentation used in our approach is reducing the mental effort for navigating a 3D environment. Nevertheless, since there are only a limited number of camera orientations that can be used to explore the virtual 3D city model, object occlusion cannot be avoided. The degree of the occlusion is dependent on the current map zoom levels (closer camera distances to objects cause larger numbers of potentially interesting objects to be occluded) and the shape of the underlying model (e.g., height of buildings and the general building density). For application cases, where occlusion caused by buildings is a problem, special rendering styles, such as semi transparent buildings or geometry cuts at a given height, could be applied to buildings occluding objects of interest for a particular application. Further, client applications could be extended to include a in detail inspection mode providing detail views for specific areas of interest. Currently we did not implement such additional interaction and presentation techniques, since these are very application specific. Since our current application focuses on a viewing scale that allows to judge the context of an real estate offered, the implemented presentation and interaction techniques were seen sufficient by users of the application to navigate through the virtual 3D city model.

Model data access plays a viable role when using 3D city model applications. Since the functionality to be implemented in client applications is kept simple (map presentation, web-based information retrieval for model features via URLs), concrete application logic can be implemented in application servers. This also enables the integration of existing web-based application into browser-based apps as well as into mobile applications running on

smartphones or tablet. This way, functionality and system components that have been created before to support work processes can be reused also for 3D city model applications with no or little adjustments regarding the layout and compatibility with touch enabled devices. Nevertheless, the approach of generating configurable client applications is still limited in terms the way existing web-based solutions integrate into such an application. It is not possible to provide a native user interface for client web services that do not implement a specific standard, e.g., a transactional WFS for viewing or editing feature data. Custom service endpoints need to be implemented as a plugin for client application, as it was done for the real-estate service used for our demonstration case. Therefore code changes are necessary to integrate new types of services into client mobile applications.

8. CONCLUSION AND FUTURE WORK

In this paper we presented a novel approach for provisioning applications based on massive virtual 3D city models on heterogeneous platforms, namely mobile devices, running Android or iOS operating systems, and web browsers. The process involves only a very limited amount of manual customization for applications dependent on the actual use case, which makes the provisioning and operation of high quality 3D visualization applications for a larger audience very cost effective and easy to handle.

The approach can be extended towards automatic derivation of applications for exploration for complex indoor models, which have not been addressed yet. Especially the exploration of complex modeled building structures, including stories, provides promising possibilities for generating, e.g., a smartMap layer per story and allowing users to select the story to show on client side. Further, these applications provide possibilities to integrate 3D city model visualization into existing systems

and workflows (e.g., for real-estate marketing) generating additional values for these systems.

The potential of additional G-Buffer layers for the interactivity and client-side rendering of map layers are not yet finally exploited. Additionally to the object ID maps, which are currently used for identifying single objects within the map, maps encoding the object type or category can be used either to implement client side highlighting effects for specific types of objects as well as providing a type sensitive interaction within the map application. Currently, the configuration for tile generation and POI-data integration is created by hand. A user interface for the overall process for application data generation could speed up the overall process and would also allow non-expert users to configure and generate application datasets.

ACKNOWLEDGMENT

This project was supported by the Berlin Partner GmbH by providing us with the virtual 3D city model of Berlin, which was funded partially by the European Fund for Regional Development. Further we thank 3D Content Logistics GmbH for their support during this project.

REFERENCES

- Akenine-Möller, T., Haines, E., & Hoffman, N. (2008). *Real-Time Rendering* (3rd ed., p. 1045). Natick, MA, USA: A. K. Peters, Ltd. doi:10.1201/b10644
- Bao, P., & Gourlay, D. (2006). A framework for remote rendering of 3-D scenes on limited mobile devices. *IEEE Transactions on Multimedia*, 8(2), 382–389. doi:10.1109/TMM.2005.864337
- Basanow, J., Neis, P., Neubauer, S., Schilling, A., & Zipf, A. (2008). Towards 3D Spatial Data Infrastructures (3D-SDI) based on open standards - experiences, results and future issues. In *Advances in 3D Geoinformation Systems* (pp. 65–86). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: doi:10.1007/978-3-540-72135-2

- Behr, J., Jung, Y., Keil, J., Drevensek, T., Zoellner, M., Eschler, P., & Fellner, D. (2010). A scalable architecture for the HTML5/X3D integration model X3DOM. In *Proceedings of the 15th International Conference on Web 3D Technology - Web3D '10* (p. 185). New York, New York, USA: ACM Press. doi:10.1145/1836049.1836077
- Boukerche, A., & Pazzi, R. (2006). Remote rendering and streaming of progressive panoramas for mobile devices. In *Proceedings of the 14th annual ACM ...* (pp. 691–694). Retrieved from <http://dl.acm.org/citation.cfm?id=1180785>
- Christen, M., Nebiker, S., & Loesch, B. (2012). Web-based large-scale 3D-geovisualisation using WebGL. *International Journal of 3-D Information Modeling*, 1(3), 16–25. doi:10.4018/ij3dim.2012070102
- De La Beaujardiere, J. (2004). OGC Web Map Service Interface. Open Geospatial Consortium. Retrieved from <http://www.opengeospatial.org/standards/wms>
- Doellner, J., Hagedorn, B., & Klimke, J. (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* (pp. 97–100). New York, NY, USA: ACM. doi:10.1145/2338714.2338729
- Gesquière, G., & Manin, A. (2012). 3D visualization of urban data based on CityGML with WebGL. *International Journal of 3-D Information Modeling*, 1(3), 1–15. doi:10.4018/ij3dim.2012070101
- Gröger, G., Kolbe, T. H., Nagel, C., & Häfele, K.-H. (2012). *OpenGIS® City Geography Markup Language (CityGML) Encoding Standard Version 2.0.0*. Retrieved from <http://www.opengeospatial.org/standards/citygml>
- Haber, R. B., & McNapp, D. A. (1990). Visualization Idioms: A Conceptual Model for Scientific Visualization Systems. In *Visualization in Scientific Computing* (pp. 74–93). IEEE.
- Hagedorn, B. (2010). Web view service discussion paper, Version 0.6. 0. *Open Geospatial Consortium Inc.* Retrieved from <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Web+View+Service+Discussion+Paper#0>
- Jackson, D. (2013). WebGL Specification. Khronos Group. Retrieved from <https://www.khronos.org/registry/webgl/specs/latest/1.0/>
- Jeschke, S., Wimmer, M., & Purgathofer, W. (2005). Image-based representations for accelerated rendering of complex scenes. *STAR Reports, Eurographics*, 1–20. Retrieved from <http://www.cg.tuwien.ac.at/research/publications/2005/jeschke-05-ISTAR/jeschke-05-ISTAR-Paper.pdf>
- Lamberti, F., & Sanna, A. (2007). A streaming-based solution for remote visualization of 3D graphics on mobile devices. *IEEE Transactions on Visualization and Computer Graphics*, 13(2), 247–260. doi:10.1109/TVCG.2007.29 PMID:17218742
- Lefebvre, S., Darbon, J., & Neyret, F. (2004). *Unified Texture Management for Arbitrary Meshes*. Retrieved from <http://hal.inria.fr/inria-00070783/>
- Mao, B. (2011). *Visualization and Generalization of 3D City Models*. Royal Institute of Technology.
- Pan, B., Zhao, Y., Guo, X., Chen, X., Chen, W., & Peng, Q. (2013). Perception-motivated visualization for 3D city scenes. *The Visual Computer*, 29(4), 277–286. doi:10.1007/s00371-012-0773-1
- Papazoglou, M. P., Traverso, P., Dustdar, S., & Leymann, F. (2007). Service-Oriented Computing: State of the Art and Research Challenges. *Computer*, 40(11), 38–45. doi:10.1109/MC.2007.400
- Prieto, I., & Izkara, J. L. (2012). Visualization of 3D city models on mobile devices. In *Proceedings of the 17th International Conference on 3D Web Technology - Web3D '12* (Vol. 1, p. 101). New York, New York, USA: ACM Press. doi:10.1145/2338714.2338731
- Saito, T., & Takahashi, T. (1990). Comprehensible Rendering of 3-D Shapes. *Computer*, 24(4), 197–206.
- Sample, J. T., & Ioup, E. (2010). *Tile-Based Geospatial Information Systems*. Boston, MA: Springer US; doi:10.1007/978-1-4419-7631-4
- Schilling, A., & Kolbe, T. H. (2010). Draft for Candidate OpenGIS® Web 3D Service Interface Standard. Open Geospatial Consortium. Retrieved from http://portal.opengeospatial.org/files/?artifact_id=36390
- Semmo, A., Trapp, M., Kyprianidis, J. E., & Döllner, J. (2012). Interactive Visualization of Generalized Virtual 3D City Models using Level-of-Abstraction Transitions. In *Computer Graphics Forum* (Vol. 31, pp. 885–894). doi: doi:10.1111/j.1467-8659.2012.03081.x
- X3D Standard. (2005). Web 3D Consortium. Retrieved from http://www.web3d.org/x3d/specifications/x3d_specification.html

ENDNOTES

- ¹ <http://sh.edushi.com>
- ² <http://bj.o.cn>
- ³ <http://osm2world.org>
- ⁴ <http://json.org>
- ⁵ http://wiki.osgeo.org/wiki/Tile_Map_Service_Specification