

Detail-In-Context Visualization for Satellite Imagery

Joachim Böttger, Martin Preiser, Michael Balzer, Oliver Deussen

Department of Computer and Information Science, University of Konstanz, Germany

Abstract

We use the complex logarithm as a transformation for the visualization and navigation of highly complex satellite and aerial imagery. The resulting depictions show details and context with greatly different scales in one seamless image while avoiding local distortions. We motivate our approach by showing its relations to the ordinary perspective views and classical map projections. We discuss how to organize and process the huge amount of imagery in realtime using modern graphics hardware with an extended clipmapping technique. Finally, we provide details and experiences concerning the interpretation of and interaction with the resulting representations.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Methodology and Techniques]: Interaction Techniques; H.5.2 [Information Interfaces and Presentation]: User Interfaces;

1. Introduction

Modern satellite and computer technologies provide data with scales ranging from the whole Earth to very small objects located on it. By using internet-based services like Google Maps and Microsoft's Virtual Earth, everybody can explore the resulting very complex datasets.

With modern graphics hardware and broad network connections, it is possible to interactively render perspective views of the globe on consumer hardware. This enables users to not only look at the Earth from straight above, but to tilt the camera, gaining the advantage of the familiar subjective perspective: Parts of the world closer to the virtual camera, in the foreground, are depicted larger than objects in the background of the resulting images. This allows for inspecting details of the Earth without losing orientation in their wider context. However, the range of scales depicted in one image is limited by the fact that the representation is getting compressed towards the horizon. Another limitation of the classical perspective is the fact that the information behind the camera can not be seen at all.

We propose a new perspective for the exploration of the Earth: Our distortion oriented approach of complex logarithmic views shows details and context with greatly differently sized objects in one single image. The mapping function we employ can represent single houses seamlessly connected to the continents they are standing on. Similar

to a panoramic perspective, this approach simultaneously presents the whole world showing all directions around the viewpoint. An additional advantage is the avoidance of anisotropic compression, preserving the shapes of the image details, which is especially important for the identification of geographic objects. By utilizing an extended clipmapping technique to organize and process the huge amount of necessary data, we are able to render our representations in realtime, allowing the user to interactively explore and navigate within very complex datasets of satellite and aerial imagery.

The paper is organized as follows: In [Section 2](#) we discuss distortion oriented detail-in-context methods and their properties. In [Section 3](#) we firstly discuss the general principle and properties of the complex logarithm. We then derive a mathematical connection between the intuitive concept of perspective and the complex logarithm, and show how to apply the transformation to satellite imagery. Then we present the implementation aspects of the data organization and the actual rendering in [Section 4](#), and describe the resulting representations in [Section 5](#). Finally, in [Section 6](#) we discuss our results, and give an outlook to future work.

2. Background

The intention of showing small details on the Earth's surface while preserving the overview of the whole world or its parts is a typical detail-in-context problem [[KR96](#), [Kea98](#)]. The

utilization of zooming and panning, or different windows for different detail levels, puts a significant mental strain on the user because of the necessary switch between the views. Rather, nonlinear magnification techniques help in the effort to show both, detail and context, in one single image.

As a side effect, these techniques, which map points in a space to other points in another space, introduce distortion that changes several properties of the depicted information. If we look at the infinitesimal elements from which the information is pieced together, the properties that are of interest to us are *size*, *orientation*, and *connectivity* between neighbors preceding the mapping.

We follow [Nee97], which describes mappings as locally linear, meaning that small circles are mapped to small ellipses. The alteration of *size* of the different parts of the information is the most critical property for their recognizability: If the scaling operations are not equal for all directions, then so-called anisotropic compression is introduced. This compression leads to degenerated shapes, which, for large differences in scaling factors, results in almost linear structures that are not recognizable anymore. Conformal mappings, on the other hand, are mappings that scale the pieces of information equally in all directions. This means that small circles are mapped to other small circles without causing distortions to local angles or introducing anisotropic compression.

A change of *orientation* of the pieces of information itself leaves intact what we perceive as the object's shape. However, it is more difficult to recognize the shape of the object after a rotation, especially if there exists a clear notion of the object's orientation—for example, the shape of the European continent is often not recognized if it is presented upside down.

The property of *connectivity* describes whether a mapping maps two infinitesimally close points to distinct positions, introducing a cut between these points. If this cut separates objects belonging together, a user has to mentally reconnect them, which increases the mental strain. However, given the case that strong magnifications are required, these cuts are preferable to approaches in which parts of the information are compressed to such an extent that they are no longer recognizable. Especially for geographic data, cuts are common when generating flat cartographic map projections.

Several algorithms were developed that allow displaying complex information by magnifying one or several parts of interest while simultaneously shrinking the surrounding space: Spence [SA99] was the first to show details in context on a computer display in one seamless visualization by using bifocal views, partially showing non-uniformly scaled content. Other approaches are the Perspective Wall [MRC91] and the Document Lens [RM93] that are using perspective projection. Another class of mapping functions are Fish-eye views [KR97], which generate symmetrical distortions around the center of interest. Furthermore, there exist many

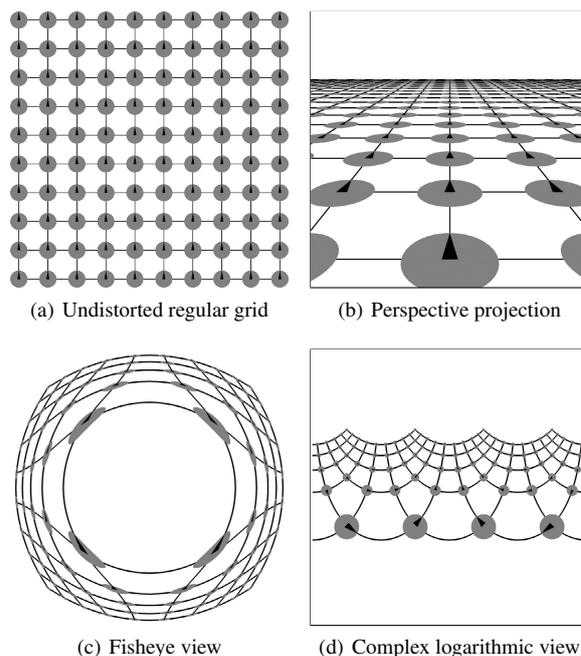


Figure 1: Examples for the distortion of a regular grid (a) using nonlinear magnification techniques (b,c,d)

other techniques, such as [KR96, Car99, RC94]. All of the mentioned detail-in-context techniques anisotropically compress parts of the information, where the shape of objects which are not in the center of interest are often unrecognizable after the transformation.

The perfect mapping that keeps all properties intact while scaling parts of a plane differently does not exist. Thus, choosing a mapping function for a certain problem always has to be a tradeoff between the different kinds of distortion it introduces. We believe that, for the exploration of complex satellite and aerial imagery over several orders of magnitude, keeping shapes intact is more important than keeping rotations constant and even justifies giving up some of the connections between points in the plane.

In [BBD06] we described a similar approach of complex logarithmic views for visualizing flat vector data in the field of information visualization, which also possesses parallels to the astronomic maps in [GJS*05]. The new contribution of this paper is founded by the differences of the method and its implementation for the application domain of satellite and aerial imagery. We present a mathematical connection to the intuitive concept of perspective to motivate and illustrate our method, and discuss how different cartographic map projections influence the resulting representations. We address relevant implementation aspects to organize and process the required massive amount of data, using an extended clipmapping approach.

3. Method

Essentially, our method is a mapping between points in \mathbb{R}^2 using a complex logarithm. In this section, after describing this complex function and its properties, we motivate our approach by deriving a connection to the ordinary central point perspective. Then, we discuss the relationship to Mercator projections and other cartographic map projections.

3.1. Complex Logarithm

The complex logarithm [Nee97] is a function that maps complex numbers onto other complex numbers. With such a complex function it is possible to describe a mapping of points within the two-dimensional plane. The complex logarithm belongs to the class of analytical functions, and thus possesses certain desirable properties: The derivative of a complex function is itself complex. Similar to the derivative of a real-valued function, the magnitude of the derivative indicates a scaling factor for a certain point in the geometrical interpretation. This scaling factor is for that point equal for all directions. This means that the parts of the mapped plane are not anisotropically compressed.

The complex logarithm $\log(z)$ maps every complex number z to another complex number with the logarithm of the original number's magnitude as its real part. The imaginary part of the transformed number is the angle of the original number in radians:

$$\log(z) = \ln(|z|) + i \arg(z)$$

Through the resulting transformation of the corresponding plane, all the points on concentric circles around the origin are mapped onto parallels with identical real value. The transformation scales parts of the corresponding plane depending on their distance from the origin. In the transformed image, the scaling factor increases or decreases exponentially by moving along the real axis.

Points on straight lines starting from the origin in the original data are mapped to parallels with identical imaginary value. The argument function is responsible for the fact that certain points connected in the original data are disconnected after the transformation: Two points with equal distance and angles of $+\pi$ and an angle slightly bigger than $-\pi$ are very close before, but on opposite sides of the defined area after the mapping. This introduces a cut along the line of points with $|\arg(z)| = \pi$. Typically, the complex logarithm would yield an image with the highest magnification on the left, and the lowest magnification on the right side. However, the chosen orientation of our representations is similar to the perspective view we are used to [Mor03], since objects far away from the center of interest are mapped more to the top of the image. The resulting mapping is further illustrated in Figure 2.

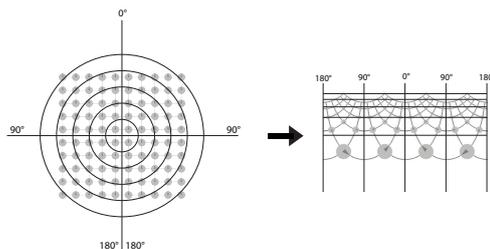


Figure 2: The complex logarithm maps concentric circles to horizontals, and directions to verticals.

3.2. From Perspective to Complex Logarithmic Views

In this section we motivate our approach by showing its mathematical and intuitive connection to the familiar perspective projection; both map a two-dimensional space, which the surface of the Earth basically is, in a way that differently scales parts of the image. Objects close to the midpoint of the complex logarithm, or to the viewpoint of the perspective projection respectively, are enlarged, while farther parts are depicted much smaller.

To show the mathematical connection between the ordinary perspective view and the complex logarithmic view, we start with the former, which maps points by projecting them along straight rays on a viewing plane, as illustrated in Figure 3. An example for a perspective view of the Metropolitan Museum in New York, can be seen in Figure 4(a). A first shortcoming of this form of perspective is, that it only shows parts of the world in the direction in which the camera is pointed. The information about what is beside and behind the camera is completely lost.

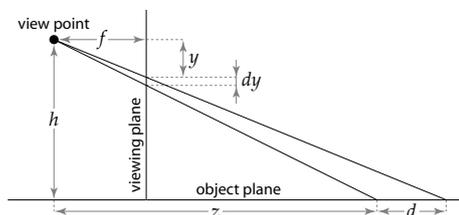


Figure 3: Side view of a perspective projection

In contrast, a panoramic perspective uses an unfolded cylinder as viewing plane, and presents information in all directions around the camera simultaneously. The result of this form of panoramic perspective view is shown in Figure 4(b). Like in the ordinary perspective view, every vertical still is the image of a line from the viewpoint outwards, but information in every direction is represented on the cylinder. The horizontals no longer correspond to lines parallel to the viewing plane, but rather to circles around the viewpoint.

Although the complete plane is depicted, the panoramic perspective view still suffers from the second shortcoming of perspective projection, namely the introduction of anisotropic compression towards the horizon. If we look at

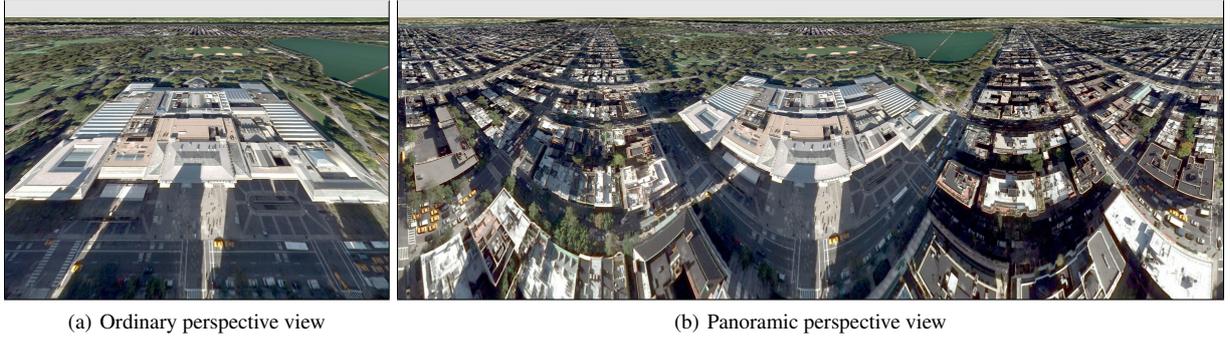


Figure 4: Perspective views of the Metropolitan Museum in New York: The information beside and behind the virtual camera is lost in the ordinary perspective view, whereas it is preserved in the panoramic perspective view. In both perspective mappings, the information far away from the viewpoint is compressed into to a singular line, the horizon.

the magnification factors for small pieces of the depicted plane in vertical and horizontal direction, we can show that the compression approaches infinity towards the horizon:

$$dx = d \cdot \frac{f}{z} \quad dy = d \cdot \frac{f \cdot h}{z^2}$$

$$c = \frac{dy}{dx} = \frac{z}{h} = \frac{f}{y}$$

Here, dx and dy are the size of the image of an infinitesimal circle with diameter d in horizontal and in vertical direction, respectively, and c is the compression factor. Larger values of c mean that the relative size of the depicted circle in the vertical direction gets smaller. The compression obviously gets arbitrarily big for large distances from the viewpoint z , squashing the information which is very distant from the viewpoint into a singular line, the horizon.

The connection to our complex logarithmic view is obvious, if we consider it as a panoramic perspective view, which simply strives to avoid any anisotropic compression towards the horizon. Figuratively speaking, it is necessary to stretch or squash every piece of perspectively mapped information just enough to compensate the compression introduced in the mapping. This means, we have to vertically scale every piece by the compression factor. This pushes every row in the depiction to a new position in the image, which we can calculate by integrating the magnification factors:

$$y_{new}(z) = \int \frac{f}{z} dz = f \cdot \ln(z) + C$$

The projection cylinder has a radius of f , which results in a projection width of $2\pi f$. The height of the projection is potentially infinite.

Summarizing, this stretched panoramic perspective view, which compensates anisotropic compression towards the horizon, maps points in \mathbb{R}^2 in a way that the horizontal position depends on the angle to the viewpoint, and the vertical position depends on the logarithm of the distance to the

viewpoint. Mathematically, this mapping is the complex logarithm from the previous section. The final result of a complex logarithmic view showing the Metropolitan Museum is given in Figure 9 at the end of the paper. The similarities of the lower part, and the differences of the upper part, between this image and the panoramic perspective are evident.

In the final result, the upper part of the image shows the whole Earth. For this result, we have to take its spherical shape into account. The complex logarithm operates on points in a two-dimensional plane, therefore we need a mapping from the sphere to the plane. These transformations are offered by cartographic map projections, which are the subject of the next section.

3.3. Cartographic Map Projections

After having analyzed different map projections, it became apparent to us, that there exists a strong relationship between complex logarithmic views and the well-known and widely used Mercator projection. Using this projection in its oblique form, is in its core a complex logarithmic view. We will therefore firstly illustrate this relationship, before generalizing to other map projections and their properties.

Relationship to the Mercator Projection

The Mercator projection [Sny87] maps the Earth's geographic coordinate system to \mathbb{R}^2 . It belongs to the group of cylindrical map projections. Beside the standard Mercator projection, where the axis of the projection cylinder runs through the poles, there also exist more general types: First, the transverse type, where the cylinder's axis is orthogonal to the axis of the standard type, and second, the oblique type with an arbitrarily angled axis, while for both types, the axis still passes the Earth's center.

The standard Mercator projection has the following properties: The scale is true only along the Equator, but reasonably correct within 15° of the Equator. The areas and

shapes of large regions are distorted, whereby the distortion increases away from the Equator, and is extreme in polar regions. The projection of the poles itself would require a cylinder of infinite height. Therefore, the cylinder is cut off for large latitude values. Usually, this latitude threshold is between 70° and 85° north or south, depending on the intended application. Mercator projections are conformal mappings in which angles and shapes within any small area are essentially true, not introducing anisotropic compression.

In our approach we aim for depicting very small details of the Earth's surface within the context of the overall world, while avoiding local distortion in a way that the shapes of the geographical objects, such as rivers, islands, or continents, remain recognizable for the user. For the special case of presenting a detail view of the north or south pole in the context of the overall world, the standard Mercator projection offers exactly such a mapping. When applying the cut of the projection cylinder at high latitude values above 85° , then the poles are extremely magnified at the top and bottom of the resulting image, and the middle of the image presents the rest of the world. This characteristic of extreme magnification at top and bottom of the Mercator projection, which is usually identified as its main drawback, can be exploited to generate detail-in-context representations of any point on the Earth's surface by utilizing oblique Mercator projections.

Given aerial imagery with sufficient resolution, a detail-in-context representation of a certain point of interest on the Earth's surface is obtained by using an oblique Mercator projection, for which the axis of the projection cylinder runs through this point of interest. To actually generate this representation, the corresponding latitude and longitude values for each point in an image can be computed by inverting this oblique Mercator projection. Then these resulting latitude and longitude values are used to look up the information in the imagery at an appropriate level of detail by applying the map projection used for the imagery.

To now generalize our method, it is important to understand that the essence of the Mercator projection is a logarithmic transformation. To be more precise, it is a concatenation of the Stereographic map projection and a complex logarithm [PL74]. In the remainder of this section, we discuss map projections that are useful for our representations.

Azimuthal Map Projections

While in general arbitrary map projections could be used, we considered the class of azimuthal map projections [Sny87], which are using a plane as projection surface, as especially applicable. The property qualifying them for complex logarithmic views is that they present true directions, but not necessarily true distances, from a chosen center point to any other point. The projections of points with equal direction yield vertical lines in the complex logarithmic view, whereas points of equal distance yield horizontal lines.

Below we discuss the results of complex logarithmic

views for the following azimuthal map projections: The aforementioned *Stereographic*, the *Azimuthal Equidistant*, and the *Orthographic* projection. These different map projections and their resulting complex logarithmic views are presented in Figure 5.

The *Stereographic* projection is a true perspective, with its point of projection being located on the surface of the sphere, opposite the point of tangency of the projection plane. It is the only true perspective or azimuthal projection that is conformal. The concatenation of the Stereographic projection and the complex logarithm is therefore a conformal mapping itself, and equivalent to the Mercator projection. Due to this relation, the resulting representation needs a potentially infinite space at the top. Since we just want to enlarge the point in the center of interest, we simply cut off the representation close to the other pole, similar to the Mercator projection.

The *Azimuthal Equidistant* projection is not a true perspective, nor is it conformal. It is constructed by plotting a given point, with a given angle to the center point, at a distance from that center proportional to its distance on the sphere. This projection is relevant in practice because its complex logarithmic view presents the whole world without a necessary cut at the top. Rather, the point opposite to the center of interest is mapped to a line at the top.

The *Orthographic* projection is a true perspective that uses a point of projection at infinite distance. It is not conformal either. In contrast to the other two projections, it does not represent the whole Earth in one image, but rather just the half that is visible from the point of projection. Hence, complex logarithmic views using the Orthographic projection reintroduce the concept of a horizon. In this respect, it is a combination of the familiar appearance of the Earth as a globe, and our detail-in-context approach.

4. Implementation

In this paper, we show how to adapt complex logarithmic views for the realtime exploration of the Earth's surface. For our prototypical implementation, we used multi-resolution tiled imagery from Microsoft's Virtual Earth, which we downloaded and cached from the Internet in realtime. The available satellite and aerial imagery is enormously complex, down from continents to single houses, which poses different problems for the organization and rendering of this massive amount of data. Due to the similarity of our mappings to perspective projections and terrain rendering, we can apply ideas from the research on large texture rendering to our problem. Hence, we adapted the frequently used clipmapping technique [TMJ98, STH*07] to our approach.

The purpose of clipmapping is to render perspective views of geometry with very large textures. This is necessary, because using an ordinary mipmap is not feasible due to memory constraints. The clipmapping method benefits from the fact that, in perspective views, not the whole mipmap is

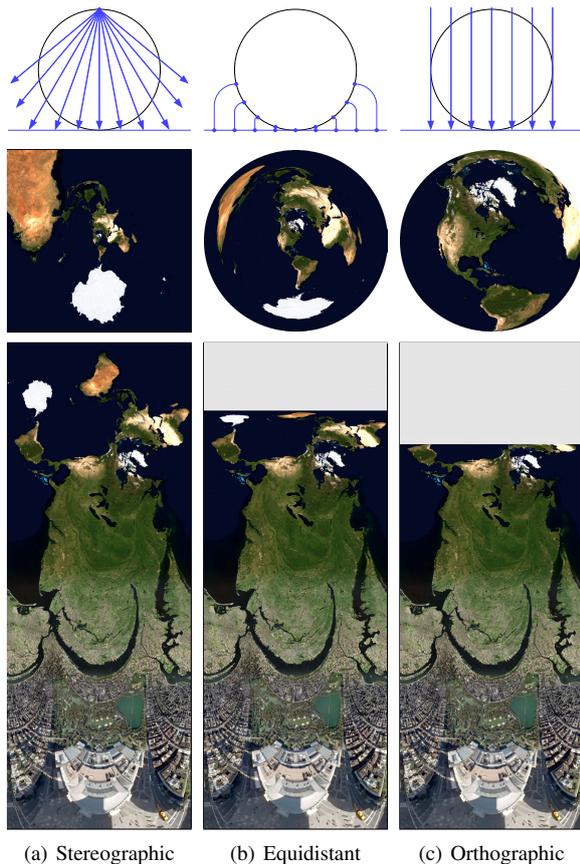


Figure 5: Complex logarithmic views (bottom) of azimuthal map projections (top, middle). The stereographic mapping is transformed to an oblique Mercator projection with the center of interest as one of the poles. It thus necessitates cutting off the second pole on top of the representation. Contrarily, the Equidistant projection is mapped to finite space, while still showing the whole Earth. The Orthographic projection only shows half of the world, introducing a virtual horizon.

needed with an equally high resolution at the same time. Rather, it is sufficient to use only a small subset of the data at any given moment. Depending on the viewpoint, the resolution of the necessary images decreases with increasing distance from that viewpoint. This means, that only nearby objects require high resolution texturing, while objects far away from the viewpoint are rendered with low resolution textures. Consequently, a clipmap is an updatable representation of a partial mipmap, in which each level of the imagery has been clipped to a specific maximum size. This results in an obelisk shape for the stack of images as opposed to the pyramid shape of mipmaps. While moving through the rendition, it is then only necessary to ensure that always the appropriate section of the mipmap for the current viewpoint is represented in the clipmap stack.

We extended the clipmapping technique with regard to our application: We took into account, that the data is sampled in tiles of a certain size, and that it does not cover the complete Earth, for each of the detail levels—for example, the resolution of the data for large cities is usually higher than for rural areas or the ocean, especially. Therefore, we introduced an additional index structure to enable a more flexible arrangement of and access to the image tiles, reducing memory consumption. Additionally, our implementation is realized by only using OpenGL, without the need of special hardware with clipmapping support.

Our rendering system consists of two different parts: The first part is responsible for loading and caching the required image tiles in the memory of the graphics hardware, and managing the index structure to allow the fast and consistent access to the clipmap stack. The second part performs the actual rendering by employing a fragment shader program, which allows for using arbitrary mappings without any geometry operations. The fragment shader calculates for each pixel the corresponding detail level in the imagery by inverting the current mapping, and then computes the pixel's color value by texture interpolation within the determined level. These two parts can be run concurrently, as long as a consistent state of the index structure is guaranteed while the actual rendering is performed.

4.1. Data Organization

Satellite and aerial imagery exists for different detail levels, where each of these levels is stored as a potentially incomplete set of small tiles. Like in clipmapping, for rendering a complex logarithmic view for a specific center of interest, we need only a small fraction of these tiles: Tiles with high resolution near the center of interest, and tiles with low resolution for parts that are far away. Therefore, we only need roughly the same small number of tiles for every detail level. Figure 6 shows a typical footprint of image tiles from different levels necessary for a certain viewpoint by an orthographic and a complex logarithmic view.

As an extension of the clipmapping approach, in order to allow a more flexible organization of the detail levels, we realized the clipmap stack using a three-dimensional index structure, as illustrated in Figure 7. Each two-dimensional layer in that structure contains a detail level, whereby the layers are sorted vertically from low resolution at the top to high resolution at the bottom. Each layer itself contains references to image tiles that are linearly organized in the memory of the graphics hardware. The neighborhood of the references in one layer is identical to the original neighborhood of the image tiles in the dataset. If an entry of a layer does not possess a valid image tile, then it refers to the corresponding entry in a higher layer that possesses a valid image tile. This enables quick access to the data with the best available resolution in the pixel shader, in case the optimal resolution is not available for a geographic location.

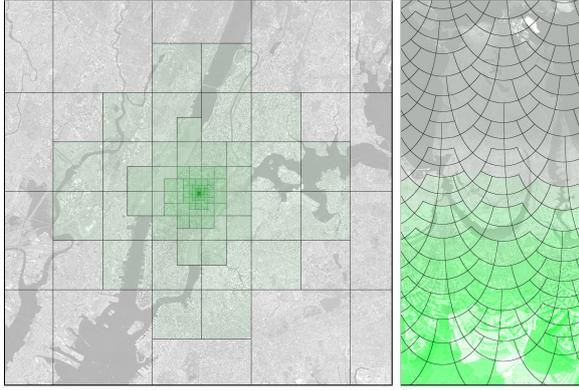


Figure 6: Zones of tiles with different size and resolution around the center of interest in the orthographic (left) and complex logarithmic view (right)—greener means higher resolution. In the transformed view, the concentric zones are mapped to horizontal stripes with almost equal size.

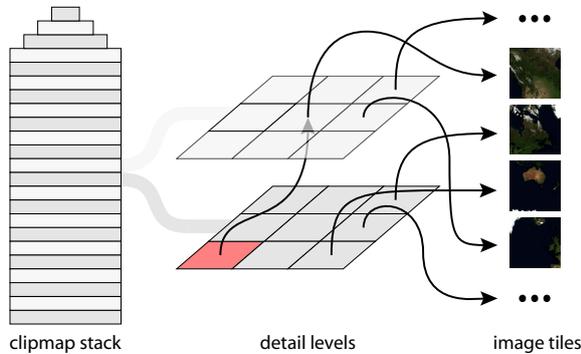


Figure 7: We use a clipmap with an index structure (left) to organize the imagery in the graphics hardware. The index structure contains one layer (middle) for each detail level in the clipmap. These layers contain references to the image tiles, which are arbitrarily arranged in the memory (right).

When the center of interest is moved, we update the references in the three-dimensional index structure, and load only these image tiles that are not yet existent in the graphics memory. The loading is performed in the order of tile importance, loading upper levels, which are offering the context information, first.

The actual size of the index structure, and thereby the number of image tiles needed in memory, depends on the number of detail levels, the size of the image tiles in each detail level, and the size of the resulting representation. For our representations we used datasets with up to 25 detail levels and up to 256 tiles per layer, with a size of 256x256 pixels per tile. Usually, the graphics memory offers more space than needed for storing the entire clipmap stack. Our technique allows to cache image tiles that have already been needed, or are presumably needed in the future, outside of the scope of the index structure.

The advantages of our approach, which realizes clipmapping via an index structure, is the more flexible updating and caching of the individual tiles, as well as the reduced memory consumption for the entire clipmap stack. The latter one is due to the fact that missing tiles, and tiles that are not in the detail levels, still occupy memory within the originally clipmap levels, while with our technique they do not.

4.2. Rendering

After having organized our data with the aforementioned extended clipmapping approach, we can actually render our complex logarithmic views. To achieve maximum flexibility, we implement the rendering algorithm for the programmable fragment shader on the graphics hardware. This allows us to use a variety of mappings without considering any complicated geometry operations, like mesh deformation and refinement or the implementation of cuts. We only draw one large rectangle that covers the entire window, and use the fragment shader to calculate the image by inverting the intended mapping, determining the optimal detail level, and sampling the corresponding image tiles for each pixel.

The detailed algorithm for computing the color value of a pixel is as follows:

1. We determine the location of the pixel as coordinates in our imagery of the Earth. For this step, we invert the concatenated mappings as shown in Figure 8: After normalizing the screen coordinates, we apply the inverse of the complex logarithm, the complex exponential function:

$$e^z = e^{x+iy} = e^x \cos(y) + ie^x \sin(y)$$

We then apply the inverted azimuthal mapping function, to yield the latitude and longitude values. Lastly, we apply the mapping function of the image data, which is, for the tiles we used, the standard Mercator projection.

2. We determine the detail level of the imagery that has the appropriate resolution by calculating the magnification factors between the original data and the pixel on the screen. This can be done analytically by differentiating the mappings from the beginning to the end. For example, for the complex exponential function, the derivative is again the exponential function. This derivative is complex, and contains the magnification factor in the magnitude of the resulting number.
3. We determine the tile for the sampling of the pixel by using the coordinates from Step 1, which describe the location of our pixel in the imagery, and the detail level from Step 2. If the tile is not loaded, we use the lower resolution tile the index structure points to.
4. We sample the pixel by using the normalized coordinates from Step 1 for the image tile from Step 3.

To prevent aliasing artifacts, we implemented trilinear interpolation by determining the two detail levels that are directly above and below the optimal resolution in Step 2.

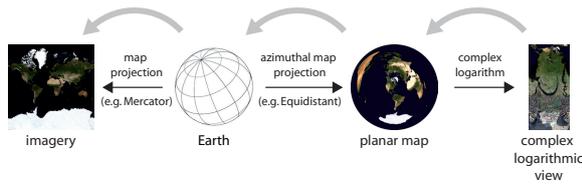


Figure 8: Our complex logarithmic views are a concatenation of an azimuthal map projection with a complex logarithm. To sample the pixels of the complex logarithmic view, we have to invert the complex logarithm and the azimuthal projection, and apply the projection of the imagery.

5. Results and Interaction

As a result, in [Figure 9](#) we present a complex logarithmic view of the Metropolitan Museum in New York. It extremely enlarges the Museum while the context of the whole Earth is preserved. Objects that are in the same direction from the center of interest are mapped to vertical lines, resembling the panoramic perspective. Objects with the same distance to the center are mapped to horizontal lines. The scale of the image varies exponentially from bottom to top. Despite the extremely different scales, small pieces of the world are left nearly undistorted, keeping their familiar shape, ranging from nearby houses, over regional objects like rivers, to coastal features, and even continents.

Due to the facts, that our representations present a view of the whole world, and modern graphics hardware allows for rendering the images at high resolutions in realtime, we can use the representations to intuitively explore and navigate on the Earth's surface. By dragging a point in the image with the mouse to the bottom, we can fluently move to any object. Since input devices only possess finite accuracy, it requires more than one mouse click, of course. But by adjusting the target while moving it closer to the bottom, we are able to correct inaccuracies during the interaction. Zooming and rotating in the original world coordinate system is changed to a one-dimensional translation in the complex logarithmic views: zooming translates the view vertically, while rotating translates it horizontally.

In [Figure 10](#) we show a series of stills from an interaction operation. Starting from a point in the Mediterranean next to Crete, a user navigates towards the Metropolitan Museum in New York. Since New York is located in North America, the user pulls this continent, which is clearly visible in the upper part of the first frame, downwards. As a result, the virtual camera follows a great circle towards the chosen target. In the second frame, after pulling North America only a few pixels closer, the camera is already close to the French Atlantic coast. In the fourth frame, the user gets close to North America and is able to locate the characteristic shape of Long Island. Flying over Long Island, Manhattan becomes visible in the seventh frame. After locating Central Park in the ninth frame, the user reaches the goal in the tenth frame.



Figure 9: Complex logarithmic view of the Metropolitan Museum in New York, in the context of the whole Earth. Up to the middle, the image is covered by the American continent, and the upper part shows the other continents.

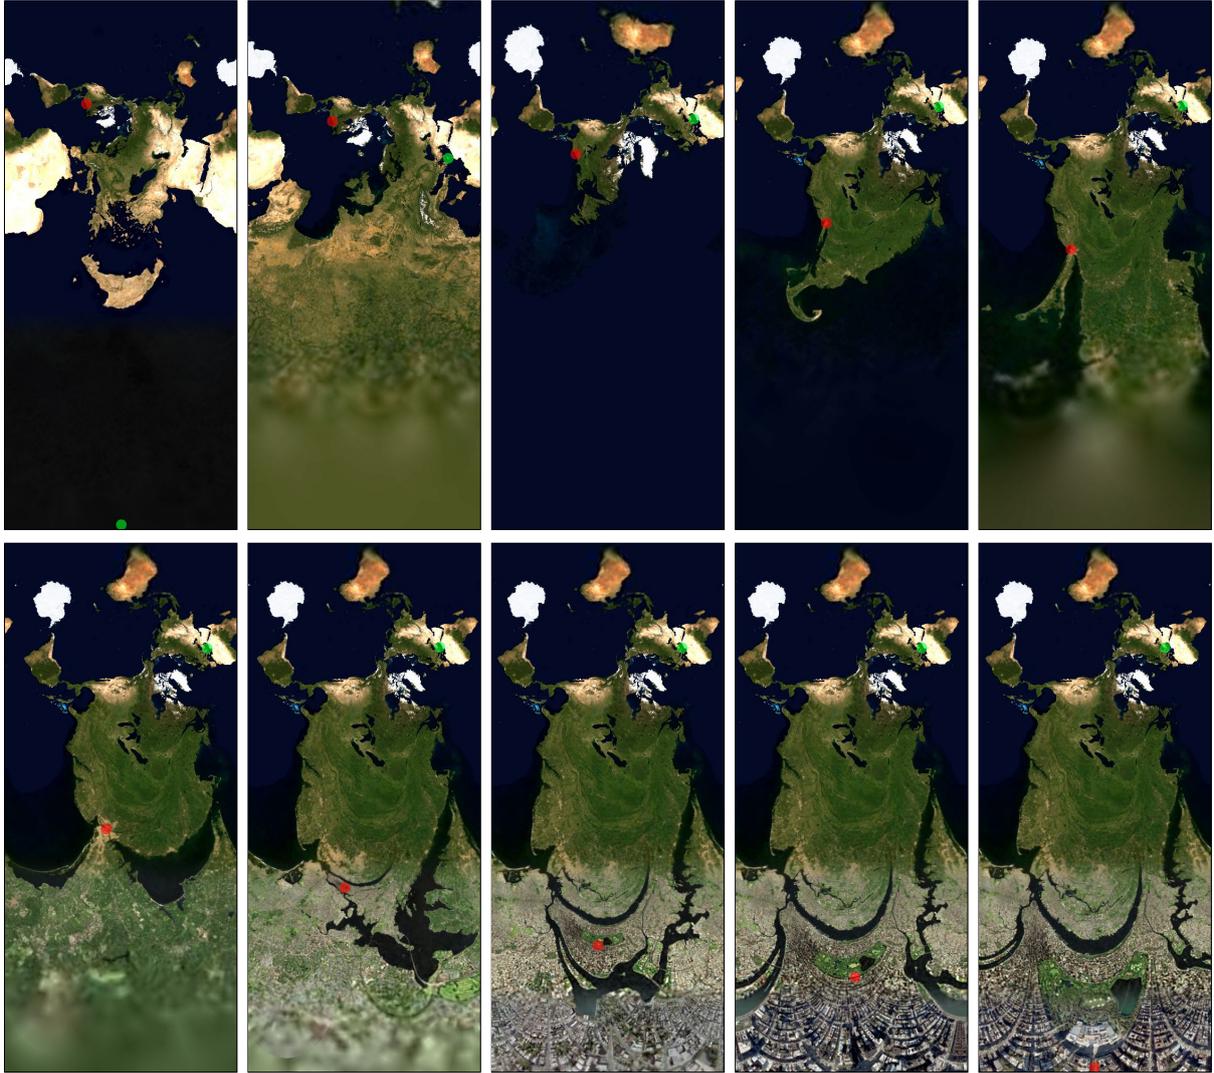


Figure 10: An example of an interaction sequence, moving from Crete to New York, by dragging it downwards. The starting point is marked green, the end point red. A detailed description is given in [Section 5](#).

6. Conclusion

We presented a detail-in-context method for very complex satellite and aerial imagery employing complex logarithmic views. We described an extended clipmapping approach for the realtime rendering of such representations. This enables the use of our method for the interactive exploration of the Earth's surface from its smallest details to the whole planet.

During informal experiments with the interactive visualization, we made a couple of interesting experiences: After the users have familiarized themselves with our representations, within a couple of minutes most of them were able to understand and navigate within the complex logarithmic views. One limitation of our approach is that it does not utilize conventional map orientation, with North towards the

top of the map. Since many users seem to strongly depend on these directional relations, some of them tended to get lost in our visualizations. Others were able to interact effectively, once they understood the similarities to a perspective projection. The fast movement in the bottom part of the visualization during interaction was seldom reported as a problem, since the users' attention was focused on the mouse pointer.

Successful navigation within complex logarithmic views depends on the familiarity with the data. The data needs to contain small, recognizable details throughout all scales, for users to be able to locate their intended targets. It is easy to find coastal cities, since the shapes of the neighboring seashore are typical, or single houses in a city, since we can follow roads and other prominent features. In contrast,

finding small cities in the middle of a country mostly covered with forest is hard, since relative directions from other features are different from the directions on ordinary maps. However, annotating the imagery with secondary map information, like labels and symbols, alleviates this problem. The distortion of linear features, like roads and rivers, seems to pose no problem, as long as the aerial data contains enough local shape information. Then, users tend to follow such features in order to reach targets as in the real world. Therefore, our approach seems to be promising for car navigation systems. The mapping makes it also easy to estimate the relations between distances to objects in different directions: The object higher in the visualization from a viewpoint is farther away in the real world.

Since every order of magnitude is mapped to a stripe with equal width, the aspect ratio of the resulting depictions is fixed by the ratio between the largest and the smallest objects depicted on the top and the bottom of the display. Our examples of single houses in the context of the whole world are more than twice as high as wide. This is a limitation on today's ordinary computer displays. However, even on these displays, the aspect ratio would be perfect for an accompanying navigational bar next to an ordinary square map. High resolution displays are very beneficial for the success of our method. They allow to represent small details with a resolution that still allows their recognition, thereby clearly showing the advantage of our method.

In future work, we plan to extend our approach to three-dimensional geographical data; mountain ranges and buildings could support the perceived similarities to ordinary perspective views. We also want to employ complex logarithmic views for car navigation systems, and to explore the possibilities of using conformal mappings for route visualizations.

Acknowledgements: This work was supported by the DFG under grant GK-1042 "Explorative Analysis and Visualization of Large Information Spaces". The authors thank Microsoft for the permission to use the Virtual Earth™ data and the support, and the reviewers for the helpful comments.

References

- [BBD06] BÖTTGER J., BALZER M., DEUSSEN O.: Complex logarithmic views for small details in large contexts. *IEEE Transactions on Visualization and Computer Graphics: IEEE Visualization Conference and IEEE Symposium on Information Visualization Proceedings 12*, 5 (September/October 2006), 845–852.
- [Car99] CARPENDALE M. S. T.: *A Framework for Elastic Presentation Space*. Ph.D. dissertation, Simon Fraser University, Burnaby, BC, Canada, 1999.
- [GJS*05] GOTT III J. R., JURIC M., SCHLEGEL D., HOYLE F., VOGLEY M., TEGMARK M., BAHCALL N., BRINKMANN J.: A map of the universe. *The Astrophysical Journal* 624 (May 2005), 463–484.
- [Kea98] KEAHEY T. A.: The generalized detail-in-context problem. In *Proceedings of the IEEE Symposium on Information Visualization* (Research Triangle Park, NC, USA, October 1998), IEEE Computer Society, pp. 44–51.
- [KR96] KEAHEY T. A., ROBERTSON E. L.: Techniques for non-linear magnification transformations. In *Proceedings of the IEEE Symposium on Information Visualization* (San Francisco, CA, USA, October 1996), IEEE Computer Society, pp. 38–45.
- [KR97] KEAHEY T. A., ROBERTSON E. L.: Nonlinear magnification fields. In *Proceedings of the IEEE Symposium on Information Visualization* (Phoenix, AZ, USA, October 1997), IEEE Computer Society, pp. 51–58.
- [Mor03] MORGAN M.: *The Space Between Our Ears*. Weidenfeld & Nicolson, London, UK, 2003.
- [MRC91] MACKINLAY J. D., ROBERTSON G. G., CARD S. K.: The perspective wall: Detail and context smoothly integrated. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA, April 1991), ACM Press, pp. 173–179.
- [Nee97] NEEDHAM T.: *Visual Complex Analysis*. Oxford University Press, 1997.
- [PL74] POLYA G., LATTA G.: *Complex Variables*. John Wiley & Sons, 1974.
- [RC94] RAO R., CARD S. K.: The table lens: merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA, April 1994), ACM Press, pp. 318–322.
- [RM93] ROBERTSON G. G., MACKINLAY J. D.: The document lens. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (Atlanta, GA, USA, November 1993), ACM Press, pp. 101–108.
- [SA99] SPENCE R., APPERLEY M.: *Data base navigation: an office environment for the professional*. Morgan Kaufmann, San Francisco, CA, USA, 1999, pp. 333–340.
- [Sny87] SNYDER J. P.: *Map Projections: A Working Manual*, vol. 1395 of *U.S. Geological Survey Professional Paper*. United States Government Printing Office, Washington, D.C., USA, 1987.
- [STH*07] SEOANE A., TAIBO J., HERNANDEZ L., LÓPEZ R., JASPE A.: Hardware-independent clipmapping. In *Proceedings of the International Conference on Computer Graphics, Visualization and Computer Vision* (Plzen, Czech Republic, January 2007), pp. 177–183.
- [TMJ98] TANNER C. C., MIGDAL C. J., JONES M. T.: The clipmap: a virtual mipmap. In *Proceedings of the annual conference on Computer graphics and interactive techniques* (Orlando, FL, USA, July 1998), ACM Press, pp. 151–158.