Complex-Logarithmic Views and Map Warping

Dissertation

zur

Erlangung des akademischen Grades

des

Doktors der Naturwissenschaften (Dr. rer. nat.)

an der

Universität Konstanz, Fachbereich Informatik und Informationswissenschaft

vorgelegt von

Joachim Böttger

Tag der mündlichen Prüfung: 24. November 2009

Referent: Prof. Dr. Oliver Deussen

Referent: Prof. Dr. Ulrik Brandes

Advisor
Prof. Dr. Oliver Deussen, University of Konstanz, Germany

Reviewers
Prof. Dr. Oliver Deussen, University of Konstanz, Germany
Prof. Dr. Ulrik Brandes, University of Konstanz, Germany

Date of Submission
May 28th, 2009

Date of Defense
November 24th, 2009

Zusammenfassung

Stetig zunehmend stehen immer komplexere Datenmengen über unsere Welt selbst, aber auch abstrakte Daten, wie beispielsweise hierarchische Filesysteme, zur Verfügung. Wenn man versucht, in solchen Datenbeständen kleine Details gemeinsam mit ihren massiven Kontexten zu verstehen, stellt der Computer ein großartiges interaktives Werkzeug dar.

Die sogenannten verzerrungs-orientierten Detail-im-Kontext Techniken nutzen vielfältige Verzerrungsfunktionen, um kleine Details im Vergleich zu ihrer Umgebung zu vergrößern, und beide gleichzeitig darzustellen. Die dabei üblicherweise verwendeten Abbildungsfunktionen, wie Fischaugenverzerrungen oder Zentralperspektive, komprimieren jedoch Teile der abgebildeten Information auf anisotrope Weise. Dies bedeutet, dass deren Formen lokal in verschiedenen Richtungen ungleichmäßig skaliert, also gestaucht oder lang gezogen werden. Hierdurch ist die enthaltene Information bei großen Unterschieden in den dargestellten Maßstäben nach einer Abbildung schwer oder gar nicht wiederzuerkennen.

Diese Arbeit untersucht zunächst die Anwendung von Verzerrungsfunktionen aus der komplexen Analysis, die frei von solcher anisotroper Kompression sind, für die Detail-im-Kontext-Problematik. Insbesondere der komplexe Logarithmus wird wegen seinem Potential für die Darstellung von extremen Größenunterschieden über viele Größenordnungen hinweg zur Entwicklung der sogenannten Komplex-Logarithmischen Views herangezogen.

Für diese Interaktionsmethode wird zunächst die Frage beantwortet, wie animierte Übergänge zwischen der komplex-logarithmischen und einer normalen euklidischen Darstellungsweise dazu beitragen können, eine Intuition für die Abbildungsfunktion zu unterstützen. Hierfür nutzt die Methode mathematische Verbindungen vom komplexen Logarithmus zu den erwähnten Verzerrungsfunktionen wie Fischaugenverzerrungen oder herkömmlicher Zentralperspektive.

Ein weiterer Beitrag besteht in der Entwicklung von Methoden zum beschleunigten Berechnen von komplex-logarithmisch verzerrten Darstellungen unter Ausnutzung von moderner Graphikhardware. Für vektorisierte geometrische Daten werden dabei die sogenannten Vertex-Shader genutzt, während für gepixelte Luftbild-Daten ein dem sogenannten Clipmapping ähnlicher Ansatz implementiert wurde, der auf der Nutzung von Fragment-Shadern beruht.

Die entwickelte Methode wurde Teil der künstlerischen Installation "Globorama", und in diesem Rahmen einer breiten Öffentlichkeit vorgestellt. Des weiteren konnten die erworbenen Kenntnisse über zweidimensionale Abbildungen zur Entwicklung von einer anderen Interaktionsmethode, dem sogenannten Map Warping, beitragen. Diese unterstützt Nutzer des öffentlichen Nahverkehrs in urbanen Umgebungen durch eine Verbindung passend verzerrter geographischer Landkarten mit schematischen Darstellungen der Netzverbindungen des Nahverkehrs.

Acknowledgements

I am grateful to everybody who contributed to the work described in this thesis.

The faculty of the research training group were a great resource of knowledge and wisdom, I am especially indebted to Ulrik Brandes, Oliver Deussen, Raouf Hamzoui, Harald Reiterer, and Dietmar Saupe for their support.

My co-authors, Michael Balzer, Werner König, Martin Preiser, Nikolaus Völzow, and Hendrik Ziezold: Thank you for the opportunity to work with you. A special thanks to Bernd Lintermann for "Globorama", which I regard as one of the most rewarding projects I have had the opportunity to be a part of.

My fellow group members and office mates, Monssef Alsweis, Carsten Colditz, Daniel Heck, Stefan Hiller, Johannes Kopf, Thomas Luft, Boris Neubert, Sören Pirk, Thomas Schlömer, Hendrik Strobelt, and Andreas Urra: It has been a pleasure and honor, thank you for your help, camaraderie and friendship.

The colleagues from the research training group, Shakeel Ahmad, Vladimir Bondarenko, Thorsten Büring, Martin Mader, Florian and Svetlana Mansmann, Martin Röder, Evghenia Stegantova, and all the others: Thank you for an inspiring research environment.

Niklas Röber, Marco Bubke, Henry Sonnet, Robert Hauschild: Thank you for your comments on the manuscript, for your support and friendship.

Alexander Abbushi: Thank you for providing a motivation to finish this dissertation.

I also thank the DFG for the support of parts of this work through the Graduiertenkolleg 1042 "Explorative Analysis and Visualization of Large Information Spaces", and my family for their support and encouragement.

Contents

Chapter 1

Introduction

We are living in a very complex world. The data we are interacting with in these days of the computer age is often very detailed. It is also often hierarchically organized, with an increasing number of hierarchical steps and orders of magnitudes available, and therefore contains very small interesting parts in very large overall contexts.

The rapid increase in computing power, while being one of the reasons for the explosion in complex data we have to deal with, also offers a great opportunity to aid us in the great undertaking of understanding our complex world. In connection with modern interface technology, especially with graphical displays, it becomes the most exciting extension of our brains to date.

This work delves into the subdiscipline of distortion oriented detail-in-context techniques, the art and science of helping an observer to cope with the richness of data by magnifying important parts while shrinking the overall overview. Common distortion techniques from this field like fisheye mappings and perspective approaches have somehow failed to be able to deal with extreme differences in scale of Euclidean information. The reason for that seems to lie in the introduction of anisotropic compression in all common mapping functions, which squishes information together until it is not recognizable anymore.

This thesis presents the results of the exploration of mapping and interaction approaches, driven amongst others by the mathematical field of complex analysis, by brain research, and the analysis of artistic examples. The remainder of this introduction contains the goals of this work, a short overview over the main results, and the organization of the thesis.

1.1 Goal

At the onset of this work stood a growing fascination with complex analysis and its application to mapping problems. Functions with complex variables can be interpreted as mappings for the geometric distortion of two-dimensional information, and as such possess certain interesting properties, the most interesting of which is the absence of anisotropic compression in the resulting depictions. The first basic question therefore was if such mappings were applicable for an improved interactive depiction of details and their context.

Towards this end, the first obvious choice of a mapping was found to be the complex logarithm, which has the potential to yield images which show single houses in the context of the whole planet we are living on. To make such depictions interactively applicable, a major question was how to employ modern graphics hardware in order to speed up the mapping process sufficiently to yield high enough frame rates. Another question was how it was possible to build an intuition for the resulting depictions, and how to interact with them. Last but not least, adequate application areas had to be identified.

During the work on this topic, the opportunity presented itself to apply the gained understanding of two-dimensional mappings towards a different exciting problem, namely schematic maps of complex public transportation systems and their embedding into the geographic space. The question here was how to bring these two navigational spaces together interactively in order to aid passengers to make informed decisions.

1.2 Main results

The main corpora of this research is the development of complex-logarithmic views for the interaction with complex data.

- After a first application of the method for diverse abstract vector data, such as hierarchical software systems, voronoi treemaps and geographical map information, the main application area is pixelated aerial imagery.

- For both kinds of data, suiting transitions between the ordinary euclidean world and complex-logarithmic layouts were found, linking this mapping to the fisheye views on the one hand, and the ordinary central perspective on the other, and thereby helping viewers to understand the properties of the mapping.

- For vector data, geometry shaders were employed in order to speed up the distortion process sufficiently for interactive framerates.

- For aerial imagery, fragment shaders in connection with an extended clipmapping method served the same goal. This method was developed within the scope of a supervised Master thesis [53] in our Computer Graphics and Media Design group.

- Interaction operations for an ordinary desktop setup were developed in order to navigate and explore the aforementioned highly complex data.

- The visualization of aerial imagery was incorporated in the interactive artistic and research installation “Globorama” in a highly successful cooperation with the Center for Art and Media in Karlsruhe and the Human-Computer Interaction group in Konstanz. The installation was presented publicly on various occasions, and made it possible for visitors to experience our world through a complex-logarithmic perspective.

The aforementioned mapping problem dealing with public transportation data was treated in a fruitful cooperation with the Research Project “Visual Navigation” and the Algorithmics Group in Konstanz, yielding an appropriate warping method as well as a novel interaction method.

1.3 Publications

Parts of this thesis were published in the following publications:

[10] Böttger, J., Brandes, U., Deussen, O., Ziezold, H., Map warping for the annotation of metro maps, IEEE Computer Graphics and Applications, Vol. 28, No. 5, pp. 56–65, 2008.

[9] Böttger, J., Balzer, M., Deussen, O., Complex logarithmic views for small details in large contexts, IEEE Transactions on Visualization and Computer Graphics, Vol. 12, No. 5, pp. 845–852, 2006.

[38] König, W., Böttger, J., Völzow, N., Reiterer, H., Laserpointer interaction between art and science, Proceedings of the 13th International Conference on Intelligent User Interfaces, Canary Islands, Spain, pp. 423–424, 2008.

[12] Böttger, J., Preiser, M., Balzer, M., Deussen, O., Detail-in-context visualization for satellite imagery, Computer Graphics Forum, Vol. 27, No. 2, pp. 587–596, 2008.

[11] Böttger, J., Brandes, U., Deussen, O., Ziezold, H., Map Warping for the Annotation of Metro Maps, Proceedings IEEE Pacific Visualization Symp., pp. 199–206, 2008.

1.4 Organization

The organization of this thesis is as follows:

In the next Chapter 2, an overview over the different motivations for this work is given. A major force driving this research is the diversity and complexity of available data in our time, ranging from geographical knowledge about the real world to abstract corpora. The methods developed in this work hereby are strongly influenced by what scientists outside of the field of computer science, artists and brain researchers do.

Chapter 3 reiterates relevant existing methods from the computer science field of human-computer interaction. While many techniques exist dealing with the relation between a focus on details and an overview, the focus here lies on the different geometrical distortion oriented detail-in-context techniques.

Chapter 4 analyzes the properties of the basic mathematical mappings which underly almost all of these techniques. The majority of distortion oriented detail-in-context techniques can be grouped together by the distortions they produce, namely either central perspective or fisheye-like distortions, both of which introduce anisotropic compression in the resulting depictions. In the second part of the chapter, we discuss conformal mappings taken from the mathematical field of complex analysis, and especially the complex logarithm, a mapping for extreme differences in scale free of anisotropic compression.

The complex logarithms first application for the interactive visualization and navigation of complex abstract two-dimensional vector data, and the implementation issues associated with them is the subject matter of Chapter 5.

Chapter 6 takes a look at the application of complex logarithmic views for the visualization of and interaction with aerial data, which is readily available over many orders of magnitude nowadays.

Chapter 7 describes the art and research collaboration “Globorama”, which employs a complex logarithmic perspective to enable viewers to browse such aerial imagery and panoramic photographs.

Chapter 8 deals with mapping methods for the navigation of complex urban public transportation systems embedded in their environment, and grew from the previously acquired knowledge. A technique to warp geographical information and their application resulting in an interactive method is explained.

This thesis ends with a conclusion and outlook to future work in Chapter 9.

Chapter 2

Motivation

The driving force behind the need for the interaction techniques for the visualization of details and their context in this work is the abundance of very complex data that we have to deal with in our world today. A prime example of this is computerized data about our world, from the microscopic dimensions of elementary particles up to the distribution of clusters of galaxies. Especially exciting is the recent abundance of digitized geographical data, where nowadays maps and aerial imagery are readily available for our whole planet, containing details like single houses and streets. This data, but also abstract data about, for example, the contents of a harddrive or complex software systems often are hierarchically organized or organizable.

Many techniques for the computerized visualization and interaction with such complex data exist, some of these techniques and their problems are the subject of the next chapter. After introducing typical examples for complex data, this chapter serves to gain inspiration by looking at how visual artists and scientists outside of the human-computer interaction community deal with visualizing such richness of details, and how the human visual system itself copes with the task of being able to discern details while still keeping an overview.

2.1 Complex Data

Since the advent of the digital computer, the exponential growth of the complexity of this machinery [46] has resulted in the availability and treatability of more and more complex corpora of data, describing our real world as well as more abstract concepts. Although understanding these data is difficult, the computer in connection with graphical displays is a great tool to aid in their comprehension. The reason for that is, that the human visual system is the broadest connection between the brain and the outside world. The fields of Visualization and Information Visualization deal with the question of how to depict real world and abstract data. This section describes typical examples with very high complexity of their subject matter.

2.1.1 Real world

One way how humans look at our universe is by thinking in hierarchies; things consist of smaller parts, which in turn consist of other, even smaller parts. Human knowledge extends to a steadily growing range of magnitudes, which grows on two fronts: The largest known objects and the smallest ones. In ancient history humans without sophisticated tools had a very limited understanding of their close surroundings, anything beyond the unaided senses was subject of speculation.

Figure 2.1: Our world from the very smallest things to its entirety.

On the macroscopic frontier, it was only necessary to understand the world for as far as an unequipped human could look or walk. Improvements in travel technology then necessitated cartography in order to expand the recording and passing on of knowledge which described our planet, culminating in a complete description of its spherical form. Telescopes, space travel and other tools such as cosmic simulations made it possible to gain knowledge which surpassed our understanding of the universe beyond this boundary, resulting in our current view of our macroscopic world.

On the other, microscopic frontier, the first humans' understanding was limited by the resolution of their visual system. The invention of the microscope lead to advances in biological knowledge of cells and previously invisible organisms. Chemistry and physics use some very large tools to get to our current understanding of atoms and smaller elementary particles, and today scratch on even more fundamental building blocks of our universe.

Analogously to the growth of knowledge, the modern world has people interacting with more and more complex data as well. In our times of globalization, we need an understanding of large scale geography as well as the very small details in our surroundings. This is reflected in the way we look at our world mathematically. Starting from ancient measuring systems, we have today arrived at a system which reflects the fact that there are many steps of recursive containment of smaller things in larger things, using orders of magnitudes to measure our world. There are many orders of magnitude between the smallest things that we deal with in close proximity and our planet.

It used to be difficult to access data of the necessary richness of detail for the whole planet, not the least due to the size of such a data base. Recently, modern mapping technology and the advent of the internet have resulted in interactive online databases like Google Earth [22] or Microsoft Virtual Earth [45]. These contain seamless imagery of all orders of magnitude from the whole planet down to the scale of single houses, and transmit only a subset of data the user is interested in at a point in time.

The available data has the property that it contains interesting and recognizable details on every level of magnification. In case that such details are not evident in the imagery itself, like in the view of a mostly uniformly green area of forest, supplementing this imagery with labels like streets and place markers makes it easily browsable.

2.1.2 Abstract data

Modern computing technology resulted in an explosion of abstract data. A classical pre-computer era example for such data describing our world is the phylogenetic tree, which classifies the biological realm of living organisms. Many other data are

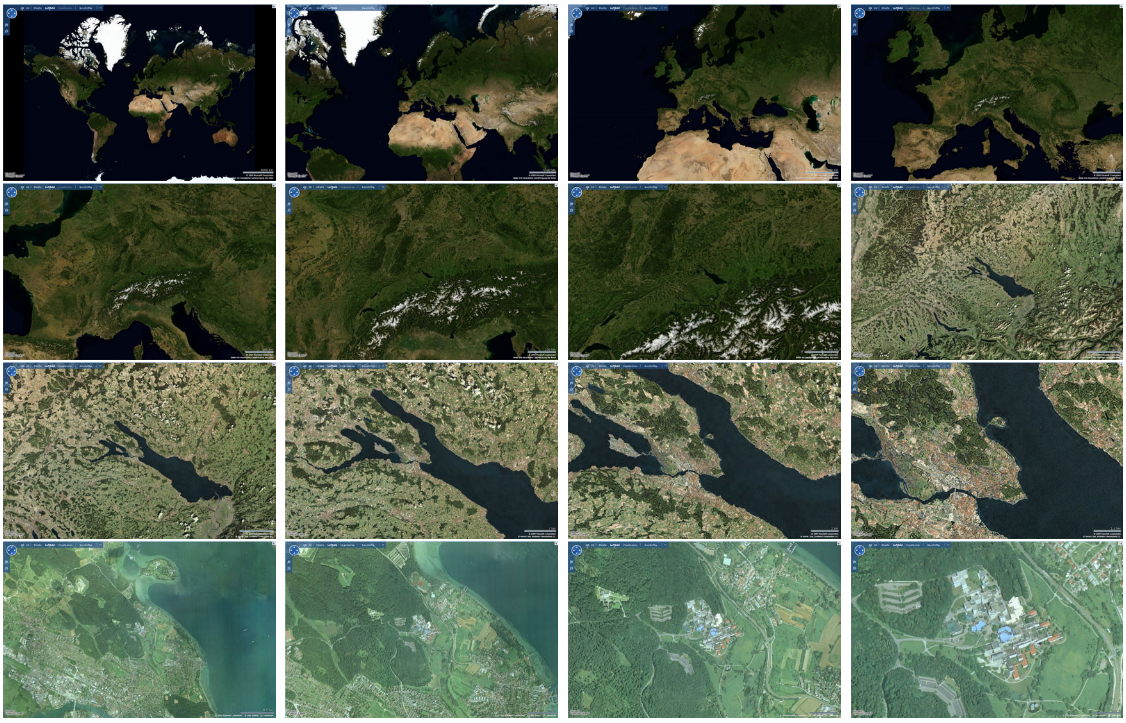

Figure 2.2: Software like Microsoft's Virtual Earth shows how complex our planet is. Starting from a view of the whole earth on the top left, every following view is magnified by a factor of two. In the last view on the bottom right, the University of Konstanz is visible (from [45]).

not even linked to any real counterpart, but result from mathematical concepts and the like.

Since our real world is organizable hierarchically, thinking in hierarchies comes natural to humans. No wonder, therefore many of the data organization schemes in computer science are hierarchical es well. Prime examples here are filesystems (see Figure 2.3 left), where folders contain other folders and files, or large object oriented software systems, which are organized in packages, subpackages, classes, subclasses, and finally single variables. If no hierarchy exists in a body of data, often one can be extracted using clustering methods. With growing amounts of data, these hierarchies are getting deeper. Like for our real world, it is often necessary to understand the relation between and the connection of the smallest and the largest entities.

For the visualization of such abstract data, often real world metaphors are used in order to find a suiting layout. For example, text corpora have been visualized as landscapes with hills and valleys, thereby evoking the association of land masses connected to each other [80] (see Figure 2.3 right). As another example, software systems have been visualized as cities and islands [4]. Such a layout then can provide desirable properties as in the real world data above, for example the right granularity of details on every level of magnitude.

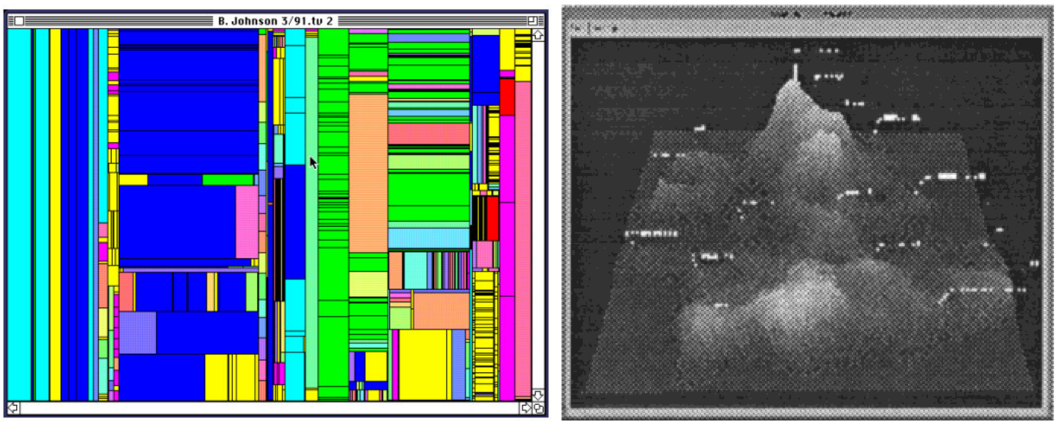

Figure 2.3: Treemap view of a complex file hierarchy (left, from [30]). Text corpora visualized as landscape (right, from [80]).

2.2 Visual Art and Science

To illustrate the relations between such very differently sized entities as described in the previous sections is a long standing concern for artists and scientists. Before the advent of computers, these illustrations were non-interactive. Many of the illustrations in this section served as the original inspiration for and echo the detail-in-context methods and interaction techniques in the next chapter. These techniques, however, suffer from several shortcomings when applied to the now available, very complex data described earlier. As we will see, to keep the connections between the different orders of magnitude intact while magnifying parts of the information is a challenge. Taking a look back here serves as a foundation to develop new improved methods.

2.2.1 Zooms

The first group of illustrations of different orders of magnitude are the zoom-like ones. These consist of a series of images of the same size, magnifying a certain point in space very differently. The first versions of this were published in the form of books, with large steps in scale, typically by a factor of ten, between the single images. With technology available, the number of images was increased, in order to yield an animation.

Kees Boeke [6] published in 1957 a book showing a series of images of a girl sitting in her garden from very different distances. On the very extremes, he showed the fringes of the then known universe, heaps of galaxies on the one end, and an atom on the other end. Leafing through the book, it is possible to mentally connect an image with its successor and predecessor, because they show the same content. Overall, the book leaves a great impression of how small or large things are compared to each other, and how complex the world we are living in is.

The Eames took the book as a guide for their famous book [48] and movie [17] “Powers of ten”, which is still often shown in school to introduce the topic of orders

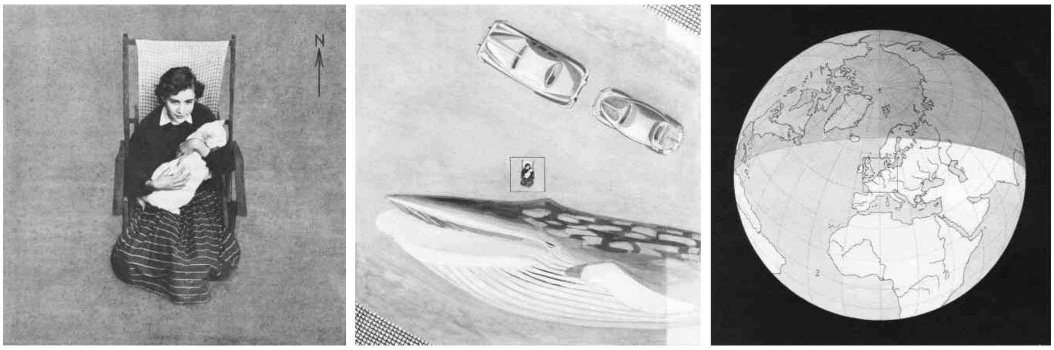

(a) (b) (c)

Figure 2.4: Boekes book shows a girl sitting in the Netherlands from very different distances (from [7]).

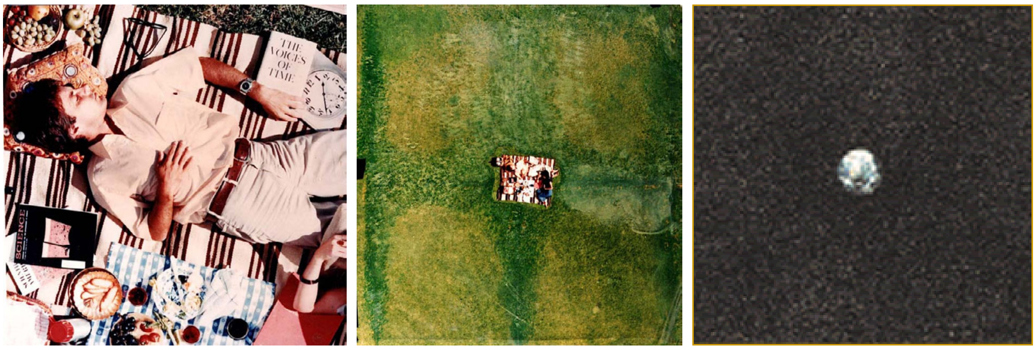

(a) (b) (c)

Figure 2.5: The Eames' video "Powers of Ten" shows an animation very similar to Boekes book (from [17]).

of magnitude in our measuring system. The animated version of Boekes vision of our world leaves an even more exhilarating impression of the complexity of our world.

With the advent of modern imaging techniques, NASA [66] constructed updated versions of parts of the Eames movie called "Great Zooms" with various well known points on our planet as starting point. The resulting animations use satellite imagery, and are especially effective if the depicted location is well known to a viewer.

2.2.2 Magnified Center of Interest

Series of images are obviously able to depict very different sizes consecutively, but a viewer has to mentally connect those images over time. Another influence on nonlinear magnification of centers of interest stems from the observation of natural phenomena which allow for the simultaneous viewing of small and large parts of our surroundings. Obviously, using magnifying glasses lead to a kind of integration of the magnified information with its surroundings, and served as archetype for later interaction techniques. But more interestingly, the observation of mirrored surfaces already lead in the Renaissance to manieristic depictions using nonlinear perspectives. The resulting nonlinear perspective from a mirrored sphere leads to fisheye views, which have interesting properties, and came to be easy to produce after the invention of photographic lenses. This view of our world lends itself nicely to portraying our world with a very enlarged center of interest, while keeping it seamlessly connected with its surroundings. It is also intuitively understandable, since the surroundings are depicted in a perspective fashion, resembling a horizon view (see Figure 2.6 left).

Another artistic attempt to depict a subjective world view stems from cartographic methods. It is easily possible to manually draw a map with very different scales for illustrative purposes. A well known example is Wallingfords “A New Yorker’s idea of the United States of America” (see Figure 2.6 right), which caricatures the view of America as seen by a New Yorker.

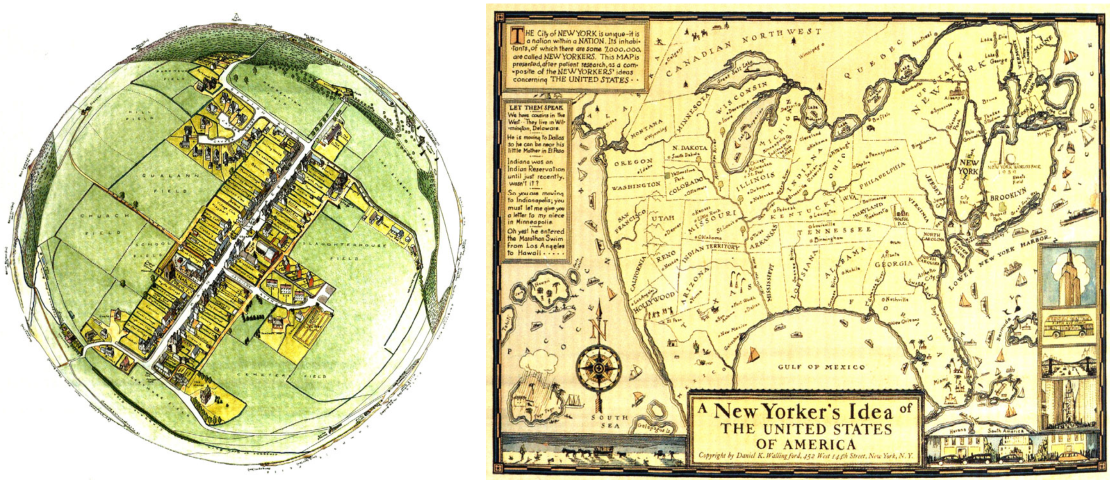

Figure 2.6: Fisheye-like view of a village with its surroundings (left). Wallingfords “A New Yorker’s idea of the United States of America” (right). Both images from [26].

2.2.3 Nontraditional Perspective

Several other artists and scientists chose a perspective-like approach in order to squeeze several orders of magnitude in a single image. It is interesting to note that most of these illustrations employ a view of the world which puts the smaller objects in the image on top of the larger ones, like they usually are e.g. in a traditional central perspective landscape. The reason for that seems to be, that we are used to this property from our natural surroundings. The placement of information is accordingly an important depth cue for our visual system [47].

Saul Steinberg uses a special form of perspective to show a New Yorker’s View of the World in his landmark illustration from 1976 [73] (see Figure 2.8). The illustration contains details like people and letterboxes as well as countries like China in one seamless view. Seemingly very similar to a central perspective, it is interesting to note, that the shown information does not align towards a single vanishing point, but rather to multiple such points, which are arranged in a vertical fashion. Steinberg used areas of uniform appearance, like the Hudson and the Pacific, for transitioning between these multiple vanishing points.

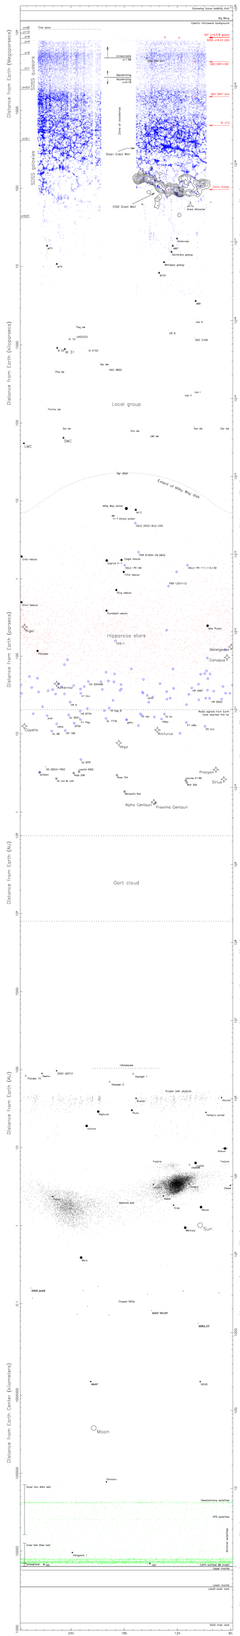

Figure 2.7: Logarithmic Map of the Universe (from [24]).

Figure 2.8: Steinbergs View of the World from 9th Avenue (from [73]).

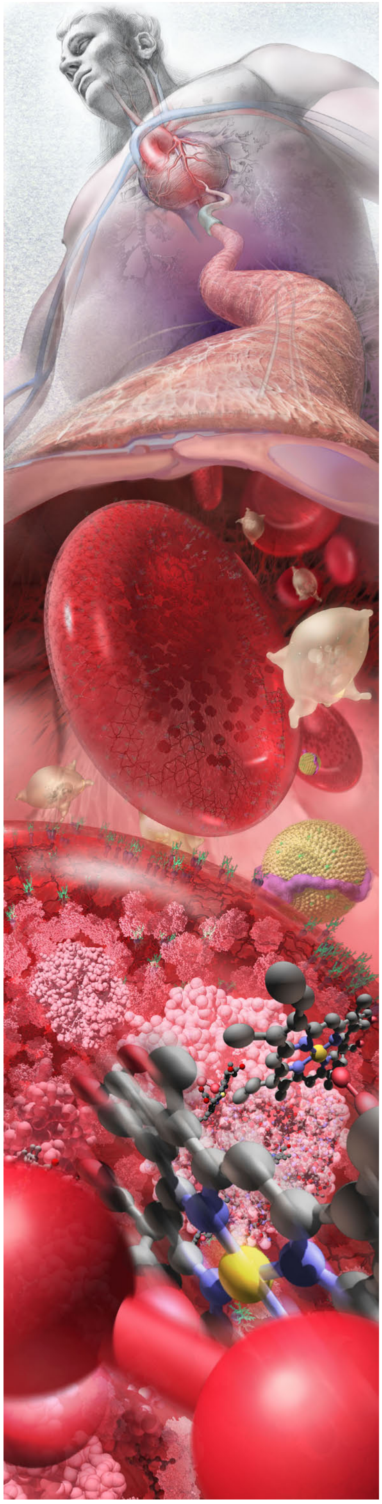

Figure 2.9: Zoom Into the Human Bloodstream (from [50]).

The Logarithmic Map of the Universe [24] (see Figure 2.7) shows our planet Earth in the context of the whole observable universe. Since there lie many different orders of magnitude in size between these two entities, the physicist Gott and his colleagues use the complex logarithm for the simultaneous depiction of everything. The resulting map is constrained to one viewpoint, and does not show any detail on the earths surface, but is otherwise very similar to snapshots of some of the visualizations in this work. It is noteworthy that the used mapping guarantees that, for example, planets are depicted as circular items, since the employed mapping is conformal.

Linda Nye and the Exploratorium Visualization Laboratory show the connection between a man and the atoms in his blood in their award-winning illustration "Zoom Into the Human Bloodstream" [50] (see Figure 2.9), connecting the different orders of magnitude in between with a nontraditional perspective. The illustration is one of a series of similar pictures. Other illustrations show atoms on the wing of a butterfly, and in a computer chip.

All these non-traditional perspectives are very high in comparison to their width; the larger the differences in scale, the more extreme this ratio gets.

2.3 Perception

Apart from how artists and scientists deal with the problem of squishing information with very different scales in one illustration, the human visual system also provides inspiration towards this subject matter. It is namely a well known fact, that our visual process employs a detail in context technique of its own in viewing our world.

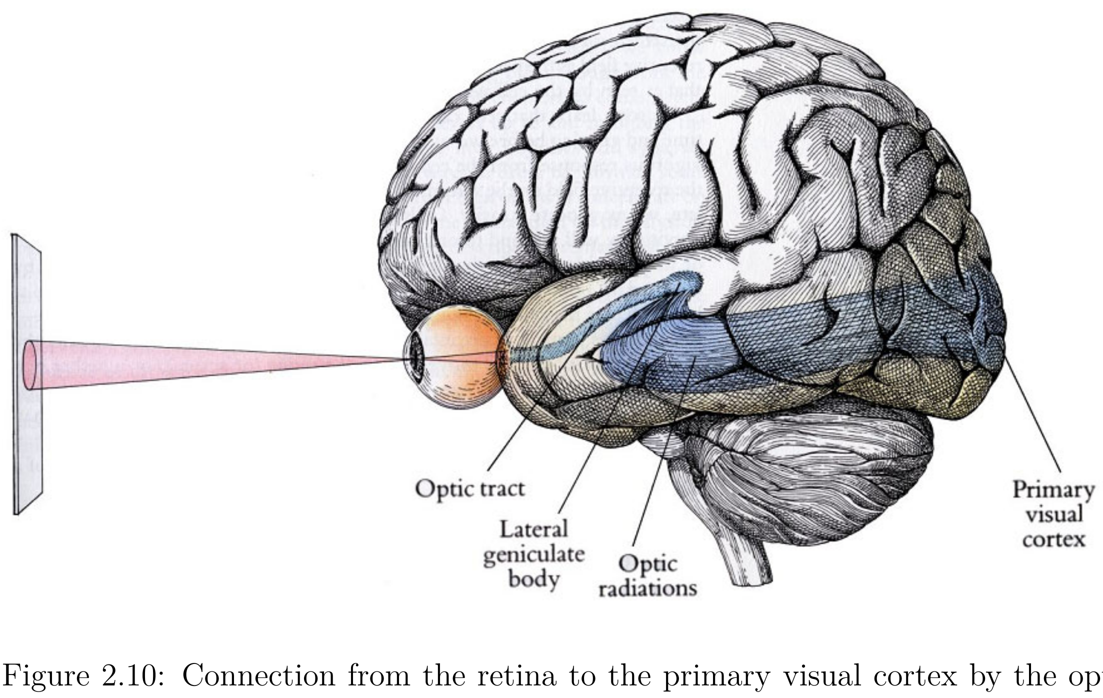

Figure 2.10: Connection from the retina to the primary visual cortex by the optic nerve, the laterate geniculate body, and the optic radiation (from [28]).

The parts that we are directly looking at are “magnified” compared with the visual periphery in our brain, which therefore commits more computing power for the things it is more concerned with, while still keeping an overview. This enables us to still notice the tiger jumping at us from out of a tree while reading our newspaper. The implementation of this technique is very interesting and relevant to the research in this work, and therefore recapitulated in this section.

2.3.1 Physiological Background

It has been long known, that we view our world by means of photosensitive receptors in our retina, which are connected to the brain by the optic nerve. Microscopic studies were able to show [28] that these receptors are not evenly spread, but much denser in the so-called fovea in the center of the retina, where the parts of our environment that we fixate are projected on. Modern physiology and neuroscience shows, that this difference in density is reflected in the primary visual cortex in the back of the brain.

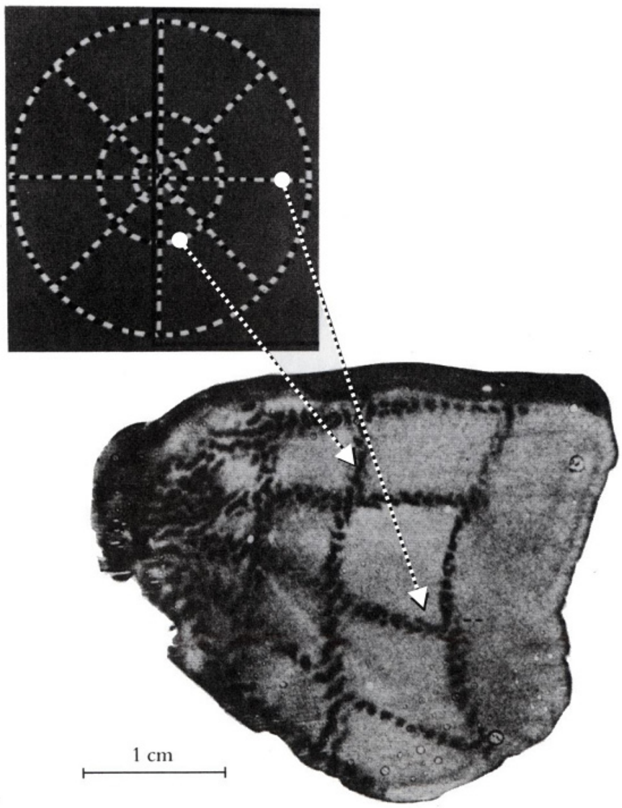

Figure 2.11: Stimulus and autoradiograph of activity patterns in the primary visual cortex of primates. Whith the fovea fixated on the middle of the stimulus pattern on the upper left, the left hemisphere of the primary visual cortex exhibits an according pattern of increased metabolism resulting in the darker areas in the lower part of the image (from [28]). The concentric circles as well as the lines from the middle outwards map towards two sets of almost parallel lines in the cortex.

The primary visual cortex is organized in a two-dimensional manner like the retina, which means that it is only a few millimeters thick, but has an area of many square centimeters, which are folded in in the back of the brain. The visual cortex is connected to the retina by the optic nerve, the lateral geniculate body and the optic radiation (see Figure 2.10). Interestingly, there exists a point to point correspondence between small areas on the retina and small areas on the primary visual cortex. Accordingly, injuries to parts of the primary visual cortex lead to a sharply defined loss of vision in the visual field. The mapping between retina and cortex is, with the small anatomically founded exception of cuts between the left and the right hemisphere, topology preserving (see Figure 2.11). This means, that neighborship relations are kept intact. However, the magnification of the mapping is far from constant. Like mentioned earlier, the magnification factor strongly depends on the distance of the part of the retina from the fovea. The fovea itself is magnified very strongly compared with the periphery, which is represented much smaller in relation (see Figure 2.12).

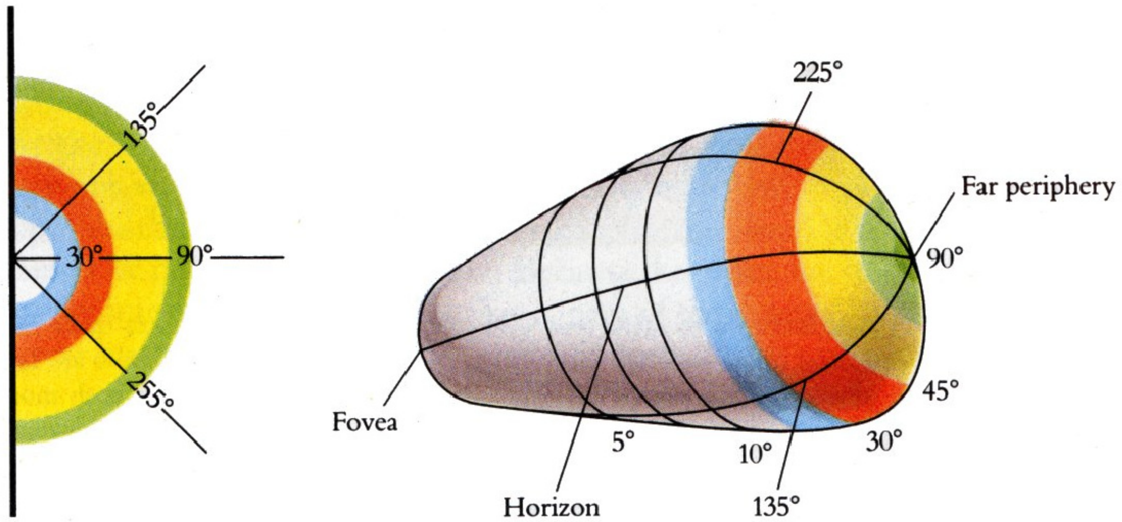

Figure 2.12: Another schematic depiction of how small the area to which the periphery is mapped is in comparison to the rest of the primary visual cortex (from [28]). The stimulus on the left is mapped to the primary visual cortex on the right. The latter is a flat structure, which can be unfolded to the shape shown here. The small bright area of the stimulus in the middle around the fovea maps to more than half of the area of the visual cortex.

2.3.2 Cortical Magnification

Daniel and Whitteridge [16] coined the term *cortical magnification* to describe this fact for their studies of the visual system of primates. The cortical magnification M is the ratio between the size of a small piece of primary visual cortex in millimeters and the corresponding angle of visual field that it represents. Daniel and Whitteridge conducted their studies using monkeys, whose visual system is very similarly organized to the humans'. Their findings showed that the cortical magnification is approximately isotropic, and therefore only depends on the distance from the fovea e as follows:

$$M(e) \sim \frac{1}{e} \tag{2.1}$$

Schwartz [63, 64] uses as approximation of the mapping a complex logarithmic mapping, which maps a complex number z to another complex number f:

$$f(z) = k \cdot \log(z + a) \tag{2.2}$$

Here, a and k are constants used to match the mapping to experimental findings. The approximation does not take into account the three-dimensional shape of the cortex. The derivative of the function is very similar to the experimentally established linear dependency of the magnification from the distance to the fovea. The mapping is also conformal, and therefore angle-preserving and free of anisotropic compression (see Figure 2.13).

This is very notable, because it hints at the fact that keeping details locally uncompressed seems to be important for the further processing of visual information in the brain. The human brain obviously works in a way which prioritizes this property even over the complete preservation of topology. This seems to be connected to a basic problem in magnifying flat information, which is reverberated throughout the following chapters.

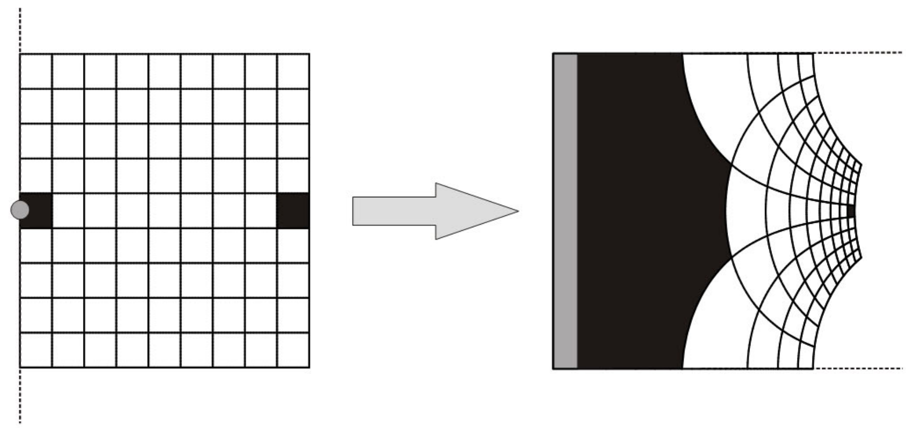

Figure 2.13: Approximation of the mapping from the retina (on the left) to the primary visual cortex (on the right) using the complex logarithm. The small black square on the retina next to the gray dot, which represents the fovea, is mapped to the large black area on the cortex, while the other black square in the periphery is mapped to an extremely small area. Note that the right angles at the crossings of the gridlines are not changed by the conformal mapping.

2.4 Conclusion

This chapter has shown the major driving forces behind this work: The abundance of diverse highly complex data, and artistic and scientific approaches to illustrate details of these data and their respective contexts. These attempts already show that there are basically two groups of mappings, the fisheye-like mappings and perspective, in order to show both in one seamless image.

Another great inspiration for this work are the findings from brain researchers, which show how the visual system deals with details in their context. As will be shown, their use of the complex logarithm and complex analysis in general to approximate the visual systems mapping from the retina to the primary visual cortex was a great influence for the methods developed in the following chapters.

Chapter 3

Detail-and-Context Approaches

This chapter deals with different approaches to cope with interaction with highly detailed data. Such data contains details which are too small to be easily readable once we try to fit the whole information space at once on a display. Examples for such data have been introduced in chapter 2.1, and will be found in later chapters, but include satellite data and abstract data sets like hierarchical graphs.

For such data, interaction paradigms like zooming and panning, and different approaches to show several disjoined views with different scale simultaneously exist. These are the subject of the first Section 3.1.

The focus of this work lies on techniques which keep the connections between details and their context intact. Therefore the manifold of distortion oriented detail in context techniques in Section 3.2 take up most of the space in this chapter.

3.1 Combinations of undistorted views

The need for these detail in context techniques follows from the fact that details have to be shown in a certain size in order to be visually recognizable. The whole information space in the same scale does usually not fit in the available space of the screen, thus it is necessary to use a smaller scale to show the whole context.

3.1.1 Zooming and Panning

The classical interaction paradigm for very complex data is zooming and panning. Here, the scale is changed over time, in order to subsequently observe details and to gain an overview. Both are not shown simultaneously.

Zooming and panning leaves shapes intact, but the context is lost when details are enlarged, and the details are too small to be recognized when the entire context is displayed.

Subsequently, when moving from one detail to another, it is necessary to zoom out, pan to another location close to the target detail, which is presumably not visible yet, and zoom in on the target. During this zooming in, the target is wandering towards the edges of the view, if it is not precisely in the center. This makes it necessary to recenter from time to time by iteratively panning and zooming.

3.1.2 Simultaneous View Combination

One widely used approach to show both, detail and context, simultaneously is to use separate windows for the detailed view as well as for the overview, thus breaking

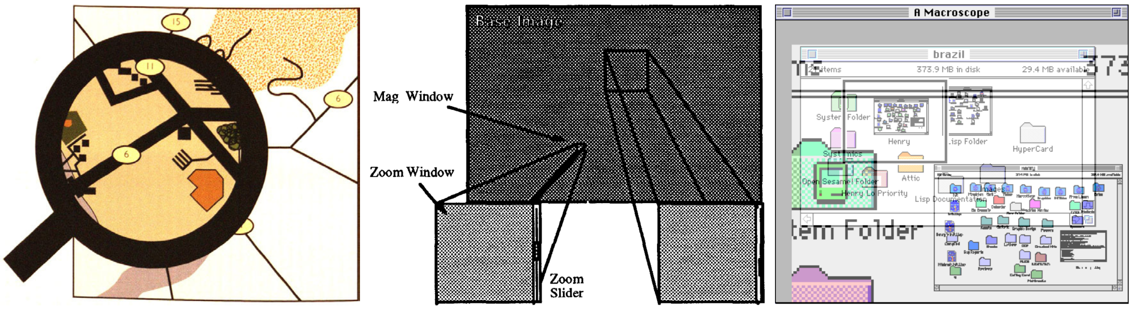

Figure 3.1: Combinations of detail and context views. Simple in-place magnification (from [71]). DragMag (from [79]). Macroscope (from [42]).

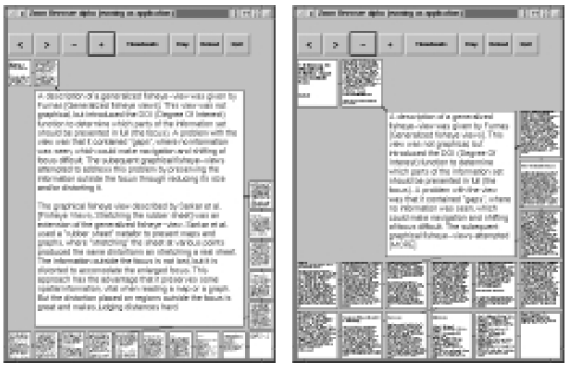

Figure 3.2: Flip Zooming (from [5]).

the connections between a detail and its context. However, the need to switch back and forth between these different views puts a mental strain on the user.

There have been several attempts to ease that burden by combining an overview and a detailed view. The basic problem, according to Spence [71], is that magnifying an area of the overview makes it impossible to simply fit that area into the context view in place. Therefore, occlusion on the fringes occurs if the magnified view is displayed in-place. Attempts to help a viewer to get over this discontinuity by interactive placement of the magnified area and graphical notation of the connections [79], or by superimposing the contextual and the magnified view with transparency [42], do not completely solve the problem of additional mental load to connect the two views convincingly (see Figure 3.1).

3.1.3 Flip Zooming

Other approaches, like Flip Zooming [5] (see Figure 3.2), use discontinuous mapping functions which leave details and their context uncompressed. However, these approaches are restricted to domains, like text display, where natural discontinuities between entities make it possible to introduce cuts in the mappings without destroying the underlying information. In the case of text display, such natural discontinuities occur between the single pages of a longer text.

The decoupled nature of such data makes it possible to optimize other criteria for display, like keeping the ordering or the relative placement of entities to a central

Figure 3.3: Perspective view in Microsoft's Virtual Earth [45]. Note that although the data is rendered using elevation data, the textures towards the horizon are anisotropically compressed.

entity intact. However, for complex data, people seem to use the relative placement of the entities and the resulting shapes for orientation. These more complex shapes tend to get destroyed by these methods, which makes finding known configurations and clusters of entities difficult.

3.2 Distortion Oriented Techniques

The distortion oriented detail in context techniques, on the other hand, employ a different approach. They are used in an effort to show both, detail and context, in one seamless visualization, by mapping them non-linearly from an original space to a display space.

The mappings used for this task are supposed to satisfy several, some mutually exclusive, properties. In particular, it is desirable to keep shapes, rotations, connections and angles constant, while changing the scale of parts of the underlying information independently. As we will see later in chapter 4, this is a mathematical impossibility, and therefore such a scale changing distortion always changes several properties of the depicted information.

There exists a variety of mappings used for the detail in context problem. The following subsections list several relevant approaches. Most of the approaches can be classified to belong to two classes of mapping functions: They either use some form of central perspective, or some form of fisheye-like mapping which leaves the connections around a magnified center of interest intact. All the common distortion oriented detail in context techniques for Euclidean spaces use mapping functions introducing anisotropic compression.

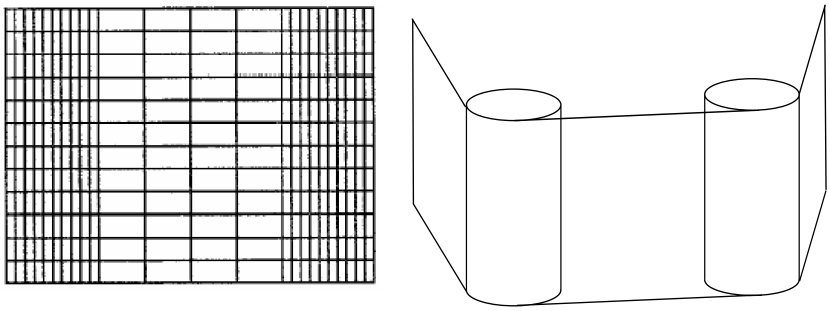

Figure 3.4: Bifocal View (left, from [72]). Original intuition behind the approach (right).

3.2.1 Central Perspective

One often used mapping in the field of satellite imagery is the ordinary central point perspective. This perspective, employed for two-dimensional data, is akin to looking at a plane from an angle.

The mapping leads to information in the foreground being displayed at a larger scale than the contextual information in the background. Although it is advantageous that the mapping is intuitively graspable, it only shows the context in one direction from the focus of interest. Everything to the sides and behind the virtual camera is cut of by the fringes of the display. Another drawback is, that only information along a single horizontal line is locally uncompressed, and that the anisotropic compression grows infinitely. Information that is very far away is mapped to a line, the horizon, as can be seen in Figure 3.3, and later more clearly in chapter 4.

3.2.2 Anisotropic Compression

Spence [72] was the first to show magnified details and demagnified context interactively on a computer in one seamless visualization. The original intuition behind his method was to imagine a long strip of paper, which contains diverse information like documents, appointments and other items.

This strip of paper is then folded over two imaginary vertical guiding posts (see Figure 3.4 right), and viewed with an orthographic perspective projection. This leads to information between the posts being displayed as usual, while information outside of this area of interest is viewed at an angle, and therefore compressed in order to fit the screen.

The first versions of the so-called Bifocal Displays (see Figure 3.4 left) accordingly divided the display in three different areas, one uncompressed for the detail view, and two areas on the left and on the right side that seamlessly connected with the detail view. The areas on the side of the display use the purest form of anisotropic compression allowing the context to fit on the screen by shrinking the underlying information in the horizontal, but using the original scale in the vertical direction.

Spence also extended the basic thought to truly two-dimensional information by applying the mapping method mentioned before in the horizontal and vertical direction simultaneously. However, using the method in two dimensions leads to

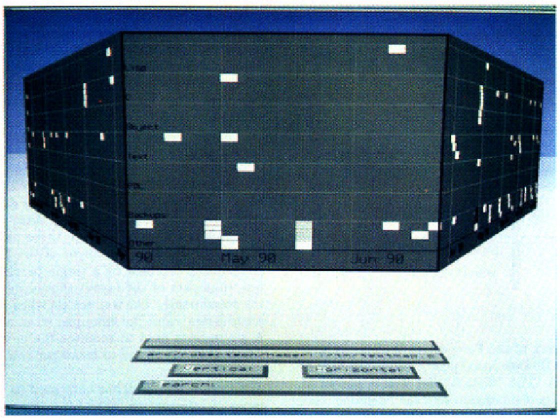

Figure 3.5: Perspective Wall (from [44]).

zones in the corners of the display which are compressed very differently compared to other details to the sides, above or below the focus. The method is therefore well suited for almost linear information, like the imagined paper strip or time dependent data, or for information layouts like tables, which are organized in a two-dimensional grid. This visualization of tables was developed further by Rao and Card [54].

It is interesting to note that Spence in 1980 already anticipated the Perspective Wall as well as some of the interactive fisheye transformations, but simply lacked the computing power to implement them on the hardware available to him at that time.

3.2.3 Perspective Wall

The Perspective Wall [44] (see Figure 3.5) uses a very similar approach to the original Bifocal Display. The method basically is a re-implementation of Spence's method. However, instead of using an easier to calculate orthographic perspective for the projection of the conceptual paper strip, the method uses true central point perspective. In the more then ten years between Spences original concept and the publication of the Perspective Wall, computing power increased enough to implement the perspective mapping, as well as other depth cues like light and shadows on an ordinary workstation.

The authors claim that this mapping alleviates the problem with compression, since the used perspective shows information placed at a distance to be smaller, in the same way as is well known from nature and is intuitively graspable for humans. Nevertheless, perspective mapping still introduces compression with growing distances. In the published examples therefore the information space usually is not very large, and the information strip is cut off before the detrimental effects of the perspective mapping are rendering the information unrecognizable.

It is possible to look at the difference between the original Bifocal Views and the Perspective Wall as a difference in distributing the anisotropic compression: While the former distributes this compression evenly over the whole context area, the latter leaves information closer to the area of interest less compressed than information items farther away.

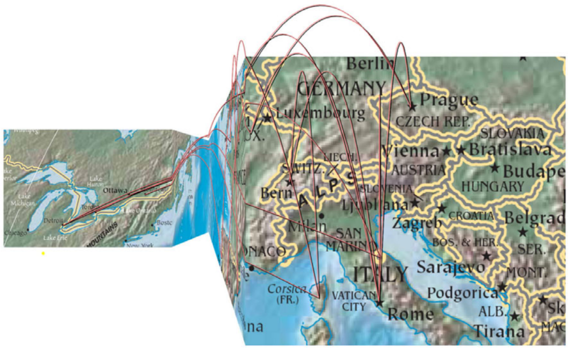

Figure 3.6: Space folding (from [18]).

Figure 3.7: Lorenz et al. (from [43])

Like in Spences original approach, showing information above or below the focus is not performed, and the method is only suited for almost linear information.

3.2.4 Space Folding

Elmqvist et al. presented another approach [18] to use perspective mappings for the interaction with complex data. Their approach is very similar to the perspective wall, but instead of perspective panels at the fringes of a focus of interest uses perspective panels to connect multiple foci (see Figure 3.6).

3.2.5 Horizonless Perspective

Lorenz et al. [43] use distorted perspectives to interact with three-dimensional geographical data. As can be seen in Figure 3.7, one of the problems improved by their approach is the compression of semi-flat information towards the horizon in a central perspective mapping. The authors counter this by tilting the geographic plane at a distance towards the camera, and connect the two zones by a transition area.

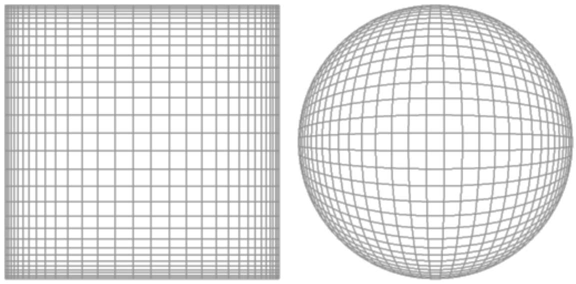

Figure 3.8: Typical Fisheye Distortions: Orthogonal (left) and radial (right) case (from [34]).

3.2.6 Fisheye Distortions

A different class of mapping functions are the so-called fisheye lenses. The cartographers Kadmon and Shlomi [31] developed mappings for the magnification of thematically interesting parts of geographic maps, which keep all the connections around a detail intact. Although implemented algorithmically, at that time, computers still were too slow for interactive adaption of the resulting depictions. Sarkar and Brown later developed interactive Graphical Fisheye Views [61] for the exploration of complex two-dimensional graph layouts. Keahey [32, 33, 34] refined fisheye mappings for the interactive display of pixelated imagery and more complex data. For a mathematical framework for this kind of transformations, see Leung and Apperleys review [41].

Fisheye mappings derive from the intuition, that for the magnification of an area of interest it is necessary to push the surrounding points outwards, away from this center, in order to create space for the enlarged information. Therefore, the fisheye mappings use a distance function that maps every point around a detail to a new point in the same relative direction. Then, using different distance metrics and functions, it is possible to generate different mappings, which resemble the effect an optical projection using a very wide angled lens would have, as well as rectangular distortions, which are more suited for example for text display.

Such fisheye mappings have to introduce compression, because they are constrained since the context around the enlarged detail has to remain closed. For radially symmetrical cases, the magnification factors along concentric circles are fixed because the circumference of the circles around the center is fixed by their radius. Figuratively speaking, the information lying on these circles literally has to stretch in order to fill the whole circle. Consequently, parts of the information cannot be scaled with the same factor in all directions, and will be anisotropically compressed (see Figure 3.8).

The differences in metrics and functions only influence how this compression is distributed. Using the Manhattan metric, for example, leads to the above mentioned rectangular distortions. Using the Euclidean distance exhibits circularly symmetrical mappings. The distance function can influence whether the compression is distributed equally around the focus of interest, or grows towards the fringes. It is also possible to leave parts outside of a certain radius uninfluenced by the map-

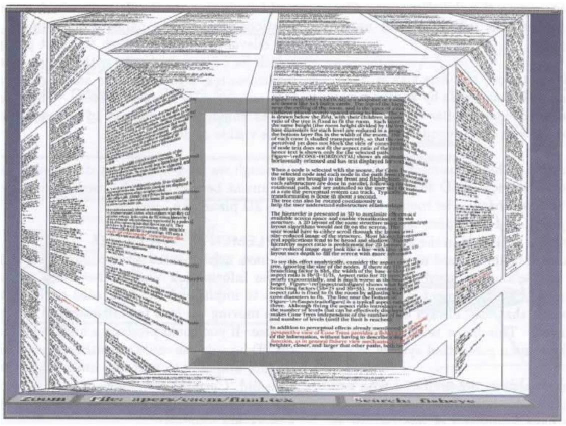

Figure 3.9: Document Lens (from [57]).

ping, and to use a distance function to smoothly integrate a magnified focus in that context.

3.2.7 Document Lens

The Document Lens [57] (see Figure 3.9) is another approach that uses perspective mapping functions: It uses a mapping of the information space to a truncated rectangular pyramid, and a subsequent perspective projection towards a viewing plane. The camera follows the focus of interest in order to avoid occlusion. The technique therefore always exhibits a rectangular magnified area of interest, and four areas connected around this uncompressed detail view showing the context. In this way, it uses the display space more efficiently than the Perspective Wall, and allows for showing the context above and below the detail area. However, the problems with the introduction of anisotropic compression in the process remain.

Beyond a certain magnification factor, the authors report that, like with the fisheye views, the information becomes unreadable, and the context display is essentially useless. This occurs when the area of interest occupies most of the viewing space.

The resulting mappings are not symmetrical. An advantage at the time of publication over some of the generalized fisheye views was the use of affine transformations for the expedient implementation of the technique.

3.2.8 Three-Dimensional Pliable Surfaces

Sheelagh Carpendale [13, 14, 15] presented a new approach employing modern graphics hardware for the presentation of details in their context. She uses the perspective transformations built into this hardware in order to produce fisheye-like transformations like Keahey. The basic intuition behind the method is to map the information on a three-dimensional sheet, and to move the areas of interest of this sheet closer to the perspective camera, in the process magnifying them. Since she stresses the continuous integration of magnified foci into their context, she chose to use gaussian functions for this three-dimensional distortion. Because moving grid points along

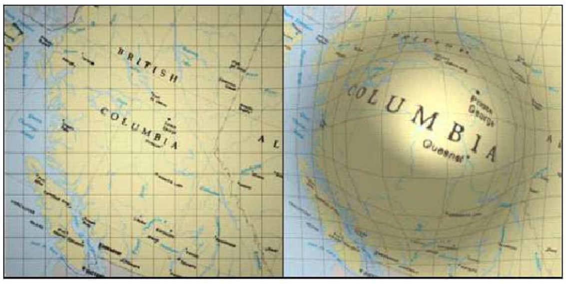

Figure 3.10: Three-Dimensional Pliable Surfaces (from [15]).

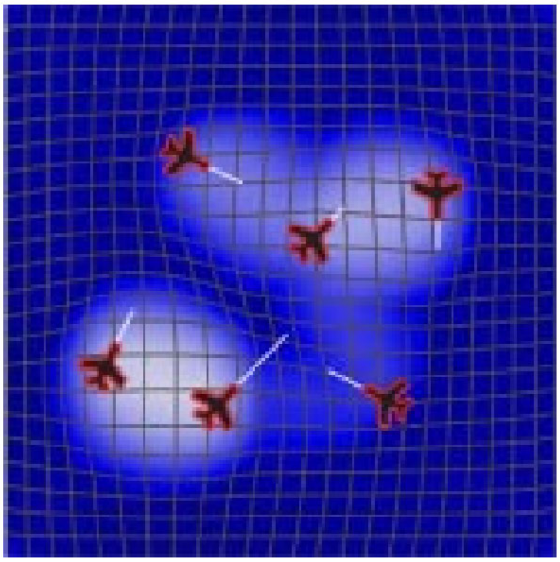

Figure 3.11: Non-linear magnification field (from [35]).

the axes towards the camera would not yield any discernible effect, the grid points are moved along the axes perpendicular to the sheet.

The method has the advantages that it is relatively easy to combine multiple foci in one distortion, as well as to add shading in order to clarify the magnification factors (see Figure 3.10). However, although the method uses perspective calculations to achieve the effect, they result in fisheye-shaped mappings. Also many properties have to be manually tuned so that no overlap occurs.

3.2.9 Non-linear Optimization

Keahey [35] also developed a technique to find a mapping function for arbitrary magnification fields by non-linear optimization. Given a magnification factor for every point in a plane, the algorithm computes the fitting mapping by minimizing deviations from this magnification. Towards this end, if the magnification deviates from the desired magnification at a certain point, neighboring points are pushed away or pulled towards that point. This allows for multiple magnification foci as well as arbitrary shapes to be magnified.

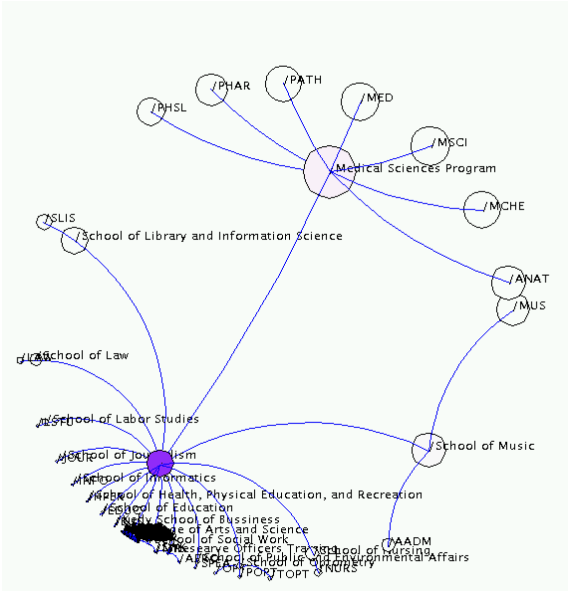

Figure 3.12: Hyperbolic tree view (from [39]).

The problem with this method lies in the used measure for magnification. Keahey uses the area magnification of the mapping, which makes the mapping not really well defined, since many combinations of magnification factors and anisotropic compression result in the same area magnification. The exclusive use of area magnification does not penalize compression of parts of the underlying data at all. The resulting mappings subsequently exhibit areas of anisotropic compression (see Figure 3.11).

3.2.10 Hyperbolic Tree View

Another relevant technique for the detail in context problem is the Hyperbolic Tree View [39, 40] (see Figure 3.12), which stresses the importance of keeping visual information free of anisotropic compression during the magnification of interesting details. This approach uses non-euclidean information as its basis, as it lays out tree data in hyperbolic space. It then maps the hyperbolic plane to a disk using a conformal mapping of hyperbolic space to the Euclidean plane, thereby not introducing anisotropic compression of the rendered shapes. This shifts the problem of presentation to the problem of finding a suitable layout for the information in hyperbolic space. It is then also not possible to show this layout on a screen in a way where all parts of the underlying information are depicted in the same scale.

The method is not directly applicable for Euclidean layouts of information, but is relevant for the problems in this work for other reasons. Using conformal mappings is founded in the desire to show configurations of entities in a consistent manner, which means to leave angles locally constant during change of the focus point. This makes it possible to find known clusters of entities easier during the interaction.

3.3 Conclusion

This chapter has shown the manifold of methods for the detail-in-context problem. These either disconnect details from their context geometrically, are only applicable for abstract data which is especially laid out in hyperbolic space, or produce mappings which are perspective- or fisheye-like in nature. However, when interacting with complex data like in section 2.1, which has extreme differences in scale, the latter two mapping approaches introduce anisotropic compression to an extent that renders the resulting depictions unsatisfactory. For the express intent to render an image which shows single houses seamlessly connected to the continents they are built upon, another approach is necessary.

In the next chapter, we will analyze central perspective and fisheye-like mappings, and introduce the complex logarithm, a conformal mapping from the field of complex analysis.

Chapter 4

Mathematical Background and Analysis

The basic work in this research stems from the analysis and application of mappings for the distortion of two-dimensional data. Such mappings map points in a plane to other points in a plane. This way, the visual information gets distorted in an arbitrary fashion.

Although some of the mappings in the previous chapter stem from three-dimensional intuition, like perspective views of planes, all these mappings depict flat information, and can therefore be described with two-dimensional mapping functions. Furthermore, the previous chapter has shown, that all the known distortion-based approaches for euclidean information can roughly be classified in two basic techniques, the perspective mappings, and the fisheye-like mappings.

Such two-dimensional mappings change different properties of the depicted information. These properties and their mathematical treatment are the subject of the first section of this chapter. In particular, it is shown that both these classes of mappings anisotropically compress parts of the depicted information. In the second section, mappings stemming from complex analysis, the conformal mappings and in particular the complex logarithm, are treated, since their application for the detail-in-context-problem is one of the main areas of research in this work.

4.1 Properties

In order to clarify the different properties associated with two-dimensional information, we use the example of a regular grid of equally sized circles as content of the distorted plane. Every circle contains a marker for one distinct direction in the plane prior to the mapping, for example, a small arrow pointing upwards (Figure 4.1).

These circles can be seen as an example of arbitrary local shapes in the plane. Following [49], every mapping can locally be described as a linear mapping, meaning that small circles are mapped to small ellipses. The transformation then can be locally described as a concatenation of a stretch in an arbitrary direction $d1$, another stretch in the direction $d2$ perpendicular to $d1$, and a translation and rotation (Figure 4.2).

The depicted shapes have several properties that can be distorted by a mapping: Primarily, their position is changed by the translation. As a result, other properties are changed as well, like magnification and local angles. If we look at the infinitesimal small elements from which that information is pieced together, the properties that are of special interest to us are size, orientation and the connections to their

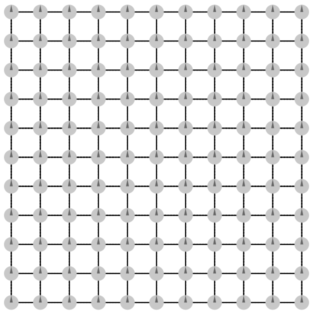

Figure 4.1: To illustrate two-dimensional mappings we use a grid of equally sized circles with markers for orientation.

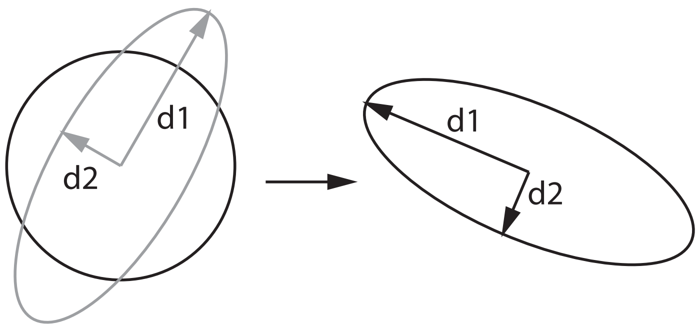

Figure 4.2: A locally linear mapping can be described as concatenation of a stretch in an arbitrary direction $d1$, a stretch in the direction $d2$ perpendicular to $d1$, and a translation and rotation.

neighbors preceding the mappings. The following sections contain a discussion and mathematical treatment of the different properties changed by two-dimensional mappings, the connections and trade-offs between, and ways to calculate them. This will help in the analysis of existing distortion oriented detail-in-context approaches, and enable the development of adequate distortions for data with extreme differences in scale between details and their context.

The commonly known mappings for detail-in-context techniques can be grouped in two classes of mappings, central perspective and fisheye-like functions. We thus use these two mappings in their purest form as examples for our analysis of the problem with anisotropic compression in common approaches.

4.1.1 Translation

Obviously, the primary property which is changed by mapping functions is the location of every single piece of data; a two-dimensional mapping is per definition a function which moves points in the plane to other points in the plane:

$$\mathbb{R}^2 : (x, y) \rightarrow \mathbb{R}^2 : (x_t, y_t) \tag{4.1}$$

All the other properties of the data are changed as a result of these translations.

Fisheyes

A typical example for one of the two basic classes of common mapping approaches, the fisheye functions, is a radially symmetrical fisheye with a root function as distance mapping. The mapping used by common fisheye techniques can be described as follows:

$$P_t = F + \frac{(P - F)}{|P - F|} a_t(|P - F|) \tag{4.2}$$

The distance function a_t changes the distance of P to the focus point F, but the relative direction α is kept the same, as can be seen in Figure 4.3. The class of root functions is widely used as distance functions.

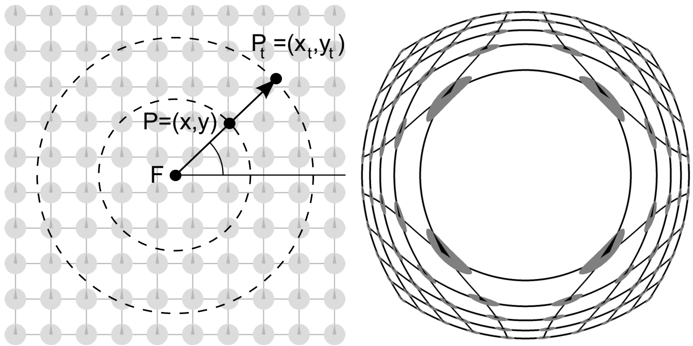

Figure 4.3: A fisheye mapping with the 5-th root function as distance function. Points are mapped outwards by that function (left). The resulting mapping magnifies the center of interest (right).

Using the nth root as an example distance function, the euclidean distance as metric, and the origin of the plane as focus point, this mapping simplifies to:

$$(x_t, y_t) = (x \cdot m, y \cdot m) \text{ with } m = \frac{\sqrt[n]{x^2 + y^2}}{x^2 + y^2} \tag{4.3}$$

The result of such a mapping can be seen in Figure 4.3.

Central Perspective

We illustrate the other class of common mapping functions, the central perspective mapping of flat information, with the projection of a horizontal plane onto a vertical viewport. An according setup and the resulting mapping can be seen in Figure 4.4.

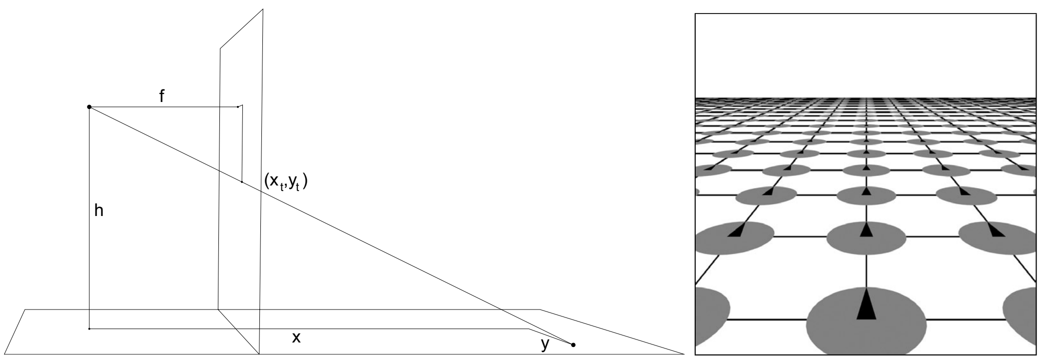

Figure 4.4: Central perspective mapping of planar information on a perpendicular viewing plane. A data plane with our example grid is mapped to a perpendicular viewing plane using a central perspective mapping (left). The resulting mapping (right).

Without loss of generality, this projection can be described as follows [20]:

$$(x_t, y_t) = (x \cdot \frac{f}{y}, h \cdot \frac{f}{y}) \tag{4.4}$$

Here, f is the distance of the view point from the view plane, and h is the height of the view point over the data plane. The resulting mapping can be seen in Figure 4.4.

Note that both these mappings change several other properties of the pieces of information different than their position implicitly as well, which is the topic of the following subsections.

4.1.2 Scaling

The scaling of the pieces of information is central for distortion-oriented detail-in-context techniques; after all, it is their goal to magnify the important parts, while shrinking the others, in order to use the available space efficiently.

Thinking about the local scale of distorted information is complicated by the fact that it is not uniquely defined; two-dimensional information can be scaled differently in different directions. Accordingly, there are different possibilities to characterize the magnification locally. To calculate a magnification in one variable, the area magnification can be used. The scaling factors in x and y direction are given by the partial derivatives of the mapping function.

The resizing of the different parts of the information is also the most critical property for their recognizability: Firstly, shapes have to have a certain size in order to be visible at all. If, for example, a shape is mapped to a space smaller than one pixel on a screen, it will be impossible to recognize it.

The problem with the so-called anisotropic compression is equally severe: If the stretching operations in different directions are not equally strong, this detrimental form of distortion is introduced in a mapping. This compression leads to shapes being crushed, which, for large differences in scaling factors, leads to them being mapped to almost linear structures that are not recognizable anymore. Another result of anisotropic compression is, that angles are locally distorted.

It is possible to calculate the largest and the smallest local magnification factors with a singular value decomposition of the Jacobian matrix, which can be approximated by using the partial derivatives. The local compression factor is the ratio between these singular values, or the condition number of the matrix. The area magnification can be calculated as the determinant of the local transformation.

Another important issue is the one of overlap: this problem occurs if the mapping function is not bijective, that is, if more than one point from the original data space are mapped to the same point. For continuous functions, this can happen if the determinant of the Jacobian matrix is negative somewhere, which means that the information at that point is "flipped".

Analyzing our two examples for detail-in-context techniques from the previous subsection, we can now make the following observations:

Fisheyes

Concerning anisotropic compression, for our example fisheye, the local scaling factors are smallest in the direction from the focus point outwards, and largest in the perpendicular directions. Mathematical analysis shows, that for root functions the anisotropic compression is constant, but proportional to the root factor. This can be shown as follows:

We look at a point in a distance d from the focus point before the mapping. After the mapping, the magnification factor at that point in the direction from the center outwards d_c is the derivative of the distance function. For the nth root function this is [8]:

$$d_c = (d^{\frac{1}{n}})' = \frac{1}{n} \cdot d^{\frac{1}{n}-1} \tag{4.5}$$

Accordingly, the magnification factor perpendicular to that d_p is

$$d_p = \frac{a(d)}{d} = \frac{d^{\frac{1}{n}}}{d} = d^{\frac{1}{n}-1} \tag{4.6}$$

since the whole circle on which the point used to lie is magnified by the remapping of their distances by that ratio. It is now easy to see, that the anisotropic compression factor c, the ratio between the two different singular values, is constant for root fisheyes:

$$c = \frac{d_p}{d_c} = n \tag{4.7}$$

This shows, that with higher roots the compression grows, and limits the applicability of such fisheye distortions for large differences in scale. However, the same is valid for other distance functions and metrics, since then the anisotropic compression is only reduced in some areas of the distorted space to the extent of other, additionally compressed areas.

Also, although our example fisheye is free of overlap, such overlap can easily occur if the distance function is not carefully designed to be continuously rising.

Central Perspective

For the central perspective mapping, the anisotropic compression factor is growing with increasing distance in the viewing direction of the information from the view point. As a matter of fact, the anisotropic compression approaches infinity there, which means that far away information is mapped to a line, namely the horizon. This can be shown by looking at the central line through the view plane shown in figure 4.4 on the right.

Along this line, we look at the magnification factors for small pieces of the depicted plane in vertical and horizontal direction (see Figure 4.5):

$$dx = d \cdot \frac{f}{z} \tag{4.8}$$

$$dy = d \cdot \frac{f \cdot h}{z^2} \tag{4.9}$$

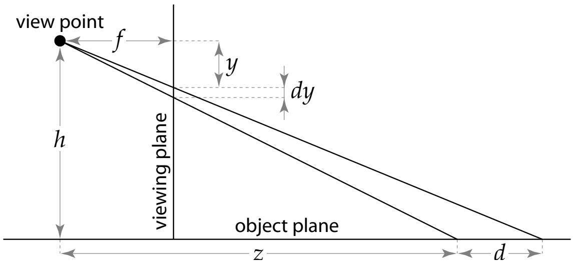

Figure 4.5: Cut along the central line of the perspective setup and according magnification factors for an infinitesimally small circle.

$$c = \frac{dy}{dx} = \frac{z}{h} = \frac{f}{y} \tag{4.10}$$

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 again is the compression factor, the ratio between these singular values. Larger values of c in this case 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 ordinary central perspective is therefore also limited in the extent of representable size differences.

4.1.3 Rotation

The rotation in the aforementioned local transformation changes the orientation of the pieces of information. Although this makes it more difficult for a user to recognize shapes, a rotation itself leaves intact what we perceive as the shape of objects.

Figure 4.6: Rotation leaves shapes intact, but might make them harder to read, especially if widely used conventions exist.

Although the shapes are intact, in special cases recognizing them might at first seem difficult in fields where usually strict conventions exist about the orientation of depicted information. One classic example is cartography, where usually maps are oriented in north-up orientation. This can be seen in Figure 4.6, where most people have difficulties recognizing Africa in south-up orientation. However, there are many ways to map our planet to a plane [70]. Some Australians, for example, promote the use of Australia-centric use south-up world maps.

Analyzing our two examples for detail-in-context techniques from the previous subsection, we can make the following observations:

The perspective leaves all the pieces of information looking almost in the same direction, and the small derivations from that are intuitively easy to interpret. This leaves us, for methods like the perspective wall or space folding, with the impression that the local orientation remains unchanged. However, perspective is often used in a way that makes it possible for the user to rotate the camera, for example when flying over the earths surface in Google Earth. The then radically changed orientations of shapes seem to pose no problem for the viewer if it is possible for him to put himself in an egocentric frame of reference. This works well for well known geographical data on the city-level, and is the display mode used in most navigation systems.

The fisheye-like mappings also change the local orientation only slightly.

4.1.4 Connectivity

Another property of mapping functions is how they deal with the connections between different parts of the mapped space: If a mapping maps two infinitesimally close points to different positions, a cut between these points is introduced. If this cut separates the details from their context, a user has to mentally reconnect them, which increases the mental strain. This work shows that cuts that still leave a connection between the currently interesting detail and its context are tolerable. Given the case that strong magnifications are required, they are preferable to common approaches, in which parts of the information are compressed to such an extent that they are no longer recognizable.

Looking at our two examples for detail-in-context mappings, the following becomes obvious: while the fisheye-mapping is leaving the connections all intact, the perspective has a singularity for points with $y = 0$. In practice, this means it can only show things in front of the camera, information under and behind it are not mappable at the same time. Additionally, the finite size of the viewport makes it necessary to cut off the mapped information at the sides, so that an opening angle of 180 degrees is not possible.

4.2 Complex Analysis

In the previous section, we have seen two archetypal example mappings which are widely used for detail-in-context distortion, but introduce anisotropic compression in that process. Conformal mappings, on the other hand, are those mappings that stretch 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.

Such mappings can be found in the field of complex analysis [25]. This mathematical discipline deals with complex valued functions, which map complex numbers to other complex numbers. They can therefore be interpreted as two-dimensional mappings. A basic insight from this analysis is, that all the functions which are analytically representable have special properties, which are useful for our application.

All these so-called analytical functions are complex differentiable, that is, they possess a derivative which is itself complex valued. This means, that the pieces of information in the corresponding mapping only get translated, rotated and isotropically magnified, leaving the shapes locally undistorted.

First and foremost, all the basic functions from ordinary analysis possess complex counterparts, like root functions, logarithms, exponentials and so forth. Apart from that, the analytic operations and combinations like addition, interpolation and so on yield a magnitude of conformal mappings.

The logarithm is a function well known for application to a single axis for plots that contain numbers with different orders of magnitude. In the following chapters, we use the complex logarithm as a two-dimensional mapping function for the purpose of showing details that are orders of magnitude smaller than their surroundings in their context. As we will later see, there exist strong mathematical connections between the complex logarithm and the fisheye and perspective mappings. Before elaborating on this in the following chapters, we first introduce the complex logarithm here.

As stated above, we want to use the complex logarithm as a two-dimensional mapping function. It maps a complex number to another complex number in the following way:

$$\log z = \ln |z| + i \arg(z) \quad \text{with} \quad -\pi < \arg(z) < \pi \tag{4.11}$$

It maps the logarithm of the magnitude of a complex number to the real value and the numbers angle to the imaginary value. Thus, points on circles around the origin (with equal distance from the origin) of the complex plane are mapped to parallel, vertical lines. Rays from the origin on outwards (with equal angle) are mapped to parallel, horizontal lines (see Figure 4.7). The origin itself is a singularity, mapped infinitely far away in the negative real direction, and magnified infinitely.

The magnification is indirectly proportional to the distance from the origin before the mapping. This means that objects close to the center of the plane are extremely enlarged. Like the one-dimensional logarithm, the complex logarithm maps distances that differ by a certain factor to equal distances: Thus, points that have distances from the center with the same order of magnitude are mapped into equally wide horizontal stripes. Consequently, enlarging the context by orders of magnitude only means enlarging its image by constant stripes.

Concerning the complex logarithms properties, we can make the following statements: since the mapping is conformal, no anisotropic compression is introduced, and there is no overlap. This makes recognition of shapes with very different scaling factors possible. However, the mapping has its drawbacks: First and foremost, the information is rotated in a, at first, counterintuitive way. The connectivity is also disrupted on an arbitrary line from the center of interest outwards. However, we

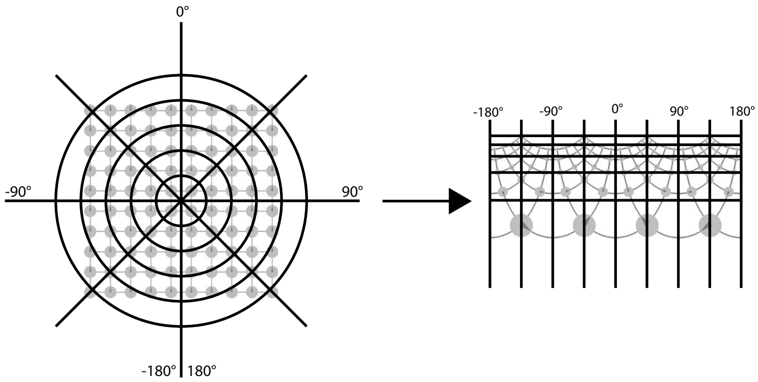

Figure 4.7: This complex logarithmic mapping maps rays from the center on outward to vertical lines, and concentric circles to horizontal lines. Note the extreme difference in distances between the horizontal lines on the right and the evenly spaced circles on the left.

will later see how building an intuition for the mapping is still possible due to the similarities to well known mappings.

4.3 Conclusion

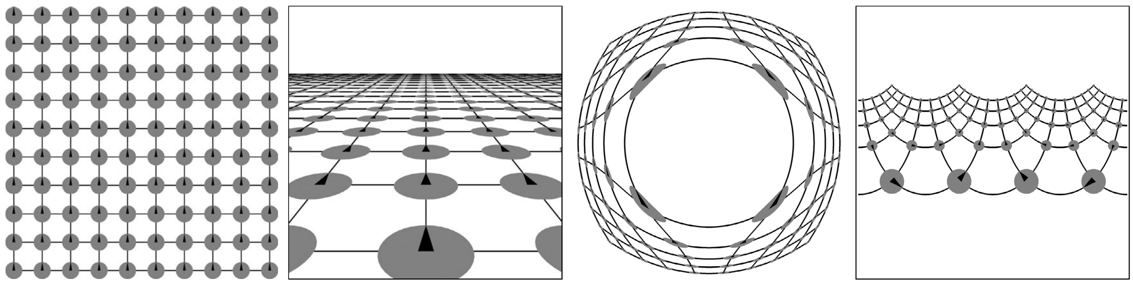

(a) (b) (c) (d)

Figure 4.8: Identical mapping (a) of a grid of small squares. Perspective (b) and fisheye (c) mapping both enlarge parts of the grid, but introduce compression, visible because circles are mapped to ellipses. The complex logarithmic mapping (d) enlarges parts of the grid without introducing compression.

All the distortion oriented detail-in-context techniques for the Euclidean plane in the previous chapter anisotropically compress parts of the information we want to show. Strong compression leads to shapes being unrecognizable after the transformation, and is inevitable with the described techniques if there are large differences in magnification factors between details and their context. The first contribution of this thesis is to map two-dimensional Euclidean data in a way that avoids compression while enlarging parts of the information in relation to other parts.

The perfect mapping that keeps all other properties intact while scaling parts of a plane differently cannot exist. Thus, choosing a mapping function for a certain problem always has to be a trade-off between the different kinds of distortion it introduces. One point made in this work is, that for the exploration of complex data over several orders of magnitude, keeping shapes uncompressed is more important than keeping rotations constant and even justifies giving up some of the connections between points in the plane in the process. Accordingly, we propose to cut the plane open along a line from the center of interest outward, which allows enlarging the center by extreme factors without introducing local compression in the mapping. The complex logarithmic mapping does exactly that.

The following chapters describe applications of that mapping for the interactive detail-in-context visualization of different data. Besides the implementation issues, a focus lies on the interaction with such a view, and the question of how to make users understand the mapping.

Chapter 5

Complex Logarithmic Views for Vector Data

This chapter describes the first attempt to employ interactive complex logarithmic views for the interaction with and navigation of abstract visual vector data. The first major application subject is a highly complex existing layout from a visualization of complex software systems.

The application demands the answering of several questions, in order to yield an interactive complex logarithmic view. Firstly, to yield useful interaction, a transition between the euclidean layout and the complex logarithmic detail-in-context layout is provided, and dragging points in order to change the center of the view is used. Secondly, on the implementation side, the rendering of the vector data employing modern graphics hardware is described. The chapter ends with application examples, and a description of our experiences with the method.

5.1 Interaction

The goal of this method is to interact with complex data, using complex logarithmic views for the magnification of a center of interest. To yield useful interaction, we first describe a transition between the euclidean layout and the complex logarithmic layout, which uses another class of complex analytic mappings, namely scaled and shifted complex root functions. We then describe moving through the data.

5.1.1 Transition

A smooth transition between complex-logarithmic and euclidean viewing modes is helpful in order to provide for seamless switching between magnified and non-magnified views, and building of an intuition for the former layout. This transition is derived from the intuition that a complex logarithmic view is similar to a fisheye view which is cut open and relaxed, and uses conformal mappings, namely root functions, which converge towards the logarithm.

The archetypal fisheye is described in 4.1. We have illustrated there how enlarging a center, but keeping everything around it connected, leads to anisotropic compression, since the distance function stretches concentric circles with constant radius more in the direction around a focus point then in the direction perpendicular to the circles. A solution is that we give up the constraint to keep everything connected in only one point on every one of those circles. We can then move everything along the circles, until the magnification factors in the two directions are equal. Then the resulting mapping is conformal, and does not compress the information or distort angles locally. For a given fisheye with the n-th root function as distance mapping, as shown in Figure 5.1(a), we can see how a regular grid of squares is distorted to circular shapes by anisotropic compression. If we move every point P_t along a circle by leaving it at the same distance from the center F, but dividing the angle α by a constant, we can observe how the stretched squares start looking like regular squares again. Once we have moved every point to another point on the same circle, but with a n-th of the angle α, we have reached conformality, and right angles remain the same after the mapping. The resulting mapping is the n-th complex root function, another function treated in the field of complex analysis.

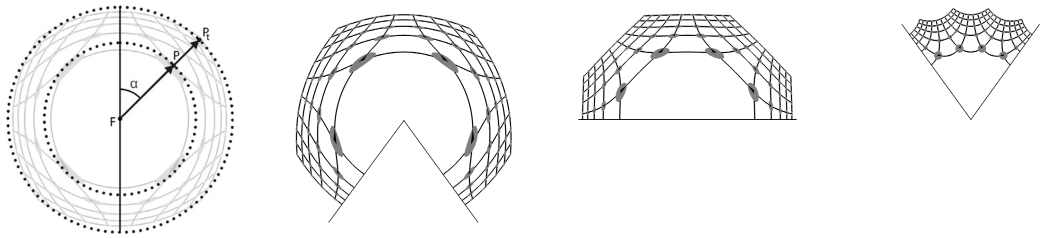

(a) (b) (c) (d)

Figure 5.1: Cutting a fisheye with the 5-th root function as distance function open by gradually dividing the angle by larger numbers. Once the angles are divided by 5, the resulting mapping is conformal, small circles are mapped to circles. The resulting mapping is a complex root function.

The complex root functions map a complex number z to another complex number in the following way:

$$\sqrt[n]{z} = \sqrt[n]{|z|} \cdot e^{i\frac{\arg(z)}{n}} \quad \text{with} \quad -\pi < arg(z) < \pi \tag{5.1}$$

Like in the fisheye described in 4.1, the n-th real root function is applied to the magnitude of a two-dimensional position. In addition, the angle is divided by the root factor n. The middle of the grid is enlarged stronger with growing n. The division of the angle leads to a cut in the grid, and to shapes that look like a slice of pie. The square root cuts the angles in half, and thus the grid only occupies half of the space. The 4th root leads to the grid occupying a quarter of the space, and so on. For larger n, the grid is mapped to a narrow wedge, and the central square takes up most of the space because it is magnified that strongly. However, the part of the mapping that shows its structure in Figure 5.1(d) looks very similar to the complex logarithm (see Figure 4.7). As a matter of fact, we will now see that shifted and scaled root functions converge towards the complex logarithm, and we can therefore use them for a smooth transition between the identical and the logarithmic mapping.

To yield the transition from the identical mapping to the complex logarithm using the root functions, we use the series expansion of the exponential function [1]:

$$e^z = \sum_{k=0}^{+\infty} \frac{z^k}{k!} \tag{5.2}$$

Transforming the base of the n-th root function

$$\sqrt[n]{z} = z^{\frac{1}{n}} = e^{\frac{1}{n}\log(z)} \tag{5.3}$$

and applying the series expansion in Equation 5.2, we can express the root as

$$\sqrt[n]{z} = 1 + \frac{1}{n}\log(z) + \frac{1}{n^2}\frac{(\log(z))^2}{2!} + \frac{1}{n^3}\frac{(\log(z))^3}{3!} + \ldots \tag{5.4}$$

If we subtract 1 from both sides, and multiply by n, this yields

$$n \cdot (\sqrt[n]{z} - 1) = \log(z) + \frac{1}{n}\frac{(\log(z))^2}{2!} + \frac{1}{n^2}\frac{(\log(z))^3}{3!} + \ldots \tag{5.5}$$

For $n \to \infty$ everything except the first summand converges to 0, and thus we yield this connection between the roots and the logarithm:

$$\lim_{n\to\infty} n \cdot (\sqrt[n]{z} - 1) = \log(z) \tag{5.6}$$

Adding a constant complex value to a complex function shifts everything in the resulting mapping by that value, and multiplication scales accordingly. We have therefore shown, that appropriately scaled and shifted root mappings converge to the complex logarithmic mapping. The resulting transition is depicted in Figure 5.2.

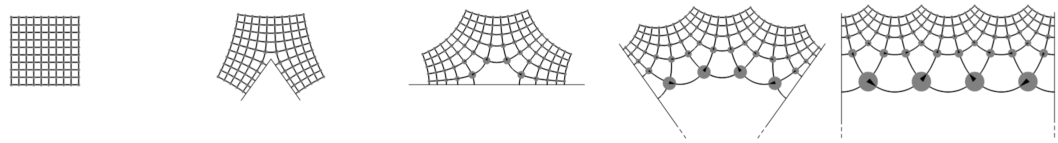

(a) $n = 1$ (b) $n = 1.25$ (c) $n = 2$ (d) $n = 5$ (e) $n = \infty$

Figure 5.2: Transition from the identical mapping (left) to the logarithmic mapping (right) using scaled and shifted complex root functions.

Using the described techniques, it is possible to conformally map two-dimensional information in a way that enlarges one point in it extremely strong, while shrinking the others in relation. In order to aid the understanding of the complex logarithmic mapping, and enable changing between logarithmic and identical mappings fluently, we employ the root functions as a transition between the two mappings: We let the user interactively manipulate a parameter $0 \leq a \leq 2\pi$, which describes the opening angle of the wedge on which the information is mapped. An angle of 2π means, that there is no cut in the mapping, and we use the identical mapping. Between 0 and 2π, we use the complex root functions. The parameter n in Equation 5.1 is calculated with

$$n = \frac{2\pi}{a} \tag{5.7}$$

to yield the desired opening angles. For small values of a, the wedge would grow very narrow, and only use the upper part of the screen, which, however, is against

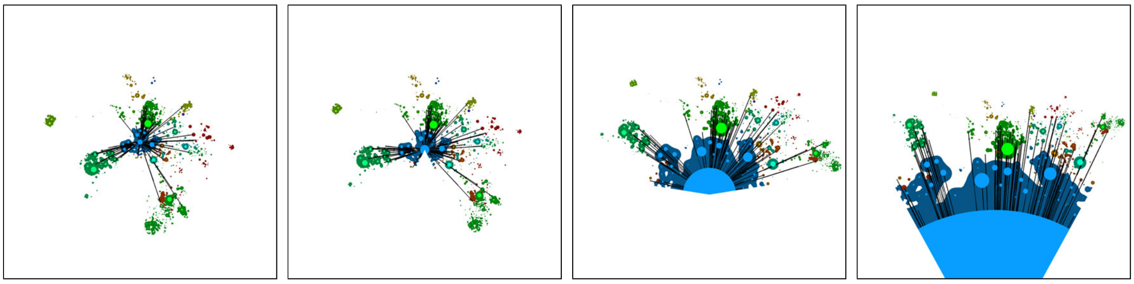

Figure 5.3: Unfolding of a complex graph layout with edges. Note how the blue node in the middle is enlarged and cut open.

our intentions. Thus, we scale and shift the wedge in the described way. The result is the transition shown in Figure 5.3. Once we are sufficiently close to $a = 0$, we switch to the complex logarithmic view, which is shown in Figure 5.9.

5.1.2 Navigation

For the navigation through the data we allow the user to drag a point of the transformed plane with the mouse to an arbitrary other point. By inverting the actual transformation, we calculate where the starting point and the end point of the dragging operation are located in the original, untransformed space. The resulting points yield a translation vector, which moves the point under the mouse pointer to the intended new position when applied to the original data. Figure 5.4 shows the according movement of the untransformed and the transformed space.

This mode of interaction is very intuitive, since it feels like grabbing and moving objects on the display, similar to Shneidermans Direct Manipulation [67]. Pulling them towards the origin of the transformation enlarges them, and pulling them away from the origin makes them smaller. After getting used to this mode of interaction, it is possible to bring very small details in focus with one dragging operation of the mouse without releasing the mouse button (see Figure 5.5).

The complex logarithmic view is only able to show a certain range of orders of magnitude, since the view size is constrained at the top and the bottom. In order to center the desired range of data, a zooming operation in the original data space is necessary. The zooming in the untransformed space translates to a simple and intuitive vertical shift in the transformed view, as can also be seen in Figure 5.4.

5.2 Implementation

Our implementation of the described technique takes geometry consisting of points, lines and triangles as input data. Arbitrary polygons can be rendered by subdividing them into triangles.

We can simply re-map point primitives by calculating their new positions. There are, however, two reasons for dealing with lines and triangles differently: Firstly, straight lines in the original data are mapped to curves in the transformed space. This requires to subdivide lines and triangles in order to avoid artifacts. The details,

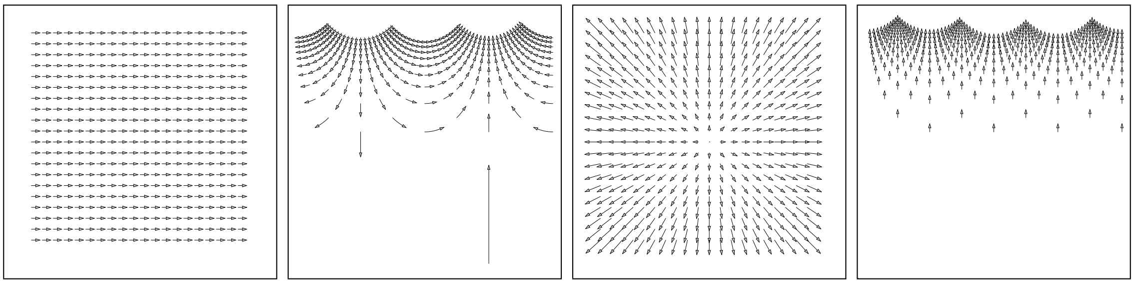

Figure 5.4: Navigation through the data using zooming and panning operations in the transformed complex-logarithmic view. Constant movement through the data (a) is transformed to the movement seen in (b): While data on the vertical corresponding to the direction of the movement moves downwards, information in the opposite direction moves upwards. Zooming in the euclidean view (c) accords to constant vertical movement in the complex logarithmic view (d).

which we view very closely, have to be divided into enough sub-primitives so that after the transformation the discretization of the curves is not too distracting. For our implementation, we subdivided the geometry empirically, until the artifacts disappeared. This leads to the use of many unnecessary primitives. However, it is possible to divide the primitives adaptively by taking their magnification factors into account, and rendering more triangles where the magnification factor is higher.

Secondly, the mappings we use introduce a cut in the mapped information. It is thus necessary to separate the vertices of line segments and subdivided triangle primitives for the rendering if they are intersected by the cut: The vertices of such primitives are mapped to opposite sides of the cut, and should no longer be connected in order to prevent artifacts. For line segments, the subdivision is straightforward. We calculate the intersection point with the cut, and replace the line segment with two line segments on each side of the cut.

For triangles, there are two different cases in which triangles can be intersected by the kind of cut resulting from the used mappings: In the first case, the start of the cut is located inside the triangle, making it necessary to cut open one of the sides of the triangle, and to replace the triangle with five smaller triangles. In the second case, the cut intersects two sides, which requires to disconnect the two halves of the triangle and to replace one half with two smaller triangles. The two different cases are shown in Figure 5.6.

The intersection tests can be performed very fast with hierarchical boundings for the primitives, since the cut only intersects a few of them at any given time. After the primitive division, we employ OpenGL [68] and vertex shaders [19] to transform the geometry. We reach interactive frame rates on common graphics hardware for data sets with hundreds of thousands of primitives.

5.3 Applications

In this section we present several example applications for complex logarithmic views, in order to show that our approach is versatilely usable. Due to the fact that

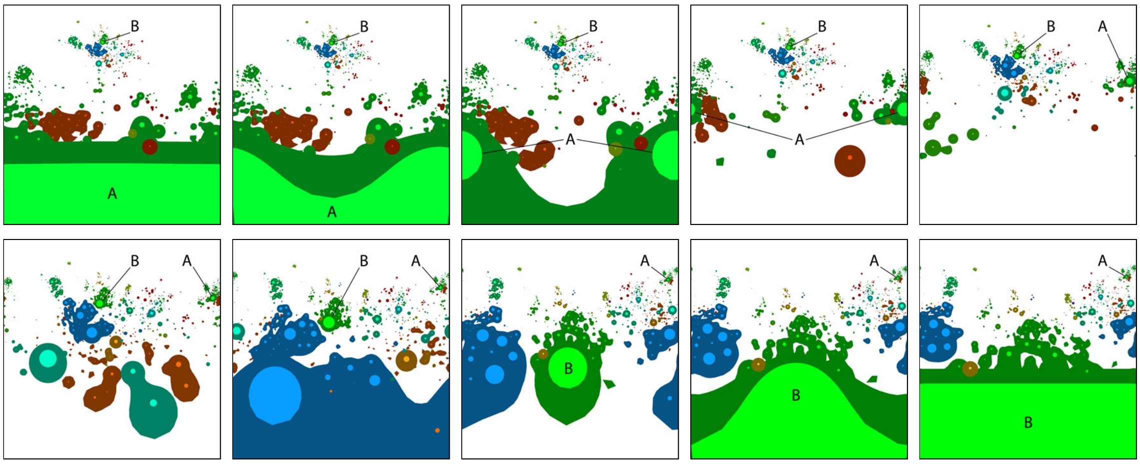

Figure 5.5: Traveling from point A to another point B: The user drags B downwards into the center of interest. At first, node A is intersected by the cut in the mapping, and divided in two parts. After the first three images, the center of the magnification has left the green area. Then, after moving through the blue area between A and B, the center of the magnification enters the area around B in the image next to the last.

we can render all sorts of geometry consisting of points, lines and triangles with our approach, the implementation of additional examples is not very time-consuming.

5.3.1 Complex hierarchical graphs

The first application of our proposed method is the exploration of complex graphs that visualize complex, hierarchically organized software packages [2]. The circles in the visualization depict the single software classes. They are surrounded by colored areas of recognizable shapes that visualize the containment within individual software packages. Edges between the circles represent inheritance relations between the corresponding classes. The visualization contains many details that are over a thousand times smaller than the entire context. The large differences between magnification factors we can show in a single image are apparent in Figure 5.7. While we need several ordinary images to visualize details in a complex graph, the complex logarithmic view simultaneously shows objects with sizes ranging over three orders of magnitude.

The original intention behind our method is to be able to explore single nodes closely, while still keeping the whole context in sight and reachable at all times. Enlarging the single nodes to a considerable size of the screen allows showing additional details for these nodes, such as source code or information about variables.

The application evidences the advantages of our approach over existing methods: The characteristic shapes that show the clustering in our graphs are well recognizable, even if they are far away from the current center of interest, because they are not compressed. This makes it possible to find known places again very quickly, even though they might be very small on the screen. Pulling these places into focus is possible with one simple drag of the mouse.

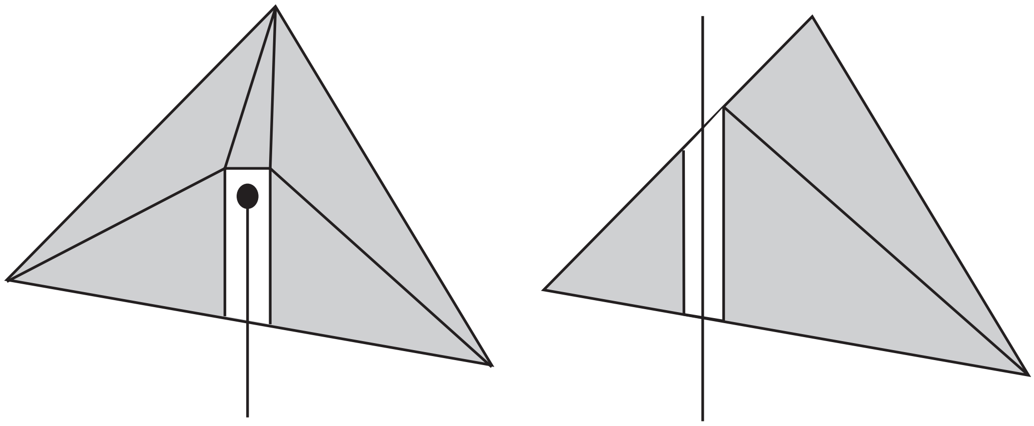

Figure 5.6: The two cases of intersection between a cut and a triangle.

In addition, the complex logarithmic mapping yields a very intuitive display for connections from a node focused in the center of the mapping to other nodes: All rays emanating from this node are mapped to vertical parallels, and thus separated from each other. This reduces the clutter between the edges starting at a focused node that is eminent in an identical mapping of the whole information space or a common fisheye (see Figure 5.9).

5.3.2 Voronoi Treemaps

Our second application example is to use our method for the enlargement of a cell of interest in Voronoi Treemaps [3]. This application also shows several notable properties of the complex logarithmic views. Since the shapes of the single cells and the borders of the hierarchies in a Voronoi Treemap are approximately circular, they are mapped to horizontal curves by the complex logarithmic view. This yields the layered structures shown in Figure 5.8.

5.3.3 Geographical information

A third domain for our method is geographical information. Figure 5.10 shows a complex logarithmic view of map data of North America centered at the Capitol in Washington. Well known features of very different sizes like the national mall, the Potomac, Florida, and the Great Lakes are easily recognizable, since they keep their shape after the mapping. In addition, the mapping does not change angles at street intersections, and thus makes it easy to follow known roads. It is also very easy to answer which places in the view are closer to the center of interest or farther away.

Our fourth example application is geographical statistical information, like the U.S. census data. This data typically contains a lot of information in populated centers, which leads to overplotting, if we want to visualize it in the larger geographical context. Figure 5.11 shows a complex logarithmic view centered at a point in the central park in Manhattan. It demonstrates that it is possible to show single street

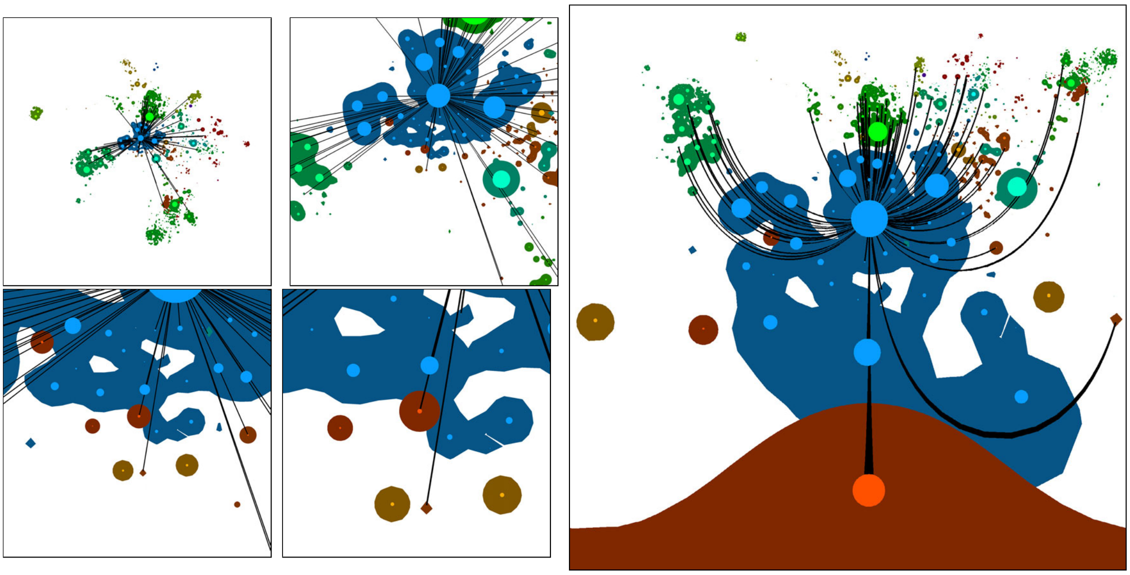

Figure 5.7: On the left side: Four views with different zoom levels centered around the same point in a two-dimensional layout of a graph. On the right side: Detail and context of the same layout within one image, using a complex logarithmic view.

blocks in the context of the whole state of New York, while not changing the shape of, for example, single counties in the countryside or long island.

5.4 Conclusion

We presented a new distortion oriented method for the detail in context problem using the complex logarithm as a mapping function. It allows to show tiny details in very large contexts in one seamless visualization, and is free of the anisotropic compression introduced in the approaches described in Section 3.2.

Although a formal evaluation was not conducted, we required several persons to use our interface and browse through the visualizations of complex software systems. Beforehand, we explained our goals, showed the visualization in the undistorted view, and explained the controls for the common zooming and panning, and the transition to the logarithmic view. After a few minutes, most users were able to move through the logarithmic views and quickly find small details they had previously seen. The strength of our method seems to lie in this task: A change from one place to another typically takes only one mouse dragging operation and is accomplished in about one second, while the same movement from one detail to another with two zooming operations, first outwards, then inwards again, takes considerably longer, since it is necessary to re-center the view between the magnification operations. The difference in speed becomes more significant with more complex data. Some users reported that it helped to imagine themselves to be at the center of the visualization, looking in all directions simultaneously, just like in a panoramic image. Then, the fact that objects become larger once you pull them towards you was reported as intuitively understandable, as well as the fact that things disappearing from the

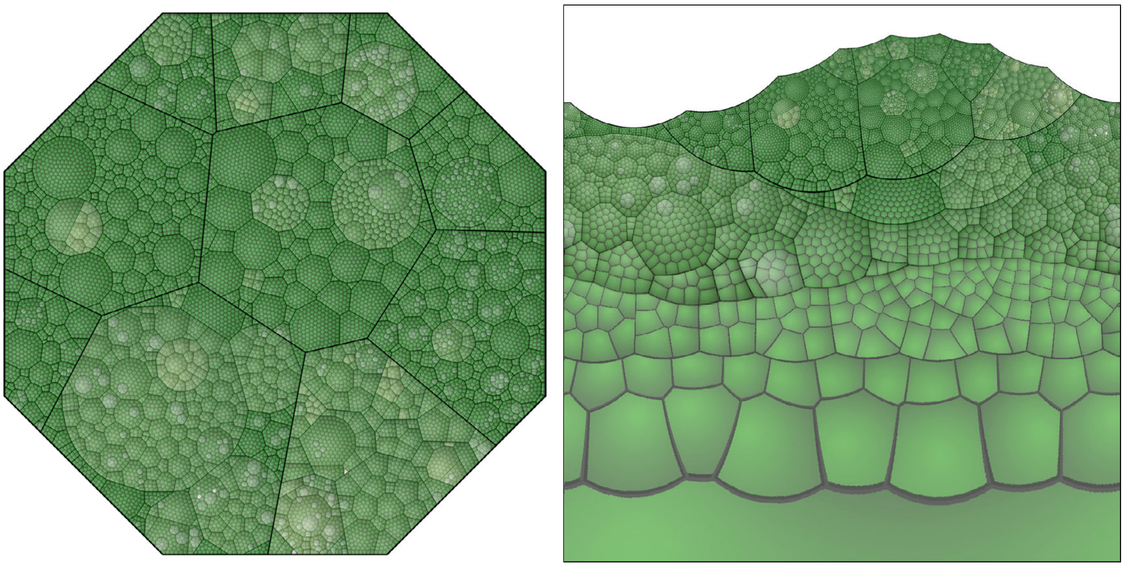

Figure 5.8: Voronoi Treemap (left) and complex logarithmic view with one cell in the Voronoi Treemap enlarged (right). The borders of the different hierarchies are transformed to almost horizontal curves.

view on one side reappear on the other. This hints at a connection to panoramic images, which is elaborated upon in the next chapter.

The cut in the mapping and the fact that the orientation of the mapped information is changed drastically, posed the most important obstacle for the interaction with complex logarithmic views. However, the degrees of freedom for the cut properties were not implemented in this first realization of complex logarithmic views. The placement of the cut is variable, if we not only translate, but also rotate the information, as will also be shown in the next chapter.

The method is suited for arbitrary two-dimensional information. It is our experience, that its advantages are especially evident if the information is containing recognizable, interconnected features throughout all the different orders of magnitude. The user has to be familiar with the shapes in the visualization, since the changes in scale and orientation make it difficult to recognize features that look similar solely by their relative positions. It is thus important to use the method with layouts that produce structures that are typical for the represented data items. Although the focus of this chapter was on abstract data, geographical data usually has these properties without extra effort. Consequently, the application of complex logarithmic views for aerial imagery is the subject matter of the next chapter.

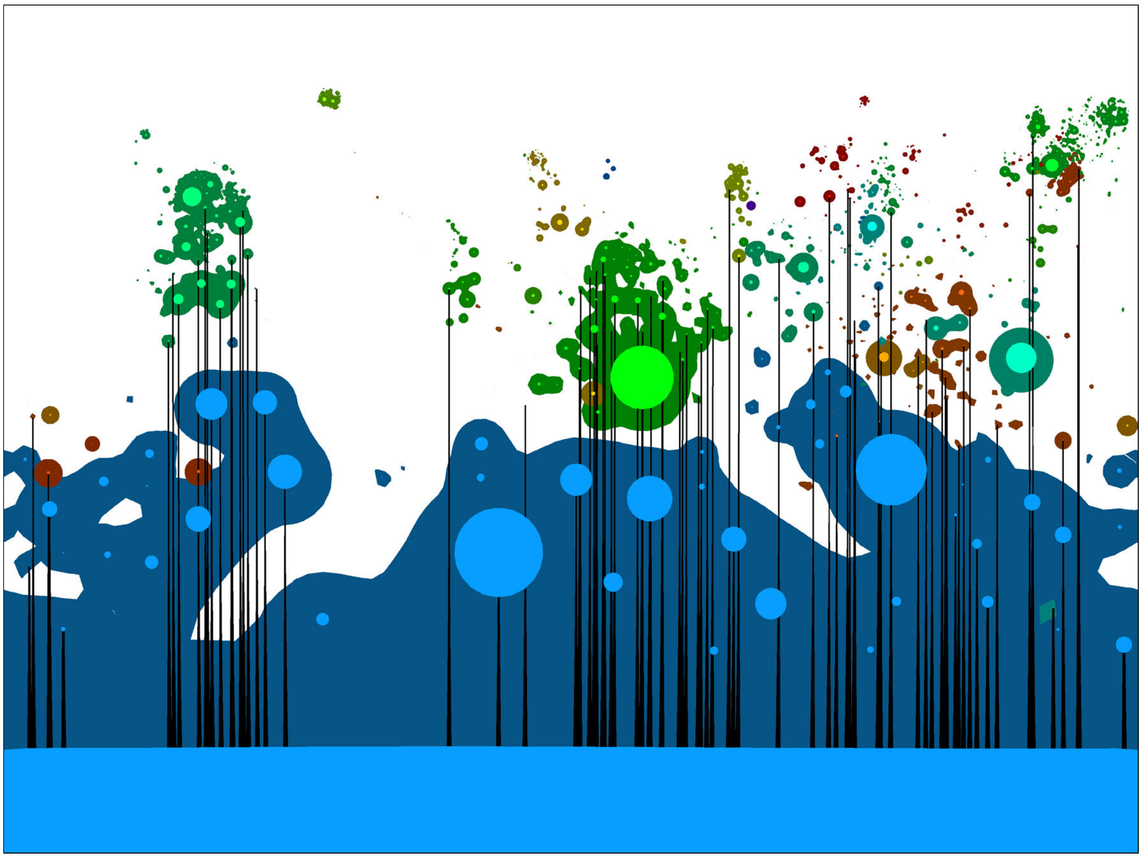

Figure 5.9: The origin of the complex logarithmic mapping is inside the node. It is cut open and mapped to the bottom of the image. Edges emanating from the node to other nodes are mapped to parallel lines.

Figure 5.10: Complex Logarithmic view (bottom) of map data (top). The view is centered at the Capitol in Washington. Points north-west of the capitol are mapped to a vertical line in the middle of the image. Points south-east are mapped to the very left and the very right.

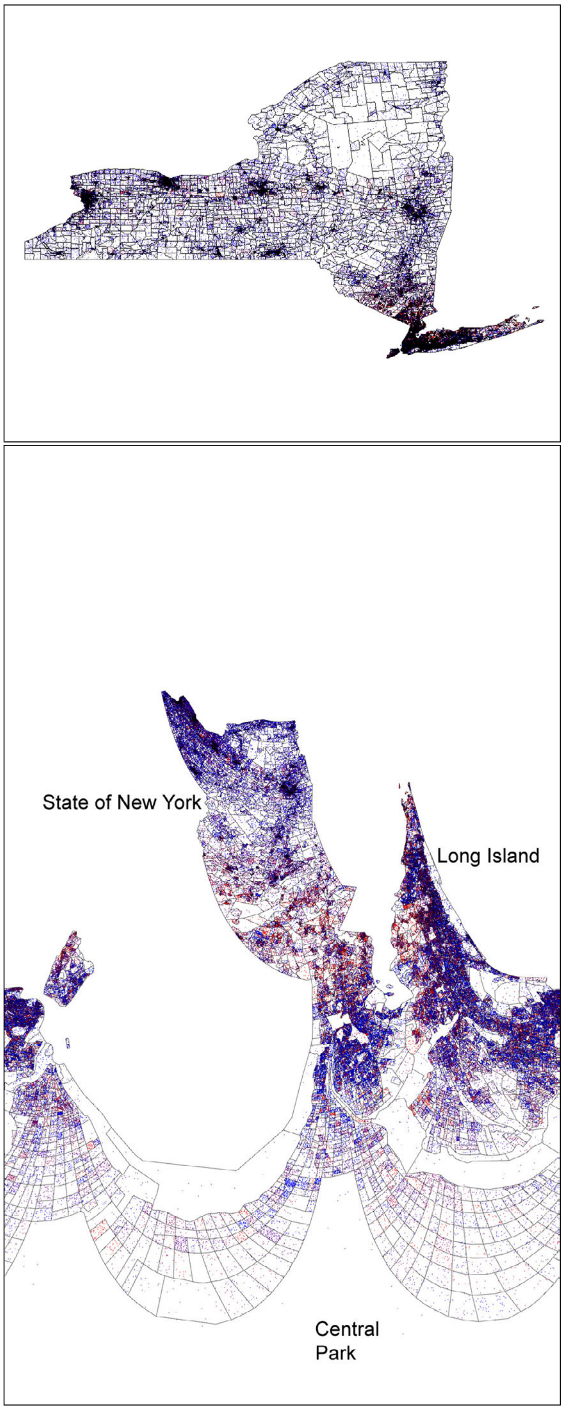

Figure 5.11: Logarithmic mapping (bottom) of census data of the state of New York (top). The view is centered at the Central Park in Manhattan. The corners of the park are still right angles, since angles are locally undistorted.

Chapter 6

Complex Logarithmic Views for Aerial Imagery

This chapter describes the application of complex logarithmic mappings for the interactive exploration of highly complex whole-world aerial imagery. This approach poses several new challenges compared to the complex logarithmic views described in the previous chapter.

In order to build an intuition for the mapping fitting this subject matter, it is prudent to view the mapping as a form of perspective, since ordinary central point perspective is a well established mapping for the browsing of aerial imagery. As we hinted at earlier, like with the fisheyes, there exists a strong connection between the two mappings, which is described in the first section of this chapter. Another challenge is the spherical shape of our planet, which has to be taken into account for the depiction of the part of the view which shows the large context containing the continents. Therefore, our method incorporates cartographic knowledge, at which we look in the following section. The implementation of real-time rendering also also needs a more sophisticated approach for the huge amounts of pixel data involved.

6.1 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 central perspective view and the complex logarithmic view, we start with the former. The central perspective maps points by projecting them along straight rays onto a viewing plane, as illustrated in Chapter 4.1. An example for a perspective view of our subject matter in this chapter, aerial imagery of the Metropolitan Museum in New York, can be seen in Figure 6.2(a). As mentioned earlier, a 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.

In contrast, a panoramic perspective uses an unfolded cylinder as viewing plane, and presents information in all directions around the camera simultaneously. The

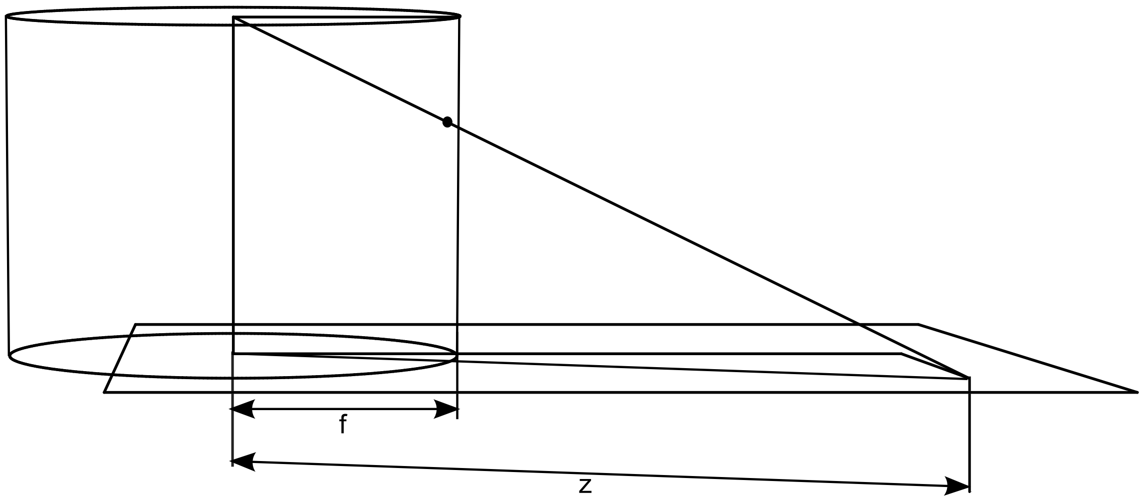

Figure 6.1: Side view of a panoramic perspective projection. Instead of on a view plane, points in the object plane are projected on a view cylinder with radius f. Note that then any cut along a plane through the view point and perpendicular to the object plane looks like Figure 4.5.

(a) Ordinary perspective view (b) Panoramic perspective view

Figure 6.2: 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.

result of this form of panoramic perspective view is shown in 6.2(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. Analogously to the central point perspective (see Section 4.1.2), the compression gets arbitrarily big for large distances from the viewpoint, 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 introduced in Section 4.2. The final result of a complex logarithmic view showing the Metropolitan Museum is shown in Figure 6.7. The similarities of the lower part, and the differences of the upper part, between this image and the panoramic perspective are evident.

In that depiction, 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.

6.2 Cartography

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.

6.2.1 Relationship to the Mercator Projection

The Mercator projection [69] 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 [52]. Using other cartographic projections for the intermediate step of flattening the earth leads to different detail-in-context mappings. In the next subsection, we discuss map projections that are useful for our representations.

6.2.2 Azimuthal Map Projections

While in general arbitrary map projections could be used, we considered the class of azimuthal map projections [69], 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.

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. While this may lead to parts of the landmasses being cut off, in practice this case is seldom, since the majority of the center of interests antipodes are located in water covered areas of the planet.

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.

Figure 6.3 shows an overview of these three map projections and the resulting complex logarithmic views.

6.3 Implementation

Here 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

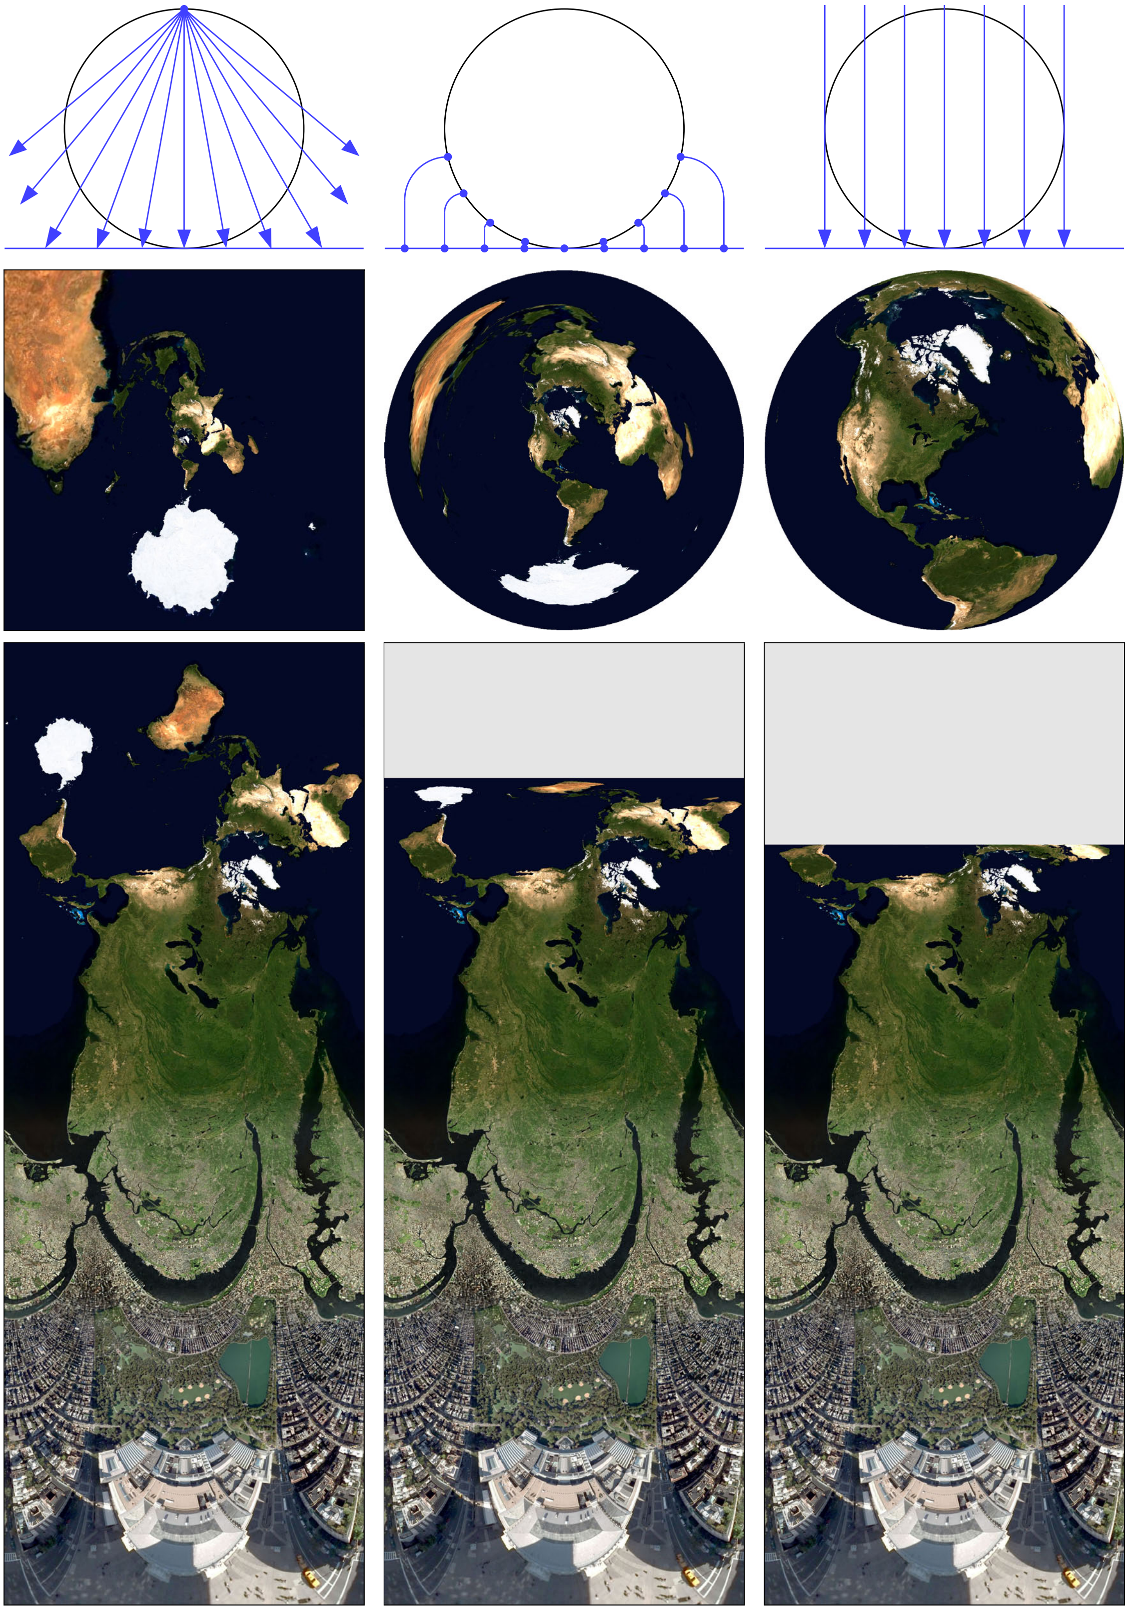

(a) Stereographic (b) Equidistant (c) Orthographic

Figure 6.3: 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.

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 [75, 65] 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 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.

6.3.1 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 6.4

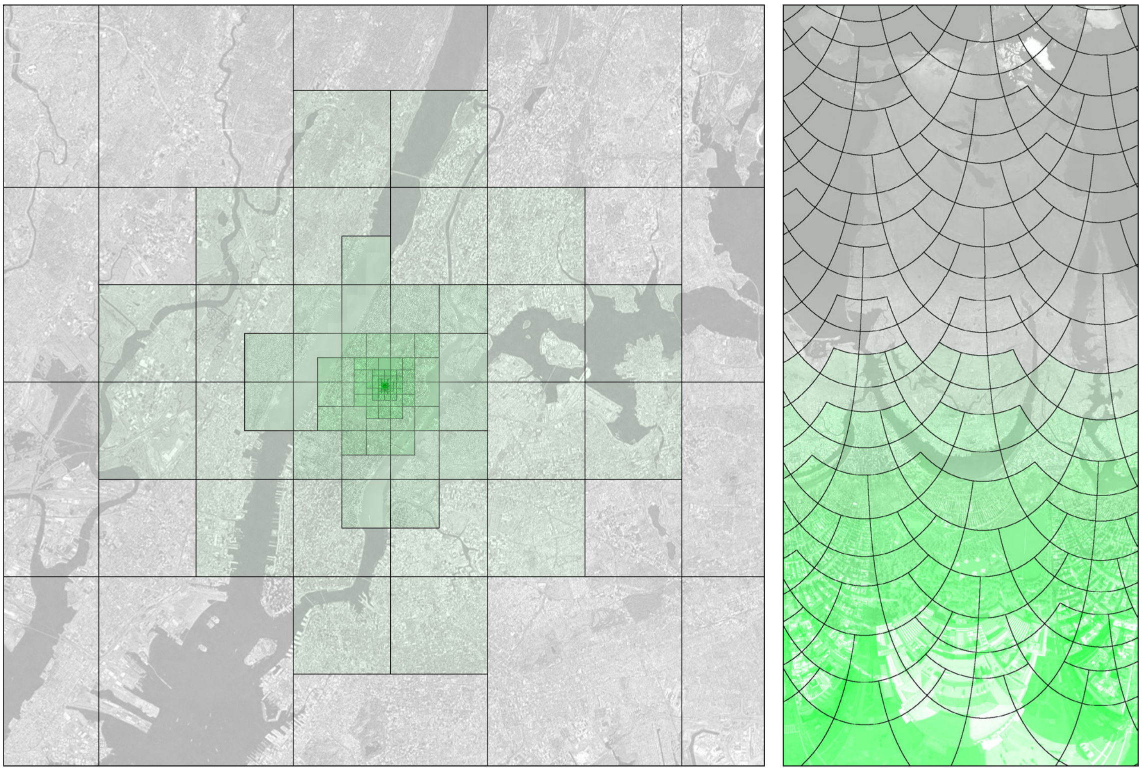

Figure 6.4: Zones of tiles with different size and resolution around the center of interest in the orthographic (left) and complex logarithmic view (right)—greener means higher resolution. In the transformed view, the concentric zones are mapped to horizontal stripes with almost equal size.

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 6.5. 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 neighborship of the references in one layer is identical to the original neighborship 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.

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

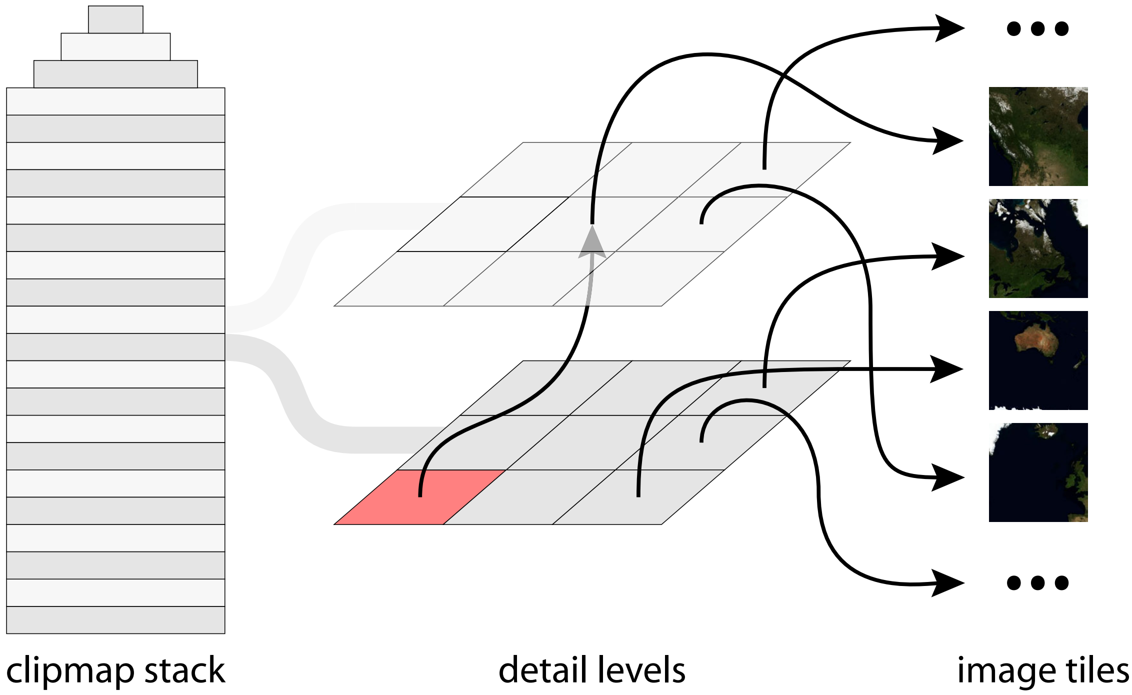

Figure 6.5: We use a clipmap with an index structure (left) to organize the imagery in the graphics hardware. The index structure contains one layer (middle) for each detail level in the clipmap. These layers contain references to the image tiles, which are arbitrarily arranged in the memory (right).

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.

6.3.2 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 [19] 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 6.6: 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.

6.4 Results and Interaction

As a result, in Figure 6.7 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

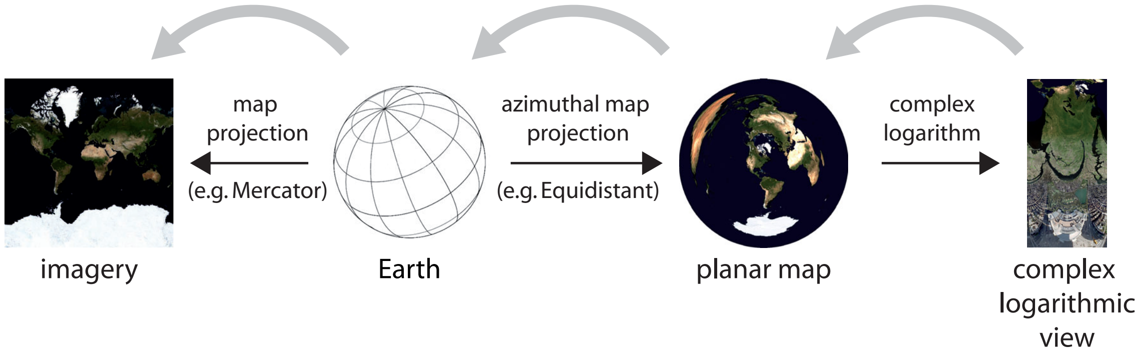

Figure 6.6: Our complex logarithmic views are a concatenation of an azimuthal map projection with a complex logarithm. To sample the pixels of the complex logarithmic view, we have to invert the complex logarithm and the azimuthal projection, and apply the projection of the imagery.

realtime, we can use the representations to intuitively explore and navigate on the Earth's surface. Like in Chapter 5, 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. Other than in the previous chapter, where the rotation of the complex logarithmic view was constant, here zooming and rotating in the original world coordinate system are changed to a one-dimensional translation in the complex logarithmic views: zooming translates the view vertically, while rotating translates it horizontally. The latter equivalence is shown in Figure 6.8.

In 6.9 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.

6.5 Conclusion

We presented the adaption of complex logarithmic views for very complex satellite and aerial imagery. In addition to an intuitive connection to perspective mappings, and the consideration of cartographic knowledge, 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.

Figure 6.7: 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.

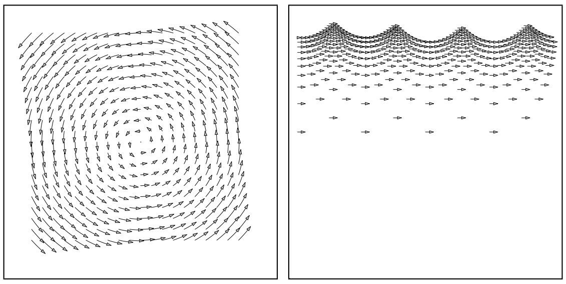

Figure 6.8: In addition to the interaction modes in Figure 5.4 in Chapter 5, the use of fragment shaders allowed for the easy implementation of rotation: rotating the data in the euclidean view (a) accords to constant horizontal movement in the complex logarithmic view (b).

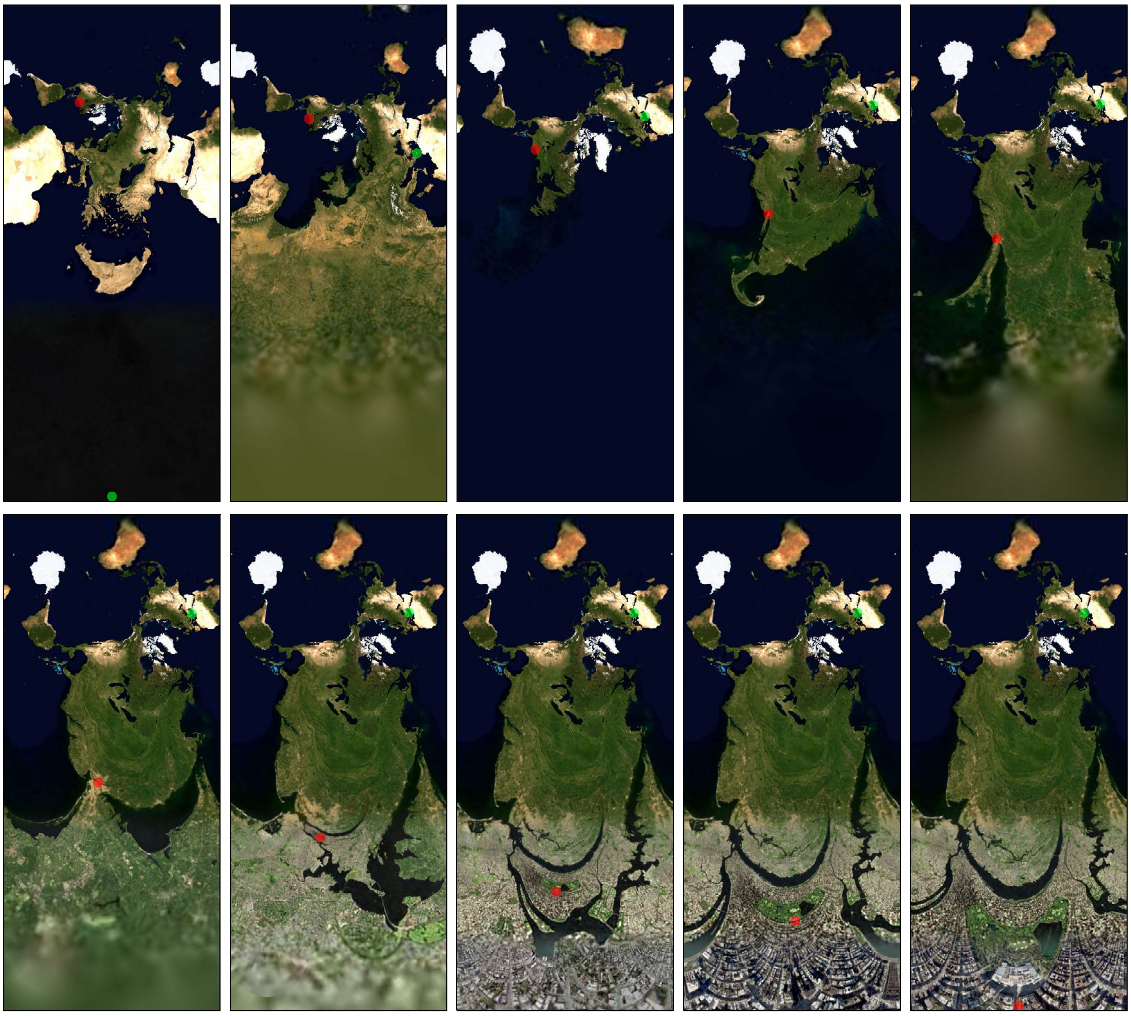

Figure 6.9: An example of an interaction sequence, moving from Crete to New York, by dragging it downwards. The starting point is marked green, the end point red. A detailed description is given in 6.4.

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.

The cut in the mapping posed less of an obstacle for successful navigation, since the additional degree of freedom of rotation made it easily possible to place it in the less important areas of the depiction. However, making the cut completely disappear would of course be preferable, but is mathematical impossible for complex logarithmic views in the plane. However, in the next chapter, we describe an application of complex logarithmic perspective in an interactive panoramic cinema which was able to display aerial imagery without that cut.

Chapter 7

Artistic Research Project "Globorama"

This chapter describes the artistic research project "Globorama", which stems from the collaboration between three groups; the Institute for Visual Media at the Center for Art and Media (ZKM) in Karlsruhe, the Computer Graphics and Media Design Group at the University of Konstanz, and the Human-Computer Interaction Group at the University of Konstanz. The project made it possible for the broad masses to explore the surface of our planet in a never before seen immersive way. Towards this end, we adapted the previously described technique for the display of satellite and aerial imagery through hugely different scales for a novel immersive display, an interactive panoramic cinema. The interaction with the resulting installation was made possible by use of advanced laserpointer interaction.

The physical manifestation of this concept described in this work was made possible, since researchers at the ZKM developed an interactive panoramic display for artistic purposes. This display has a very high resolution through use of six video-projectors, and is able to render distorted aerial imagery in realtime employing a computer with three modern graphics boards. Because of the above mentioned panorama-like properties of complex logarithmic views, it seemed natural to connect that perspective with this exciting new display. This way, one of the shortcomings of complex logarithmic views on ordinary flat screens, namely that the world is cut open along an imaginary line behind the viewer, ceases to exist. The resulting experience while moving through the world in the panoramic display is very immersive and intuitively graspable.

To make interaction and navigation with the installation possible, we use modern laserpointer-interaction developed by the Human-Computer Interaction Group at the University of Konstanz [38]. We use a specifically developed laserpointer module and computer-vision algorithms adapted to the panoramic screen in order to move through the world, and to navigate all the different scales. Our experiences have shown, that this mode of interaction needs little or no explanation, and is very helpful for the, mostly very brief, activation of people in exhibitions.

This chapter serves to describe the adaptions made for the interactive display of aerial imagery in a panoramic display, the additions to the basic display algorithms for the installation, and the user experience.

7.1 Installation

The physical installation consists of a large interactive panoramic display, and an additional laserpointer interaction device.

Figure 7.1: The custom laser pointer interaction device (from [38]).

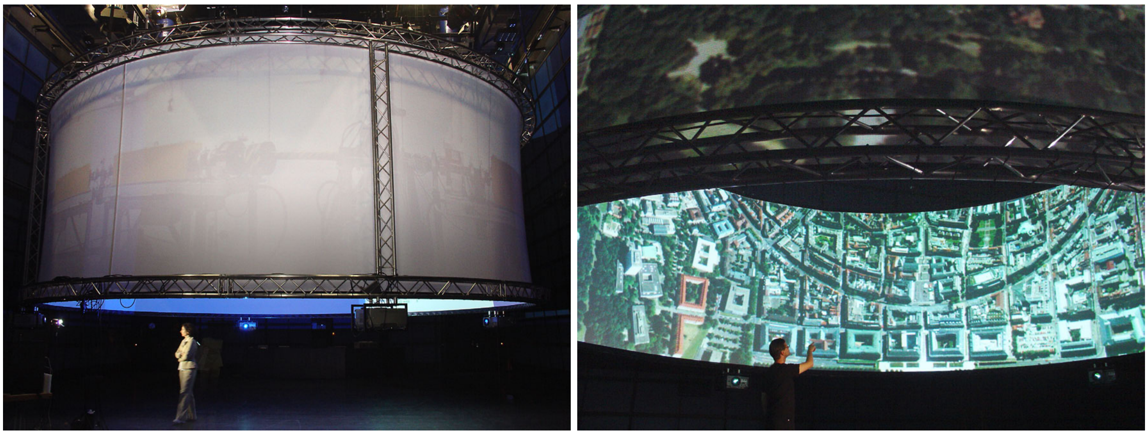

Figure 7.2: The panoramic display developed at the ZKM in Karlsruhe. Six projectors display a seamless image with 8192 pixels width. Outside view of the panoramic display (left). Inside view of the display hosting Globorama (right). The center of this visualization is the city of Karlsruhe (images: Bernd Lintermann, ZKM).

The panoramic display was developed at the Center for Art and Media (ZKM) in Karlsruhe for artistic installations. It frontally projects a panoramic image on a screen with ten meters diameter and a height of three meters, using six video projectors. These projectors' images are calibrated and distorted in a way that the installation yields one seamless virtual projection screen with a resolution of 8192x928 pixels. The installation is driven by one display server containing three graphics cards. Customized software, the so-called Panorama-Player, enables the transparent development of OpenGL-accelerated applications.

The interaction device of the installation, which was developed by the Human-Computer Interaction Group in Konstanz, employed a customized laserpointer with additional accelerometers and tactile feedback, and used specially adapted computer vision algorithms [37]. The resulting device enables easily understandable pointing operations, and possesses several buttons for the steering through the visualization.

The installation was amended with a directional sound system, which played aural background that was tailored in real time to the visualization.

7.2 Interaction

Adapting the complex logarithmic perspective to the panoramic screen consisted of two parts; the adaption of the display algorithm is detailed in the next section. The

Figure 7.3: Changing the level of magnification in Globorama. Zooming in and out of the visualization is akin to selecting a different height for the stripe to be displayed (images: Bernd Lintermann, ZKM).

interaction in the panoramic display environment also had to be adapted compared with the interaction employing an ordinary computer workstation with a mouse.

Due to the very wide aspect ratio, not all the orders of magnitude fit on the screen at the same time. Since zooming in the panoramic screen equates a vertical shift (see Section 5.1.2), we chose to display only a horizontal stripe of the context around the current viewpoint, and to enable a viewer to change the zoom factor interactively. This movement in the panoramic display (see Figures 7.3 and 7.4 left) feels very much like rising or falling, the comparison to flying up- and downwards in a hot air balloon describes the experience very aptly.

Moving through the environment, i.e. changing the viewpoint, is also very interesting: Since things on the one side of the screen move downwards, or towards the viewer, while they move upwards, or away from the user on the other side (see Figure 7.4 right), it feels like the landscape tilts during movement. The similarities to the ordinary perspective here help a lot to grasp the ongoings during interaction.

We chose not to implement a facility for rotating of the visualization, since a user can very easily rotate himself, turn around in the display.

To make the interaction easy and intuitively graspable, we chose the following interaction method: The first button on the laserpointer device had the function of “I want to go there”. Pressing this button shifted the center of interest towards the chosen point. In order to be able to fluently change the magnification of the visualization with the same button, the magnification was simultaneously changed if the chosen point lied close to the upper or the lower margin of the screen. Choosing a point in one of these two zones therefore magnified a new center of interest automatically if it threatened to get too close to the viewer to be seen, or demagnified it respectively. This yielded a very easy point and click navigation, which needed little explanation.

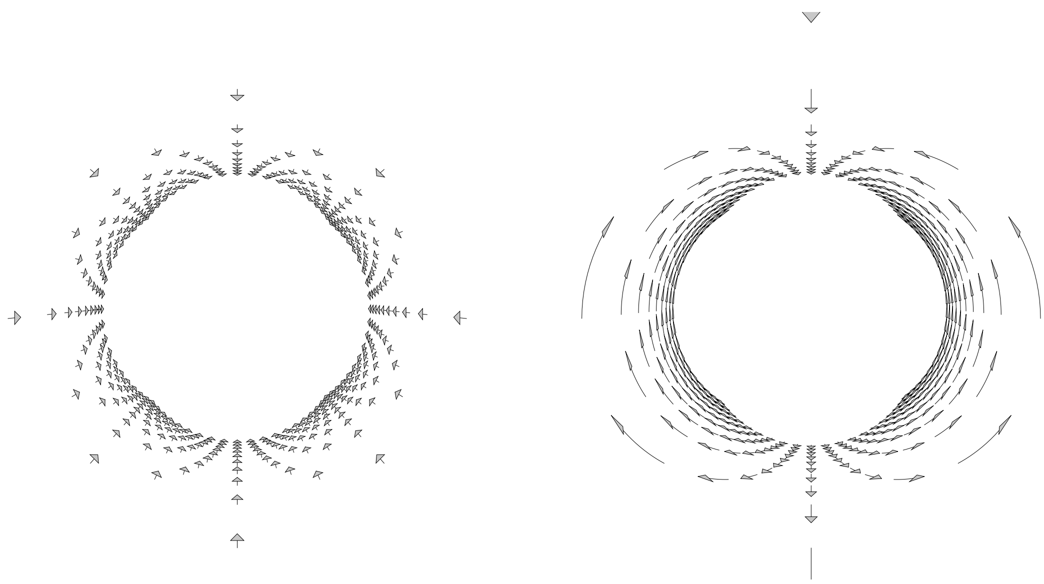

Figure 7.4: Views from the bottom upwards into the panoramic cinema: Zooming in Globorama (left). Movement through the visualization (right). Zooming raises or lowers the content of the screen uniformly, whereas changing the center of interest magnifies parts of the surroundings on one side of the screen while the opposite parts shrink.

The second button served as a way to quickly change the center of interest to several predefined points on the globe by superimposing a world map as a menu. The third button helped people who got completely lost by returning them to a predefined home, for example the current real world position of the installation in a level of magnification that made quick recognition and reorientation possible.

The visualization was further enriched by several data in addition to the aerial imagery. Textual labels were superimposed from a geographic database of geocoded names. On several locations, georeferenced panoramic photos could be activated by selecting specific markers. Then, these photos were displayed on the screen until another button was pressed. Hovering over the markers gave users tactile feedback by vibration of the laserpointer device. Similarly, it was possible to choose realtime webcam feeds from all over the world.

7.3 Adaption of display algorithm

The display algorithms had to be improved sufficiently for the extremely high resolution. The display algorithms for distorted imagery described in the previous chapter proved flexible and fast enough to work for the high resolutions required in this installation, since they are based on fragment shaders, and therefore can fully employ the three graphics cards in the installation. Several optimizations were possible due to the fact that only a relatively narrow stripe of the complex logarithmic view was displayed at any given time, and that the change in magnification always took place relatively slowly.

The caching algorithms were adapted on several levels in order to speed up data throughput. The required data was fetched from the internet and saved on harddisk on a separate computer, which was connected to the display computer by local area network. After extended use by thousands of visitors, it showed that relatively few requests for new data were necessary, since most of the interesting points were visited previously. The flexible index structure described in Section 6.3.1 made it possible to load only data from the required orders of magnitude, which allowed for an efficient use of the memory available on the graphics cards.

The use of six video projectors was made possible by one display server with three high-end NVidia graphics cards. These cards posses many parallel Fragment Shader units and a lot of memory, and each of them is able to drive two projectors. Since all the distortion functionality is encapsulated in the most local of units, the fragments, parallelization was no issue, and the installation was able to derive the full speed-up that was to be expected by adding more graphics boards. Therefore, although the overall resolution was 8192x928 pixels, every graphics card only had to distort the imagery for a little more than a third of this number of pixels.

The installation showed that very extreme distortions like the complex logarithmic views are feasible for high resolutions using modern graphics hardware. The visualization with additional textual and label display was fast enough to run with 60hz for extended periods of time.

7.4 User experience

The experience of exhibiting an installation which employs complex logarithmic perspective made it possible to evaluate how inexperienced users cope with the concept. Globorama was shown to several thousand people during four weeks at the Panorama Festival at the ZKM, and to hundreds of thousands of viewers at the Thyssen-Krupp science park. It was generally well received.

Watching the users, it became clear that not everybody immediately coped with the concept. Most of the people, however, where able to understand the mapping, and used the installation to explore places well known to them without any instructions other than regarding the function of the different buttons. This was possible thanks to the very intuitive interaction device as well as because the mapping depicted local shapes of our world in the same way that mapping applications like Google Earth do. There were several interesting comments from users, which showed, that the employed perspective was graspable if they realized, that they had to imagine to stand at a point on the surface of the Earth and to look in different directions. For example, several users intuitively used term like “behind” or “in front of” very naturally, like in “you see, there behind the alps you can see Italy, which is shaped like a boot”, or “look, there is our hotel, right in front of the beach”.

Within the framework of an evaluation of the laserpointer interaction device, the Human-Computer Interaction Group at the University of Konstanz handed out questionnaires, amongst others asking the visitors about their experiences with the employed complex logarithmic perspective. Although that mode of evaluation has a more informal character and does not make definite statements, the retrieved 73 questionnaires seem to suggest that a majority of visitors enjoyed the different worldview, and were able to intuitively understand the mapping concept and to reach their navigational goals.

7.5 Conclusion

This chapter described the adaption of the complex logarithmic views for an interactive artistic and research installation in an panoramic display. It thereby offered the opportunity to show this perspective of our world to a broad public, and is evidence that the display algorithms are robust and fast enough even for high resolutions.

The reactions of the public to the installation were a real enrichment for the research in this work, and the cooperation with the Centrum for Art and Media in Karlsruhe and the Human-Computer Interaction Group in Konstanz was a great experience.

Chapter 8

Map Warping for the Annotation of Metro Maps

This chapter describes the application of mapping techniques for a specific problem: to aid in the navigation of metropolitan transportation systems by interactively linking schematic metro maps with geographic data. The work described in this chapter is the result of a collaboration with the Research Project "Visual Navigation" and the Algorithmics Group at the University of Konstanz, who came up with the basic idea, and approached the author of this thesis with the, from a mapping and computer graphics standpoint, interesting problem.

While the mapping solution yields results which are very different from complex logarithmic views, there are similarities: due to the visualized geographic data, both mapping techniques need to fulfill similar constraints. The mapping has to be as conformal as possible and free of overlap, in order to enable recognition of geographical features. The technique also automatically results in a form of detail-in-context visualization, since vastly different scales are implied by schematized metro maps.

After the basic problem definition, our technique for warping geographic information in order to fit a schematic map with as little anisotropic compression as possible is described in this chapter. The remainder of the chapter describes the resulting interactive technique.

8.1 Motivation

The mapping problem described in this chapter arises from the existence of two different types of maps that serve as navigational aids in metropolitan areas. While maps of public transportation systems are designed to effectively and efficiently convey possible itineraries, street-level maps usually serve to convey relative positions and distances of a wealth of locations. These very different purposes have led to likewise differences in the design of such maps.

8.1.1 Street Level Maps

One goal of detailed street maps is to minimize distortion, which means showing the real world in a way which is geometrically similar to what we would see from a vantage point high above a city. Accordingly, street intersections have the same angles like in reality, and features like parks and rivers have the same distinct shapes they have in the real world. This makes it easy to mentally put ourselves on the map and to autonomously navigate through the city, since we can use the shapes to find out where we are.

The street-level maps contain a lot more information than only the network of streets, such as landmarks, public facilities and many other aspects of the surrounding environment. By representing geographical information in addition with this very high level of detail the street-level maps show a large variety of relations between locations. This abundance of detail is necessary, since a city map is used for many different tasks, most of which require navigational decisions on a much smaller scale than the decisions that have to be made while traveling in a public transportation network.

8.1.2 Metro Maps

Schematic transportation maps are designed to clearly show the navigational information of the transportation system on a preferably small map. The pioneer of schematic public transportation maps, Harry Beck, conceived his well-known Tube Map of London – which is considered a design landmark and forms the basis for schematic public transportation maps today – in 1931 [21]. To achieve an expedient representation of the underground map he found it was necessary to make the central area appear larger, since the stations were closely crowded there. Drawing a map of the whole area in a limited space mapped the stations in the center too close to each other to leave space to make their connections distinguishable. Beck imagined he was using a convex lens to ensure readability in the center and in the periphery at the same time. He formulated the general design principle for metro maps to place all the stations at equal distances, although their geographical distances are very different. This also reflects the fact that the traveling time of a metro is approximately independent from the distance of the stations in the real world, since it takes a train a relatively long time to drop off and pick up passengers, and to accelerate and decelerate.

Another typical schematization requirement for metro maps is straightening the route lines by placing the stations of a line on straight lines, if possible. The overall shape is further simplified by restricting the positions of the stations to be only at a few discrete angles relative to each other, which, for example, leads to route lines only being mapped to verticals, horizontals or diagonals. This makes it easy to mentally connect stations belonging to the same line.

The resulting metro maps avoid intersections with small angles, and are generally easily and intuitively readable. The automatic layout of metro maps has recently grown into an active field of research [27, 51, 74, 81]. However, most of the metro maps in use are still manually fabricated by designers, who tweak the maps until they look just right, and do not always strictly adhere to the aforementioned design principles. We use manually fabricated layouts as input for our method. One argument for the use of existing manually fabricated layouts is that the inhabitants of a city are already familiar with them. We assume that over time people have adapted their mental map of the whole city to these layouts. Subsequently adapting to a different layout imposes mental strain on them. Nevertheless, our method works with automatically generated metro maps as well.

For finding a way from one station to another station within a public transportation system the user only requires an overview of the relations between the stations concerning connectivity. Therefore, the type of relations represented in a schematic map of a transportation system is very different from that of a street-level map; the easy perception of the existence of services and connections takes precedence over geographic accuracy. Spatial relations are preserved on an ordinal scale, if at all.

8.1.3 Combined maps

Many city guides contain street-level maps that are annotated with the stations and lines of the local transportation system. These annotated geographic maps are suitable for many purposes, but since all advantages of schematization are lost, metropolitan areas around the world rely on schematized maps of their transportation systems.

One approach for combining schematic transportation maps with street level information are spider maps [60, 77], which are transfer guides for metropolitan areas. A Spider Map is a schematic transportation map that centers on one station and displays the local area surrounding the station as a geographic street map, aiding the user when changing bus and metro lines. However, this makes it necessary to adjust the schematic layout for providing display space for the station of interest and its surrounding streets. Additionally, starting at an arbitrary point in the city and having to chose one of the nearby stations to walk to, a street network map centered at this specific starting point is more useful than a street map centered at one station.

To alleviate the above mentioned disadvantages, we propose to annotate schematic maps with all the information usually incorporated into city maps without modifying their design. Thus, in a way, our approach is opposite to the annotation of geographic maps.

Designers sometimes include real world items such as coast lines of the sea or large rivers, but apart from that, most of the features of the real world are not shown, partly because their placement in the schematic map is not trivial. Due to the lack of distance relations in the schematic map the expedient positioning of street-level details requires an extrapolation of the deformations caused by the schematization. Today, many schematized metro maps contain no detailed information other than the stations and their connections, but rather restrict themselves to describing the navigational space of using the transportation system as well as possible. Evidently, such maps are strongly specialized, since there is no reliable possibility of reading out any geographical accurate information about the transportation system or its surroundings.

To analyze these changes in map layouts, Jenny [29] used MapAnalyst, a tool originally developed for the visualization of geographical errors in historical maps. He annotates the schematic map with visual hints to aid the understanding of the implied distortion. Applying this to the London tube map, Jenny observes the typical features of schematic transportation map layouts. For example, in his scale isoline visualization, the fisheye character of that map is clearly noticeable. Klippel and Kulik [36] introduce another approach to visualize the distortion due to schematization by applying it to the commonly known grid squares of a city plan.

The essential feature a street level map and a schematic metro map have in common is that they contain position indicators for station locations. Since schematic maps are designed to preserve relative positions of stations, it seems natural to use them for aligning the street-level map with the schematic map.

It is important to note that our approach is not about morphing one map into another, but about a smooth transition between two very different visual navigation aids, one accurately representing distances and local details and one simplifying the use of a transportation system by abstracting from irrelevant details. This distinction becomes evident when comparing with an example used by Reilly and Inkpen [55, 56], where a slider controls the morphing between a street-level map and a London tube map. Coupled with alpha-blending, *map morphing* is used to make correspondences obvious. Facilitating an intuitive understanding of correspondences is a prerequisite in our approach, just like the avoidance of fold-overs and occlusions. The integration of high-resolution, detailed street-level information and high-level, schematized network diagrams, however, cannot be accomplished in this way, but requires combination with the appropriate level-of-detail and zooming techniques.

8.2 Warping

In order to merge the two maps, we use mappings from the field of image warping. Ruprecht et al. [58, 59] describe different methods used for the distortion of two-dimensional information. All these methods solve the basic problem of warping: Given two-dimensional information and a set of control points in this information, the goal is a mapping function moving these control points from their starting positions to arbitrary end positions. The mapping function should have several properties for a satisfactory warping: It is supposed to be interpolating, which means that the starting positions of the control points are precisely mapped to their end positions. Furthermore, the mapping should be smooth, that is, it should not introduce discontinuities between the control points. Ideally, the mapping should also be free of overlap. In contrast to the above mentioned distortion analysis and visualization by MapAnalyst [29], for warping geographic information to support navigation tasks, it is very important to avoid overlap, since otherwise parts of the information completely disappear.

For the automatic integration of two corresponding maps, each one optimized for adequately displaying its respective navigational information, triangulation-based methods are inapplicable, because they suffer from foldover and other discontinuities, which are not easily solved. Therefore, we chose a warping method using scattered data interpolation that produces smooth and interpolating mapping functions. In addition, we want to keep angles in the distorted map as similar to the angles in the real world as possible, since this keeps the shapes of real world features recognizable.

8.2.1 MLS

Schaefer et al. [62] describe a moving least squares algorithm that interpolates a similarity transformation between the control points. This way, angles are less distorted compared with only interpolating with general affine transformations.

Given a set of control points p, their position after the warping q, and an arbitrary single point v, Schaefer et al. solve for the optimal affine transformation l_v that minimizes

$$\sum_i w_i |l_v(p_i) - q_i|^2. \tag{8.1}$$

The method is called a Moving Least Squares minimization, because the weights w_i depend on the point v:

$$w_i = \frac{1}{|p_i - v|^{2\alpha}} \tag{8.2}$$

The parameter α controls the decay-profile for the distance, and should be larger than 1. For our examples, we experimentally chose it to yield satisfying results, and a typical value was 1.5.

This leads to a different transformation $l_v(x)$ for each single point v. Restricting the allowed transformations to similarity transformations, Schaefer et al. find the following optimal mapping functions for the single points v:

$$l_v(x) = (x - p_*) \frac{1}{\mu_s} \sum_i w_i \begin{pmatrix} \hat{p}_i \\ -\hat{p}_i^{\perp} \end{pmatrix} \left(\hat{q}_i^T - \hat{q}_i^{\perp T} \right) + q_* \tag{8.3}$$

Here, p_* and q_* denote the weighted centroids:

$$p_* = \frac{\sum_i w_i p_i}{\sum_i w_i} \tag{8.4}$$

$$q_* = \frac{\sum_i w_i q_i}{\sum_i w_i} \tag{8.5}$$

Furthermore, $\hat{p}_i = p_i - p_*$, $\hat{q}_i = q_i - q_*$, $\mu_s = \sum_i w_i \hat{p}_i \hat{p}_i^T$, and $\perp$ is an operator which maps a vector (x, y) to $(-y, x)$.

We apply these mapping functions for single points individually to the points in our geographical datasets. As Schaefer et al. point out, the mappings still suffer from overlap. A simple example clarifying the resulting overall mapping functions by applying them to a regular grid can be seen in Figure 8.1(a). In Figure 8.1(b), the overlapping parts of the resulting 2D mapping function are clearly visible. Thus, we combined the above mapping method with the overlap control described by Tiddeman et al. [76] to achieve overlap-free mapping functions.

8.2.2 Overlap Control

Tiddeman et al. describe a general method to avoid overlap problems. One key observation of the method is, that for any given mapping function, another mapping function can be derived by scaling the mapping, i.e. interpolating it with the identical transformation. Such a scaling operation with a scaling factor s yields, for our case, the following mapping function:

$$l_s(v, s) = (1 - s)v + s l_v(v) \tag{8.6}$$

The other key observation of the method is, that overlap occurs at any point in a given mapping function if the determinant of its Jacobian changes signs. It is therefore necessary to restrict this determinant J to be at least positive. Since values of J close to 0 mean, that the mapping at that point compresses the warped information very strongly, Tiddeman et al. restrict J further by requiring it to be larger than a minimal value J_{min}.

J can be calculated using estimates of the partial derivatives by mapping two points close to a point v as follows:

$$\left(\frac{\partial f}{\partial x}, \frac{\partial g}{\partial x}\right) \approx \frac{l_v(v) - l_v(v + (\delta, 0))}{\delta} \tag{8.7}$$

$$\left(\frac{\partial f}{\partial y}, \frac{\partial g}{\partial y}\right) \approx \frac{l_v(v) - l_v(v + (0, \delta))}{\delta} \tag{8.8}$$

$$J = \frac{\partial f}{\partial x}\frac{\partial g}{\partial y} - \frac{\partial f}{\partial y}\frac{\partial g}{\partial x} \tag{8.9}$$

Here, δ is some small value. For several scaling factors s, with $0 < s < 1$, it is guaranteed that the resulting mapping function is free of overlap. To find an optimal scaling factor, it is necessary to solve the quadratic equation

$$J = \left(\left(s\frac{\partial f}{\partial x} + 1\right)\left(s\frac{\partial g}{\partial y}\right) + 1\right) - s^2\frac{\partial f}{\partial y}\frac{\partial g}{\partial x} = J_{min} \tag{8.10}$$

for the Jacobian determinant to be J_{min}. Solving a quadratic equation yields between two and zero roots. Since the Jacobian determinant is always equal to 1 at $s = 0$ and only gets smaller than the minimal value J_{min} at the roots of the equation, the mapping is locally free of overlap or strong compression for all scaling factors larger than 0, but smaller or equal to the smallest root in the interval between 0 and 1. To ensure quick convergence, the method uses this root as scaling factor. Since the control points should not overshoot their destination, 1 is used if there is no such root.

It would be necessary to solve the equation for all points in the mapping function in order to find an overall optimal scaling factor. Since this is not possible, the equation is usually solved at discrete positions on a regular grid. We solve it for every single point we map individually. Then, the overall best scaling factor is the minimum of all the locally optimal factors.

Scaling the whole mapping with the derived scaling factor yields a new mapping, which does not fulfill the constraints of the warp, but already brings the control points some portion of the way closer to their destinations, as can be seen in Figure 8.1(c). Iteration of the process and concatenation of the partial mappings brings the control points arbitrarily close to their destinations. A drawback of this method is, that the convergence is not guaranteed for all cases. Also, choosing J_{min} too small leads to unnecessarily strong compression, while choosing it too big prevents quick convergence. A typical value we used was 0.5. We found that with this value the overlap control worked very well for several different examples, which converged within 5-15 iterations. The result of the iterative process for our simple example can be seen in Figure 8.1(d).

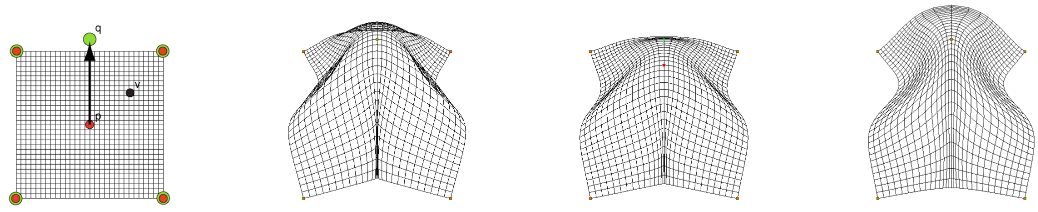

(a) (b) (c) (d)

Figure 8.1: Undistorted grid (a) with fixed control points at the corners and one control point moving the middle of the grid p to a position outside of the grid q. The MLS method results in overlap in the 2D mapping function (b). Scaling the mapping yields a mapping function (c) which moves the control point closer to its destination position. Iterating this process and concatenating the partial mappings results in a mapping function (d) fulfilling the constraints without overlap. Note how the angles at the corner are still right angles after the mapping.

8.3 Interactive Method

We augment schematic maps of transportation systems by superimposing them on street-level maps that are fitted using image warping techniques. Schematic transportation maps usually contain little or no detail describing the environment of stations or their embedding in the surrounding area. The annotation of a distorted city map alleviates this deficiency and improves further the usability of schematic transportation maps by merging two different navigational spaces. We obtain an easily readable transportation network map on which we can show all the typical city map features such as rivers, streets, and parks without compromising on the schematization. Furthermore, we present two interaction techniques: we couple zooming with warping and control over the level of detail in what we call *Warping Zoom*, and adapt a fisheye technique for the exploration of geographical details in the schematic context.

When people use a city's public transportation system, they are faced with a seemingly simple task: They start at one point somewhere in the city, want to get to a nearby station, look for the best connection to another station close to where they want to go, and finally want to reach that destination itself. Usually people use two maps to accomplish that task: On the one hand, an ordinary street-level city map, and on the other, a schematic map of the system of public transportation in the area.

The reason for that is, that both of these representations of our world have their advantages and disadvantages: The ordinary map is very well suited for gaining detailed information down to every single street, but for several reasons struggles to give a fast overview over the network of public transportation, even if that network is included in the map. A schematic map, on the other hand, is optimized for the readability of information concerning the connections and structure of the transportation network. However, it seldom shows the different stations in their surroundings and fails to deliver the needed contextual information for the task mentioned above.

In this section, we describe a method to produce a compound map containing both, network and detailed street information, by warping the street map information to fit a schematized map of a public transportation network. For that warp, we use the earlier introduced mapping from the field of image distortion, which is especially well suited for geographical information. We also introduce a *Warping Zoom* and an adaption of fisheye views which yield in dynamic interactive maps applicable for both, street-level navigation as well as navigation in the public transportation system.

8.3.1 Merging of two data spaces

In our method we use the positions of the stations in both of these maps to merge them in one compound depiction by warping the street-level map to fit the schematic map. Towards this end, we use the corresponding pairs of positions as control points in a warping technique from the field of image warping. The positions in the detailed map serve as starting positions, and the positions in the schematic map as end positions for this warp. This yields a mapping function which, when applied to the geographically correct map, shifts the stations to their positions in the metro map, and distributes all the other features of the real world smoothly between them. We then use this warped map to augment the schematic metro map, in order to support navigation and orientation in the parts of the real world between the metro stations.

8.3.2 Prototype Implementation

As proof of concept, we implemented a prototypical system that generates combined schematic and geographic maps.

Schematic maps are the eminent representations of public transportation services and therefore available for most larger cities.

For the geographic information, we use U.S. Census TIGER map data [78], which contain vector data of detailed street information and landmarks such as water surfaces, parks, airports and public institutions. An advantage of these vector data is the possibility to provide good quality of rendered maps over a wide range of resolutions. They also allow to transform the topography independent from, e.g., textual and symbolic labels to ensure readability. Moreover, these particular data are in the public domain and sufficiently detailed to demonstrate the potential of our approach for actual city plans. We manually annotated the data with the geographic positions of metro stations, compiling this information from other publicly available sources like GoogleMaps [23]. Figure 8.2 shows a geographic map of the Washington Metropolitan Area annotated with the geographically correct positions of the metro stations, the corresponding schematic metro map, and the metro map annotated with the warped geographic map.

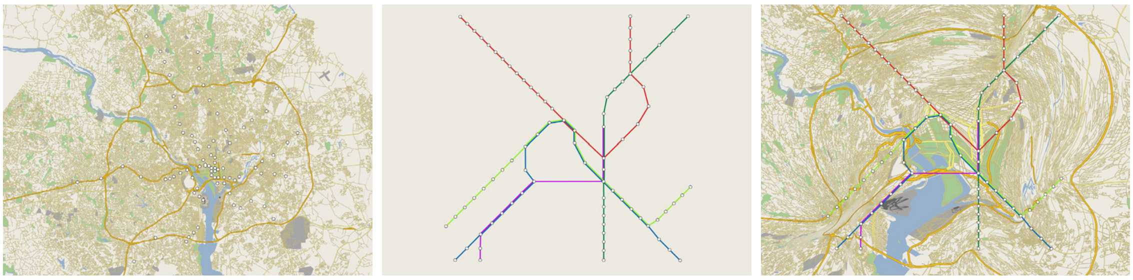

Figure 8.2: Geographic map of the Washington Metropolitan Area with positions of metro network stations superimposed (left). A metro map layout of the same area optimized for readability (middle). In our compound map (right), the metro map is annotated with the warped geographic map.

To warp the geographic map, long lines in the original data are first sampled sufficiently fine to avoid artifacts when rendering them, and the polygons between them. This is necessary because, although the mapping functions described earlier are smooth, straight lines are mapped to curves. Note that mapping only the start and end of a line, and connecting these in the warped image again by a straight line, does not yield the desired result of smoothly deformed curves in general. After subdivision, we evaluate the mapping function for every point of the geographic vector data as described in Section 8.2. We neither use a grid with fixed cell size nor rasterize our data beforehand, like it is usually done in the process of applying image warping functions. The calculation of the mappings itself is time-intensive; our examples took around 1 hour to process on an ordinary desktop computer. However, the mapping has to be calculated and saved only once for every set of control points. To preserve quality, we render the distorted street-level data consisting of lines and polygons after the warp using OpenGL [68] and GLUT, reaching interactive frame rates. In the end, the metro stations, which were manually drawn into the street-level map, have the same positions as the stations in the schematic metro map. We thus obtain a compound map showing topological and topographic information – the schematic metro map annotated with the distorted geographic map.

A particularly nice feature of our warping approach is that it allows to interpolate the mapping between exact geography and schematization. Placing stations in a convex combination of their geographic positions and their positions in the schematized map yields a compromise between geography and schematization. This compromise can be extended to the entire compound map by linearly interpolating between the geographic map and its distortion based on the schematic map. We will sketch an important application of this feature in the next section.

8.3.3 Examples and Use Cases

The main purpose of our work is to ease the transition between schematic maps, which are useful for navigating in a transportation system, and geographic maps, which are better suited for autonomous navigation and locating sites.

We can envision three different use cases for our method: Firstly, on large, static depictions, like wall maps at metro stations, it would be possible to show more detail where it is needed compared with an ordinary map.

Secondly, with limited space, a static overview over the compound map yields annotation of the schematic map with a focus on large streets and landmarks, but can still aid rough orientation in the city. These two cases already make it obvious that level of detail needs to be addressed.

The most interesting use case is the interactive application of our method for small displays, like PDAs. Here, our method can really demonstrate the advantages of linking the two navigational spaces, as we will describe in Section 8.3.5.

As mentioned earlier, we require as input a street-level map of an arbitrary city annotated with the stations of its transportation system and a schematic transportation map of the respective city. We applied our method to maps of the Washington and the Boston area. We chose these two cities because they contain typical features like airports, lakes, rivers, coastlines, islands, harbors, parks as well as a fairly complex transportation system with nontrivial graph structure. Exemplarily we present geographic and compound maps of the transportation systems of the cities in Figure 8.3. In the Boston case, the center of the area is greatly magnified compared to the surrounding area, similar to a fisheye lens. This effect is even more clearly visible in the compound map of the Washington map. The mapping functions manage to keep the areas close to the stations relatively undistorted, while areas between the stations are more strongly stretched.

8.3.4 Level of Detail

Addressing level of detail turned out to be an important issue for our technique. Since showing all the small details on a limited space can lead to indistinguishable visual clutter, it was necessary to consider the local magnification and compression for the depiction of the different features of the geographic map. During the iterative mapping, we calculated for each point an estimation of the partial derivatives at that point for overlap control. We can use these estimations for level-of-detail control as well, since the determinant of the Jacobian yields the local area magnification, and its condition number is proportional to the local compression.

We found it helpful to modify the thickness of linear features like streets directly proportional to the local magnification, and indirectly proportional to the compression. This way, the density of features is evenly distributed over the whole depiction.

8.3.5 Interactive Warping Zoom

For the task of street-level navigation, the typically small map size of schematic metro maps is inapplicable. In order to read the navigational information of the street level, the annotated metro map needs to be enlarged to the size of a regular street-level map or the compound map needs to be enhanced by an interactive zooming technique.

(a) (b) (c) (d)

Figure 8.3: Geographic maps of Washington (top left) and Boston (bottom left). On the right, the maps are fitted to the respective schematic metro maps. Note that it is now possible to discern details in the cities' centers, which are not visible on the left, due to the fisheye-like character of the implied mapping functions.

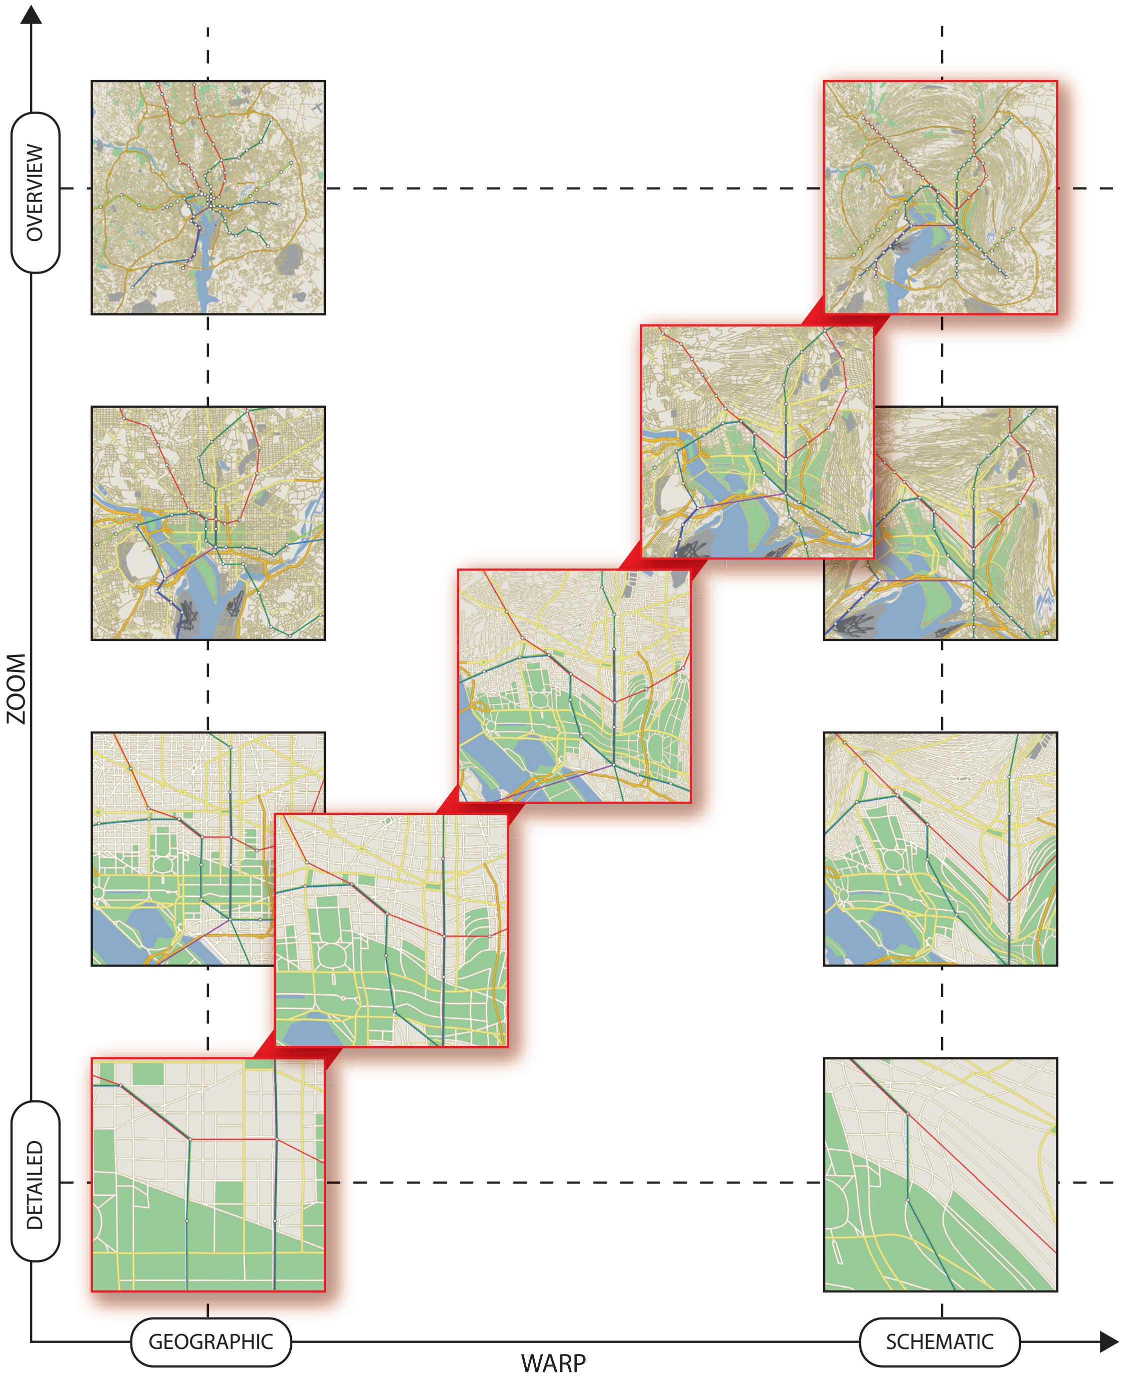

Figure 8.4: The Warping Zoom (here shown on the diagonal in red) is a combination of zooming and at the same time warping between the schematized and the geographically correct map: This makes it possible to employ the schematized layout for an overview, and to employ a detailed geographical layout for localization and street-level navigation. Note that it is not possible to discern the connections in the center of Washington in the geographic overview (top left) in this resolution.

The enlargement of the compound map contradicts the main advantage of schematic maps, which is to give a quick overview over the transportation system on a preferably small map. So we chose to implement a zooming technique, which couples scaling of the viewport with a transition between our schematic compound map and the geographically correct compound map, which equals the undistorted street-level map annotated with a geographical transportation map. To achieve the transition between the two maps we take advantage of the gradual distortion technique introduced in Section 8.3.2. While zooming, we interpolate between the distorted and the undistorted map and simultaneously translate the map in a way that keeps the center of the map at a constant position on the screen. This technique we call *Warping Zoom*.

The effect of our *Warping Zoom* technique is shown in Figure 8.4. The resulting dynamic compound map is especially well applicable on mobile devices: because their display size is usually very small, dynamic maps with zooming functionalities are generally favored. While zoomed out of an interactive general city map, the user wants to get a quick overview of the city. When zooming in, the user wants to get detailed information about a specific region or point or even wants to read navigational information of the street level.

Contrarily, in case of navigation within the city with use of public transportation, the destination is reached approximately by public transportation. This navigation step is aided adequately by a schematic transportation map. So, when zooming out, the user gets a quick overview of the transportation system – which is naturally a schematic layout of the transportation system. In our implementation this schematic transportation map is annotated by warped street-level information in an adequate level of detail. The stations of the transportation system are the interfaces between two navigational spaces – the transportation system and the street-level space. Leaving the transportation system at a specific station, the user has to navigate on the street level to reach the destination exactly. Therefore, the user requires geographical accurate information about the surroundings rather than a quick overview of the transportation system he just left. Additionally, since just a few stations are displayed on the zoomed section, the advantages of schematization are invalidated. Thus, when zoomed in, the compound map has to be displayed in an undistorted/unschematized layout.

We found that, in order to make the interpolation between the distorted and the undistorted map feel intuitive, we had to define start and end scaling factors for the transition considering the configuration of the different maps and display sizes. We display the undistorted compound map when only a few stations are visible. The other extreme is defined once the whole schematized metro map just fits into the viewport.

The level-of-detail-control makes all steps of the transition readable: On the coarsest level, the schematic map is mainly annotated with the most prominent features of the city, like parks, rivers and large roads. The smaller streets appear more clearly when the user zooms in and wants to navigate in the street network.

Figure 8.5: Geographically correct shape of an airport in Washington (left). Ordinary fisheye view applied to the schematic map (middle). Adapted fisheye view (right).

8.3.6 Adapted Fisheye Views

Although the Warping Zoom technique makes the examination of small geographic details possible, it is also desirable to show magnified geographic detail and schematized network information in one seamless view, like in the spider-maps referenced earlier. However, mixing the two spaces by simple spatial interpolation potentially leads to overlapping problems, since the absence of overlap is guaranteed only for concatenation of overlap-free mappings. Therefore, to guarantee a seamless transition between detail and context, we chose a different approach: We apply fisheye mappings to the schematic space.

Using ordinary fisheyes would only make the local compression inherent in the warped information more visible. To counter this compression, we calculate the compression factor using the estimations of the partial derivatives at a center of interest. These define a locally linear transformation, and using a singular value decomposition, it is possible to estimate the magnification factors in the different directions. If these are not equal, the information at that point is compressed. To counter that compression, before applying an ordinary fisheye mapping like Keahey and Robertson [34], we apply a mapping which stretches the area around the center of interest. We then let the stretching decay depending from the distance of that center. At a certain radius from the center, the schematically distorted information is left unchanged by our mappings.

The resulting mappings are free of overlap, and stretch the warped information just enough to locally guarantee angular faithfulness. We illustrate our method by an example in Figure 8.5. On the left, the geographically correct shape of an airport in Washington can be seen. Magnifying the respective detail in the schematized view using an ordinary fisheye view (middle) exhibits a compressed shape with changed aspect ratio. The right cut-out shows the result of our technique: The shape of the airport is enlarged, and it is close to its original shape. In order to fit the geographical detail into the schematic map, it is rotated. The difference between the ordinary fisheye and our method is illustrated by the red circle in the middle cut-out, which corresponds to the ellipse on the right.

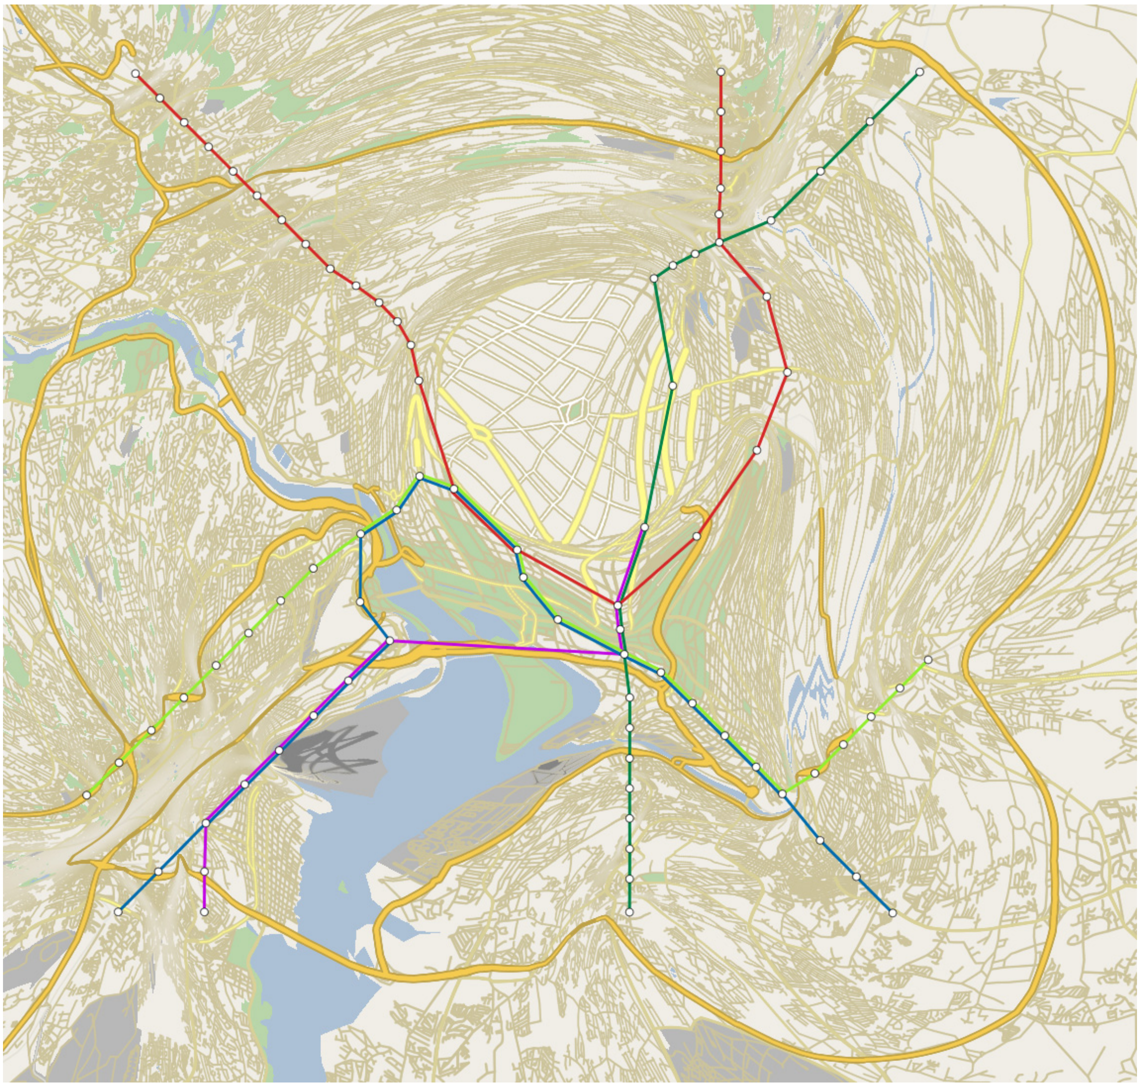

Figure 8.6: Compound map of Washington D.C. in a schematic layout with an area around a point of interest, magnified and stretched by our fisheye technique.

Figure 8.6 shows the whole compound map of Washington D.C. in a schematic layout. The area around a point of interest is magnified with our fisheye technique. Such an interactive map can be used for combined navigational tasks, similar to the Spider Maps mentioned earlier.

8.3.7 Distance Information

To aid in the understanding of the geographical distance relations in the schematic view, we found it helpful to annotate the metro map with isolines at certain distances from the closest stations. During the iterative mapping, we distort the points of a regular grid in addition to the street-level data. To this end, we apply the mapping function to the single grid points, which results in a distorted grid as can be seen in Figure 8.7.

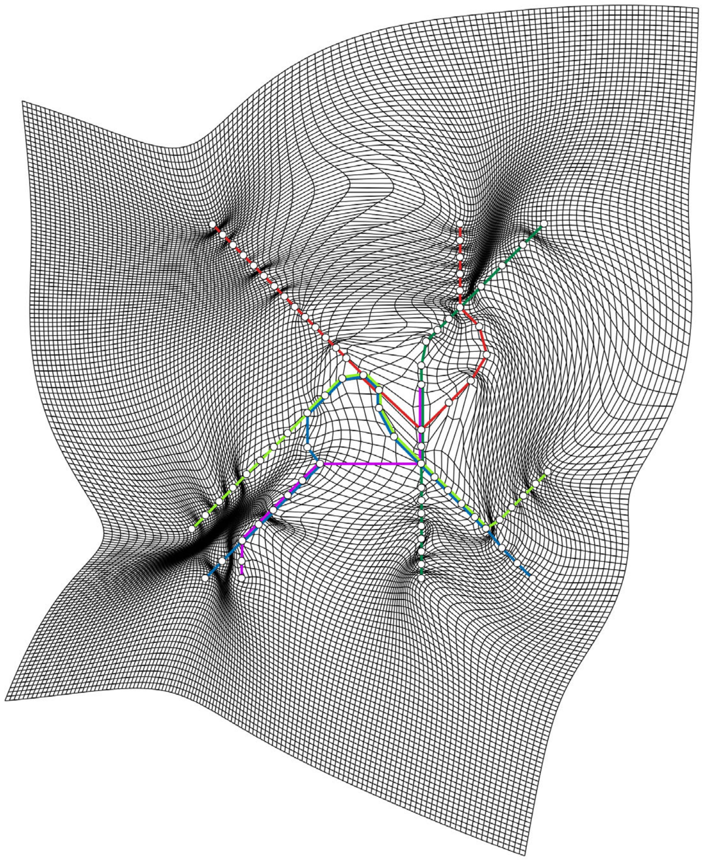

Figure 8.7: A regular grid distorted with the same mapping as the one used for warping the Washington data in Figure 8.3. The grid is used to render isolines around the stations in order to aid in the understanding of the geographical distance relations.

We can then visualize the real world distances from points in the map to the next station by first calculating the distance of every regular grid point to the closest station. Then, applying a marching squares algorithm to the distorted grid yields isolines denoting equal distances to the closest station in the real world.

While for the undistorted grid, these isolines consist of circular shapes centered around each station, the shapes are more complex after the distortion, as can be seen in Figure 8.8. For example, two stations, which are close to each other, are connected by an hourglass-like structure. Moreover, a gap between these structures indicates large distances between the corresponding stations.

This way, for example, the appropriate selection of the nearest station to a specific destination can be supported.

8.4 Conclusion

We presented a method to annotate schematic transportation maps with street-level information of the respective city. With the introduced *Warping Zoom* we attain a dynamic map that is as applicable for street-level navigation as for navigation in a

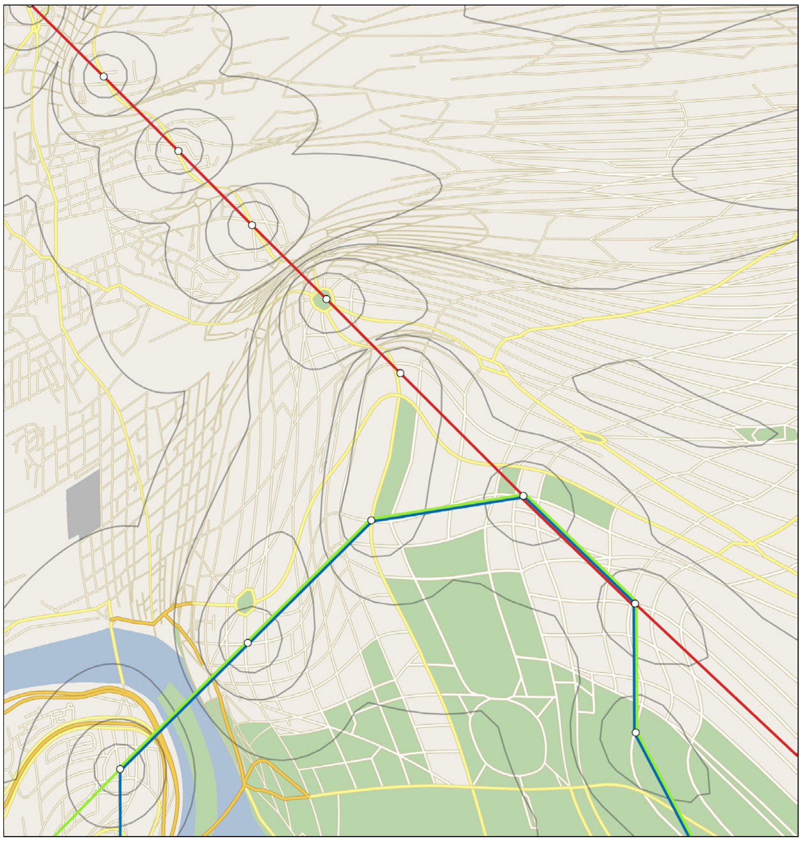

Figure 8.8: Isolines around stations. The lines are at constant geographic distance to the closest station, so that nearby stations are connected by blob-like shapes, while gaps between these shapes indicate large distances.

public transportation system. When the user switches navigational spaces by leaving the station of a public transportation system, switching map layouts is advisable as well. The reason is, that the user needs to read out disparate relations by switching navigational spaces: connected stations versus geographical relations of streets and city details. Our zooming technique implements this switching of layouts by continuous warping between the distorted schematic compound map and the geographical compound map. Only applying an ad-hoc map switch apart from zooming would require the user to find the old position on the new map. In addition the user might lose orientation. The in-between layouts, which are partially distorted/schematized, can support more complex navigation scenarios, for example, in a situation where a trade-off between the readability of walking distances and connectivity between metro stations has to be made. We also presented an adaption of fisheye views to achieve a detail-in-context visualization for combined navigational tasks.

Chapter 9

Conclusion and Outlook

This work describes new methods to map two-dimensional euclidean information to interactive displays in order to help to ameliorate the detail-in-context problem. The major influences for this undertaking, which are described in Chapter 2, are the growing body of very detailed information, artists' and scientists' concern with and methods for the illustration of such very complex data, and also brain science, which tells us how our own visual system deals with detail and context. Chapter 3 summarized a manifold of interactive methods developed to interact with detailed data, especially the distortion-oriented detail-in-context approaches.

After a careful mathematical analysis in Chapter 4, it was shown that a fundamental problem inherent in the strong magnification of details in their context is the introduction of anisotropic compression, which can render local shapes unrecognizable. One of the techniques to solve this problem is the application of the complex logarithm as a mapping, a function which is analytic and therefore introduces no anisotropic compression.

The challenges in this mapping are the questions of how to make it intuitively understandable for viewers, how to interact with the renditions, and how to speed up the rendering process sufficiently for interactive display employing modern graphics hardware. For the first application, vectorized geometric data, these questions are answered in Chapter 5, amongst others with the presentation of a smooth transition between complex logarithmic and euclidean layouts via complex root functions, direct mouse interaction by dragging points, and the use of vertex shaders. For the second application, pixelated aerial imagery of the whole earth, another intuitive connection between complex logarithmic views and central perspective, improved mouse interaction, and a sophisticated clipmapping-like rendering approach using fragment shaders were used, as is detailed in Chapter 6.

The developed new perspective of our world was incorporated into an interactive artistic and research installation in collaboration with the Center for Art and Media in Karlsruhe and the Human-Computer Interaction Group in Konstanz. The installation, which is described in Chapter 7 was successfully shown to the broad public on several events, and was a great experience for the author of this thesis.

Another successful cooperation with the Algorithmics Group and the Center for Junior Research Fellows at the University of Konstanz, which is described in Chapter 8 lead to the development of warping techniques for the navigation of complex public transportation networks. The connection lies in the properties of the underlying data, and therefore similar challenges like for complex logarithmic views.

The work described in this thesis was very rewarding, because it draws on a broad spectrum of influences; not only on art and scientific visualization, but also on cartography, perception and computer graphics. The potential of the acquired knowledge is by no means exhausted. Consequently, in the next and last section of this thesis, open questions and directions for future development of the described techniques are shown up.

9.1 Future Work

The developed method of map warping certainly magnifies several areas, but in the process introduces anisotropic compression. Complex logarithmic views can magnify only one center of interest without that flaw. It is an open question how to apply similar conformal mappings for two or more focus points, or even along lines of interest, such as for route descriptions. The potential of complex analysis in this respect has only begun to be explored.

Another promising extension would be the inclusion of three-dimensional data for the renditions. The use of geographic geometry data seems relatively straightforward, but the access to adequate data for the whole planet is still very restrictive due to the huge cost associated with the acquisition of such data. Geographic data is also more or less two-and-a-half dimensional, organized in a planar fashion with an additional elevation. The use of very complex fully three-dimensional data would need different mapping methods, but would open up very different application subjects, for example from biology, electronics design, and the more abstract visualization layouts in information visualization.

For the extension to the third dimension, it is interesting to note that modern graphics hardware not only continuously gets faster and faster, but also offers new capabilities with every next generation. The rather new geometry shaders, which nowadays are common in consumer hardware, offer exciting potential for the acceleration of the necessary refinement and distortion operations for quick rendering. They also would improve an implementation of the method in Chapter 5, enabling the rendering of much more complex data.

The methods in this work also might profit from a careful evaluation. The tasks for such an evaluation should be quick, successive navigation to different very small details in a large context for the complex logarithmic views, and making decisions about the fastest way to get from one point in a city to another using public transportation for map warping.

Concerning the latter, it is not clear whether the used mapping functions are optimally suited for our perception of geographical information. They stress the importance of keeping angles locally intact, maybe at the expense of readability for extreme distortions.

Although we think existing handmade metro map layouts are still clearly superior to automatically generated ones, merging our method with a method for the automatic generation of metro map layouts seems promising. This way, it might be possible to find an even better compromise between detail and schematization, avoiding extreme distortions in the street-level map and hard to read configurations in the superimposed schematic representation.

Last but not least, the application of labeling algorithms could make both approaches much more usable.

Bibliography

[1] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publ., New York, 9. dover pr. edition, 1972.

[2] M. Balzer and O. Deussen. Exploring relations within software systems using treemap enhanced hierarchical graphs. In *Proceedings of the 3rd IEEE International Workshop on Visualizing Software for Understanding and Analysis 2005 (VISSOFT 2005)*, pages 89–94. IEEE Computer Society, 2005.

[3] M. Balzer and O. Deussen. Voronoi treemaps. In *Proceedings of the IEEE Symposium on Information Visualization*. IEEE Computer Society, 2005.

[4] M. Balzer, A. Noack, O. Deussen, and C. Lewerentz. Software landscapes: Visualizing the structure of large software systems. In *Joint Eurographics and IEEE TCVG Symposium on Visualization*, pages 261–266, 2004.

[5] S. Björk. Flip zooming - the development of an information visualization technique. In *Gothenburg Studies in Informatics, Report 19, October 2000, ISSN 1400741X*, 2000.

[6] K. Boeke. *Cosmic View - The Universe in 40 Jumps*. John Day Company, Inc., 1957.

[7] K. Boeke. *Zoom in 40 Schritten durch den Kosmos - Ein Bilderbuch der Größenverhältnisse vom Atomkern bis zur Unendlichkeit*. WeltZeit-Verlag, 1982.

[8] I. N. Bronstein, K. A. Semendjajew, G. Musiol, and H. Mühlig. *Taschenbuch der Mathematik*. Verlag Harri Deutsch, Thun, Frankfurt am Main, 1993.

[9] J. Böttger, M. Balzer, and O. Deussen. 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):845–852, September/October 2006.

[10] J. Böttger, U. Brandes, O. Deussen, and H. Ziezold. Map warping for the annotation of metro maps. *IEEE Computer Graphics and Applications*, 28(5):56–65, 2008.

[11] J. Böttger, U. Brandes, O. Deussen, and H. Ziezold. Map warping for the annotation of metro maps. In *Proceedings VGTC Pacific Visualization Symposium*, pages 199–206, March 2008.

[12] J. Böttger, M. Preiser, M. Balzer, and O. Deussen. Detail-in-context visualization for satellite imagery. *Computer Graphics Forum*, 27(2):587–596, 2008.

[13] M. S. T. Carpendale. *A Framework for Elastic Presentation Space.* Ph.D. dissertation, Simon Fraser University, Burnaby, BC, Canada, 1999.

[14] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. 3-dimensional pliable surfaces: For the effective presentation of visual information. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, Information Navigation, pages 217–226, 1995.

[15] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. Multi-scale viewing. In *ACM SIGGRAPH 96 Visual Proceedings: The art and interdisciplinary programs of SIGGRAPH '96*, pages 149–149. ACM Press, 1996.

[16] P. M. Daniel and D. Whitteridge. The representation of the visual field on the cerebral cortex in monkeys. *Journal of Physiology*, 1961.

[17] C. Eames and R. Eames. Powers of ten, 1977. Pyramid Films.

[18] N. Elmqvist, N. Henry, Y. Riche, and J.-D. Fekete. Melange: space folding for multi-focus interaction. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1333–1342, New York, NY, USA, 2008. ACM.

[19] R. Fernando and M. J. Kilgard. *The Cg Tutorial: The Definitive Guide to Programmable Real-Time Graphics.* Addison-Wesley Professional, Boston, Mass., 2003.

[20] J. Foley, A. van Dam, S. Feiner, and J. Hughes. *Computer Graphics: Principles and Practice, Second Edition in C.* Addison-Wesley, 2nd edition, 1996.

[21] K. Garland. *Mr Beck's Underground Map.* Capital Transport Publishing, 38 Long Elmes, Harrow Eeald, Middlesex, 1994.

[22] Google Inc. Google Earth. http://earth.google.com/.

[23] Google Inc. Google Maps. http://maps.google.com/.

[24] J. R. Gott III, M. Juric, D. Schlegel, F. Hoyle, M. Vogeley, M. Tegmark, N. Bahcall, and J. Brinkmann. A map of the universe. *The Astrophysical Journal*, 624:463–484, May 2005.

[25] O. Greuel and H. Kadner. *Komplexe Funktionen und konforme Abbildungen*, volume 9 of *Mathematik für Ingenieure, Naturwissenschaftler, Ökonomen und Landwirte.* BSB B.G. Teubner Verlagsgesellschaft, 1982.

[26] K. Harmon. *You are here: personal geographies and other maps of the imagination.* Princeton Architectural Press, New York, 2004.

[27] S.-H. Hong, D. Merrick, and H. A. do Nascimento. Automatic visualisation of metro maps. In *Journal of Visual Languages and Computing*, volume 17, pages 203–224. ELSEVIER, 2006.

[28] D. H. Hubel. *Eye, Brain, and Vision*. W. H. Freeman, New York, 1988.

[29] B. Jenny. Geometric distortion of schematic network maps. *Bulletin of the Society of Cartographers*, 40:15–18, 2006.

[30] B. Johnson and B. Shneiderman. Tree-maps: a space-filling approach to the visualization of hierarchical information structures. In *Visualization, 1991. Visualization '91, Proceedings., IEEE Conference on*, pages 284–291, 1991.

[31] N. Kadmon and E. Shlomi. A polyfocal projection for statistical surfaces. *The Cartographic Journal*, 15(1):36–41, June 1978.

[32] T. A. Keahey. *Nonlinear Magnification*. PhD thesis, Indiana University Computer Science, Dec. 1997.

[33] T. A. Keahey. The generalized detail-in-context problem. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 44–51, Research Triangle Park, NC, USA, October 1998. IEEE Computer Society.

[34] T. A. Keahey and E. L. Robertson. Techniques for non-linear magnification transformations. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 38–45, San Francisco, CA, USA, October 1996. IEEE Computer Society.

[35] T. A. Keahey and E. L. Robertson. Nonlinear magnification fields. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 51–58, Phoenix, AZ, USA, October 1997. IEEE Computer Society.

[36] A. Klippel and L. Kulik. Using Grids in Maps. *Theory and Application of Diagrams: First International Conference, Diagrams 2000, Edinburgh, Scotland, Uk, September 2000: Proceedings*, 2000.

[37] W. A. König, H.-J. Bieg, T. Schmidt, and H. Reiterer. Position-independent interaction for large high-resolution displays. In *IADIS International Conference on Interfaces and Human Computer Interaction 2007.*, pages 117–125. IADIS Press, 2007.

[38] W. A. König, J. Böttger, N. Völzow, and H. Reiterer. Laserpointer-interaction between art and science. In *13th international Conference on Intelligent User Interfaces. IUI '08.*, pages 423–424, New York, NY, USA, 2008. ACM.

[39] J. Lamping and R. Rao. Laying out and visualizing large trees using a hyperbolic space. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, Visualization I, pages 13–14, 1994. TechNote.

[40] J. Lamping, R. Rao, and P. Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 401–408, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

[41] Y. K. Leung and M. D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. *ACM Trans. Comput.-Hum. Interact.*, 1(2):126–160, 1994.

[42] H. Lieberman. A multi-scale, multi-layer, translucent virtual space. In *International Conference on Information Visualization*. IEEE, September 1997.

[43] H. Lorenz, M. Trapp, M. Jobst, and J. Döllner. Interactive multi-perspective views of virtual 3d landscape and city models. In *11th AGILE International Conference on GI Science*. SPRINGER, May 2008.

[44] J. D. Mackinlay, G. G. Robertson, and S. K. Card. The perspective wall: Detail and context smoothly integrated. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 173–179, New Orleans, LA, USA, April 1991. ACM Press.

[45] Microsoft Corporation. Virtual Earth. http://www.microsoft.com/virtualearth/.

[46] G. E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine*, 1965.

[47] M. Morgan. *The Space Between Our Ears*. Weidenfeld & Nicolson, London, UK, 2003.

[48] P. Morrison, P. Morrison, C. Eames, and R. Eames. *Powers of Ten: Dimensions between Quarks and Galaxies*. Heidelberg: Spekt. Wiss., 1984.

[49] T. Needham. *Visual Complex Analysis*. Oxford University Press, 1997.

[50] L. Nye and the Exploratorium Visualization Laboratory. Zoom into the human bloodstream. published online, 2008. available at http://www.nisenet.org/viz_lab/illustrations.

[51] M. Nöllenburg and A. Wolff. A mixed-integer program for drawing high-quality metro maps. In P. Healy and N. S. Nikolov, editors, *Graph Drawing, Limerick, Ireland, September 12-14, 2005*, pages 321–333. Springer, 2006.

[52] G. Polya and G. Latta. *Complex Variables*. John Wiley & Sons, 1974.

[53] M. Preiser. Interaktive kartographische Abbildungen mit starken Größenunterschieden. Master's thesis, University of Konstanz, March 2008.

[54] R. Rao and S. K. Card. 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*, pages 318–322, Boston, MA, USA, April 1994. ACM Press.

[55] D. F. Reilly and K. M. Inkpen. Map Morphing: Visualizing relationships between map views. *Proceedings of Graphics Interface*, 2004.

[56] D. F. Reilly and K. M. Inkpen. White rooms and morphing don't mix: setting and the evaluation of visualization techniques. *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 111–120, 2007.

[57] G. G. Robertson and J. D. Mackinlay. The document lens. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 101–108, Atlanta, GA, USA, November 1993. ACM Press.

[58] D. Ruprecht and H. Müller. Image warping with scattered data interpolation methods. Technical Report 443, FB Informatik LS VII, 1992.

[59] D. Ruprecht, R. Nagel, and H. Müller. Spatial free-form deformation with scattered data interpolation methods. *Computers and Graphics*, 19(1):63–71, 1995.

[60] San Francisco Cityscape. Spider Map. http://sfcityscape.com/maps/spider.html.

[61] M. Sarkar and M. H. Brown. Graphical fisheye views. *Commun. ACM*, 37(12):73–83, December 1994.

[62] S. Schaefer, T. McPhail, and J. Warren. Image deformation using moving least squares. *ACM Transactions on Graphics*, 25(3):533–540, 2006.

[63] E. L. Schwartz. *Visual Science and Engineering: Models and Applications*. Marcel Dekker, 1993.

[64] E. L. Schwartz. Topographic mapping in primate visual cortex: Anatomical and computational approaches. In D. Kelly, editor, *Visual Science and Engineering: Models and Applications*, volume 43 of *Optical Engineering*. New York : Marcel Dekker, 1994.

[65] A. Seoane, J. Taibo, L. Hernandez, R. López, and A. Jaspe. Hardware-independent clipmapping. In *Proceedings of the International Conference on Computer Graphics, Visualization and Computer Vision*, pages 177–183, Plzen, Czech Republic, January 2007.

[66] G. W. Shirah and H. G. Mitchell. Nasa's great zooms: A case study. In *VIS '02: Proceedings of the conference on Visualization '02*, Washington, DC, USA, 2002. IEEE Computer Society.

[67] B. Shneiderman. Direct manipulation: A step beyond programming languages. *IEEE Computer*, 16(8):57–69, 1983.

[68] D. Shreiner. *OpenGL Reference Manual*. Addison Wesley, 1999.

[69] J. P. Snyder. *Map Projections: A Working Manual*, volume 1395 of *U.S. Geological Survey Professional Paper*. United States Government Printing Office, Washington, D.C., USA, 1987.

[70] J. P. Snyder. *Flattening the Earth*. The University of Chicago Press, Chicago and London, 1993.

[71] R. Spence. *Information Visualization*. Addison-Wesley, Harlow,Munich, 2005.

[72] R. Spence and M. Apperley. *Data base navigation: an office environment for the professional*, pages 333–340. Morgan Kaufmann, San Francisco, CA, USA, 1999.

[73] S. Steinberg. View of the world from 9th avenue. The New Yorker Magazine Cover, March 29th 1976.

[74] J. Stott and P. Rodgers. Metro Map Layout Using Multicriteria Optimization. In *Proceedings 8th International Conference on Information Visualisation*, pages 355–362. IEEE, July 2004.

[75] C. C. Tanner, C. J. Migdal, and M. T. Jones. The clipmap: a virtual mipmap. In *Proceedings of the annual conference on Computer graphics and interactive techniques*, pages 151–158, Orlando, FL, USA, July 1998. ACM Press.

[76] B. Tiddeman, N. Duffy, and G. Rabey. A general method for overlap control in image warping. *Computers and Graphics*, 25(1):59–66, 2001.

[77] Transport for London. Bus Maps. http://www.tfl.gov.uk/tfl/gettingaround/maps/buses.

[78] U.S. Census Bureau. Topologically Integrated Geographic Encoding and Referencing system (TIGER). http://tiger.census.gov/.

[79] C. Ware and M. Lewis. The dragmag image magnifier. In *CHI '95: Conference companion on Human factors in computing systems*, pages 407–408, New York, NY, USA, 1995. ACM.

[80] J. A. Wise. The ecological approach to text visualization. *J. Am. Soc. Inf. Sci.*, 50(13):1224–1233, 1999.

[81] A. Wolff. Drawing subway maps: A survey. *Informatik – Forschung & Entwicklung*, 22(1):23–44, 2007.