Selected Trends in Scientific Visualization

Lukas Mroz and Helwig Löffelmann and Eduard Gröller

Institute of Computer Graphics, Vienna University of Technology,
Karlsplatz 13/186/2, A-1040 Wien, Austria

ABSTRACT

Visualization became an important branch of scientific research during the past decade. Some topics of visualization, e.g., flow visualization and volume visualization, already provide lots of high-quality solutions to the most important problems. Others such as information visualization, are rather young disciplines which continuously provide lots of new advances. One such topic is the visualization over the Internet. Quite a number of approaches emerged during the last two years which facilitate this idea and embed the World Wide Web into the visualization process. Another trend in visualization is collaborative visualization within an augmented reality setup. Users exploit real 3D interaction and augmented reality to combine communication and investigation during research or education. Thirdly, the visualization of multi-dimensional and multi-modal data is a rather hot challenge to current scientific research. Coping with the restriction of just a few dimensions for representing visualization results is a big open problem in visualization and currently subject to eager research.

1. INTRODUCTION

When computer graphics is applied to scientific data for purposes of gaining insight and testing hypotheses, we speak of scientific visualization. As a field of science, scientific visualization has produced a large number of algorithms and mathematical results to which this activity gives rise.

Generally speaking, visualization comprises methods and techniques to generate images out of numbers.[1,2] While computers are well-suited to manipulate and process large amounts of data, the results of such calculations often have to be transformed into visual information to be easily accessible to the human analyzer. Visualization of scientific data proceeds in several stages, that together form the so called visualization pipeline (see Fig. 1).

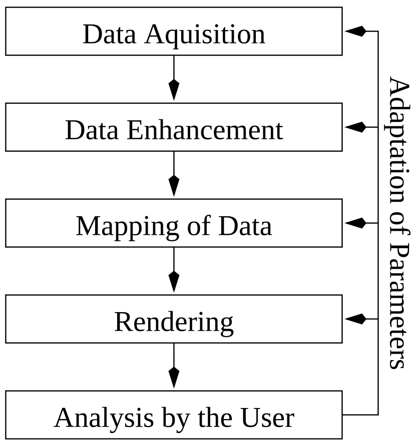

Figure 1. The visualization pipeline

At the beginning of the visualization pipeline is the data acquisition procedure. Data may originate from experiments, observations, or numerical simulations. Data enrichment and enhancement is used on the initial raw data to produce well-behaved data sets or derived data as required by an application. These operations may include interpolation, noise reduction, filtering, smoothing, and gradient calculation. Afterwards a mapping of the scientific data onto the

E-mail: { mroz | helwig | groeller} @cg.tuwien.ac.at – URL: http://www.cg.tuwien.ac.at/home/

attributes of a visual representation is performed. Attributes of this visual representation can be geometry, color, transparency, time, and material properties. Rendering, i.e., image generation by, for example, projecting the visual representation onto an image plane, is another step of the visualization pipeline. After analyzing an image the user may want to change various parameters in the different stages of the visualization pipeline to further increase the insight into the properties of a given scientific data set.

This paper gives an overview and some new ideas for three rapidly developing areas in scientific visualization research. First, existing and new ideas for the visualization of multi-dimensional data are explained. These also enable the exploration of high-dimensional dynamical systems. Then we will show how the Internet can be used to enable people to use visualization tools at remote sites, without the need to invest their money in high grade hardware, or care about the implementation of specific software. As third part we will demonstrate the impact of virtual reality techniques to the paradigms of scientific visualization collaboration. See-through HMDs enable people to cooperate on virtual visualization results in 3D space.

2. VISUALIZATION OF MULTI-DIMENSIONAL DATA

Multi-parameter data sets are becoming more and more common in many scientific disciplines. Depending on their origin, several different aspects of the data can be distinguished. Data sets can be discrete (like measured data) or continuous (like dynamical systems given by differential equations), they can be spatially coherent (such as medical images) or spatially incoherent (like census data). Data sets may consist of a collection of sampled data. Each sample is an n-dimensional data item. For scalar values, n equals 1, if each sample consists of three different variables, n equals 3. The m-dimensional space where samples are taken is described by an m-dimensional lattice. High-dimensional data structures L_n^m may have a large number $n > 3$ of sampled variables, or they can be sampled within a high-dimensional ($m > 3$) space.[3]

Two important goals of visualization techniques are the identification of individual parameters (e.g., what is the value of some data within a specific region?), and the detection of regions and correlation of variables (e.g., where do data points of a specific value reside?). Five different sets of visualization methods for high-dimensional data can be distinguished and are briefly discussed below. Visualization of high-dimensional data often uses combinations of these methods with interactivity as a way to introduce time as an additional dimension into the space where the data is explored.

2.1. Attribute Mapping

Attribute mapping is one of the most common methods to visualize high-dimensional data. This method uses one- or two-dimensional lattices to define some simple geometric primitives, e.g., contours or planes. The attributes of these geometric primitives can be used to visualize the remaining variables. The most often used attribute is the color of the geometric primitive. A major advantage of color coding is the easy calculation and interpretation. As this technique is very often used (e.g., weather forecast maps), many users are familiar with this kind of visualization. A disadvantage is that colors do not have a unique order, so color coded images have to show a color legend to allow an exact interpretation. Although at most three variables can be simultaneously encoded into color, this leads to images where different variables are not distinguishable any more. In realistic applications color coding is restricted to two variables. Another problem is the fact that about eight percent of the population suffers from some kind of color blindness.

2.2. Geometric Coding

Geometric coding, originally proposed by Pickett,[4] is used for displaying high-dimensional data on a low-dimensional lattice by displaying distinct geometric objects within the lattice and mapping the high-dimensional data to some geometric features or attributes of these objects. The most well knows examples for this technique are:

Glyphs are often utilized for interactive exploration of data sets. A glyph is a generic term describing a graphical entity whose shape or appearance is modified by mapping data values to some of its graphical attributes. An interactively positioned glyph adapts its appearance according to the underlying data.[5,6]

Icons are a generalization of a single pixel to higher dimensions having multiple perceivable features and attributes. The fact that shape and color are perceptually separable features is used for the display of color icons. Icons represent separable features by using color, shape and texture perception to code multiple variables. Levkowitz,[7] for example, proposed an icon which encodes six different parameters by color coding six different lines within a square icon.Individual variables are not recognizable any more, but correlation patterns appear.

2.3. Sonification

Sonification is another method of making more than three dimensions perceptible for a researcher. A sound is produced according to the mapped parameters.[8] Variables can be mapped to the loudness, the pitch, and even to the orchestration of the sound. A disadvantage of these methods is that the parameters of the sound are not independent, i.e., changing the pitch also influences the perceived loudness. Nevertheless, sonification is a good tool for visualizing high-dimensional data sets because it stimulates a different organ and thus may avoid overloading the human visual system.

2.4. Reduction of Dimension

Focusing techniques may be involved to select subsets, reduce dimension by projection, or to perform more general manipulation of the layout of information on the screen. Examples for subset selection techniques are panning, zooming,[9] and slicing.[10] Reduction of dimension can be achieved by simply projecting high-dimensional spaces along some axes into a low-dimensional space. Techniques for more general layout manipulation include methods for adapting to a user's point of interest such as fish eye views[11] and rooms.[12]

One consequence of focusing is that each view will only convey partial information about the data. This can be compensated by linking several focused visualizations. Linking can be done by sequencing several visualizations over time (guided tour) or by showing them in parallel simultaneously. This can be done using separate views or within one single image by using parallel coordinates[13] (see also Sect. 2.5) or dimensional stacking.[14] Dimensional stacking is a recursive projection method, where two dimensions are mapped on the horizontal and vertical axis, creating a discrete grid. Within each cell of this grid the process is applied in a recursive way until all dimensions are assigned. Another form of linking is the use of hierarchical axis[15] where n independent variables are hierarchically stacked on the x-axis and the analysis of one dependent variable is done on the y-axis. The pixel constraint on the x-axis is overcome by a set of hierarchical (color coded) symbols that represent the data not only at data points, but along whole data lines (or subspaces).

2.5. Parallel Coordinates

Parallel coordinates[13] represent dimensions on parallel axes which are arranged orthogonal to a horizontal line and uniformly spaced on the display. Each point of the data set corresponds to a polyline that intersects the parallel coordinate axes at the coordinate values of the data point. Characteristics of the data can be detected by looking at the patterns that are produced by the polylines. Linear dependencies between two coordinates of a set of points in n-space correspond for example to all polylines intersecting in a single point between the axes. By interactively brushing through the data set statistical characteristics like outliers and clusters can be recognized very easily.

2.6. High-Dimensional Dynamical Systems and Visualization

A lot of natural phenomena can be approximately described by differential equations. Scientists in many fields, e.g., in chemistry, physics, biology, economy, medical research, and other fields, investigate models called dynamical systems. Each differential equation describes the change of one state variable, thus a set of n differential equations defines the behavior of n state variables describing an n-dimensional dynamical system. For each "sampled" set of state variables the differential equations give an n-dimensional vector describing the direction of the flow at the specific point. Thus the discretized flow described by n differential equations forms a vector field of dimension n, where each vector itself is of dimension n. This can be interpreted as an L_n^n data set.

The main difference to measured "real world" data is that a dynamical system typically describes a (maybe) complex but smooth flow. The behavior of this flow is entirely determined by its topology. Visualization methods have to reveal this information to the scientist. Fig. 2 shows the visualization methods described and some advantages and disadvantages. As can be seen all the methods focus on the representation of statistical quantities and do not give any information on the underlying flow topology.

	attribute mapping	glyphs	parallel coordinates	reduction of dimension
dimension	$n \leq 3$	$n > 3$	$n > 3$	$n \geq 3$
identification of individual parameter	only for one variable	only for few glyphs	good for low quantity of data	depends on visualization method
discontinuities, trends, outliers	good	good	good	good
local extrema	possible (depends on color coding scheme)	possible (e.g., shape coding)	good	possible
correlations	possible	possible	good (for neighbouring axes)	

Figure 2. Methods for the visualization of high-dimensional data (similar to a scheme by Schumann[16])

2.7. Visualizing High-Dimensional Dynamical Systems

In the next sections we present three rather new techniques for the visualization of high-dimensional dynamical systems.[17]

Extruded Parallel Coordinates

Extruded parallel coordinates are based on parallel coordinates. A trajectory is sampled during its evolution at discrete points in time and its coordinates are inserted as polylines in a parallel coordinate system. Instead of using the same coordinate system for each sample, we now move the parallel coordinate system along the third spatial axis. The polylines of the samples can be viewed as cross sections of a moving plane with a complex surface which defines the trajectory. The geometry of the surface can be generated and modified fast and easily, allowing an interactive exploration of trajectories in the dynamical system. All exploration methods used by parallel coordinates can be used as well since rotating the surface and parallel projecting it reveals exactly the parallel coordinate representation. This can be used as a starting point for the exploration of the trajectory. Clustering and correlation can be visually detected.[13] Rotating the surface a little bit reveals the evolution of the trajectory over time without any animation methods that would have to be used for parallel coordinates. Convergence or divergence can be observed by varying the starting coordinates of the trajectory slightly (cf. Fig. 3). The changing shape of the surface shows for each dimension if its attracted or repelled by some topological structure.

Linking with Wings

This idea is based on a new method of linking data. Two arbitrary dimensions of the high-dimensional system are selected and displayed as a two-dimensional trajectory within a base plane (the high-dimensional trajectory is projected into a two-dimensional subspace). The third dimension (along the z-axis) can now be used to display a third variable over the base trajectory. If the resulting three-dimensional trajectory is connected with the base trajectory this connection can be thought of as a wing on the base trajectory. This wing can be tilted at each point

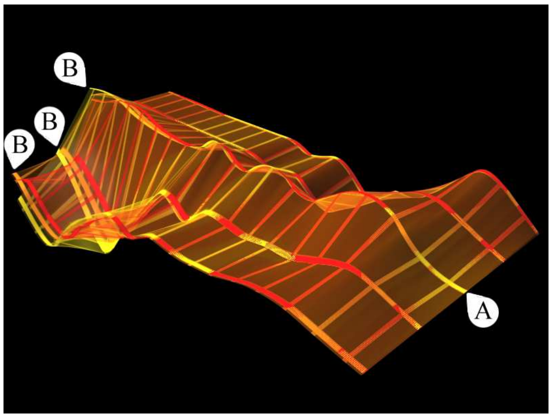

Figure 3. Small jitter at the starting point (A) of a chaotic attractor yields large effects after some time period (B)

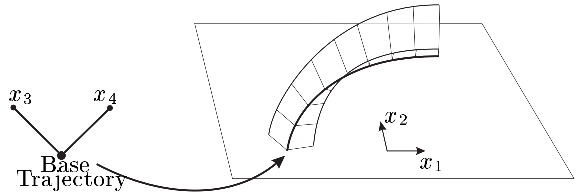

Figure 4. Two-dimensional base trajectory with two wings for the third and fourth dimension linked to it.

within a plane normal to the base trajectory. When different tilting angles are used several additional dimensions can be linked to the base trajectory on separate wings (see Fig. 4). Theoretically any number of wings can be added to display high-dimensional trajectories. As the number of wings increases occlusion might become a severe problem. To avoid this the wings can be rendered transparently with opaque trajectories at the top.

Three-dimensional Parallel Coordinates

Three-dimensional parallel coordinates are again based on the parallel coordinate method. As described in Sect. 2.5 the basic idea of parallel coordinates is to depict each variable on a one-dimensional space. All these one-dimensional spaces are put together within a two-dimensional space and linked with one-dimensional polylines. All information is packed in the two-dimensional space. Since the visualization of three-dimensional structures poses no problem we increased each dimension of the parallel coordinate method. The basic information now resides in separate two-dimensional spaces (planes) where two-dimensional trajectories are shown. These planes are combined within three-space and linked by surfaces which connect the separated projections of trajectories (see left part of Fig. 5).

The positioning of the planes is more flexible in comparison to the parallel lines of the parallel coordinate method. The planes can be moved and rotated within three-space to avoid occlusion in different regions of the structure. The right side of Fig. 5 shows two coinciding planes, where the connecting surface is bent normal to the planes to give a better overview of the linking. Using this technique, an arbitrary number of dimensions can be depicted.

If the structure of the trajectory is more complex, as for example in the case of a trajectory of a chaotic attractor, the resulting visualization can be rather crowded. Since the rendering of the presented structures is fast, interactive brushing methods can be used to overcome this problem. For instance the linking surface can be rendered transparently, and only a small temporal interval of the linkage is rendered opaquely. When the interval is moved corresponding parts of the trajectories can then be detected (see Fig. 6).

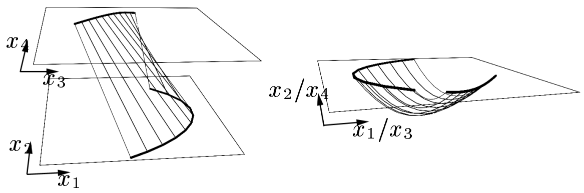

Figure 5. Linking parallel planes instead of lines extends the idea of parallel coordinates by one dimension (left); these parallel planes can coincide, with three-dimensional linking surfaces (right)

3. VISUALIZATION OVER THE INTERNET

One of the most impressing developments in communication technology during the past decade has been the evolution of the World Wide Web (WWW) based on the Internet. The combination of easy-to-use web browsers on multiple platforms and distributed servers provides a powerful facility to distribute information to a large number of users. We see the Internet as an excellent foundation to connect new developments in the field of scientific visualization and a large community of users, almost regardless of their hardware and software resources. Especially, we want to exploit the portable and transparent computing opportunities of the Internet and the ease-of-use of the WWW to

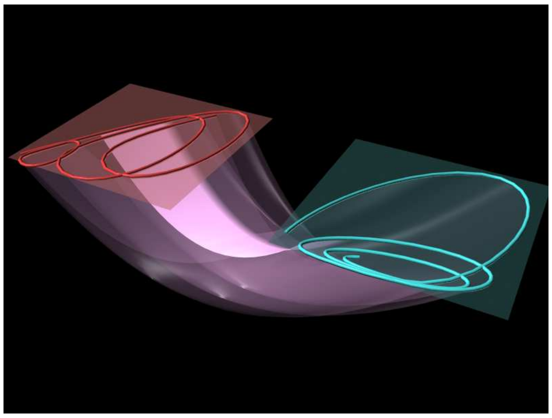

Figure 6. Complex linkage with highlighted time interval for the trajectory of a chaotic attractor

bring the benefits of advanced visualization techniques to users working with dynamical systems without access to expensive hardware and software resources. To be broadly accepted by the user community, a good solution must include several important features: A visualization system should be portable and run on various platforms ranging from UNIX workstations to desktop PCs. Java and VRML provide the power to realize these features. The system should offer a set of predefined visualization configurations covering common applications and presents the same intuitive and easy-to-learn interface on each platform. The system should be flexible, allowing the user to visualize his or her data in a customized way with a minimum amount of work. Furthermore, the investigation of the user data should be interactive, almost regardless of the user's desktop computer. Computation tasks may be distributed between client and server.

Both the users and the developers benefit from this approach. The users apply improved visualization techniques to investigate their problems without big expenses using a simple and easy-to-learn interface. The scientists developing new visualization techniques receive rapid feedback about the usefulness of their methods from a broad field of applications.

3.1. Existing Internet Applets

The visualization applets which can be found on the Internet can be subdivided into the following three classes:

Limited front-ends for existing systems

The client is used as a front-end to enter parameters and data only, and to display the results of the visualization as an image or via a VRML viewer. The actual visualization is performed on the server, in most cases using a general purpose visualization package, like AVS,[18] the IBM Data Explorer,[19] or the IRIS Explorer.[20] Although this solution is appropriate for providing special services to a limited number of researchers, it requires a powerful visualization server to cope with multiple requests for computationally expensive visualization procedures. An example for Internet visualization using this technique is the Air Quality Data Visualization,[21,22] which provides a visualization of levels of different pollutants at various places over Great Britain. The visualization service is realized as a combination of HTML forms, CGI scripts, and the IRIS Explorer as a visualization engine. The results are displayed at the client using a VRML viewer.

Applets implementing certain visualization techniques for demonstration or educational purposes

All calculations involved in the visualization process are performed on the client side. Some examples for this rather often used approach are the Experimental Volume Visualization Environment (EVOLVE),[23] or VizWiz.[24] The Java applet EVOLVE demonstrates volume rendering methods facilitating wavelet-based progressive refinement techniques. The user selects from some precomputed (wavelet compressed) data sets and parameters defining a visualization technique. The calculation of the visualization results is then performed by the client software. VizWiz implements a set of visualization techniques for volumetric and elevation datasets into a Java applet. Several of the techniques can be applied simultaneously to visualize example or user data.

VRML as a distribution medium for pre-calculated visualization results

The actual visualization is performed off-line, generating a set of geometric representations, which are converted into VRML and placed on the web. The Normal Mode Visualization[25] for example, visualizes the normal modes of molecular vibration for a sample molecule. The user selects an activating frequency band from a 2D plot in the upper part of the page and watches the resulting vibration mode of the molecule as an animated VRML 2.0 model in the lower part of the page.

3.2. An Advanced Approach to Bridge the Gap

We are currently developing a concept study for a platform independent visualization system for dynamical systems using Java and Java3D to achieve portability. The system should be easy-to-learn and -use, and enable low cost interactive exploration of dynamical systems on low- and high-end platforms. The user should be able to create visualization results and to explore his or her data, finally composing presentations of the research results. The viewer of the results should be able to interact with them by changing viewing parameters and/or additional parameters of the visualization, which have been specified by the creator. To increase the flexibility, the visualization tasks will be distributed between server and client. This allows the user to experiment with the system, without having to download it first, or to run computationally expensive modules on a high performance platform. As interactivity is crucial for efficient work, interactive response to parameter changes will be achieved by adapting the quality of the visualization during interactions to the available computational resources, using for example progressive refinement or levels of complexity techniques.

4. COLLABORATIVE VISUALIZATION USING AN AUGMENTED REALITY SETUP

As a highly interdisciplinary field, scientific visualization frequently requires experts with different background to cooperate closely. Many valuable insights only occur in face-to-face discussions over the relevant data.

Early adopters of virtual reality (VR) systems soon realized that one of the immediately useful applications comes from the field of scientific visualization, where scientists try to understand complex data sets and can benefit from true 3D, stereoscopy, and interactive exploration, e.g., GROPE,[26] the nanomanipulator,[27] and the virtual wind tunnel.[28] The need to support collaboration of human users lead in two directions: remote collaboration[29] and collaborative virtual environments (VEs) where users come together in one place, and can interact and communicate in a natural way.

In the latter category, two very successful approaches have been developed: The CAVE,[30] and the workbench, with two variants - Responsive Workbench[31] and Virtual Workbench (VW).[32] Both systems present stereoscopic images to the user via large display screens and LCD shutter glasses, and both systems support multiple users.

The **CAVE** is a small room composed of three projection walls and a projection floor, on which computer-generated images are displayed. Among the CAVE's advantages are high resolution, wide field of view, insensitivity towards lag for rotational head movements, and a strong feeling of immersion.

The **Workbench** is essentially a table on which computer generated images are projected, resulting in a typical setup used by, e.g., surgeons, engineers, and architects. The resource requirements are less demanding than those of the CAVE, and the horizontal workspace is very useful for manipulation with hand-held tools.

Both systems suffer from the drawback that true stereoscopic images usually can only be rendered for one "leading" user wearing the head tracker - the users have to remain close to the leading user, because distortions increase proportional to their distance to the tracked point of view. Applications in which users surround an object do not work in the CAVE and are only possible for two participants in the latest version of the Workbench.

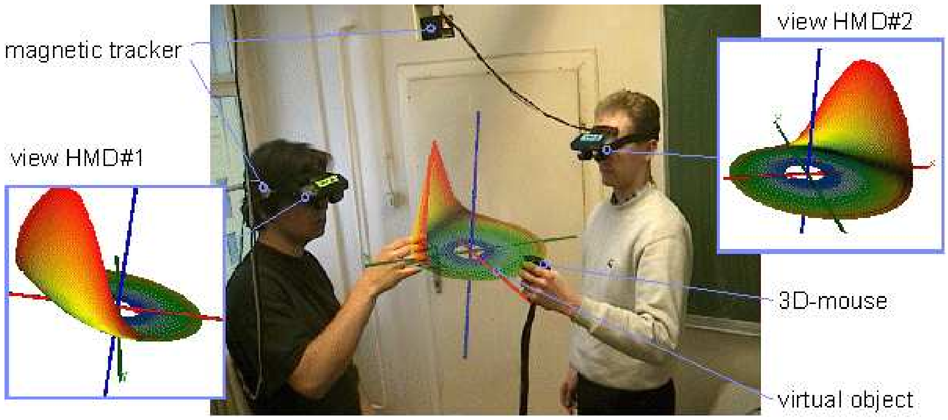

Figure 7. Hardware configuration of Studierstube (displaying Rössler attractor)

4.1. The "Studierstube" Approach

We propose a system that allows multiple collaborating users to simultaneously study three-dimensional scientific visualizations in a "study room" (German: "Studierstube", inspired by the classic play "Faust" by Goethe). Each participant wears an individually head-tracked see-through HMD providing a stereoscopic real-time display (Fig. 7).

The use of individual displays and head-tracking for each participant allows stereoscopic, undistorted images to be presented to everyone. There are no constraints regarding the viewpoint, e.g., unlike the Workbench, users may sit on opposite sides of the table, which in combination with the see-through property of the HMDs allow users to see each other. The HMDs we use - Virtual I-O i-glasses - are very lightweight and unobtrusive, but only of limited resolution and small (30°) field of view. Rendering separate images for each user gives great flexibility in the choice of the presented image, but also makes the rendering effort proportional to the number of users, whereas CAVE and (single-user) Workbench require only a constant rendering effort. The defining features of Studierstube are:

Augmented props – we exploit the capabilities of AR to construct a three-dimensional user interface needed for controlling the presentation and possibly the simulation by introducing tracked real-world objects that combine physical items and overlaid computer graphics, such as the Personal Interaction Panel (PIP).[33] The PIP is used as a device for handling virtual objects or as a cutting plane.

Customized views – in addition to individual choice of viewpoint, customized views of the data are possible, e.g., one user may want to see stream lines added to the basic image, while another may not. Two users in the same room may see different aspects of the same object at the same time.

Usage of space – the space in the Studierstube can be used similarly to a CAVE (multiple users standing around), but also allows a workbench setup (users gathering at a blank table with free choice of seating).

Organisational advantages – while the cost of the Studierstube's hardware components are certainly higher than a conventional desktop visualization station consisting of only a graphics workstation, they are very conservative compared to a setup like the CAVE. This is particularly important as the potential users of the Studierstube - research groups - are typically operating on a tight budget. Furthermore, the setup consumes little space and is relatively easy transportable.

4.2. Visualization in Studierstube

To combine augmented reality and scientific visualization, a new integrated solution could be developed, but employing an existing, general-purpose scientific visualization system (in our case, AVS) allows a wider range of applications and eases development. Since this desktop-based system is not designed for the real-time requirements of AR, we use what Shaw et al.[34] call a decoupled simulation: The visualization system and the AR user interface (called display server in an analogy to X-Windows) run as completely independent processes, typically executing on separate machines and communicating with each other over a network.[35]

The system is composed of two loops: the display loop, a tightly coupled human-in-the-loop component, where real-time response is essential, and a loose coupling between display server and visualization application for the exchange of visual information. The display server continues to serve real-time graphics to the user while the application may work in background, delays of several seconds for recomputation are quite acceptable.

4.3. Evaluation

Our evaluation showed that investigating dynamical systems in the Studierstube has some significant advantages:

Rapid feedback for visualization mapping. It has been seen that a fast feedback loop allowing to trigger new calculations of the visualization mapping, e.g., the integration of a stream surface within the AR setup, speeds up and enhances the visualization and understanding process. The process of choosing an initial visualization mapping, then refining it converges much faster as researchers are better able to estimate their requirements.

Reduced abstraction. Abstract mathematical constructs such as phase space, attractors etc. can be deterring in their complexity, even for experts. The intuitive 3D user interface of Studierstube greatly simplifies the exploration of these structures and allows the user to get familiar with the 3D representation of the dynamical systems much faster than via a 2D display.

Educational settings. The collaborative setting of Studierstube is of great utility for explaining dynamical systems. Tutorial situations are strongly supported by the capability to freely choose one's viewpoint and the unhindered natural communication.

Customized views. When multiple researchers investigate dynamical systems together, Studierstube's capability of separating customized views allows individual markers and visualization icons to be added. These icons can be kept private to avoid display clutter, or they can be shared with others.

5. SUMMARY AND CONCLUSIONS

This paper give an overview over some rapidly developing areas in scientific visualization research. New visualization methods for multi-dimensional data are explained. They enable the exploration of high-dimensional dynamical systems. It was shown how the Internet can be used to enable people to use visualization tools at remote sites, needing neither expensive hardware nor any complicated visualization software. Finally, the impact of virtual reality techniques to the possibilities of scientific visualization collaboration have been demonstrated. See-through HMDs enable people to cooperate on virtual visualizations.

The still accelerating development of fast computers and new interaction devices will certainly inspire the creation of further visualization techniques and interfaces and will help make this branch of computer graphics to one of the most useful application fields. New innovative ideas and user friendly interaction paradigms are still to be invented. Scientific visualization is very promising, but still in its childhood.

ACKNOWLEDGEMENTS

The authors thank Rainer Wegenkittl, Anton Fuhrmann, and Dieter Schmalstieg for their contributions to this article. Parts of this work have been supported by the Austrian Science Foundation (FWF) under project no. P-12074-MAT.

REFERENCES

1. H. Hagen, H. Müller, and G. M. Nielson, eds., *Focus on Scientific Visualization*, Springer, 1993.
2. L. Rosenblum *et al.*, eds., *Scientific Visualization, Advances and Challenges*, Academic Press, 1994.
3. R. D. Bergeron and G. Grinstein, "A reference model for the visualization of mult-dimensional data," *Proceedings of Eurographics '89* , 1989.
4. R. M. Pickett, "Visual analysis of texture in the detection and recognition of objects," *Picture Processing and Psycho-Pictories* , 1970.
5. G. D. Kerlick, "Moving iconic objects in scientific visualization," *Proc. IEEE Vis. '90* , pp. 124–129, 1990.

6. W. C. de Leeuw and J. J. van Wijk, "A probe for local flow field visualization," *Proceedings IEEE Visualization '93* , pp. 117–123, October 1993.
7. H. Levkowitz, "Color icons: Merging color and texture perception for integrated visualization," *Proceedings of IEEE Visualization '90* , pp. 164–170, 1990.
8. G. Grinstein and S. Smith, "The perceptualization of scientific data," *Proceedings of SPIE '90* , 1990.
9. A. Buja, C. Hurley, and J. A. McDonald, "Elements of a viewing pipeline for data analysis," in *Dynamic Graphics for Statistics*, Wadsworth and Brooks, 1988.
10. G. W. Furnas, "Dimensionality constraints on projection and section views of high dimensional loci," *Computer Science and Statistics: Proceedings of the 20th Symposium on the Interface* , 1988.
11. G. W. Furnas, "Generalized fisheye views," *Proceedings CHI '86, Human Factors in Computing Systems* , 1986.
12. D. A. Henderson and S. K. Card, "Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface," *ACM Transactions on Graphics* **5**(3), pp. 211–243, 1986.
13. A. Inselberg and B. Dimsdale, "Parallel coordinates: a tool for visualizing multidimensional geometry," in *Proceedings IEEE Visualization '90*, pp. 361–378, July 1990.
14. J. LeBlanc, M. O. Ward, and N. Wittels, "Exploring n-dimensional databases," *Proceedings of IEEE Visualization '90* , pp. 230–239, 1990.
15. T. Mihalisin, J. Timlin, and J. Schwegler, "Visualization and analysis of multi-variate data: A technique for all fields," *Proceedings of IEEE Visualization '91* , pp. 171–178, 1991.
16. H. Schumann, "Visualisierung wissenschaftlicher Daten," 1993. Lecture Notes.
17. R. Wegenkittl, H. Löffelmann, and E. Gröller, "Visualizing the behavior of higher dimensional dynamical systems," in *Proceedings IEEE Visualization '97*, October 1997.
18. C. Upson, T. Faulhaber, Jr., D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam, "The application visualization system: A computational environment for scientific visualization," *IEEE Computer Graphics and Applications* **9**, pp. 30–42, July 1989.
19. B. Lucas, G. D. Abram, N. S. Collins, D. A. Epstein, D. L. Gresh, and K. P. McAuliffe, "An architecture for a scientific visualization system," in *Proceedings IEEE Visualization '92*, pp. 107–114, Oct. 1992.
20. D. Foulser, "Iris explorer: A framework for investigation," *Computer Graphics* **29**(2), pp. 13–16, 1995.
21. J. D. Wood, K. W. Brodlie, and H. Wright, "Visualization over the world wide web and its application to environmental data," in *Proceedings IEEE Visualization '96*, pp. 81–86, 1996.
22. "Air quality data visualization." `http://csindy12.leeds.ac.uk:8010 /jason/html/pollution.html`.
23. "Experimental volume visualization environment." `http://www.inf.ethz.ch /personal/lippert/EVOLVE/`.
24. C. K. Michaels and M. J. Bailey, "Vizwiz: A java applet for interactive 3D scientific visualization on the web," in *Proceedings Visualization '97*, pp. 261–267, October 1997.
25. "Normal mode visualization." `http://www.pc.chemie.th-darmstadt.de/vrml/vib/`.
26. F. Brooks Jr, M. Ouh-Young, J. J. Batter, and P. J. Kilpatrick, "Project GROPE - haptic displays for scientific visualization," *Proceedings of SIGGRAPH'90* **24**(4), pp. 177–185, 1990.
27. R. M. Taylor *et al.*, "The nanomanipulator: A virtual reality interface for a scanning tunneling microscope," *Proceedings of SIGGRAPH'93* **27**(3), pp. 127–134, 1993.
28. S. Bryson, "The virtual wind tunnel," *Proceedings of IEEE Visualization'91* (43), pp. 17–25, 1991.
29. S. Bryson, "The distributed virtual wind tunnel," *SIGGRAPH'93 Course* (43), pp. 3.1–3.10, 1993.
30. C. Cruz-Neira *et al.*, "Scientists in wonderland: a report on visualization applications in the cave virtual reality environment," *Proceedings of IEEE 1993 Symposium on Research Frontiers in Virtual Reality* , pp. 59–67, 1993.
31. W. Krüger, C. Bohn, B. Fröhlich, H. Schüth, W. Strauss, and G. Wesche, "The responsive workbench: A virtual work environment," *IEEE Computer* **28**(7), pp. 42–48, 1995.
32. U. Obeysekare *et al.*, "Virtual workbench – a non-immersive virtual environment for visualizing and interacting with 3D objects for scientific visualization," *Proceedings of IEEE Visualization'96* , pp. 345–350, 1996.
33. Z. Szalavári and M. Gervautz, "The personal interaction panel - a two handed interface for augmented reality," *Computer Graphics Forum (Proceedings of EUROGRAPHICS'97)* **16**(3), pp. 335–346, 1997.
34. C. Shaw, J. Liang, M. Green, and Y. Sun, "The decoupled simulation model for VR systems," *Proceedings of SIGCHI* , pp. 321–328, 1992.
35. A. Fuhrmann, H. Löffelmann, and D. Schmalstieg, "Collaborative augmented reality: Exploring dynamical systems," in *Proceedings IEEE Visualization '97*, October 1997.