

Adaptive Compression for Remote Visualization

Zoran Constantinescu^{*}, Monica Vlădoiu^{**}

^{*} ZealSoft Ltd., str. Tg. Neamț, nr. 11, Bucharest
e-mail: zoran@unde.ro

^{**} Petroleum –Gas University of Ploiești, Informatics Department, Bd. București, 39, Ploiești
e-mail: mmvladoiu@acm.org

Abstract

Remote visualization techniques that use client-server environments allow users to access large datasets. One possible solution for remote visualization is the use of compression techniques, in which images are generated and compressed at the servers' side and then the encoded images are transferred over a data network, decompressed and displayed at the clients' side. In this paper we propose an adaptive algorithm based on reinforcement learning for choosing one of the available compression methods in order to maximize the frame rate. Our experiments show that such an algorithm can work in a dynamic and uncertain environment, consisting of a visualization server, a visualization client, and a network for transferring the compressed images between the server and the client.

Key words: *adaptive compression, reinforcement learning, remote visualization*

Adaptive Compression for Remote Visualization

Remote visualization techniques that use client-server environments allow users to access large datasets. One possible solution for remote visualization is the use of compression techniques, in which images are generated and compressed at the servers' side and then the encoded images are transferred over a data network, decompressed and displayed at the clients' side. One of the problems in remote visualization is how to increase the frame rate for the user [7, 11]. A possible solution is to reduce the amount of data transferred over the network, in our case to choose an efficient compression algorithm. However, the better a compression algorithm is, the more computing is necessary for both compression and decompression, increasing the time needed to process the image. The choice of the most efficient compression depends also on the content of the image [14]. Therefore we propose an innovative intelligent adaptive compression method for selecting different image compression algorithms for remote visualization. The selection is made based on the performance of previously compressed frames and network transfer delays. A reinforcement learning technique to select the compression algorithm for each individual frame is used. The algorithm was tested using the SGI OpenGL Vizserver [25], but it can be easily adapted to other remote visualization systems.

Transferring the full data set to the researcher's desktop for visualization purposes is most of the time impossible, due to the lack of memory and storage space of local desktop computers [1, 2, 5]. Scientific Visualization research applies a client-server approach to this problem [10]. Remote visualization can be done using different strategies. In a first scenario, the server

renders the images and streams them to the client. In a second scenario, the server is doing some of the rendering calculations, such as geometry transformations or visibility determination, while the client is doing the final rendering. Another scenario has the client doing all the rendering computations. Each of these scenarios has tradeoffs. For example, performing the rendering completely on the client side requires high-end desktop computers, not always available to researchers. Performing some of the rendering on the server can greatly improve the visualization, however the low-end client resources may not provide sufficient power to finish the rendering in time [12].

In the rest of this paper we consider only the first scenario, where the server is doing all the computations including the rendering, and the client is responsible only with the display of the final image. Image streaming makes possible remote visualization using low-end desktop computers (*thin clients*), and can be made independently from any visualization algorithm used. However, image streaming can require significant network bandwidth. For example, if the resolution of displayed image is 640x512 pixels with 4 bytes for the RGB colors and the alpha channel, then the size of such an image is 1.25 MBytes. The maximum theoretical frame rate, which can be obtained using a 100 MBps bandwidth network, is 10 frames per second, considering that the full bandwidth could be used. If we consider remote visualization over a wide area network, then the achievable frame rate will be much lower.

One possible solution to this problem, when using image streaming over the network, is to use compression algorithms. The server renders the image, then compresses that image and sends it over the network. The client is then responsible for decompressing the encoded image and displaying it. Using different compression techniques for the images, the amount of data transferred over the network can be significantly reduced. How much an image can be compressed depends essentially on the image content [13, 21]. This means that by using different compression algorithm for the same image, different compression rates can be obtained. The problem is to find an automatic way of selecting the right compression algorithm. Different methods that are based on analyzing the image and on using the compression algorithm that gives the best-compressed size are presented in the literature [14, 15, 16, 17]. However, most of these methods can be used only for certain types of images.

In this work we present an algorithm for selecting the compression methods during a remote visualization session without analyzing the content of the image. We propose an adaptive algorithm for dynamically selecting one of the compression algorithms to be used for each individual frame. The selection is done using a reinforcement-learning algorithm, and it is based on different performance measures from the environment: past and present frame rates, compressed image sizes, compression times, estimated bandwidth. The compression method that increases the overall frame rate is chosen. However, from time to time, other compression methods are also used for short periods of time, in order to estimate the potential benefit of selecting them.

Reinforcement learning is a computational approach to learning whereby an agent tries to maximize the total amount of reward (the frame rate in our situation) it receives when acting with a complex, uncertain environment [18]. As opposed to other machine learning methods, in this method the learner is not told which actions to take, but instead must discover which actions yield the most reward by trying them. In many cases, the actions may affect not only the immediate reward, but also the next situation and, through that, all subsequent rewards.

Related Work

Renderer implementations exploiting image compression have mostly adopted relatively simple lossless schemes, which rely on frame differencing and run-length encoding. While these techniques can deliver acceptable frame rates over LANs, their compression ratios are highly

dependent on image content, and are insufficient in slower networks. SGI OpenGL Vizserver is a product developed by Silicon Graphics, Inc., to enable remote-visualization applications [26].

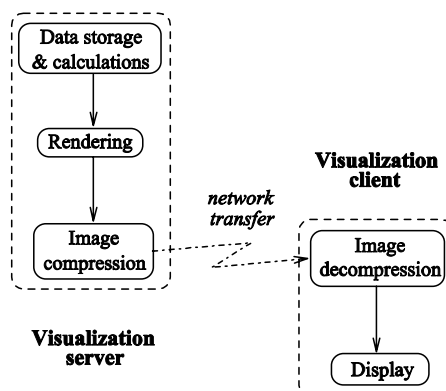


Fig. 1. Remote Visualization

Specifically, OpenGL Vizserver is designed to provide users remote access to graphics pipelines of Onyx2 Infinite Reality machines so that they may view rendered output from visualization applications at geographically remote locations while utilizing the powerful pipeline and memory of an Onyx2 machine located at a some centralized place. OpenGL Vizserver uses programmable compression modules to compress and decompress frames of the rendered scene. It comes with five standard modules (CCC, ICC, SCC, SICC, LCC) and an API that provides the capability to develop new modules with user defined functionality. Each compression module has the capability of taking advantage of frame-to-frame coherency inherent in most visualizations by implementing an inter-frame compression scheme where only the changing portions of each frame are compressed and sent to the clients. The CCC, ICC, SCC, and SICC compression modules implement lossy compression algorithms. These four schemes are derived from the Block Truncation Coding (BTC) algorithm that compresses a 4x4 pixel block down to two colors plus a 4x4 pixel mask. In addition to lossy compressors, there is also a lossless compression module called LCC. This preserves the original image quality while still saving bandwidth. In many cases the savings are as high as 4x without any reduction in image quality.

A similar framework exists which provides remote control to Open Inventor or Cosmo3D based visualization applications [4]. It allows transparent access to remote visualization capabilities and allows sharing of expensive resources. A visualization server distributes a visualization session to Java based clients by transmitting compressed images from the server frame buffer. Visualization parameters and GUI events from the clients are applied to the server application by sending CORBA (Common Object Request Broker Architecture) requests. Both of these two solutions require the user to explicitly select the compression algorithm to be used. In most of the situations, the user does not have any knowledge about the compression algorithm.

An adaptive compression algorithm for medical images was presented in [6]. The adaptive algorithm presented is based on a classification of digital images into three classes and followed by the compression of the image by a suitable compression algorithm. The content of the image is analyzed based on a validation of the relative number and absolute values of the wavelet coefficients. A comparison between the original image and the decoded image will be done by a difference criteria calculated by the wavelet coefficients of the original image and the decoded of the first and second iteration step of the wavelet transform.

Compression of images was used in [8, 9] for visualizing time varying volume data over a wide area network. The rendering was done on a remote parallel computer and compression of the images was used for significantly reducing the cost of transferring output images from the parallel computer to the local display. They used lossy compression methods combined with lossless compression methods, which were capable of providing acceptable image quality for

many applications, while retaining desirable properties such as efficient parallel compression and fast decompression. They experimented with different combinations of the JPEG, BZIP and LZO compression algorithms, and then selected the combination of JPEG and LZO as giving the best frame rates for their system.

Image Compression

The use of image compression algorithms can significantly improve the amount of data transmitted over the network. All compression algorithms are based on the same principle: compressing data by removing redundancy from the original data. Any non-random collection data has some structure, and this structure can be exploited to achieve a smaller representation of the data, where no structure is discernible. This is the case of using lossless compression algorithms. An important feature of image compression is that in many situations it can be lossy, being acceptable to lose image features to which the human eye is not sensitive. Images can be loss compressed by removing irrelevant information even if the original image does not have any redundancy [15, 23, 24].

Different image compression algorithms can be used for different types of images. Each type of image may feature redundancy, but they are redundant in different way. This is why any given compression method may not perform well for all images, and why different methods are needed to compress the different image types. The choice of the best algorithm is not trivial, most of the time requiring a certain experience with the algorithms. During a visualization session, the type of image can also change, making even more difficult to choose the appropriate algorithm. One important factor, which is important in choosing the compression algorithm, is the amount of computation needed for both compressing and decompressing the image. More efficient algorithms, capable of generating smaller compressed images are usually requiring more CPU power. This becomes very critical, especially for high-resolution images. There is a tradeoff between the amount of computation time needed to generate the compressed image and the amount of time used to transfer it over the network.

There are cases when an investment in a more efficient compression algorithm can result in a higher frame rate, especially when the remote visualization is done over low bandwidth networks. In many situations, the actual network bandwidth available, which can be used, is less than the maximum bandwidth. This is the case when the remote visualization is done without having a dedicated network connection between the visualization server and client, especially when using wide area networks for visualization over long distance [20]. An additional problem is that this available network bandwidth can change significantly during a remote visualization session. This can be due to other data traffic in the network.

For our study, we used four lossless compression algorithms. The choice was mainly made based on the performance of these algorithms for general image compression and the availability of optimal implementations as software libraries. The first algorithm (ZLIB) is the so-called "deflation" algorithm, which is used in the popular programs zip and gzip. This is a dictionary based compression method: it selects strings of symbols and encodes each string as a token using a dictionary. It is based on the LZ77 compression method combined with static Huffman encoding. The compression time and image sizes are pretty good, however for certain image type compression can be very poor. The second algorithm called Lempel-Ziv-Oberhumer (LZO), an optimized dictionary based method, which is more suited for real-time compression-decompression. It offers pretty fast compression and very fast decompression, however it favors speed over compression ratio. The resulting compressed images can be very large, thus increasing the transfer time over the network. The third algorithm used (BZIP2) is based on the Burrows-Wheeler method, which is a compression method using block sorting. The input stream is read block by block and each block is encoded separately as one string. The main idea

is to start with a string S of “ n ” symbols and to scramble (permute) them into another string L , which satisfies: (1) any area of L will tend to have a concentration of just a few symbols; (2) it is possible to reconstruct the original string S from L . The method is a general-purpose method, which works well on images and can achieve very high compression ratios. The disadvantage of this algorithm is that it requires a lot of computing, both compression and decompression being slow. Since the algorithm is compressing individual blocks independently, it is possible to use a parallel version of the compression to reduce the time. The last algorithm we used is a simple Run Length Encoder (RLE). The idea behind this approach is the following: if a data item d occurs n consecutive times in the input stream, replace the n occurrences with the single pair nd . This is well suited for certain types of images, with large areas containing the same pixel value. The size of the compressed stream depends on the complexity of the image. The more detail we have, the worse the compression is. The algorithm being extremely simple, very efficient implementations could be implemented. It is also well suited for a parallel encoding.

Reinforcement Learning

Reinforcement learning is a computational approach for goal directed learning from interaction [18]. The learner is not told which actions to take, but instead must discover which actions yield the most reward by trying them. Reinforcement learning is different from supervised learning, the kind of learning from examples provided by a knowledgeable external supervisor. In interactive problems it is often impractical to obtain examples of desired behavior that are both correct and representative of all the situations. In uncharted situations, where one would expect learning to be most beneficial, an agent must be able to learn from its own experience.

In reinforcement learning, the learner and decision maker is called *the agent*. The thing it interacts with, comprising everything outside the agent, is called *the environment*. These interact continually, the agent selecting actions and the environment responding to those actions and presenting new situations to the agent. The environment also gives rise to rewards, special numerical values that the agent tries to maximize over time. More specifically, the agent and environment interact at each of a sequence of discrete time steps, t . At each time step t , the agent receives some representation of the environment’s state, st , and on that basis selects an action, at . One time step later, in part as a consequence of its action, the agent receives a numerical reward, $rt+1$, and finds itself in a new state, $st+1$. At each time step, the agent implements a mapping from states to probabilities of selecting each possible action.

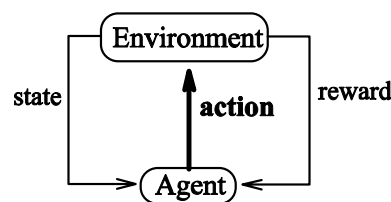


Fig. 2. The agent environment interaction

This mapping is called *the agent’s policy*. Reinforcement learning methods specify how the agent changes its policy as a result of its experience. The agent’s goal, roughly speaking, is to maximize the total amount of reward it receives over the long run. One of the challenges that arise in reinforcement learning is the trade-off between exploration and exploitation. To obtain a lot of reward, a reinforcement agent must prefer actions that it has tried in the past and found to be effective in producing reward. But to discover such actions, it has to try actions that it has not selected before. The agent has to exploit what it already knows in order to obtain reward, but it also has to explore in order to make better action selections in the future. The agent must try a variety of actions and progressively favor those that appear to be best.

Another key feature of reinforcement learning is that it explicitly considers the whole problem of a goal-directed agent interacting with an uncertain environment. All reinforcement-learning agents have explicit goals, can sense aspects of their environments, and can choose actions to influence their environments. It is usually assumed that the agent has to operate despite significant uncertainty about the environment it faces.

Adaptive Compression

The adaptive compression algorithm that we propose uses a reinforcement algorithm as presented in the previous section. We consider the frame rates as the rewards for each time step. An example of frame rate variation during a typical visualization session is shown in Figure 3. The figure shows the current and average frame rates obtained by using the RLE (Run Length Encoding) compression algorithm for two situations: one 100 MBps and one 10 MBps network connection of the client to the LAN. There are large variations in the current frame rate, especially when there is enough available network bandwidth (in the left and right regions of the figure). The average is done using the last ten frame rates.

Due to these large variations, the algorithm is making the selection of the compression algorithm based on these average values. For each selected compression method, at least 10 frames will be rendered using this method, providing this way a better estimate of performance of the algorithm.

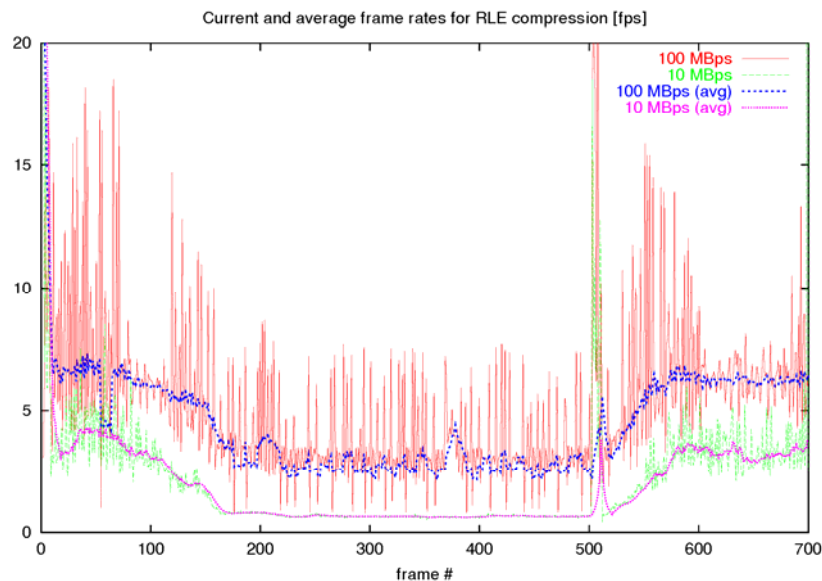


Fig. 3. Frame rate variations and average

The adaptive algorithm works as follows (see Figure 4): it starts with one of the compression methods (LZO in our case) and it uses it for the next 10 frames to get an estimate of its performance. After that, it tries in a similar way the other compression methods, and when all the methods are tested it chooses the best of the algorithms. From time to time, another compression method, different from the current one, is selected randomly and evaluated. If the new method is providing a better performance, i.e. increased frame rate, then it is selected as the

next compression method. We used an interval of 50 frames between trying another compression method.

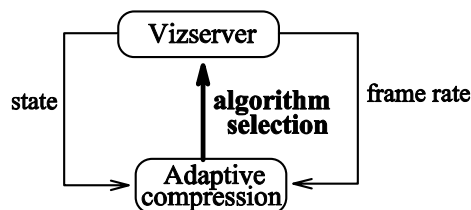


Fig. 4. The adaptive algorithm

Experimental Results

We conducted tests using an SGI Onyx2 2400 parallel computer as the remote visualization server. This computer consists of 32 R12000 RISC processors at 300 MHz, with a total memory of 16 GBytes and two Infinite Reality3 graphic pipelines. For the local visualization client we used a desktop PC with a Pentium 3 processor, running Linux (with a 2.4.19 kernel) at 500 MHz, with 256 MBytes of memory. As a remote visualization system, we used the SGI OpenGL Vizserver software, which allows remote rendering of the images on the SGI server, which are then compressed and sent over the network to the client for display (Figure 5). The SGI Vizserver offers an API for writing additional compression modules. We implemented four compression modules using the four lossless methods described in section 3. These modules are basically wrappers for existing software libraries, which implement the compression methods. The modules give a simple interface to both the compression and decompression, which is used by the adaptive algorithm [3]. This is implemented as a compression module for the SGI Vizserver using the development API provided with the software. The adaptive algorithm was implemented using the C++ programming language.

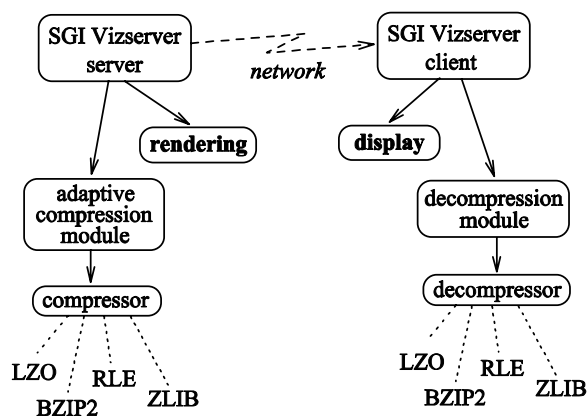


Fig. 5. Vizserver architecture

We have chosen the SGI Vizserver for several reasons. Firstly, because we have had access to an SGI parallel visualization server which had it available, and second, because the API used for the compression modules is very simple, making it very easy the implementation of different compression techniques. Another reason was that the use of the Vizserver is transparent to the applications. There was however some problems we experienced. One of them is that the version we have been using (3.1 beta) is quite unstable. We have experienced many crashes of the server software while developing and experimenting with different compression algorithms. One of the disadvantages in using SGI Vizserver is that the server hardware must be an SGI

computer. However, both the algorithm for the adaptive compression and the four compression modules that we propose are very easy to adapt to other similar remote visualization systems, due to the modularity of their implementation. One possible useful parameter we did not have access to while using the SGI Vizserver framework was the effective time required for sending each of the compressed frames over the network. The only available parameters we could use were the compression time for the frames and the time between two consecutive calls for the frame compression algorithm.

In our experiment, the size of each frame was 640x512 pixels with 4 bytes per pixel (RGB plus alpha channels). We used the Volview program for visualizing a volume data set of 256x256x77 voxels of a CT scan. The Volview is part of the SGI Volumizer2 software, and uses hardware accelerated 3D texturing for volume visualization. This is a direct data visualization technique that uses textured data slices, which are combined in a specific order using a blending operator. This technique takes advantage of graphics hardware and resources by using OpenGL 3D-texture rendering, allowing applications to obtain high interactive performances. The experiments were conducted using two different network connections between the client and the 100 MBps LAN containing the server. In the first situation, we connected the client using a 100 MBps network card to the LAN. In the second situation, we used a 10 MBps network card for connecting the client. Using a modified version of the Volview program, we recorded the translation and rotation vectors of the volume data for each frame generated during a typical interactive visualization session. We then played back the same session using the four different compression methods and then the adaptive algorithm. Frame rate averages for the two network connection situations, for all five situations are presented in Figure 6 and Figure 7. In both situations, the adaptive algorithm (*adap*) is searching for the best algorithm in the beginning, thus giving low frame rates. However, when it finds the best algorithm, it keeps it for the rest of the visualization session.

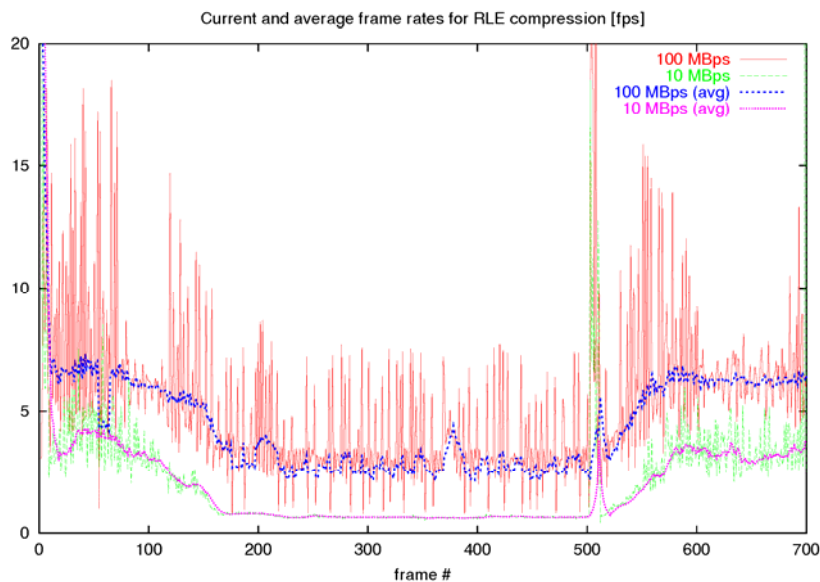


Fig. 6. Average frame rate - 100 MBps network

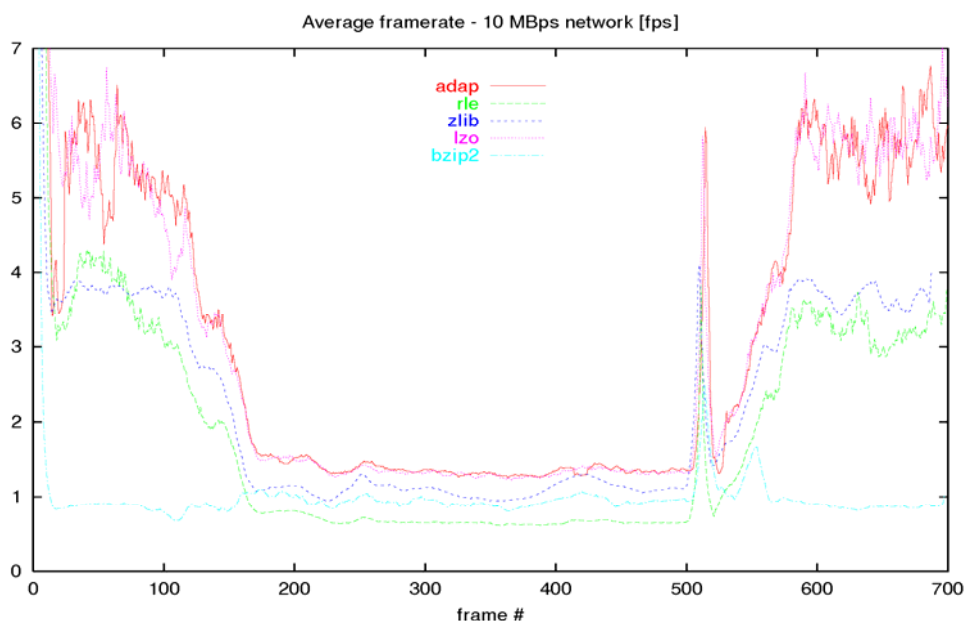


Fig. 7. Average frame rate - 10 MBps network

Conclusion

In this paper we presented an adaptive algorithm based on reinforcement learning for choosing one of the available compression methods in order to maximize the frame rate. Our experiments show that such an algorithm can work in a dynamic and uncertain environment, consisting of a visualization server, a visualization client, and a network for transferring the compressed images between the server and the client. One of the problems we experience with the current algorithm is that, in certain situations, one of the compression methods, which are evaluated by the adaptive algorithm, is giving really poor frame rates. This affects the interactive responsiveness of the application. One possible improvement of the algorithm would be to use a different selection algorithm for evaluating the next possible method, by making actions, which give small rewards to be less likely to occur. In this way, compression methods, which give poor frame rates, will be less probable to be selected in the future.

The modules for the compression methods and the adaptive algorithm are available for download, both as source code and binary at the following web site: <http://www.unde.ro/zoran/vizserver/>

References

1. Bederson, B., Shneiderman, B. - *The craft of information visualization: readings and reflections*, Amsterdam, Boston, Morgan Kaufmann, 2003
2. Bonneau, G.-P., Ertl, T., Nielson, G.M. - *Scientific Visualization: the visual extraction of knowledge from data*, Berlin, Springer-Verlag, 2006
3. Constantinescu, Z. - *A Desktop Grid Computing Approach for Scientific Computing and Visualization*, PhD Thesis, Norwegian Univ. of Science and Technology, Trondheim, Norway, 2008

4. Engel, K., Sommer, O., Ertl, T. - A Framework for Interactive Hardware Accelerated Remote 3D-Visualization, *Proceedings of TCVG Symposium on Visualization – VisSym*, 2000
5. Hansen, C.D., Johnson, C.R. - *The Visualization Handbook*, Boston, Elsevier-Butterworth Heinemann, 2005
6. Hludov, S., Schröter, C., Meinel C. - Adaptive Compression of Image Data, *Proceedings of SYBEN 1998*, Broadband Networks, Zürich, Switzerland, 1998
7. Kaufman, A., Nielson, G.M. - *Visualization '92: Proceedings*, October, Boston, Massachusetts, IEEE Computer Society Press, 1992
8. Ma, K.-L. - Visualizing Time-Varying Volume Data, *Computing in Science and Engineering*, vol. 5, no. 2, pp. 34-42, Mar./Apr. 2003
9. Ma, K.-L., Camp, D.M. - High performance visualization of time-varying volume data over a wide-area network status, *Proceedings of the 2000 ACM/IEEE conference on Supercomputing*
10. Nielson, G.M. - Visualization in Scientific and Engineering Computing, *Computer*, 21, pp. 58-66, 1991
11. Nielson, G.M., Bergeron, D. - *Visualization '93: Proceedings*, San Jose, California, IEEE Computer Society Press, 1993
12. Nielson, G.M., Hagen, H., Müller, H. - *Scientific Visualization: Overviews, Methodologies, and Techniques*, Los Alamitos, California, IEEE Computer Society, 1997
13. Patrikalakis, N.M. - *Scientific Visualization of Physical Phenomena*, Tokyo, New York, Springer-Verlag, 1991
14. Reed, T.R. - *Digital Image Sequence Processing: Compression and Analysis*, CRC, 2004
15. Salomon, D., Motta, G. - *Handbook of Data Compression*, 5th ed., ISBN: 978-1-84882-902-2, 2010
16. Sayood, K. - *Introduction to Data Compression*, 3rd Edition, Morgan Kaufmann Series in Multimedia Information and Systems, 2005
17. Shi, Y.Q., Sun, H. - *Image and Video Compression for Multimedia Engineering: Fundamentals, Algorithms, and Standards*, 2nd Edition, CRC, 2008
18. Sutton, R.S., Barto, A.G. - *Reinforcement Learning: an Introduction*, Cambridge, Massachusetts, MIT Press, 1998
19. Utnes, T., Brors, B. - Numerical Modelling of 3-D Circulation in Restricted Waters, *Applied Mathematical Modelling*, 17, pp. 522-535, 1993
20. Ware, C. - *Information Visualization: Perception for Design*, San Francisco, Morgan Kaufman, 2004
21. Wright, H. - *Introduction to Scientific Visualization*, New York, Springer, 2007
22. Thalmann, D. - *Scientific Visualization and Graphics Simulation*, Chichester England, New York, Wiley, 1990
23. Tufte, E.R. - *Visual Explanations Images And Quantities, Evidence And Narrative*, Cheshire, Connecticut, Graphics Press, 1997
24. Tufte, E.R. - *The Visual Display Of Quantitative Information*, Cheshire, Connecticut, Graphics Press, 2001
25. *** - *Silicon Graphics*, <http://www.sgi.com>, accessed 2009

Compresie adaptivă pentru vizualizare la distanță

Rezumat

Tehnicile de vizualizare la distanță care folosesc medii client-server permit utilizatorilor să acceseze seturi de date de mari dimensiuni. O soluție posibilă pentru vizualizarea la distanță este folosirea tehnicilor de compresie, în care imaginile sunt generate și comprimate pe server și apoi imaginile codificate sunt transferate prin rețea, decomprimate și afișate la clienți. În această lucrare propunem un algoritm adaptiv bazat pe învățare prin întărire pentru a alege una dintre metodele de compresie disponibile pentru a maximiza numărul de cadre pe secundă. Experimentele efectuate arată că un astfel de algoritm poate lucra într-un mediu necunoscut, dinamic, constând dintr-un server de vizualizare, un client de vizualizare și o rețea pentru a transfera imaginile comprimate dintre server și client.