

Virtual Desktop Display Acceleration Technology: RVEC

● Kazuki Matsui ● Kenichi Horio ● Yuichi Sato ● Shinichi Sazawa

Virtual desktop services have been attracting attention due to their ability to improve data security and reduce the operation and management costs of personal computers. A virtual desktop is a desktop environment virtualized in a cloud that can be accessed remotely and used in the same way as a conventional desktop environment. It thereby eliminates the need to store information on a client terminal. As smartphones grow in sophistication and popularity, the need for using virtual desktop services in a mobile environment is increasing, and smooth operability is required as well. Fujitsu Laboratories has developed Remote Virtual Environment Computing (RVEC) technology, which improves the operational responsiveness of user terminals. It works by decreasing the amount of data transfer for videos and high-definition images. It reduces the amount of data transferred to one-tenth that with conventional techniques. This technology supports virtual desktop services for handling graphics processing of computer-aided design (CAD) as well as other such applications in a mobile environment with smartphones and other devices.

1. Introduction

Virtual desktop services have been attracting attention in recent years as a means of improving data security and reducing the costs of operating and managing personal computers. These services create a virtual desktop environment in a cloud without storing information on the client terminal and enable that desktop to be used the same as in an ordinary terminal environment through remote access. The need for accessing virtual desktop services from a mobile environment is also growing as high-function smartphones continue to spread. In either scenario, smooth operability is a prime requirement.

Fujitsu Laboratories has developed Remote Virtual Environment Computing (RVEC) as a display acceleration technology for improving terminal operating response for the user. This technology creates a virtual desktop environment on a cloud and decreases the amount of data

transferred when handling videos and high-definition images remotely from a client terminal to about one-tenth that with conventional techniques. Fujitsu has adopted RVEC as a core technology of its Engineering Cloud.¹⁾

This paper provides an overview of the virtual desktop display acceleration technology^{2),3)} developed by Fujitsu Laboratories.

2. Methods for achieving a thin client system

A virtual desktop created on a cloud and operated remotely via a network is one type of thin client system.

We here explain the operating principle of a thin client system. Such a system outputs drawing results from an application executed on the server side and the contents of the desktop screen to a frame buffer and sends the screen data in that frame buffer to the client side in real time. This makes it possible to immediately

reproduce a screen identical to the server's frame buffer on the client side.

It is therefore important in a thin client system to quickly detect any changes in the desktop screen and to deliver those screen changes to the client side using the least amount of data possible. There are two types of methods for achieving this, as shown in **Figure 1**. The "screen transfer method" shown on the left sends screen changes to the client in the form of drawing-command or still-image data. The "video transfer method" shown on the right sends the entire desktop to the client in the form of a video, in a manner similar to that used for television.

3. Technical issues

With the screen transfer method, if an application is being used that generates screen changes frequently (as when processing video, animation, or computer graphics), commands or still images must be transmitted every time the screen changes, resulting in a high data transfer rate. If the available network transmission bandwidth is insufficient, the user will see a drop in operating response.

In contrast, the video transfer method allows for smooth viewing of rich content in the

form of video but imposes a large processing load on the server since it must convert the entire desktop to video data. In general, dedicated hardware must be used to achieve an acceptable level of performance.

When comparing the data transfer amounts among thin client systems, one must also take into account the frame rate governing the screen update interval. This frame rate affects how quickly screen changes should be sent to the client.

For a user of a thin client system, a frame rate of 30 frames per second (fps) generally provides for smooth operation with virtually the same operating response as when executing an application locally on a personal computer. A reduction in the frame rate results in a corresponding reduction in the data transfer rate, leading to a longer screen update interval and deterioration in the operating response from the user's perspective. It is therefore important to reduce the amount of data transferred while refraining as much as possible from reducing the frame rate.

To solve these technical issues, we have developed a display acceleration technology and an image compression technique as introduced in the following section.

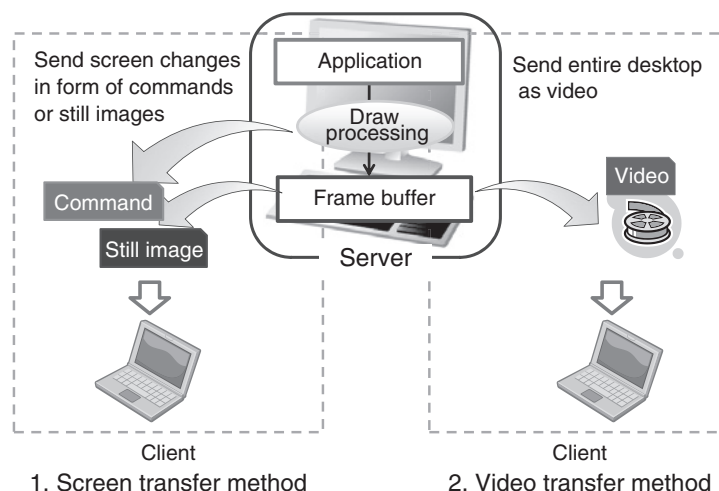


Figure 1
Methods for delivering screen changes to client side.

4. Solutions

4.1 Display acceleration technology

The RVEC virtual desktop acceleration technology we developed solves the technical issues discussed above. It uses a hybrid system that converts a region within the screen having many updates to video while sending all other differences in screen updates as still images (Figure 2).

Given a desktop screen divided into small sub-regions, RVEC calculates at regular intervals the number of screen updates in each sub-region and extracts those sub-regions having a high number. It then generates a rectangular region to contain all such extracted sub-regions and treats that region as a video region while considering the remaining region of the screen to be a still-image region. RVEC transfers these two types of regions to the client terminal using a system appropriate for each (RVEC original protocol). The numbers in each sub-region in the upper-right portion of Figure 2 indicate

the number of times that the screen has been updated in that sub-region within a fixed period of time. The larger the number, the larger the number of updates in that sub-region and the greater the amount of data sent to the client for screen updating.

Since RVEC measures the number of screen updates for each fixed period of time corresponding to the drawing frame rate and is constantly adjusting the video region and still-image region, it can optimally control the amount of transferred data in accordance with the screen display. In addition, RVEC can perform its operations solely on the basis of software by concentrating the video region in one portion of the desktop screen to improve processing efficiency. This approach negates the need for introducing dedicated hardware.

As described above, RVEC divides the desktop screen into a video region and a still-image region for sending screen updates in accordance with the number of updates that have

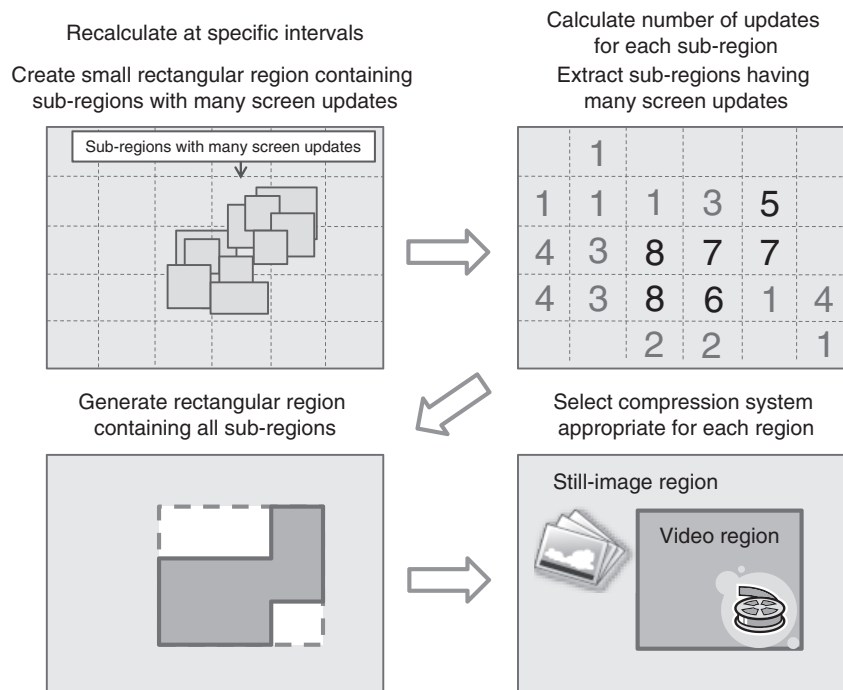


Figure 2
Determining video region.

occurred on the virtual desktop. While there are various systems like MPEG for coding video, each performs compression processing using the screen data for the previous frame. These systems normally achieve a high compression rate compared to still-image compression systems, which perform compression processing using only the data for a single image.

Since a reduction in the still-image compression rate effectively increases the total amount of data to be transferred, there is a need to improve still-image compression efficiency. Such improvement was achieved by having RVEC select an appropriate compression system in accordance with the properties of the display screen.

The RVEC architecture is sketched in **Figure 3**. First, the regions with screen updates are extracted, and regions to be sent as video are determined. (The remaining regions are sent as still images.) Appropriate compression systems are then selected. For regions to be sent as still images, a still-image compression system is

selected in accordance with the properties of the display screen.

With RVEC, we have been able to reduce the amount of data transferred when regenerating high-definition (HD) video (1280×720 pixels) in our virtual desktop test environment to about one-tenth (0.93 Mb/s) that with existing remote-desktop screen transfer systems.

4.2 Image compression technique

Selection of a still-image compression algorithm can greatly affect the data transfer rate. For example, the properties of a photo image of a green valley used as the standard desktop wallpaper are vastly different from those of a wireframe model used in computer-aided design (CAD). The compression rate of still-image data can thus be significantly improved by selecting a still-image compression algorithm appropriate for the image in question. For example, when displaying a wireframe model commonly used in CAD, screens having a large number of lines must be transferred at high

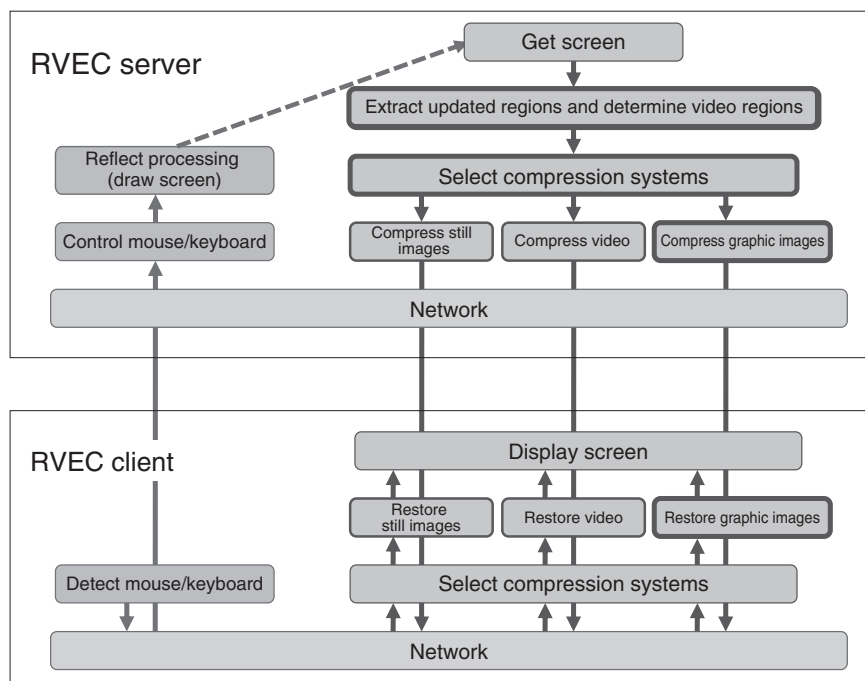


Figure 3
RVEC architecture.

speed with a high compression rate. In this case, using lossy compression like JPEG compression would make the lines blurry, thereby degrading image quality and hindering engineering design work (**Figure 4**). Using general-purpose lossless compression like zlib compression can improve the compression rate if the compression-level parameter is set appropriately. Its use, however, would greatly lengthen the processing time, making it impossible to complete compression processing at the frame rate of 30 fps needed to prevent visual discomfort. For this reason, conventional virtual desktop technologies have used the simple Hextile⁴⁾ technique, which divides an image into 16×16 small tiles and treats the image as single-color rectangles within that set of tiles. Hextile processing can be performed at extremely high speed owing to its simple algorithm. Nevertheless, when applied to images commonly found in CAD processing, it delays the display of those images on remote terminals because of its poor compression performance, which results in high bandwidth usage.

To address this problem, we developed an original lossless image-compression technique

that simultaneously solves the processing-speed and compression-rate problems. Its use enables smooth remote desktop operations even for complicated graphical images. This compression technique for computer graphics extends the Hextile technique so that a sequence of single-color rectangles in a certain direction is approximated as a vector to be matched with pixel data. However, actual pixel data do not strictly agree with this sequence of rectangles of fixed size, as shown in **Figure 5**. To reduce the amount of this small-difference data between the vector data and the actual pixel data, we use Wyle encoding,⁵⁾ which is remarkably effective when the data distribution is concentrated on a small quantity. This combination of vector approximation and Wyle encoding compresses the images remarkably, especially for images having a large quantity of line drawings, which is common in engineering work. Moreover, this technique maintains the high-speed characteristics of the Hextile technique. Using typical CAD images in our virtual desktop test environment, we compared our compression technology with a variety of other techniques in terms of compression rate and compression

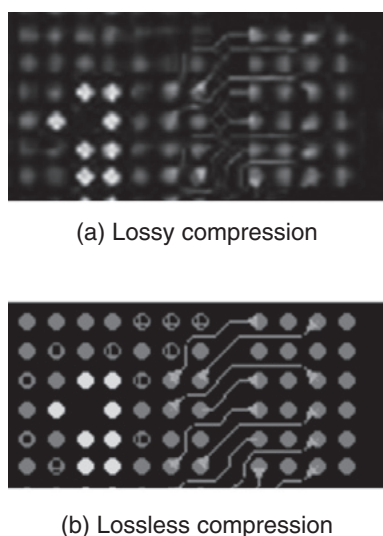


Figure 4
Compression quality.

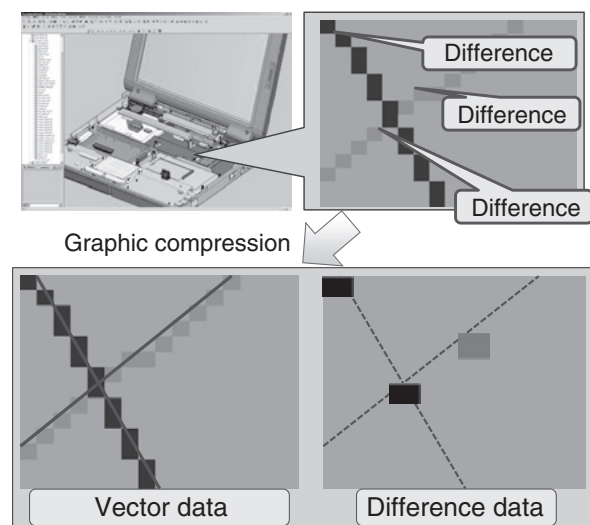


Figure 5
Compression of computer graphics.

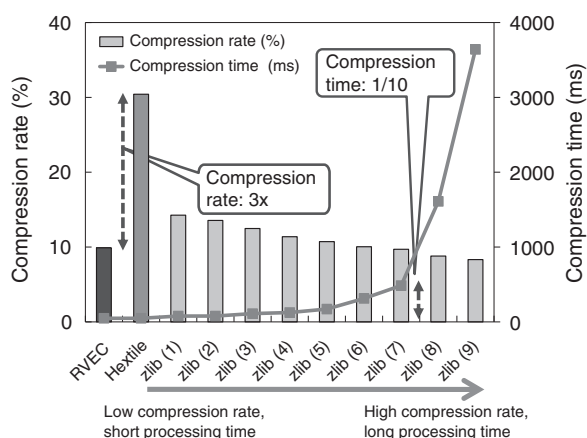


Figure 6
Evaluation of compression performance.

speed. As shown in **Figure 6**, the compression rate with RVEC was about three times that with the Hextile technique at nearly the same processing speed. The processing speed with RVEC was about ten times faster than that of zlib (level 7) at nearly the same compression rate.

5. Future activities

Fujitsu Laboratories is researching and developing means of resolving several technical issues that arise when applying RVEC acceleration technology to cloud services.

1) Efficient use of GPUs in a virtual environment

In a virtual environment, it is essential that the cost of providing technology that applies graphics processing units (GPUs) and general-purpose computing on graphics processing units (GPGPUs) be reduced by improving the degree of consolidation of guest OSs in the cloud. It is also vital that processing performance be improved so that applications using a large quantity of graphics as in CAD can be executed at high speeds.

With RVEC, we will pursue the effective use of GPUs and GPGPUs in the cloud, means of improving response on thin-client terminals of CAD applications executed on the cloud, and techniques for further reducing the amount of data transferred.

2) Performance improvement on networks with long delay

We are currently working on the application of the “random parity stream” (RPS)^{6)–8)} high-speed data transfer technology to RVEC. The objective of this technology is to minimize response latency of cloud services in various international and domestic network environments having long network delay.

3) Optimization on mobile networks

The use of thin clients is spreading in mobile environments as reflected, for example, by the expanded use of tablet computers running Android OS. However, there is still a need to improve thin-client response and make data transfer more efficient by taking into account the individual characteristics of different types of mobile networks such as wireless LAN, mobile WiMAX, and LTE.

6. Conclusion

Fujitsu Laboratories has developed virtual desktop display acceleration technology for improving the operating response of thin client systems. This technology is currently undergoing in-house trials for two-dimensional CAD. From the first half of fiscal year 2012, three-dimensional CAD has been applied to commercial cloud services. Studies are also underway on the application of this technology to a variety of mobile solutions with smartphones in mind such as the secure access of a virtual desktop from a mobile environment.

References

- 1) S. Saito et al.: Engineering Cloud: Flexible and Integrated Development Environment. *Fujitsu Sci. Tech. J.*, Vol. 47, No. 4, pp. 408–417 (2011).
- 2) Fujitsu Laboratories: High-speed Display Technology for Improving Virtual Desktop Operating Response. (in Japanese). <http://pr.fujitsu.com/jp/news/2011/05/2.html>
- 3) Fujitsu Journal: Virtual Desktop High-speed Display Technology. (in Japanese). <http://jp.fujitsu.com/journal/strength/technologies/201106.html>
- 4) The RFB Protocol, Version 3.8. <http://www.realvnc.com/docs/rfbproto.pdf>

- 5) H. Wyle et al.: Reduced-Time Facsimile Transmission by Digital Coding. *IRE Trans.on Communications Systems*, Vol. 9, pp. 215–222, Sep. 1961.
- 6) H. Kameyama et al.: High-speed File Transfer Solution for Cloud Computing. (in Japanese), *FUJITSU*, Vol. 60, No. 5, pp. 470–475 (2009).
- 7) Fujitsu Journal: High-speed File-transfer Technology 20 Times Faster than FTP: BI.DAN-GUN. (in Japanese).
<http://jp.fujitsu.com/journal/strength/technologies/200906-02.html>
- 8) High-speed File-transfer Solution BI.DAN-GUN.
<http://www.fujitsu.com/global/services/software/filetransfer/>



Kazuki Matsui

Fujitsu Laboratories Ltd.

Mr. Matsui is engaged in the research and development of mobile-services platform technology and virtual desktop technology.



Yuichi Sato

Fujitsu Laboratories Ltd.

Mr. Sato is engaged in the research and development of a total simulation platform with an emphasis on monozukuri and financial applications.



Kenichi Horio

Fujitsu Laboratories Ltd.

Mr. Horio is engaged in the research and development of mobile-services platform technology.



Shinichi Sazawa

Fujitsu Laboratories Ltd.

Mr. Sazawa is engaged in the research and development of cloud technology and parallel simulation technology.