

BASS Application Sharing System*

Omer Boyaci and Henning Schulzrinne
Department of Computer Science
Columbia University
{boyaci,hgs}@cs.columbia.edu

Abstract

Application and desktop sharing allows sharing of any application with one or more people over the Internet. The participants receive the screen-view of the shared application from the server. Their mouse and keyboard events are delivered and regenerated at the server. Application and desktop sharing enables collaborative work, software tutoring, and e-learning over the Internet. We have developed an application and desktop sharing platform called BASS which is efficient, reliable, independent of the operating system, scales well via heterogenous multicast, supports all applications, and features true application sharing.

1. Introduction

Application and desktop sharing allows two or more people to collaborate on a single document, drawing or project in real-time. Some applications like Netbeans and Google Docs are collaboration-aware and allow more than one person to work on the same document at the same time [6, 3]. However, most applications are not collaboration-aware. Fortunately, collaboration features can be added to these applications transparently. There are two models for collaboration-transparent application sharing: application-specific and generic. An application-specific solution allows to share a specific application such as Microsoft Word [26], while a generic one allows to share any application. Application-specific solutions are expensive in terms of engineering cost and they may not allow to use all features of the application [26]. Also, to participate in a sharing session, all participants must have a copy of the shared application. In the generic model, the application can be anything such as a word processor, browser, CAD/CAM, Powerpoint or movie editor. Also, the participants do not need to install the application. One disadvantage of generic application sharing is that its generic

nature makes it less efficient as compared to the application-specific model in certain scenarios. We have developed an application and desktop sharing system, BASS, based on the generic model.

Application sharing differs from desktop sharing. In desktop sharing, a server distributes any screen update. In application sharing, the server distributes screen updates if and only if they belong to the shared application's windows.

The main challenges of application and desktop sharing are scalability, reliability, true application sharing, operating system independence, and performance. BASS scales quite well via reliable multicast as discussed in Section 6.6. The sharing system should be efficient in the sense that it should transmit only the changed parts of the screen, and it should not consume all the bandwidth and CPU resources while doing this. BASS uses the most efficient technique, a mirror driver, to detect changed regions of the screen. BASS uses different encodings for different regions of the screen which yields high performance for movies in terms of both bandwidth and frame rate.

Section 2 discusses related work. The system architecture is discussed in Section 3, and the details of the BASS protocol are discussed in Section 4. Client and server architectures are explained in Sections 5 and 6, respectively. Section 7 compares the performance of BASS to other systems in terms of bandwidth and frame rate. Finally, Section 8 summarizes our work.

2. Related work

The X Window System is a network-transparent window system. Several X protocol multiplexors such as SharedX, DMX, XMX, and CCFX have been developed [15]. These multiplexors fail to support heterogeneous X servers. Also, it is not possible to support late joiners.

Microsoft provides Windows Meeting Space for Windows Vista and Netmeeting for Windows XP. Netmeeting was released in 1999 for Windows 98; in our tests, it failed to display pop-ups and menus. Windows Vista introduces application sharing feature as part of Windows Meeting

*This work was supported in part by a grant from FirstHand Technologies.

Space, but all the attendees must use Windows Vista.

VNC [24] is a cross-platform open source desktop sharing system but it supports only screen sharing. VNC uses a client-pull based transmission mechanism which performs poorly compared with server-push based transmissions under high round-trip time (RTT).

SharedAppVnc [25] supports true application sharing, but the delay is on the order of seconds. It uses a lossy codec and does not support multicast.

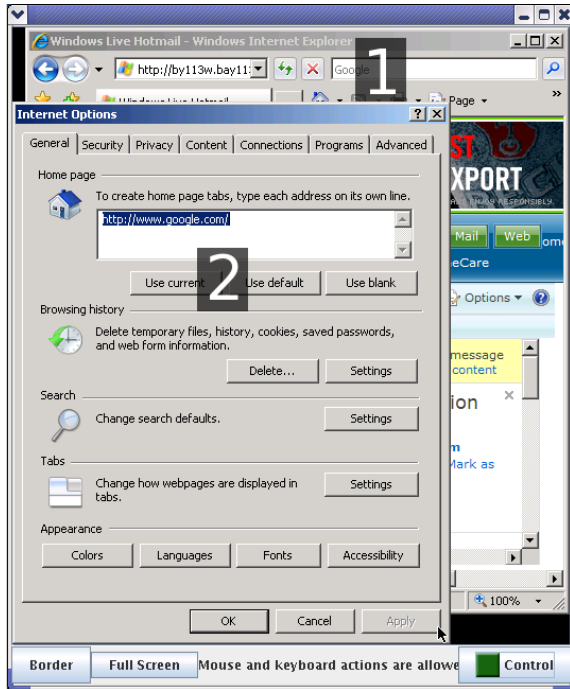


Figure 1. BASS client view

Some sharing systems such as UltraVNC [12] and Multicast Application Sharing Tool (MAST) [22] claim application sharing support. However, they distribute all pieces within the boundary of the shared window. This may expose information from other non-shared applications. Shared application may open new child windows such as those for selecting options or fonts. UltraVNC and MAST failed to share child windows. A true application sharing system must blank all the non-shared windows and must transfer all the child windows of the shared application. For example, if a user wants to share only the “Internet Explorer” application, which has the title “Windows Live Hotmail - Windows Internet Explorer”, from the desktop seen in (Figure 2), then the participants should only see the main and the “Internet Options” windows. BASS (Figure 1) displays only these two windows with a correct size while blocking the desktop background and the non-shared windows. In our test MAST, discussed in the following section, did not display the shared application in correct size and did not

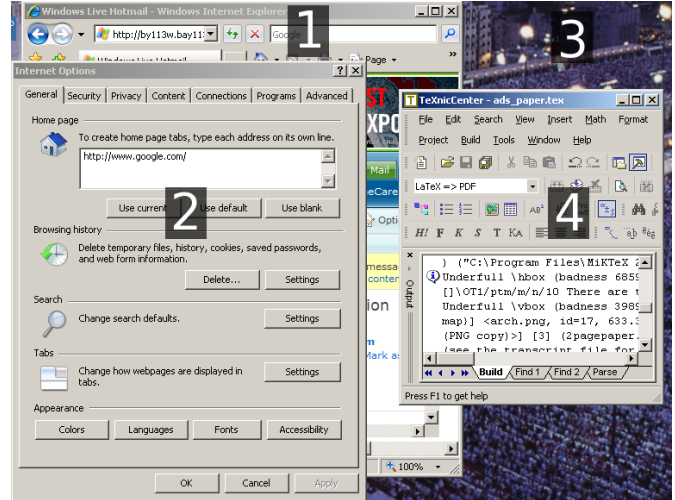


Figure 2. Desktop with overlapping windows




Figure 3. MAST client view

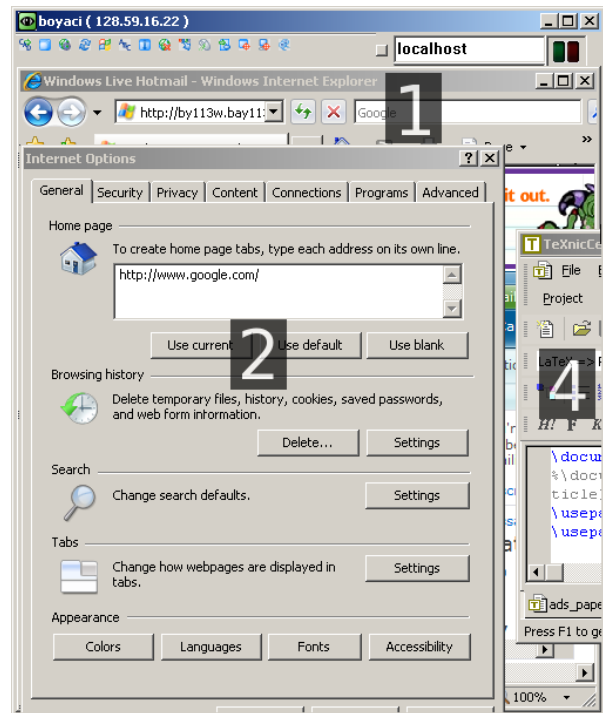


Figure 4. UltraVNC client view

	OS Independent	Scalable	Application Sharing	Remote Control	Recording
VNC	+	-	-	+	+
Windows Meeting Space	-	-	+	+	-
Leopard Screen Sharing	-	-	-	+	-
TeleTeachingTool	+	+	-	-	+
Mast	+	+	-	+	-
BASS	+	+	+	+	+

Table 1. Comparison of related work

block the non-shared application and desktop background (Figure 3). Similarly, UltraVNC failed in our tests due to following problems: the cursor position did not match, windows belonging to unshared applications are shared, new windows belonging to same application are not included and long menus are not shown properly (Figure 4). Current sharing solutions perform poorly if the user wants to share photos or movies. They use the same encoding for text, computer-generated images, movies, and photographic images. Lossless encodings give poor performance for movies and photographic images. While lossy encodings generate visual artifacts around texts and computer-generated images such as straight lines. THINC [16] and RDP [10] can play full motion movies if the bandwidth between the user and participant is tens of Mb/s [12, 10]. Due to their high bandwidth requirements, they do not scale well, and they do not perform well for realistic bandwidth conditions. BASS is the only system which uses different encodings for different regions of the screen. BASS uses the Theora [11] video codec to stream movies, JPEG [13] to transmit images, and PNG [9] for the rest.

VNC, Netmeeting, Windows Meeting Space and Mac OS X Leopard Screen Sharing all rely on unicast only, so they do not scale well to larger groups. Sharing an application via unicast increases the bandwidth usage linearly with the number of participants. For instance, Microsoft suggests Windows Meeting Space for groups of no more than 10 users. The TeleTeachingTool [27] and MAST use multicast to overcome this limitation. The TeleTeachingTool adds multicast support to VNC servers; however, it is developed just for online teaching, so it does not allow participants to control the shared desktop. Also, it does not support true application sharing due to its underlying VNC system. MAST allows remote users to participate via their keyboard and mouse, but its screen capture model is based on polling which is very primitive and not comparable to current state-of-the-art capturing methods like mirror drivers, discussed in Section 6.1. Although both TeleTeachingTool and MAST use multicast for scalability, they do not address the unreliable nature of UDP transmissions. Even if the packets are delivered, they may be out of order. In order to compensate for packet loss, the TeleTeachingTool

and MAST periodically transmit the whole screen which increases the bandwidth and CPU usage. Table 1 compares the sharing systems discussed so far. Nieh et al compared sharing systems in detail [20].

3. System architecture

BASS is based on a client-server architecture. The server is the computer which runs the shared application. The participants use a lightweight Java [5] client application for connecting to the server, and they do not need the shared application. Clients receive screen updates from the server and send keyboard and mouse events to the server.

The Java client works in almost every operating system. The server could not be written in Java because Java does not have OS level windowing information and can not learn screen updates from the OS. Therefore, there should be a server for each operating system and we have developed a Windows XP server and mirror driver. A mirror driver is the best known technique for capturing screen update events and will be discussed in detail in Section 7.1. The mirror driver runs in kernel mode and notifies the user mode server when it detects changes in the GUI of the shared application. The server then prepares an RTP packet containing encoded image of the updated region. RTP allows the clients to re-order the packets, recognize missing packets and synchronize application sharing with other media types like audio and video. The screen updates can be encoded with PNG [9], JPEG [13] or Theora according to their characteristics. PNG is an open image format which uses a lossless compression algorithm [19, 18] and more suitable for computer-generated images. JPEG is lossy, but more suitable for photographic images. Theora is an open source video codec comparable to H.264 and suitable for movie encoding.

The server supports both multicast and unicast transmissions. For unicast connections, either UDP or TCP can be used. Since TCP provides reliable communication and flow control, it is more suitable for unicast sessions. Multiple TCP clients sharing a single application may have different bandwidths, so we have developed an algorithm which sends the updates at the link speed of each client. For UDP

clients, the server controls the transmission rate because UDP does not provide flow and congestion control. Several simultaneous multicast sessions with different transmission rates can be created at the server. The server can share an application to TCP clients, UDP clients, and several multicast addresses in the same sharing session.

Participants can join a sharing session anytime, and they need the full screen buffer before receiving partial updates. Therefore, they send a RCTP-based feedback message, Full Intra-frame Request (FIR) [23], after joining the session. The server prepares and transmits the image of the whole shared region after receiving an FIR message. Preparing a full screen update is costly in terms of CPU, so the sharing server stores the generated image for some time.

Although multiple users can receive screen updates simultaneously, clearly only one of them can manipulate the application via keyboard and mouse events. BASS uses the Binary Floor Control Protocol (BFCP) [17] to restrict the control of the application to a single user. BFCP receives floor request and floor release messages from clients, and then it grants the floor to the appropriate client for a period of time while keeping the requests from other clients in a FIFO queue. All BFCP messages, keyboard and mouse events are transmitted directly to the server using TCP. Java's key-codes are used for mouse and keyboard events because these events are captured from a Java client and regenerated by a Java component at the server. These key-codes are publicly available [4].

We have also added a recording feature to BASS. Participants can record the sharing session to a file. This file may be used for watching the session locally or streaming to multiple receivers simultaneously. This feature is very useful for preparing lectures or software tutorials for future use.

4. The BASS protocol

BASS clients and servers communicate using the BASS protocol. The protocol defines five messages from server to client and two messages in the reverse direction. The "Open new window" message instructs the client application to open a new window. For TCP clients, the server transmits this message right after establishing the connection. UDP-based clients send a "full intra-frame request" to the server when they join the session. Receiving this FIR message, the server transmits the "Open new window" message. The server transmits "Region update" messages for updated regions. Whenever the shared window resizes, the server sends a "Window size update" message. "Move rectangle" instructs the client to move a region from one place to another, which is efficient for some drawing operations like scrolls.

Each protocol message consists of an RTP header, mes-

sage type identifier, window identifier, and message specific payload. The message specific payload structure and length are determined by the message type. Neither TCP nor RTP declares the length of the RTP packet. Therefore, RTP framing [21] is used to split RTP packets within the TCP byte stream.

5. Client architecture

The BASS client is very simple and lightweight compared with the server. It receives screen updates from the server and displays them to the participant. It is completely stateless in the sense that it can disconnect and reconnect to the server. Due to its simplicity, clients for different platforms can be easily developed. A Java BASS client has been developed in our lab. BASS clients can listen for or initiate connections. Participants of a sharing session can be view-only or they can request input control from the server. To request input control from server, a user presses the "control" button in the GUI of the client application (Figure 1). The client sends a floor request message, and then the server responds with a "granted" or "request queued" message. The server grants the floor immediately if nobody else currently controls the floor. Otherwise, the request will be queued in a FIFO queue, and the floor will be granted to the requesters one-by-one automatically. Users can release the floor by pressing the "control" button again. The floor is automatically released after a period of inactivity or after the client holding the floor leaves the session. After the server grants the floor, the client captures all keyboard and mouse events locally and transmits them to the server via RTP messages.

6. Windows XP server architecture

The Windows server allows Windows XP users to share an application with other participants. The Windows server has two main components, a kernel-mode mirror driver and user-mode sharing server. The mirror driver tracks the updated regions of the screen, and notifies the user-mode sharing server about these updates through a shared memory region. The user-mode sharing server learns the updated regions, prepares region updates, and transmits these updates to the participants (Figure 5). The sharing server also receives and regenerates mouse and keyboard events from participants. These two components are examined in detail in the following sections.

6.1. Mirror driver

A mirror driver is a display driver that mirrors the drawing operations of a physical display driver. The Windows OS calls the physical and mirror display drivers with the

same GDI (graphics device interface) commands. It is the most efficient way of learning screen updates because the operating system provides the exact coordinates of screen updates. We had to develop our own mirror driver for the sharing server because there is no free and open source mirror driver. Remote Desktop Connection and VNC also use their own mirror drivers to efficiently learn the screen updates. Stability and correctness are very important for kernel-mode components because they may easily cause restart or blue screen. Our mirror driver is completely stable such that we have not observed any crashes at all for the last two years.

We have used shared memory to establish a communication channel between the mirror driver and the sharing server. Both mirror driver and the sharing server map the same region of the memory to their own address spaces. The shared memory consists of a frame buffer to keep the screen state and a ring buffer which is used by the mirror driver to insert update commands and coordinates. There are two types of commands, BitBlt and MoveRect. The Windows OS notifies the mirror driver for an update, and then mirror driver inserts this update to the ring buffer with command type and the coordinates of the region. In case of application sharing, the BASS server computes the bounding rectangle of the shared application windows and informs the mirror driver about the tracking region. The mirror driver only tracks this specific region instead of the whole desktop, decreasing CPU overhead.

6.2. Server architecture

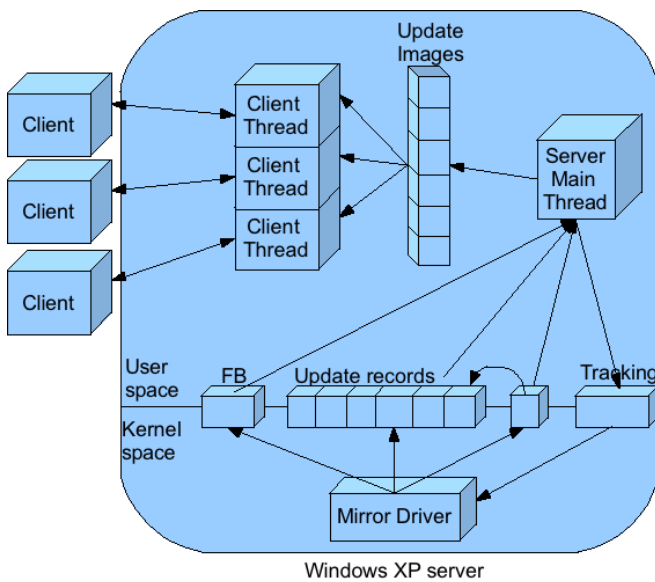


Figure 5. Windows XP server architecture

The Windows XP sharing server is a user-mode process

which complements the mirror driver. While the mirror driver keeps track of the frame buffer and the list of updated regions, the sharing server handles connection establishment, process keyboard and mouse events, and optimizes, compresses and transmits screen updates (Figure 5). The multi-threaded sharing server can serve multiple clients simultaneously. The server can wait for incoming connections and it can also connect to clients directly if instructed by the user. The sharing server has been designed considering the following challenges. Participants may have different bandwidths and they can join in anytime. UDP-based multicast and unicast sessions should be reliable even though UDP does not provide reliability. Some regions or windows may require different encoding for better performance.

The sharing server has one main thread, one manager thread, and a number of client threads. The main thread periodically checks for updated regions and prepares encoded images of these regions. These images are inserted to the image ring buffer which stores them until they are transmitted by client threads.

6.3. Serving window updates to users with different bandwidth

Encoding screen updates is a CPU-intensive operation, so there is no point in generating lots of updates if the clients do not have enough bandwidth to display them. The manager thread observes each client's bandwidth and throttles the main thread according to the highest bandwidth client. This technique can be easily explained with a movie playing example. Assume that the user shares a movie with remote participants. The same region of the screen is updated 24-30 times per second. Initially, the server tries to generate as many updates per second as possible. If at least one of the clients has enough bandwidth to receive these updates, the manager thread allows the main thread to continue this pace. But if none of the clients has enough bandwidth to deal with that update rate, then the manager thread slows the main thread by forcing it to sleep between update generations. Therefore, the main thread will generate a single update by combining several updates into one. The manager thread tries to equalize the update generation rate to the fastest client's bandwidth speed. This technique prevents unnecessary CPU usage in the sharing server.

The effective bandwidths of clients may change during the sharing session. This is properly handled by the manager thread because it checks clients' effective bandwidths periodically. A similar technique is utilized by the low bandwidth client threads. The fastest client thread transmits all the generated updates, however other client threads may not transmit all the generated updates due to their low bandwidths. These low bandwidth client threads skip some

of the region updates if there are newer updates for these regions. Going back to video player example, the fastest client thread may transmit 12 frames per second, whereas the other client threads may transmit some of these generated frames permitted by their bandwidths.

6.4. Encodings

The updates are distributed as PNG images except for movies or photos. PNG is very suitable and efficient for computer-generated images. However, its lossless nature results in a large increase in the compressed size of photographic images with negligible gain in quality, compared with JPEG and Theora which are specifically designed for photographic images. But the server does not know whether an updated region contains photographic or computer-generated content, because the mirror driver runs at the frame buffer level and at that level, there are only pixels and no metadata. Fortunately, detecting movie playing is very easy due to its specific characteristics. Different from other applications, movies generate 24-30 updates per second in a specific region of the screen. Benefitting from this characteristic, we have developed an algorithm to detect movie playing in order to use JPEG or Theora encoding for this region. Consecutive updates to a specific region trigger the detection. The Detection algorithm encodes the region using JPEG and compares the compressed image size between JPEG and PNG. If the JPEG size is less than a quarter of the PNG size, the server switches the default algorithm for this region to Theora and stores this result in a lookup table for subsequent updates. Theora is the default encoding for movies because it is four times more bandwidth efficient than JPEG. However, encoding Theora is costlier than JPEG, especially for high resolution movies. JPEG uses approximately four times more bandwidth than Theora but can generate 1.5 times more frames. The user can switch the server's default encoding for movies from Theora to JPEG if all the participants have enough bandwidth. If the compression ratio of JPEG or Theora regions drops below 12:1 compared to raw bits during the session, the server deletes the stored encoding information for this region from the lookup table. Similarly, regions which are marked as PNG regions are periodically rechecked.

6.5. Minimizing the effect of packet loss for UDP clients

With UDP, packets can get lost. Region updates may require several kilobytes or even megabytes. Unless designed carefully, a single packet loss may destroy the region update completely. In order to minimize the effect of packet loss, we have developed an algorithm which generates several small PNG images for a given update region. Blindly

generating a PNG image for each scan line may increase the bandwidth usage because a new zlib compressor object should be created for each new PNG. Creating a new zlib compressor decreases the compression ratio, so the bandwidth usage increases. Our algorithm tries to maximize the number of scan lines included in a single UDP packet while trying to keep the packet size below 1500 bytes, which the MTU for Ethernet. Due to its adaptive nature, it may feed tens of lines for a text document whereas it may feed only a single line for a photographic image. We have observed an increase of approximately twenty percent in bandwidth due to small PNGs.

Transmitting self-contained UDP packets minimizes the effect of packet loss. Instead of losing the complete region update, participants may lose only a few scan lines in case of a packet loss. They may end up with imperfect frame buffer due to packet losses, but they can continue to participate to the session. The retransmission mechanism discussed in the next section helps to restore the frame buffer state.

6.6. Reliable multicast

In the previous section, we described our algorithm which minimizes the effect of packet loss. In this section, we explain another supplementary technique which retransmits missing packets [14]. When a packet loss occurs, the participant views the rest of the image except missing scanlines. The client application sends a negative acknowledgement (NACK) for this missing packet [23]. Receiving this NACK, the server retransmits the requested packet. We have modified the oRTP library [8] which is used in the server side and extended our own Java RTP library on the client side to support this feature. The client/server applications do not deal with retransmission and buffering of packets, as these are handled by the RTP libraries.

The retransmission mechanism accommodates malicious or corrupted client behavior and NACK storms. In order to protect itself from malicious clients, the server does not retransmit a packet more than three times. If a packet failed to reach to several multicast clients, they all send a NACK back to the server causing a NACK storm. Instead of sending a NACK right after detecting a packet loss, clients wait for a random amount of time (0-100 ms). If a client observes a multicast NACK from another client while waiting, it suppresses its own multicast NACK request.

6.7. Streaming movie support

BASS is able to detect the regions of screen where a movie is playing, and it uses Theora or JPEG encoding for these regions. Encoding in real-time is computationally expensive. Although a Pentium 4 can encode 426x320 movie in full motion, it can only encode 6-10 frames per second for

a 852x480 movie. Therefore, we implemented another feature into BASS which enables to stream full motion movies to participants regardless of the resolution. The user copies the movie file into a BASS's specific directory. BASS automatically detects the movie and displays it in the GUI. If the movie is not encoded in Theora, BASS transcodes the movie into Theora using `ffmpeg2theora` [2]. This preprocessing takes 30 seconds for a 20 seconds 852x480 MPEG-4 movie. After the transcoding user can stream the movie to participants using negligible CPU power.

7. Performance results

We compared the bandwidth usage of sharing systems for web browsing. We also compared them for playing movies in terms of both bandwidth usage and frame rate. All sharing systems use 24 bits per pixel except RDP, which uses 16 bits per pixel. If RDP used 24 bits, it would consume fifty percent more bandwidth. For all tests, we used a Pentium 4 3GHz CPU and 1GB memory as the server and a Athlon XP 2600+ CPU and 1GB memory as the client. Server and client are connected over an 100 Mb/s LAN. For the movie playing comparison over a low bandwidth measurement, we restricted the bandwidth of the client to 3 Mb/s using `NetLimiter` [7]. To count the frame rate, we used a `Canopus TwinPact100` scan-line converter. This box takes the RGB output of the client as input, and it outputs a digital movie stream via firewire cable. We recorded this movie stream using the `iMovie` application of a `Macbook pro`. We then counted the individual frames to find the actual frame rate.




Figure 7. Web browsing performance

Figure 7 compares RDP, VNC and BASS for web browsing. During the measurement, the server automatically visited the home pages of the twenty most popular webpages according to `alexa.com`. We developed and used a Java-based automated testing application, available at [1], which visits each website for ten seconds. Some of these websites have animations and advertisements. Therefore, the bandwidth usage depends on how eagerly a particular sharing system transmits updates. We can say that all systems consume almost the same bandwidth, around 1 Mb/s. Although VNC and BASS use similar compression techniques, VNC

consumes less bandwidth because it uses a single compressor during the session, while BASS uses a separate compressor for each update. Using a new a compressor for each update allows BASS to compress each update only once regardless of the number of participants. However, VNC has to compress the same update for each participant because each participant has a different compressor. In case of more than one participant, VNC consumes more CPU, while the CPU usage of BASS remains constant.

We measured the multimedia performances of sharing systems by playing a movie over both an unlimited bandwidth link and a 3 Mb/s bandwidth link. The movie is a 20 seconds soundless 852x480 24 fps MPEG-4 encoded trailer of Warren Miller's *Higher Ground*. The BASS server can be configured by the user to use JPEG or Theora for movies. BASS-T and BASS-J represent BASS systems which use Theora and JPEG for movies, respectively. BASS-M represents BASS's Theora streaming feature, discussed in Section 6.7, instead of playing them in default media player. Figure 6 compares sharing systems over an unlimited bandwidth link. BASS-M and THINC are able to play the movie in full motion, however THINC consumes 112 Mb/s, while BASS-M consumes only 1.6 Mb/s. RDP gives the second highest frame rate, but consumes 45 Mb/s. BASS-J gives 9 fps consuming just 2 Mb/s, and BASS-T gives 6 fps consuming less than 1 Mb/s. VNC is the worst performer in terms of frame rate. In conclusion, BASS gives acceptable frame rate while using less than 2 Mb/s.

Comparing sharing systems in unlimited bandwidth environments is not very realistic because some participants may have low bandwidths. We repeated the same experiments over a 3 Mb/s link (Figure 8). Frame rates of all sharing systems dropped less than a frame per second except BASS whose frame rate remained the same. BASS-M is able to play full motion movies over an 1.6 Mb/s link. In conclusion, over low bandwidth links, all three BASS configurations yield a frame rate that is at least six times than the other sharing systems.

8. Conclusion

We have developed an application and desktop sharing system which is scalable, efficient, and independent of the operating system. BASS supports all applications due to its generic model, and transmits only the shared application and its child windows. We have used industry standards like RTP, BFCP, and PNG. BASS uses very little CPU thanks to the mirror driver. BASS, VNC and RDP consume roughly the same bandwidth for web browsing. However, BASS uses several times less bandwidth than the others for playing movies. Others can develop clients and servers compatible with BASS by implementing its open protocol.

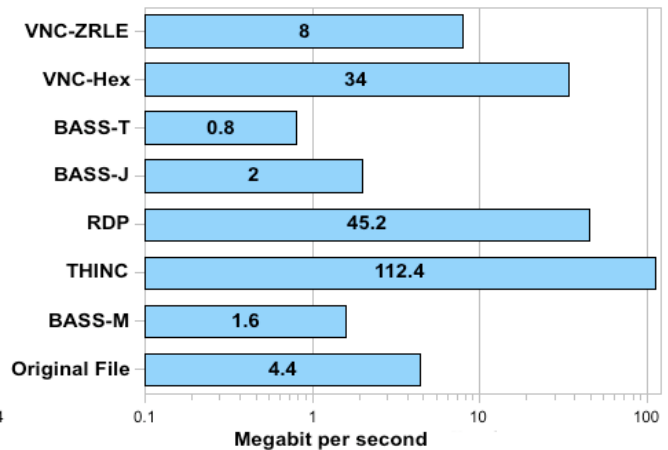



Figure 6. Comparison of sharing systems in terms of movie performance (unlimited bandwidth)




Figure 8. Movie performance comparison of sharing systems (network is limited to 3Mb/s)

References

- [1] BASS Application Sharing System. <http://www.cs.columbia.edu/~boyaci/bass/>.
- [2] ffmpeg2theora. <http://v2v.cc/~j/ffmpeg2theora>.
- [3] Google docs. <http://docs.google.com>.
- [4] Java keycodes. <http://java.sun.com/javase/6/docs/api/java/awt/event/KeyEvent.html>.
- [5] Java technology. <http://java.sun.com/>.
- [6] Netbeans. <http://www.netbeans.org/>.
- [7] Netlimiter. <http://www.netlimiter.com/>.
- [8] oRTP. <http://freshmeat.net/projects/ortp>.
- [9] PNG. <http://www.libpng.org/pub/png/>.
- [10] RDP. www.microsoft.com/windowsserver2003/techinfo/overview/termserv.mspx.
- [11] Theora open video codec. <http://www.theora.org/>.
- [12] UltraVNC. <http://www.ultravnc.com>.
- [13] Joint photographic experts group. <http://www.jpeg.org>, 1992. ISO 10918-1.
- [14] B. Adamson, C. Bormann, M. Handley, and J. Macker. Negative-acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Protocol. RFC 3940, Nov. 2004.
- [15] J. E. Baldeschwieler, T. Gutekunst, and B. Plattner. A survey of X protocol multiplexors. *SIGCOMM Comput. Commun. Rev.*, 23(2):16–24, 1993.
- [16] R. A. Baratto, L. N. Kim, and J. Nieh. Thinc: a virtual display architecture for thin-client computing. *SIGOPS Oper. Syst.*, 2005.
- [17] G. Camarillo, J. Ott, and K. Drage. The Binary Floor Control Protocol (BFCP). RFC 4582, Nov. 2006.
- [18] P. Deutsch. DEFLATE Compressed Data Format Specification version 1.3. RFC 1951, May 1996.
- [19] P. Deutsch and J.-L. Gailly. ZLIB Compressed Data Format Specification version 3.3. RFC 1950, May 1996.
- [20] A. M. Lai and J. Nieh. On the performance of wide-area thin-client computing. *ACM Trans. Comput. Syst.*, 2006.
- [21] J. Lazzaro. Framing Real-time Transport Protocol (RTP) and RTP Control Protocol (RTCP) Packets over Connection-Oriented Transport. RFC 4571, July 2006.
- [22] G. Lewis, S. M. Hasan, V. N. Alexandrov, and M. T. Dove. Facilitating Collaboration and Application Sharing with MAST and the Access Grid Development Infrastructures. In *E-SCIENCE '06*, 2006.
- [23] J. Ott, S. Wenger, N. Sato, C. Burmeister, and J. Rey. Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF). RFC 4585, 2006.
- [24] T. Richardson, Q. Stafford-Fraser, K. R. Wood, and A. Hopper. Virtual network computing. *IEEE Internet Computing*, 1998.
- [25] G. Wallace and K. Li. Virtually shared displays and user input devices. In *2007 USENIX Annual Technical Conference*. S. Xia, D. Sun, C. Sun, D. Chen, and H. Shen. Leveraging single-user applications for multi-user collaboration: the cword approach. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*. P. Ziewer and H. Seidl. Transparent teleteaching. In *AS-CILITE*, pages 749–758. Auckland, New Zealand, 2002.