

Improving Video Performance in VNC Under Latency Conditions

Cynthia Taylor and Joe Pasquale
Computer Science and Engineering
University of California, San Diego
La Jolla, CA

Email: {cbtaylor,pasquale}@cs.ucsd.edu

Abstract—Virtual Network Computing, or VNC, is a popular thin client application used to access files and applications on remote computers. However, VNC can suffer from significant losses in throughput when there is high latency between the client and server. These losses become especially apparent in the case of video, where updates are both large and frequent.

In this work, we present a *Message Accelerator* proxy for VNC. This simple but highly effective solution mitigates high latency network effects for video performance while maintaining the advantages of a client-pull system. By operating near/on the server, it can send updates to the client at a rate corresponding to proxy-server interactions which are faster than client-server interactions. When testing using video, our Message Accelerator design results in frame rates an order of magnitude higher than plain VNC when running under high latency conditions.

I. INTRODUCTION

Thin client systems, with their ability to connect multiple users to the same computer system, are a valuable tool for supporting collaborative systems. Users can connect to multiple desktop sessions on the same computer, allowing for shared data and applications, or can have multiple connections to the same desktop, allowing them to collaborate by actually sharing the same virtual screen. VNC has been studied as a tool for both synchronous and asynchronous collaboration [1], [2]. VNC has also been used collaborative tools such as wall-sized and tabletop displays [3]–[6]. The ability to run multiple sessions on a single server offers other more system-related advantages as well. Having all of the data and computation running on a central server means that a business can just update software on one computer instead of hundreds or thousands of PCs running on employees’ desks, saving money on maintenance. Keeping all data on a central server provides security from lost laptops or disgruntled employees. Even the electricity costs of running thin clients can be cheaper than that of running PCs. As a result, many businesses have been moving over to thin client systems [7], [8].

However, one problem with thin client systems is that they may suffer from poor performance, resulting from a number of factors. There is the computational overhead on the server side of tracking and encoding display updates, and of decoding and rendering them on the client side. There is the limiting factor of bandwidth, or how much data can be sent across the network in a given time period. Perhaps the most significant limiting factor in thin client systems is network latency [9]. Latency

affects every message, regardless of its size, sent between the client and server. These performance effects are especially noticeable in media-intensive applications, such as video.

While performance is important to all tasks, it is especially important for video. Highly interactive tasks usually have lower update rates than video, as they are governed by the user’s reaction time. Schmidt et al. performed a characterization of interactive applications that showed that input rates are usually much lower than monitor refresh rates, indicating a machine-independent upper bound on how fast humans manufacture input events, and that seventy percent of events occur with more than one hundred milliseconds between events [10]. He also showed that most user input events actually result in very few pixels needing to be updated. These applications, with their infrequent screen updates, are less impacted by high network latency. In contrast, video applications require frame updates every 30 to 60 milliseconds, and updates frequently involve change in a large number of pixels. Because of its frequent updates, video is especially impacted by high latency conditions.

In this work, we focus on the thin client system VNC, or Virtual Network Computing [11]. VNC is one of the most widely used thin client systems. It has an open protocol, known as the Remote Frame Buffer or RFB protocol, and there are many open source implementations of both the client and server. [11]–[16]. Due to its many implementations, which are available for all major operating systems, it is possible to use VNC across different computing platforms, allowing users to access a Linux machine from a device running Windows, or vice versa. Because it is open source, VNC has been one of the more popular systems for thin client research. Research on VNC has included adapting it for high resolution tiled displays, adapting it to control home appliances, and adapting it for optimal viewing on cellular phones, as well as trying new data compression schemes [3], [4], [17]–[20]. With its use in both real world and academic settings, VNC is one of the most ubiquitous thin client systems.

In this work, we offer a way to mitigate VNC’s poor video performance under high latency conditions. The solution that we present is a very general and simple proxy, called a *Message Accelerator*, that works with VNC to mitigate the effects of network latency. The Message Accelerator runs on or near the server machine and requests updates from the server

in a client-pull fashion, thus mimicking the VNC client. It then forwards these updates to the client as soon as possible, at the faster rate of proxy-server interactions rather than the slower rate of client-server interactions. Because the Message Accelerator does not wait to receive a request from the client before sending, it is not affected by high network latency in the same way that traditional client-pull systems are affected.

The Message Accelerator requires no changes to the existing client or server, making it easy for a user to install (essentially “plug and play”), and eliminating parallel code maintenance issues. Using a Message Accelerator with VNC can result in the client receiving updates an order of magnitude faster in high latency situations. Our solution is specifically designed for cases such as video, where all video updates are being generated on the server, without depending on user input.

We will first provide some background and related work in Section 2. In Section 3, we will go into more detail about VNC. In Section 4, we will describe our changes to VNC. In Section 5, we will discuss our experiments and results. In Section 6 we describe how our system evolved from our original design to the system we present here. We will briefly discuss future works in Section 7, and then conclude in Section 8.

II. BACKGROUND AND RELATED WORK

This work is an improvement to the Virtual Network Computing or VNC system [11]. The VNC system works as a user-space application that requires no modifications to either the operating system or user applications. The VNC server continually scrapes the framebuffer, an area of memory in which color values for every on-screen pixel are stored. The VNC client sends a request about a specific on-screen rectangle, and the server responds with an update consisting of an encoding of the changes between now and the last time the client requested information about that rectangle. We will go into more detail about the operation of VNC in Section 3.

Thin client systems fall into two general categories in terms of when the server sends an update to the client [21]. In *server-push* systems, updates are sent to the client as soon as the server generates them. For example, if the user is playing video, in a server-push system, when a new frame of video is played, the server will immediately send an update to the client with the changes to the framebuffer. In *client-pull* systems, the server waits to send an update until it receives an explicit request from the client. If a user is playing video on a client-pull system, when a new frame of video is available or is generated, the server stores the changes it makes to the framebuffer in an internal representation, and when the server receives a request from the client, it sends an update containing all the changes to the framebuffer since the last update. VNC is a client-pull system.

The SLIM system offers a contrast to VNC [10]. SLIM uses a very similar message protocol to VNC, but differs in other ways. Instead of scraping the framebuffer to detect updates, SLIM works as a virtual device driver and offers a software library with display options for higher-level applications. In




Fig. 1. The VNC server scans the framebuffer and sends updates to the client.

SLIM’s implementation, X-Windows is modified to use the SLIM virtual device driver so that applications can either explicitly invoke SLIM display commands through its library, or pass their display commands to X-windows and have it forward them to SLIM. In this way, applications with specific display needs such as multimedia applications can send commands to SLIM that include more semantic information about how they should be processed. However, applications must be specifically altered to take advantage of this. Even if no user-level applications are modified, at the very least the window manager must be modified, and the modified applications must be kept updated in parallel with the original applications. The system is also limited to operating systems with available source code because of this. SLIM uses a server-push update system, and requires a network and machines capable of handling server-push updates with no loss of data.

THINC also acts as a virtual device driver, but at the abstract hardware level so no changes need to be made to applications or operating system [22]. THINC also uses a similar message protocol to SLIM and VNC. Device driver commands are tracked in an intermediate encoding until they are finally translated into messages and sent to the client. Commands are highly processed while in their intermediate encoding to avoid sending any information that is fully or partially overwritten by later updates. THINC is a server-push system, but rather than sending messages immediately, before sending a message to the client it checks the socket to see if the write will block, and waits until the buffer is next flushed on blocking. THINC uses special handling for video, using a stream-based system to avoid jitter between frames. The special handling of video adds even more complexity to an already quite complex system. In addition, since device drivers are OS-specific, there must be a virtual device driver written for each operating system on which THINC is used.

This work was also informed by the papers of Nieh et al, which describe and compare many currently existing thin client systems [9], [23], [24]. Nieh stresses that the performance of thin client systems is especially affected by latency, rather than bandwidth, and that client-pull systems are affected more than others due to the fact that they must send two messages for each update.

III. DETAILS OF VNC

In this section, we will describe the basic workings of a particular implementation of VNC, TightVNC [15]. These

details can easily be generalized to other VNC systems. We explain the concept of client-pull, the core operating feature of this and all VNC systems, in the context of a real working system that is publicly available and in wide use. We will focus on the implementation of the RFB protocol and the sending of messages between the client and server [12]. The system is illustrated in Figure 1.

A. The VNC Client

The VNC client is implemented as a very simple program; it is not multithreaded, and it blocks on reads and writes. After a brief initialization phase in which it exchanges setup information with the server, it falls into a `while(true)` loop. In this loop, it first waits to receive an update from the server. It processes the update and redraws the display. Then it issues a request for a new update.

The client intersperses read calls throughout the processing of the update, rather than reading the entire update and then processing it. The structure of the update encourages this: it contains a general header containing information about the message, and then a series of rectangles, each with a rectangle header that contains the dimensions of the data following it and its encoding. Due to this message structure, it is impossible to tell how many bytes an update is without processing a significant portion of the message. An additional advantage is that by processing the message as it is being read, the client can get as much processing as possible done before it has received the message in its entirety.

There are times when the client has processed as much of the update as it has received, but cannot read any more from the server. When this happens, the client uses the idle time to collect user input, such as mouse movement and typing, and sends the appropriate messages to the server. This is the only time when it processes user input.

B. The VNC Server

The VNC server is more complex than the client. As the server runs, it notes when the framebuffer is changed. It stores an internal representation of the area modified, and the new modifications made to it. This is called the *modified region*. The server continues to add any modifications made to the modified region until it receives a request from the client.

The requests the client sends are very short, containing just the dimensions of a rectangle, and a bit specifying whether the update is incremental or not. The rectangle is always the dimensions of the client's display window and rarely changes within a session. The first request the client sends is not incremental, and all requests afterwards are incremental. If the request the server receives is not incremental, the server immediately sends all of the framebuffer information for the rectangle.

If the request is incremental, the server compares the rectangle the client requested to the modified region. If they overlap, the server sends an update with the modifications within the overlapping area, and clears those modifications from the modified region. If the areas do not overlap, as in

the case where no changes have been made to the framebuffer since the last request, the server saves the rectangle as the *requested region*. The next time the framebuffer is modified, the server checks to see if the modification falls within the requested region. If it does, the server sends the client an update with the new modifications and deletes the requested region. If it does not, the server adds the modifications to the modified region and keeps the requested region. This means that the server can send an update in response to a request from the client any time after it receives the request. If the server receives more than one request before there are any modifications to the requested region, it will just send one update in response to all requests. We emphasize though that the server will only send an update if it receives a request from the client, and no sooner than having received the request.

C. VNC as a Client-Pull System

As described above, VNC is a client-pull system. This has both advantages and disadvantages in terms of system performance. Having the server wait until it receives a request to send an update provides a very simple flow-control mechanism to prevent the client from being overwhelmed by updates. The client will never receive an update that it is not ready to process, since the server sends updates only after the client specifically requests them. This is especially good for the types of highly resource-limited clients that motivate this work.

However, the sending of requests means that each update will accumulate at least twice as much network latency as updates in the server push system. In the rest of this paper, we will describe how we modify the VNC system to retain the advantages of a client-pull system, while reducing the problems caused by this additional latency.

IV. OUR IMPROVEMENTS TO VNC

A. Goals

The goal of our improvements to VNC is to create better video performance in the case of high latency between the server and client. We cannot change the time between when an action (such as a mouse movement or keystroke) is taken on the client machine and when the client receives an update that reflects that action. What we can change is the number of updates the client receives between these two events. The more updates the client receives, the less jumpiness the user will experience. This is especially important for video applications, where the number of updates per second the client receives becomes the number of frames per second the user sees.

We would like our changes to be as "plug and play" as possible, involving little to no effort on the part of the user. With that in mind, we have designed our system so that our improvements do not require changing any of the existing client or server code, but only the insertion of a proxy. With these design choices, we avoid the issues of parallel code maintenance, as any VNC system that adheres to the standard RFB protocol will be compatible with our proxy [12].




Fig. 2. The Message Accelerator quickly acquires an update from the server, then pipelines sending the update to the client.

B. Design

Our changes to VNC take the form of adding a *Message Accelerator* proxy between the client and server. We run the Message Accelerator on the same machine as the server, but it could also run on a separate machine (but close to the server to minimize network delay) if the server is not amenable to running the proxy. The Message Accelerator requests updates from the server, and forwards them to the client. It is logically invisible to both client and server.

Since the Message Accelerator is on the same machine as the server and communication does not have to travel over the network, it can make requests in a client-pull manner at a much faster speed than the actual client. It requests new updates at the rate that it can send them to the client. Since the client is single-threaded and performs reads interspersed with processing, it will never read at a faster rate than it can process. The Message Accelerator checks to make sure it can write to the client's socket without blocking before it sends the update, so it will never send at a rate faster than the network can handle. As a result of all of these factors, the Message Accelerator sends updates at the client's natural processing rate, regardless of network latency.

The Message Accelerator works by pipelining updates to the client. Since the Message Accelerator is running on or near the server, it can very quickly send requests to the server and receive updates in response. The Message Accelerator can duplicate and resend client requests, without needing to wait for a new request from the client. It then sends the updates from the server to the client at a constant rate, regardless of network latency. This is illustrated in Figure 2.

C. Implementation

The VNC session starts up with an initialization period during which the Message Accelerator simply forwards everything sent to it by the client to the server, and everything sent to it from the server to the client. This period is where the client sends the server information such as the user's password, what encoding the client would like to use, and other parameters. It is necessary for the Message Accelerator to simply forward this information, since there is no way for it to know these parameters ahead of time.

Once the initialization process is over, the Message Accelerator settles into a loop, illustrated in the algorithm below.

```

while (True) {
    select ();

    if (can read server) {
        updates[received] = read(server);
        received = received + 1;
        unansweredRequest = False;
    }

    if (can write client) {
        if (sent < received) {
            write(client, updates[sent]);
            sent = sent + 1;
        }
    }

    if (can read client) {
        request = read(client);
    }

    if (can write server) {
        if ((!unansweredRequest)
            && (received - sent < 1)) {
            write(server, request);
            unansweredRequest = True;
        }
    }
}

```

This loop constantly polls to check if either the server or client are available for reading or writing. The client sends short ten byte requests that the Message Accelerator saves and sends to the server. The server sends large updates (an average size for typical desktop video is around 0.9 MB) that the Message Accelerator forwards to the client. The Message Accelerator will often receive partial updates from the server, and will have to assemble an update over many reads.

If the Message Accelerator can read from the server, it reads everything available, and then parses it to see if it is a complete update. If the update is complete, it marks that the server is ready to receive the next request, since in a client-pull system we do not want to send a request until we have fully received an update. If the update is not complete, it saves what it has received so far and waits for more. After it has read and parsed the available information from the server, the Message Accelerator updates variables that keep track of the number of complete updates and bytes of partial updates it has received.

If the Message Accelerator can write to the client, it writes as much as it can, and then records how much it has written. When the Message Accelerator can read from the client, it reads a request, and saves it to send to the server. Since the client always sends a single complete request, the Message Accelerator does not have to do any sort of book-keeping for it. If the client's display dimensions change, the client will change the rectangle dimensions in the request it sends, and all the updates the server sends after it receives this request will reflect that change.

When the server is available for writing, the Message Accelerator first checks to make sure that it has received a

complete update since the last time it sent a request. It then compares the number of updates it has received from the server to the number of updates it has forwarded to the client. If it has sent all of the updates it has received, it sends a request to the server. Otherwise, it does nothing.

The Message Accelerator sends to the client whenever the client will not block its write (similar to the send mechanism in the THINC system [22]). Because we check to make sure that the write will not block before we send it, the rate at which the Message Accelerator can send to the client is effectively controlled by how fast the client can process the updates. However, because the Message Accelerator sends requests to the server even if it has not yet received a new request from the client, it is not affected by network latency in the same way as a simple client-pull system.

The Message Accelerator sends requests to the server at the same rate that it is capable of sending updates to the client. It must keep updates buffered to make sure it always has something ready to send to the client. By buffering the bare minimum of updates, we minimize the effects of changing send rates. If the Message Accelerator was buffering a large number of updates and the speed at which it was sending to the client changed, these buffered updates would appear in fast motion (i.e., fast forward fashion) if the sending rate decreased, and in slow motion if the sending rate increased. In addition, perception of any user input would be delayed until all of the buffered updates had been sent. However, in order for maximum performance results, we do want to always have an update prepared to send to the client as soon as we receive a request. To satisfy both of these constraints, the maximum number of updates that the Message Accelerator should have buffered at any time in order to always have something ready to send to the client is $\lceil \frac{\text{server processing time}}{\text{accelerators send time}} \rceil$. Because it generally takes the server less time to process an update and send it to the Message Accelerator than it takes the Message Accelerator to send that update to the client, this number is typically one.

V. PERFORMANCE EVALUATION

A. Experimental Design

As described above, after initialization the VNC Client goes into a loop where it reads and processes an update, sends a request to the server, and then waits for the next update. To compare the performance of a standard (i.e., no Message Accelerator) and an improved (i.e., with Message Accelerator) system, we instrumented the client to measure how long it takes for the client to go through each iteration of this loop. Since the client does this for each update it displays, the faster it goes through this process the faster it displays updates. Since this measurement is being taken in the client, it is blind to whether or not the system is using the Message Accelerator.

To test the systems, we need a media intensive application with many framebuffer updates per second. We used an episode of the television show *Angel*, playing on a loop in the VLC media player. Our testing was done with the client running on a Dell Vostro 200 machine with the Ubuntu




Fig. 3. The graphs show the median video update rate for both unmodified system and message-accelerated system, as a function of network latency. As the network latency increases, the update rate rapidly decreases in the unmodified system, but remains high and fairly constant with the Message Accelerator.

Intrepid Ixex OS, and the server and Message Accelerator running on a Dell Optiplex 755 with the Fedora OS. All TCP settings were at the default.

To simulate network latency, we use the *netem* Network Emulator to insert artificial network delay [25]. To generate measurements, we leave the server running, and run a script that runs the client connected directly to the server for one minute, and then runs the client connecting through the Message Accelerator for one minute. It goes through this loop ten times. All of our experiments were run using the raw encoding of both systems. (We used the raw encoding because it was the easiest to implement. In a production setting, better performance could be achieved by using different encodings.) By running first unmodified VNC and then VNC with the Message Accelerator, any network or processing effects outside of our system will affect both systems. All reported times are in milliseconds (ms).

B. Results

Figure 3 displays the median rate of both systems in updates per second. The unmodified VNC system is drastically impacted by network latency. With network latency of just 20 ms, the display rate at the client has almost halved. At network latency of 150 ms, the display rate is ten times lower than it was with no network latency. In contrast, the results from VNC with the Message Accelerator show that there is no significant impact on the median update time with increasing network latency. While the update rate begins to dip starting at a latency of 50 ms, it never gets worse than a reduction by 10%, even up to 300 ms. The lowest point in the dip in rate that occurs around 150 ms was a phenomena that occurred consistently in testing. Preliminary experiments indicate that altering the TCP window size moved or eliminated this drop in rate, so at this time we believe it is an artifact of our default TCP window

TABLE I

EFFECTS OF NETWORK LATENCY ON CLIENT INTER-UPDATE TIME. THE MEDIAN INTER-UPDATE TIME STAYS FAIRLY CONSTANT WITH THE MESSAGE ACCELERATOR, BUT RAPIDLY GROWS IN THE UNMODIFIED SYSTEM. DUE TO OUTLIERS AND OTHER EFFECTS, THE MEAN AND STANDARD DEVIATIONS GROW IN BOTH SYSTEMS, BUT MUCH LESS SO WITH THE MESSAGE ACCELERATOR.

Latency	Unmodified System			Message Accelerator		
	Median	Mean	St Dev	Median	Mean	St Dev
0	76.9	77.3	4.5	76.0	76.1	2.6
10	96.8	101.5	26.7	75.9	76.3	7.6
20	138.6	150.5	69.3	75.9	77.2	14.5
30	198.7	225.3	122.1	75.9	76.8	22.0
40	258.4	301.4	174.9	75.9	79.6	41.9
50	318.6	382.4	235.0	76.0	80.2	45.3
100	620.5	925.7	1117.9	83.1	111.8	172.0
150	920.6	1544.6	1243.3	86.4	128.3	187.1
200	1220.2	2313.1	1861.8	79.5	157.6	317.9
250	1520.8	3183.6	2500.7	79.2	182.4	366.5
300	1820.9	3723.3	2842.7	80.4	252.7	694.2

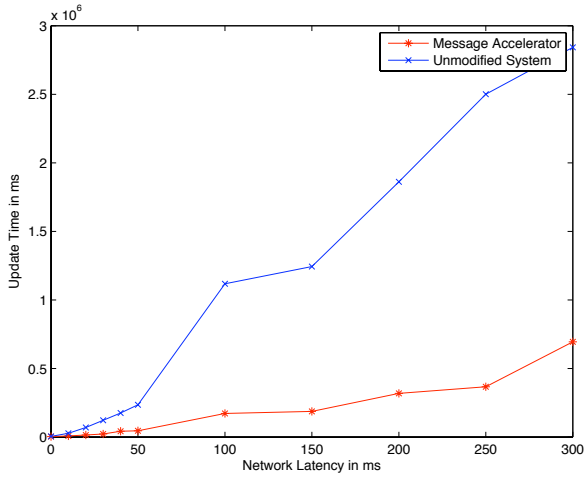


Fig. 4. Standard deviation of client update times in ms. The standard deviation quickly grows in the unmodified system, reflecting the uneven rate at which frames of video are displayed.

size. Even with this dip at 150 ms, compared to the unmodified system, the Message Accelerator system is outperforming the unmodified system by an order of magnitude.

Table I shows the inter-update times (reciprocal of the update rate) for both the unmodified system and message-accelerated system. For both systems, we see that the mean inter-update time is significantly higher than the median. This is due in part to some extreme outliers, where the update takes significantly longer than normal. Some of these outliers are the first update sent in a session. The first update takes a long time to send because display data is not yet being buffered, and is much larger than the later updates because it contains the entirety of the frame buffer. There are also outliers due to packets being lost or sent out of order, since packets re-sent by the TCP protocol are affected by the network latency.

On a qualitative and subjective note, while watching the videos, the unmodified system becomes noticeably jittery with just 20 ms of latency, and becomes completely unwatchable with 100 ms of latency. At 100 ms, the unmodified VNC system demonstrates user observable delay while rendering




Fig. 5. In our original Design, the message accelerometer was on a separate machine.

single frames, and the video becomes a series of semi-related frames with no illusion of movement between them. In the VNC system with the addition of our Message Accelerator, the video is still very watchable even at 300 ms, with very little noticeable lag or jitter.

As shown in Figure 4 (as well as Table I), we also observe that the standard deviation rises much more sharply in the unmodified system, quickly becoming ten times higher than that of the system with the Message Accelerator. This demonstrates that the difference in client display times between frames is much larger in the unmodified system. Disparate rendering times such as this mean that video plays at a jerky, uneven rate, alternately appearing sped up and slowed down.

Overall, we see that adding the Message Accelerator to VNC drastically improves the frame rate under network latency, even when the latency is not that large. With the Message Accelerator, the time it takes an update to be displayed remains consistent over increased network latency, even when the latency becomes quite high. By both qualitative and quantitative measures, the system is much improved by the addition of the Message Accelerator.

VI. THE EVOLUTION OF THE MESSAGE ACCELERATOR

Over the course of designing our system, we made many changes from our original design. Our system evolved as we learned more about the intricacies both of VNC and TCP. We offer this section on the evolution of our Message Accelerator to provide some insight on the inner workings of VNC and the lessons we learned.

A. The Original System

In our original design for the Message Accelerator, illustrated in Figure 5, it would have resided on its own machine, one that was approximately one WiFi hop away from the client machine. In this position, the latency between the Message Accelerator and client would be small, e.g., around 3 ms in a typical wireless-access network environment. With its ability to quickly communicate with the client, the Message Accelerator would measure the rate at which the client was capable of processing updates, and then send requests to the server at this rate. Because the Message Accelerator was not waiting to receive a new update before it sent the next request, it could send requests and receive updates at the rate at which the client was processing them, regardless of network latency. The Message Accelerator would then buffer the updates from the server, and dispense them to the client in a client-pull fashion, sending a complete update in response to a request.

B. The Message Accelerator without Update Delimitation

One of the first problems we discovered with this system was that waiting until the Message Accelerator had received the complete update from the server before we sent it to the client was causing a huge performance hit. The Message Accelerator was reading the update in from the server one packet at a time and waiting until the Message Accelerator had gotten a complete update before sending it to the client. It then sent it to the client, which then read it in a little bit at a time. This was causing a huge amount of overhead. With no network latency, it was taking an average of 103 ms for the client to receive and process an update with the Message Accelerator, as compared to 77 ms without the Message Accelerator. We changed our Message Accelerator design so that instead of collecting a complete update and sending it in response to a client request, the Message Accelerator forwarded everything it received from the server to the client immediately. This meant that the client could immediately process as much of the update as it had received, taking advantage of its interleaving of network reads and update processing.

The Message Accelerator no longer had to parse the updates as it read them in, or have any internal representation of an update. It could simply read data in from the server and immediately forward it to the client. This also meant that our system was no longer a true client-pull system, and no longer had the built-in client-pull safeguards. To avoid overloading the client, we were relying on a combination of sending requests at the same rate the client was processing, and checking to make sure writes to the server would not block. This change significantly increased performance, bringing the time it took the client per update down to an average of 87 ms.

C. The Server Accelerator and Client Accelerator

We next discovered that sending requests to the server in a non-client pull fashion was causing delays. The server is designed so that if it receives multiple requests before it is ready with an update, it will just send one update in response.

When the Message Accelerator sent requests at a very fast rate, the server would frequently process and then ignore requests that came too close together, so it was doing all of the work of reading and parsing a request without sending more frequent updates in response. We expected there would be a performance advantage in sending requests to the server in a client-pull fashion, but with the Message Accelerator on a machine next to the client, we could not do that without suffering from network latency in the same manner as a system without the Message Accelerator.

So we added another accelerator, one that was placed on the same machine as the server. This new *Server Accelerator* read requests from the *Client Accelerator*, and forwarded them to the server in a client-pull system, then forwarded the updates it got from the server back to the Client Accelerator. When the Server Accelerator was able to write to the server, it checked if it had received a new request from the client, and if it had been sent a full update in response to the last sent request. If both those checks passed, it sent a new request to the server. We were still depending on the Client Accelerator to set the rate at which the requests were sent.

D. Automatically Detecting Request Rate

Our original design for the Message Accelerator automatically determined the amount of time it took the client to process the request, completely separate from the time it took the update to get to the client over the network. The Message Accelerator would then send requests to the server at that rate. In this way, the system would be able to run at optimal speed, regardless of network latency. If the rate at which the client received requests was reflected in the measurement of how fast it was processing, we could suffer from a drift effect in the feedback loop, where the rate at which the Message Accelerator sent requests got slower in response to the client receiving and processing messages more slowly, which in turn was due to the slowing Message Accelerator request rate.

Exactly how to measure pure processing was unclear, especially given the interleaving of network reads and processing. We once more had to add knowledge and processing of the updates into the Client Accelerator, after taking them out when the Client Accelerator was simply forwarding all data from the server. However, the Client Accelerator continued to forward data from the server immediately after receiving it, sending partial updates to the client, in order to keep the performance advantages.

Measuring the time from when the Client Accelerator began sending the update to when it received the next request from the client would completely capture the time it took the client to process the update. However, it would also be highly affected by the time it took to send the update to the client, especially now that the Client Accelerator was forwarding partial updates as it received them. We tried measuring the time from after the Client Accelerator had finished sending the update to when it received a new request from the client, but that returned very small times that did not capture the amount

of time the accelerator was spending processing, again due to the interleaved reads.

In experiments with the Server Accelerator, we had observed that the server sent updates to the Server Accelerator much faster than the Server Accelerator could send them to the client. This resulted in a slow-motion effect when the client displayed the updates, since they were being shown on the client at a much slower rate than they had been generated on the server. To limit this effect, we put in checks so that the Server Accelerator would only prefetch a limited number of updates, and would wait to send requests to the server until it had sent these updates to the client. We noticed that with higher network latencies, the Server Accelerator was sending requests to the server at the rate at which it could send updates to the client, since this was higher than the rate at which the Client Accelerator was sending requests. Because we were checking that the client would not block before we wrote to it, the Server Accelerator could not send faster than the network could handle. Because of the interleaving of network read and processing, the Server Accelerator could not send faster than the client could process. Since the Server Accelerator is always sending as much data as it can without blocking regardless of how long it takes to hear back from the client, it is not affected by network latency in the same way as a purely client-pull system.

We removed the Client Accelerator, and had the Server Accelerator send to the client as fast as it could without blocking. When the Server Accelerator can write to the server, it checks to see if it has received a full update in response to the previous request, and if it has already sent all of the updates its received to the client. If both checks pass, it sends a request to the server.

With this simple mechanism, we no longer had to worry about separating processing time from network time. We also no longer had to calculate an explicit rate at which to send updates. By limiting the number of updates the Message Accelerator prefetched, we ensured that the updates were being generated by the server at the same rate they were being received and processed by the client, and thus we avoided any slow motion effects.

E. TCP Window Size

In all versions of our accelerator that featured the Client Accelerator, we found that making changes to TCP window size had a significant effect on performance. We experimented with very small and very large window sizes, and discovered that the optimal performance was when window sizes were large enough to allow a number of updates within a TCP window, but small enough that dropped packets were discovered relatively quickly. The TCP defaults resulted in our system running with a median time of 257 ms per update with a 100 ms network delay. Changing the TCP window size to 32 MB resulted in a median time of 96 ms with a network latency of 100 ms, and additionally setting the Selective Acknowledgment flag got our median time down to 85 ms.

However, when we removed the Client Accelerator and once more had just the server and client machine, we found that changing the TCP window size and flags had very little effect on performance. With a network latency of 100 ms, the Message Accelerator had a median time of 79 ms with the default settings, a median time of 81 ms with a window size of 32 MB, and a median time of 79 ms with the higher window size and the Selective Acknowledgment flag.

VII. FUTURE WORK

Using an adaptive message-accelerating proxy to improve the performance of client-pull systems may be an approach that can be extended to other systems. The message modifying program offers a simple way to improve performance or add additional features, is easy to deploy, and works with existing binaries. It does not suffer from problems of parallel code maintenance, and will continue to work as long as the message format between the client and server remains the same. The proxy can even be deployed to a different machine from the server, and still offer performance advantages.

There are additional ways that a modifying proxy could be used with thin client systems. The proxy could dynamically tightly compress updates when network speeds were low, and uncompress when the client device had low batteries or other computational issues. With an additional client application, the server application could encrypt updates, and the client application could decrypt them either on the client machine, or on a machine with a trusted network connection to the client. The server application could also perform machine vision tasks such as object detection or face recognition.

Clearly, adding an adaptive proxy to a client-pull system offers any number of ways to improve performance or transform data. Updates can be buffered, information can be cached, or messages can be modified in a wide variety of ways. While some systems have taken some of these approaches, we are not aware of any other system that uses a proxy for message acceleration.

VIII. CONCLUSION

Adding a Message Accelerator proxy to a VNC system is a very simple but highly effective way of improving video performance with VNC under high latency conditions tenfold. Even with small amounts of network latency, video performance is as good or better than an unmodified VNC system. It is easy to install, requiring no recompilation of client or server code. The Message Accelerator virtually erases frame-rate reduction due to the negative effects of network latency on client-pull systems. A video displayed using the Message Accelerator system will look almost the same, in terms of frame rate, with a 300 ms network latency as it does with a 3 ms network latency, while a video playing on an unmodified system will take ten times as long to display updates when latency increases to 100 ms as it would at 3 ms.

REFERENCES

- [1] S. Li, Q. Stafford-Fraser, and A. Hopper, "Integrating synchronous and asynchronous collaboration with virtual network computing," *IEEE Internet Computing*, vol. 4, no. 3, pp. 26–33, 2000.
- [2] A. Adler, J. Nash, and S. No "el, "TellTable: a server for collaborative office applications," in *Proc. CSCW*, 2004, pp. 6–10.
- [3] Y. Liu and O. J. Anshus, "Improving the performance of vnc for high-resolution display walls," *Collaborative Technologies and Systems, International Symposium on*, vol. 0, pp. 376–383, 2009.
- [4] D. Stødle, J. M. Bjørndalen, and O. J. Anshus, "De-Centralizing the VNC Model for Improved Performance on Wall-Sized, High-Resolution Tiled Displays," in *NIK '07: Norsk Informatikkonferanse*. tapir akademisk forlag, Nov. 2007, pp. 53–64.
- [5] L. Renambot, B. Jeong, R. Jagodic, A. Johnson, and J. Leigh, "Collaborative Visualization using High-Resolution Tiled Displays," *ACM CHI Workshop on Information Visualization Interaction Techniques for Collaboration Across Multiple Displays*, 2006.
- [6] K. Nakashima, T. Machida, H. Takemura, and K. Kiyokawa, "A 2D3D integrated tabletop environment for multi-user collaboration," *Computer Animation and Virtual Worlds*, vol. 18, no. 1, pp. 39–56, 2007.
- [7] C. Lawton, "'Dumb Terminals' Can Be a Smart Move," *The Wall Street Journal*, January 2007.
- [8] D. Tynan, "Think Thin," *InfoWorld*, July 2005.
- [9] A. Lai and J. Niegh, "Limits of Wide-Area Thin-Client Computing," *Proc. SIGMETRICS 2002*, pp. 228–239, 2002.
- [10] B. K. Schmidt, M. S. Lam, and J. D. Northcutt, "The interactive performance of slim: a stateless, thin-client architecture," in *Proc. SOSP '99*. ACM Press, 1999, pp. 32–47.
- [11] T. Richardson, Q. Stafford-Fraser, K. Wood, and A. Hopper, "Virtual network computing," *Internet Computing*, vol. 2, no. 1, pp. 33–38, 1998.
- [12] T. Richardson, "The RFB Protocol," RealVNC Ltd, Tech. Rep., 2007.
- [13] "RealVNC," <http://www.realvnc.com/>. [Online]. Available: <http://www.realvnc.com/>
- [14] "Chicken of the VNC," <http://sourceforge.net/projects/cotvnc/>. [Online]. Available: <http://sourceforge.net/projects/cotvnc/>
- [15] "Tight VNC," <http://www.tightvnc.com/>. [Online]. Available: <http://www.tightvnc.com/>
- [16] "UltraVNC," <http://www.uvnc.com/>. [Online]. Available: <http://www.uvnc.com/>
- [17] T. Haraikawa, T. Sakamoto, T. Hase, T. Mizuno, and A. Togashi, "μvnc over plc: a framework for gui-based remote operation of home appliances through power-line communication," *Consumer Electronics, IEEE Transactions on*, vol. 48, no. 4, pp. 1067–1074, Nov 2002.
- [18] B. Shizuki, M. Nakasu, and J. Tanaka, "VNC-based access to remote computers from cellular phones," in *Proceedings of the IASTED International Conference on Communication Systems and Networks (CSN 2002)*. Citeseer, 2002, pp. 74–79.
- [19] X. Zhang and H. Takahashi, "A Hybrid Data Compression Scheme for Improved VNC," *Systemics, Cybernetics and Informatics*, vol. 5, no. 2, pp. 1–4, 2007.
- [20] K. Kaplinsky, "VNC tight encoder-data compression for VNC," in *Modern Techniques and Technology, 2001. MTT 2001. Proceedings of the 7th International Scientific and Practical Conference of Students, Post-graduates and Young Scientists*, 2001, pp. 155–157.
- [21] S. Rao, H. Vin, and A. Tarafdar, "Comparative Evaluation of Server-push and Client-pull Architectures for Multimedia Servers," in *Proceedings of NOSSDAV*, pp. 45–48, 1996.
- [22] R. A. Baratto, J. Nieh, and L. Kim, "THINC: A Remote Display Architecture for Thin-Client Computing," Department of Computer Science, Columbia University, Computing Science Technical Report CUCS-027-04, 2004.
- [23] J. Nieh, S. J. Yang, and N. Novik, "A Comparison of Thin-Client Computing Architectures," Department of Computer Science, Columbia University, Computing Science Technical Report CUCS-022-00, 2000.
- [24] —, "Measuring thin-client performance using slow-motion benchmarking," *ACM Trans. Comput. Syst.*, vol. 21, no. 1, pp. 87–115, 2003.
- [25] "NetEm," <http://www.linux-foundation.org/en/Net:Netem>. [Online]. Available: <http://www.linux-foundation.org/en/Net:Netem>