# Performance Comparison of Windows-based Thin-Client Architectures

Daniel Schlosser, Andreas Binzenhöfer, and Barbara Staehle
University of Würzburg, Institute of Computer Science
Würzburg, Germany
Email: [schlosser—binzenhoefer—bstaehle]@informatik.uni-wuerzburg.de

***Abstract*—The basic idea behind thin-client architectures is to run applications on a central server instead of installing them separately on each client. The Windows Remote Desktop Protocol (RDP) and the Citrix Presentation Server are two well known approaches to separate the location of where the user input is processed from the computer he is actually working on. While both alternatives solve the same problem, they rely on significantly different mechanisms to handle the exchange of user input and screen updates between client and server.**

**In this paper we therefore compare the performance of both protocols under different aspects. In particular, we study the load caused on network layer as well as the satisfaction of the end user with the service quality achieved by the different terminal services. As this performance heavily depends on the current network conditions, we emulate realistic scenarios in a controlled testbed environment and measure the time required for typical office tasks on application layer. As a result, we quantify the Quality-of-Experience (QoE) perceived by the end-user, compare the overhead required by the different available protocols, and unveil their advantages and disadvantages. Our results can be used to decide which protocol to use in which scenario.**

## I. Introduction

A significant part of the expenditures of companies goes into purchasing, maintaining, and updating their IT infrastructure as well as the software needed by their employees. In this context, the thin-client paradigm offers a great alternative to traditional approaches. Instead of installing and maintaining the required applications separately on each client, all computing resources are centralized on one single machine or a server farm. The thin-client itself is a small low-cost computer mainly used as a display which also receives and forwards the input of the user. This process of sharing centralized hardware brings along many benefits like a more effective utilization of computer resources, easier maintenance of software versions, and a significant reduction of the total cost of ownership.

There are various choices of how to realize a thin-client architecture. Microsoft's Remote Desktop Protocol (RDP) [1] comes with all current versions of Windows and offers the possibility to connect to a remote computer, typically a Terminal Server, and use it as if sitting in front of it. The Citrix Presentation Server [2] is a remote application publishing tool which extends the basic Windows Terminal Services (WTS) platform. It relies on the Independent Computing Architecture (ICA) protocol which transmits user input to the server and high-level window display information to the client. From a technical point of view, the considered architectures differ in the amount of overhead created, the robustness against network failures, and the responsiveness to user input. However, independent of such technical aspects, it is the satisfaction of the user with the service which counts at the end of the day. This satisfaction might, e.g., be quantified by how much additional time a user needs to do his work in a thin-client environment as compared to working on a local PC.

In this paper, we therefore address the need to obtain an objective comparison of the different terminal services. We expose the two considered protocols to controlled network conditions in a realistic testbed environment and compare their performance in terms of bandwidth and user satisfaction. The Quality-of-Experience (QoE) perceived by the end-user heavily depends on Quality-of-Service (QoS) parameters such as packet loss (PL) and network round trip time (RTT). We therefore vary those parameters in our testbed and measure the time it takes an end-user to perform typical tasks like typing text, navigating through menu-entries, or scrolling through a document. The results reveal the advantages and disadvantages of the different approaches and can therefore be used as a much needed guideline for which protocol to use in which situation.

The remainder of this paper is organized as follows. Section II gives a brief overview of related work. The measurement setup as well as the measurement methodology are explained in detail in Section III. We discuss the results and their implications in Section IV, while Section V finally concludes this work.

## II. Related Work

The number of providers hosting office applications for remote users is growing, but the percentage of bandwidth consumed by this type of service remains negligible. Therefore, little work has been dedicated to analyze the characteristics of traffic caused by thin-client based applications. One of the few studies in this area has been performed by the Tolly Group [3] who evaluated the usability of Microsoft PowerPoint via WAN. They examined the consumed bandwidth and completion time of a common PowerPoint operation executed on a machine running Citrix MetaFrame XP client software accessing a server hosting the corresponding Presentation Server. They report "on machine" experience for high bandwidth links like Fast Ethernet, Ethernet, or WLAN and up to twice as long completion times for slower links.

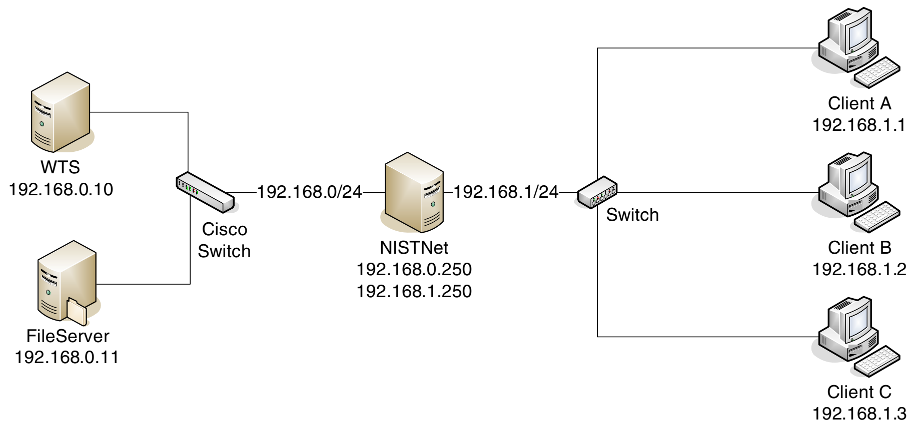


Fig. 1. Overview of the testbed environment

Schmidt et al. [4] examined a stateless thin-client architecture in order to show that LAN environments are able to support real time interactions with programs running on a remote server. Tolia et al. [5] evaluated VNC performance, focusing on the operation response times for image processing, presentation creating, and text processing. Their findings show that highly interactive applications are more sensitive to a higher network round trip time than simpler applications. More general measurements are reported by Lai and Nieh [6], who evaluated the performance of thin-client computing in a WAN. For their work, they used both measurements in the Internet2 and simulations in a testbed. One of the investigated characteristics was the completion time of typical office tasks like typing, scrolling, or image downloading. To analyze the difference between the considered platforms, they compared the amount of data transferred before the completion of each test, the required bandwidth, and the test duration. The fact that the typing task was completed faster, and required less bandwidth than scrolling already gave a first insight into the structure of traffic caused by office applications. Finally, the issue of benchmarking the user perceived performance on a thin-client system is addressed in [7]. The authors use a so called slow-motion benchmark, which relies on monitoring network traffic corresponding to a user's action, in order to measure the performance of several thin-client platforms.

In this work, we compare the traffic generated by the different terminal services and investigate in how far different network conditions influence the usability of office applications. In particular, we distinguish between the client and the server side, different delay and packet loss schemes as well as their influence on the performance of Microsoft Office applications.

## III. Measurement Setup and Methodology

In order to emulate a standard environment for thin-client computing, we set up a testbed as depicted in Figure 1. For the server side we use two 3.4 GHz Intel Xeon servers with 3.5 GB RAM each, running on Windows 2003 Server standard edition with Service Pack 1. The Windows Terminal Server (WTS) in Figure 1 is responsible for hosting the server side applications of the tested terminal services. The second server is set up as a file server and used to store user data. In order to emulate varying network conditions we use a NISTNet [8] machine in the middle of the communication channel. For this task we use a Dual Pentium III 500Mhz computer with 512 MB RAM running OpenSuSE 10.0 and NISTNet 2.1012.c. Note, that both the processing power as well as the memory of this machine are well above the minimum requirements for NISTNet. On the client side we use a Pentium IV 2.6 GHz machine with 1 GB RAM running Windows XP with Service Pack 2. All hosts are connected using 100 Mbit. We dimensioned the hosts and the network in such a way, that none of these components is a bottleneck and the performance of the applications is only affected by the used terminal service solution as well as by the emulated network conditions.

As software applications we regarded Word and Excel, the two most popular products of the Microsoft Office family. Both applications were used to compare the performance of the standard Remote Desktop Protocol (RDP) client included with Microsoft Windows XP in version 5.1.2600 and the Citrix Metaframe Presentation Server 4.0 in combination with the ICA Client in version 9.237. For RDP, the color depth was set to 16 bit and the persistent bitmap cache was enabled as preset in standard mode. For the evaluation of the Citrix Metaframe environment, data compression and session reliability were enabled and the color depth was also set to 16 bit. We did not use any other options, which might influence the performance of the client like Speed Screen Latency Reduction or local buffer for mouse and keyboard.

In order to expose both architectures to the same user behavior, the entire user input is done automatically. At first, the client starts the corresponding application. Then, a test consisting of some typical office tasks is performed several times. For the evaluation of Word we consider three different user tasks. The first task is to enter some standard text in English language. In order to capture the performance of the typing test, we continuously check if the typed text is actually displayed on the screen. To do so, the test waits until some predefined pixels turn to a specific color. The second task is to scroll through the text. Thereby, we also verify, that the scrolling is done correctly. The last task in the Word test is to select some options in the menu. Again, it is verified that the menu shows up and the correct entry is selected. For the Excel measurements, we also emulate a user performing three typical tasks. First, some values are entered into an empty Excel spreadsheet. Then, these values are selected using the mouse pointer and finally a bar chart is created. As in the Word test, all steps of each task are verified before the next task is executed. We used the open-source tool AutoHotkey [9] to reproduce the same input behavior, to check whether the changes have been applied properly, and to measure the duration of each task.

To obtain credible measurement results, we repeatedly performed each test for the duration of an hour. Within this hour the network emulation settings on the NISTNet machine remain unchanged. During the entire time we record all network traffic using WinDump [10]. Once the measurement is over, we collect all data and reset the testbed. The tests that were done in the first five minutes and the last five minutes of the hour are discarded in oder to prevent incorrect data caused by a transient time which might occur at the NISTNet machine.

## IV. Results

In this section we present the results of the measurements performed in our testbed environment. We compare the performance of Citrix and RDP under two main aspects, the processes on network layer as well as the user perceived quality on application layer. That is, we take a closer look at the network traffic produced by the two protocols and investigate in how far network parameters like delay and packet loss influence the service quality as perceived by the end-user. Note, that both protocols use TCP/IP connections between the thin-client and the corresponding server in order to establish a reliable connection.

### A. Bandwidth Usage on Client Side

At first we take a closer look at the bandwidth used by the clients during our measurements. To do so, we repeatedly performed the three Word tests scrolling, menu, and typing with both the Citrix client and the RDP client. Figure 2 shows a 20 s moving average of the bandwidth sent by the client in uplink direction during an emulation time of 100 s. Packets sent from the client to the server mainly encode user input like keystrokes or mouse movements. Note, that since the tasks have periodically been repeated during the emulation, the bandwidth consumption remains nearly constant over time. Thereby, there is almost no difference between the two protocols when performing the typing test, which requires approximately 4 kbps in both cases. The Citrix client, however, requires significantly more bandwidth during

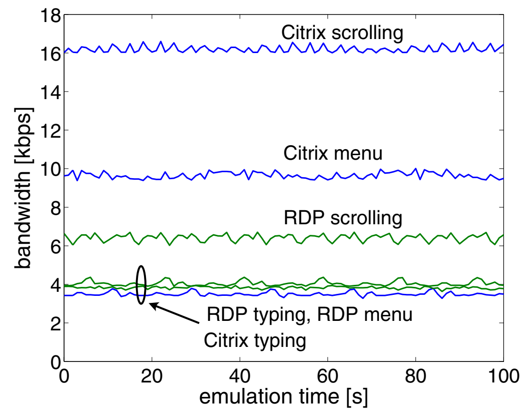


Fig. 2. Bandwidth used by the client for different MS Word tasks.

the other two tests. The reason for this is that both tests involve mouse movement which forces Citrix to send more packets than in the case of merely encoding and transmitting simple keystrokes.

To illustrate this in more detail, we take a closer look at the activities on packet layer during the Word scrolling test. Figure 3 visualizes the packet sizes sent by both the Citrix and the RDP client during a randomly selected 5 second interval. In both cases there are almost no TCP acknowledgments

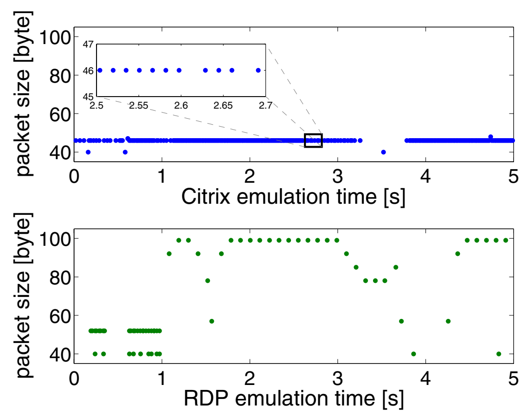


Fig. 3. Packet sizes sent by Citrix and RDP during the Word scrolling test.

(packets of size 40 byte), since our automated end-user continuously produces user input and packets sent by the server can therefore be acknowledged using regular data packets. In terms of payload, the RDP client sends packets of varying size in order to inform the server about updates on client side. The Citrix client, however, encodes all user input using a fixed payload size of 6 byte and immediately transmits it over the network. That is, the 16 kbps used during the scrolling test (cf. Figure 2) are not realized by accumulating information and sending larger packets but by sending small packets of fixed size in very short intervals. This can also nicely be seen within the zoom area in Figure 3 which shows the packets sent during a time period of 200 ms. This result is in agreement with previous work [11], which showed that in the design of Citrix responsiveness had played a more important role than protocol overhead.

### *B. Bandwidth Usage on Server Side*

Since the differences between Citrix and RDP show a similar behavior for both Word and Excel, we now concentrate on the latter on the server side. Figure 4 shows a 20 s moving average of the bandwidth sent by the server for two different Excel tasks: selecting some values and inserting a chart. The values in brackets indicate the average amount of bandwidth which was used by TCP acknowledgments. That is, the Citrix server mainly sends acknowledgments and only a small portion of the overall bandwidth is actually used to transmit information like screen updates from the server to the client. When neglecting the acknowledgments and concentrating on actually transmitted payload, the Citrix server requires roughly the same amount of bandwidth as the RDP server during the selecting test. Regarding the insertion of a chart, however, the RDP server requires significantly more bandwidth than the Citrix server. A possible reason is that Citrix is optimized for office applications and probably merely transmits some meta information while RDP needs to refresh the entire affected area on the screen of the client. We repeatedly observed this behavior during other measurements which involved graphic-intensive tasks like the insertion of a picture into a Word document or PowerPoint animations.

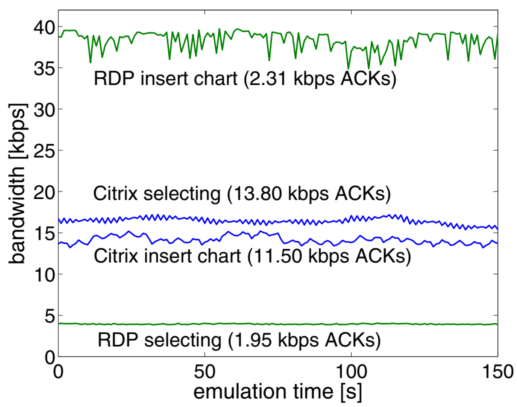


Fig. 4. Bandwidth used by the server for different Excel tasks.

To understand this phenomenon in more detail, we take a closer look at the packets sent by the server during the insertion of a chart into an Excel spreadsheet. Figure 5 shows the empirical cumulative distribution function (CDF) of the packet sizes sent by the Citrix server as well as by the RDP server while the user inserts the chart. As mentioned above, the Citrix server mainly sends acknowledgments (roughly 90% of all packets) to inform the client that it received the input of the user. It only needs a few larger packets to transmit the actual screen updates and other additional information. The RDP server on the other hand requires significantly more bandwidth to visualize the same screen refreshes on the client side. In particular, it only sends about 45% acknowledgments, while more than 10% of all packets have the maximum possible size. Note, that in our case the MTU was set to 1394 byte since we used an additional VPN header.

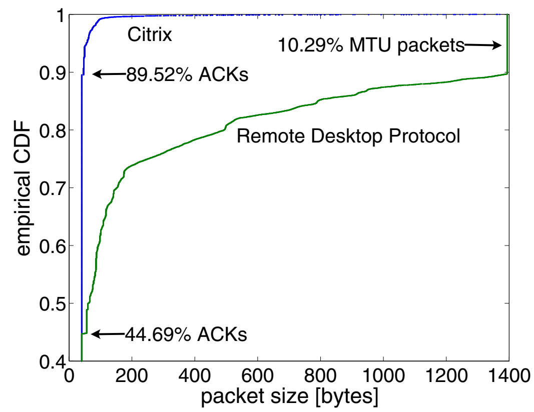


Fig. 5. Empirical CDF of the packet sizes sent during the Excel chart test.

### *C. Influence of Network Parameters*

So far we concentrated on thin-client traffic under perfect network conditions. In this section we study in how far network parameters like packet loss and delay influence the bandwidth used by the different thin-clients. To gain a better overview, we use a simple approach which is based on methods used in the design of experiments. That is, we perform the same measurements for a low Round Trip Time (RTT) of 0 ms and a high RTT of 500 ms as well as for a low packet loss (PL) of 0% and a high PL of 2%. From this we then derive what factors mainly influence the traffic sent by thin-clients. Figure 6 shows the bandwidth used by a

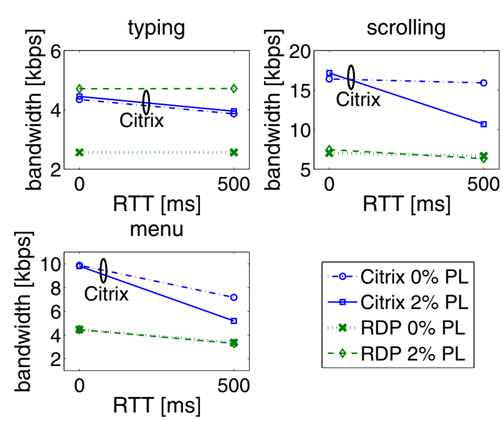


Fig. 6. Influence of network parameters on the bandwidth usage

Citrix client and an RDP client for different Word tasks under different network conditions. For all three tasks, the bandwidth sent by the RDP client is almost independent of the RTT in the network. This is especially true for the typing test, which, however, is the only test in the RDP setup which is heavily influenced by packet loss. That is, when faced with packet loss, the RDP client uses almost twice as much bandwidth by sending more and larger packets. This is probably due to the fact, that the most important information sent from the client to the server are mouse clicks and keystrokes. While the scrolling test and the menu test mainly involve mouse movements, the typing test solely consists of keystrokes and therefore probably forces the RDP client to quickly repeat its transmission on application layer.

With the exception of the typing task under packet loss, the Citrix client always requires more bandwidth than the RDP client for the same tasks. Unlike the RDP client, the bandwidth sent by the Citrix client also depends on the current RTT in the network. This is especially true for the scrolling test under packet loss. In this case, we observed that if server sent packets get lost the Citrix client occasionally stops the transmission of new packets. We believe that the Citrix client waits for an acknowledgment on application level before it accepts and forwards any additional user input. The following section therefore investigates how these packet loss and packet delay scenarios translate into user-perceived quality.

### D. Quality-of-Experience for Different Network Conditions

In this section we compare the Quality-of-Experience (QoE) of the two different thin-client architectures and analyze in how far it is influenced by Quality-of-Service (QoS) parameters like the packet loss or the RTT. To do so, we measure the duration of the Word and Excel tests described in Section III under different network conditions. Thereby, the test duration is defined as the time it takes our automated end-user to successfully perform all three Word or Excel tasks, respectively. Also note, that in the following we use the median instead of the mean of all measured test durations as it is statistically robust against outliers.

Figure 7 shows the median of the duration of the Word tests for both Citrix and RDP in dependence of the current RTT in the network. The error bars represent the interquartile range, i.e. the difference between the third and first quartiles, as a measure of statistical dispersion. If there is no delay in the network, both protocols achieve approximately the same performance, with Citrix having a slight edge over RDP. Moreover, the test duration is almost deterministic in this case as indicated by the very small interquartile range. Even a packet loss of 2% does not influence the test duration in this case. As long as there is no packet loss in the network, there is not much difference between Citrix and RDP. The test durations increase with an increasing RTT but still remain relatively constant for a given RTT. Regarding the curves for 2% packet loss, however, RDP clearly outperforms Citrix while the test durations for both protocols vary significantly.

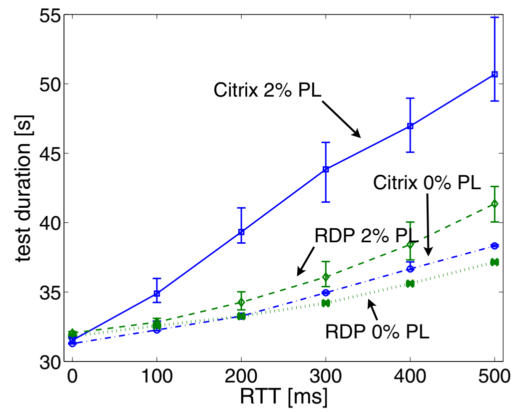


Fig. 7. Test durations using Word and different network delays

That is, the test durations vary during packet loss depending on how much lost information needs to be retransmitted between client and server. In any case, when using Word, RDP is able to handle packet loss better than Citrix.

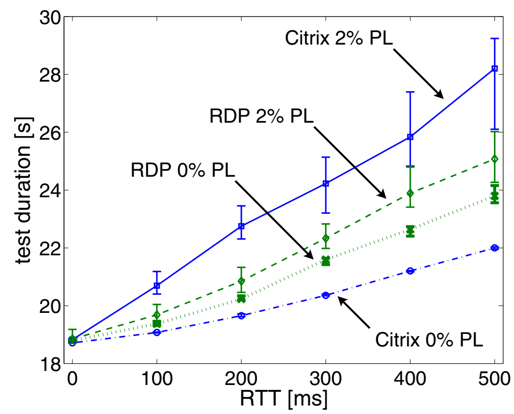


Fig. 8. Test durations using Excel and different network delays

We obtain similar results for the durations of the Excel tests as shown in Figure 8. Again, the test durations are almost deterministic as long as there is no packet loss in the network. This time, however, Citrix achieves noticeably shorter test durations in scenarios without packet loss. This is mainly due to the insertion of the Excel chart which takes longer under RDP than under Citrix. In exchange, RDP is hardly influenced by the 2% packet loss, while the test durations for Citrix increase considerably under packet loss and additionally show a higher variance.

In order to analyze the influence of packet loss on the performance of the two thin-client approaches in more detail, we take a closer look at the test durations for a fixed network delay and varying packet loss. Figure 9 shows the median test duration as well as the interquartile range for the Word tests in such scenarios. The packet loss was varied between 0% and 3% in steps of 0.5%, the RTT was set to 0 ms and 200 ms,

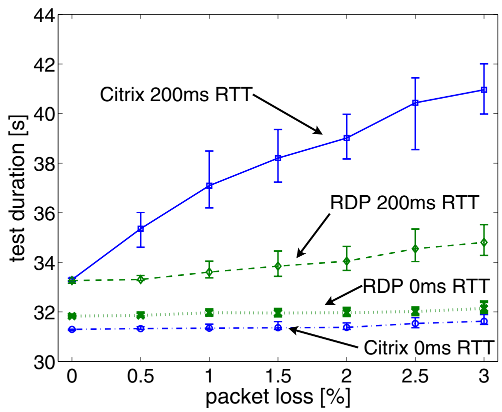


Fig. 9. Test duration using Word and varying packet loss

respectively. The figure again illustrates that if the network delay is very low, packet loss does not have much influence. As soon as we introduce some network delay, however, Citrix is again much more sensitive to packet loss than RDP. As the packet loss increases from 0% to 3%, the median duration of the tests increases by 4.5% for RDP and by 23.1% for Citrix. We obtain a very similar behavior for Excel and therefore omit the corresponding graphs. Using Excel, the test durations are slightly closer to each other, whereas the relative increase of the test duration for an RTT of 200 ms is 4.3% for RDP and 22.8% for Citrix.

## V. Conclusion

Thin-client architectures are becoming increasingly popular in company networks as they offer a more efficient utilization of computer resources, easier maintenance of software versions, and a significant reduction of the total cost of ownership. However, there are only very few research studies which address this topic. In this paper we therefore compared the performance of two popular Windows-based thin-client architectures, the Citrix Presentation Server and Microsoft's Remote Desktop Protocol. For both approaches, we performed extensive measurements in a realistic testbed environment using emulated network conditions. In particular, we analyzed the traffic sent on network layer as well as the service quality perceived by the end-user. The latter was realized by comparing the time needed to perform typical tasks in Microsoft Word and Excel.

Since Citrix is optimized for Microsoft Office applications, it achieves a slightly better Quality-of-Experience under perfect network conditions. This is especially true for graphic-intensive tasks like the creation of an Excel chart or the insertion of a picture into a Word document. Typical network parameters like packet loss and delay influence the user perceived QoE to a different extent. An increase of the round trip time in the network affects both considered thin-client approaches in a similar way. The test duration increases with an increasing RTT but remains relatively constant for a given RTT. As soon as there is packet loss in the network, however, the duration of the individual test becomes more variant. The reason is that packet loss occurs randomly and thus each test run is affected to a different degree by packet loss. We also observed that RDP clearly outperforms Citrix under packet loss. That is, the individual office tasks take significantly longer using the Citrix client as compared to the RDP client when there is packet loss in the network.

Based on our measurements, we conclude that Citrix should be prefered in Local Area Networks as it offers a better user experience under perfect network conditions. For Wide Area Networks or other connections which are more vulnerable to packet loss, however, the RDP client might be chosen as it is more robust against short interruptions in the packet flow. In future work we will analyze additional applications like video streaming and further improve our understanding of the relation between QoS in the network and the QoE perceived by the user when working with thin-client architectures.

## Acknowledgments

The authors would like to thank Phuoc Tran-Gia for the insightful discussions during the course of this work and Björn Boder for his helping comments as well as for providing the hardware necessary to conduct the studies. Further acknowledgments go to Markus Weiß for assisting during the course of the measurements.

## References

[1] Microsoft, "Remote Desktop Protocol," http://www.microsoft.com/technet/prodtechnol/Win2KTS/evaluate/featfunc/rdpfperf.mspx, 2007.

[2] Citrix Systems, Inc., "Citrix Presentation Server," http://www.citrix.com/English/ps2/products/product.asp?contentID=186, 2007.

[3] The Tolly Group, "Citrix Systems, Inc. MetaFrame XP Presentation Server Windows-based Application Access Performance and Functionality," 2003.

[4] B. K. Schmidt, M. S. Lam, and J. D. Northcutt, "The Interactive Performance of SLIM: a Stateless, Thin-Client Architecture," in *SOSP '99: Proceedings of the seventeenth ACM symposium on Operating systems principles*, Charleston, South Carolina, United States, 1999.

[5] N. Tolia, D. G. Andersen, and M. Satyanarayanan, "Quantifying Interactive User Experience on Thin Clients," *IEEE Computer Society Press Computer*, vol. 39, no. 3, 2006.

[6] A. M. Lai and J. Nieh, "On the Performance of Wide-Area Thin-Client Computing," *ACM Transactions on Computer Systems*, vol. 24, no. 2, 2006.

[7] J. Nieh, S. J. Yang, and N. Novik, "Measuring Thin-Client Performance using Slow-Motion Benchmarking," *ACM Transactions on Computer Systems*, vol. 21, no. 1, 2003.

[8] M. Carson and D. Santay, "NistNet: a Linux-based network emulation tool," *SIGCOMM Computer Communication Review*, vol. 33, no. 3, 2003.

[9] AutoHotkey, http://www.autohotkey.com/.

[10] WinPcap, http://www.winpcap.org/.

[11] B. Emmert, A. Binzenhöfer, D. Schlosser, and M. Weiß, "Source Traffic Characterization for Thin Client Based Office Applications," in *13th EUNICE and IFIP TC6.6 Workshop on Dependable and Adaptable Networks and Services*, University of Twente, the Netherlands, July 2007.