

The Effect of Latency on User Performance in Warcraft III^{®*}

Nathan Sheldon, Eric Girard, Seth Borg, Mark Claypool, Emmanuel Agu
{claypool,emmanuel}@cs.wpi.edu
CS Dept., Worcester Polytechnic Institute
100 Institute Road,
Worcester, MA 01609, USA

Abstract

Latency on the Internet is a well-known problem for interactive applications. With the increase in interactive network games comes the increased importance of understanding the effects of latency on user performance. Classes of network games such as First Person Shooters (FPS) and Real Time Strategy (RTS) differ in their user interaction model and hence susceptibility to latency. While previous work has measured the effects of latency on FPS games, there has been no systematic investigation of the effects of latency on RTS games. In this work, we design and conduct user studies that measure the impact of latency on user performance in Warcraft III, a popular RTS game. As a foundation for the research, we separated typical Warcraft III user interactions into the basic components of explore, build and combat, and analyzed each individually. We find modest statistical correlations between user performance and latency for exploration, but very weak correlations for building and combat. Overall, the effect of even very high latency, while noticeable to users, has a negligible effect on the outcome of the game. We attribute this somewhat surprising result to the nature of RTS game-play that clearly favors strategy over the real-time aspects.

*Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. NetGames'03, May 22-23, 2003, Redwood City, CA, USA. Copyright 2003 ACM 1-58113-734-6/03/0005...\$5.00

1 Introduction

Over the past decade, the Internet has grown in popularity and capability at exceptional rates. In 1997, there were 36.6 million homes with computers and only 18 million of them had Internet access [6]. By the year 2000, the number of homes with computers had grown to 51 million, 41.5 million of which had Internet access, and many with broadband Internet connections such as cable modems and DSL lines.

This growth in popularity and capability of the Internet has led to an increasingly diverse set of applications with varying network behaviors and requirements. Key metrics of latency and throughput are useful in characterizing application behavior. Traditional applications such as file transfer, Usenet news and email are primarily concerned with throughput and can tolerate delays on the order of minutes. Web browsers are also concerned about throughput, but the interactive nature of browsing requires latencies on the order of seconds or at most tens of seconds [4]. Emerging real-time applications such as IP telephony and networked games typically have the lowest throughput requirements but are even less tolerant of latency than other applications. Knowing how these real-time applications react to latency and loss is the crucial first step in designing the next generation network hardware and software that will support their requirements. In addition, classifications of real-time applications according to latency tolerance will enable designers, developers and engineers to make informed decisions on appropriate quality for classes under such architectures as Diff-Serv [5].

The most popular real-time applications are

multi-player network computer games that can make up around half of the top 25 types of non-traditional traffic for some Internet links [13] and are predicted to make up over 25 percent of Local Area Network (LAN) traffic by the year 2010. In 2000, the U.S. economy only grew 7.4% while the computer and video game industry grew by 14.9%, out-pacing growth in other high-tech industries and even Hollywood over the previous five years [10]. In 2002, over 221 million computer and video games were sold, or almost two games for every household in America.¹ Knowledge of how network related issues, such as latency and packet loss, affect the usability of games can be of great use to the companies that make these games, network software and equipment manufacturers, Internet Service Providers (ISPs), and the research community at large. In particular, if established latency requirements and any associated trade-offs were known, ISPs could establish tariffs based on customers' indicated maximum delays, requested Quality of service (QoS) and the ISP's ability to meet these demands.

Two of the most popular categories of real-time network games are *First Person Shooter* (FPS) games and *Real Time Strategy* (RTS) games. FPS games, first made popular by *Doom*,² have the player view the world through the eyes of a character (the first person). They then move around slaying monsters and other players, with an amalgamation of ranged weaponry (the shooter). RTS games, first made popular by *Dune 2*,³ are generally characterized by resource collection, unit construction, and battles that consist of large numbers of animated soldiers going through a repetitive, animated attack.

While there has been research qualitatively characterizing the effects of latency for some popular FPS games [2, 9] as well as a general awareness of latency issues [3, 11, 14], quantitative studies of the effects of latency on network games have been lacking. Moreover, it is unlikely that FPS games have the same network requirements as do RTS games. In FPS games, exact positioning and timing is re-

quired, because, for example, a target must still be at the location where the player aimed in order for the shot to hit. In RTS games, the positioning and timing is more forgiving because, for instance, a command can be issued to attack a unit, regardless of its current location or its direction and time of movement.

This work studies the effect of latency on user performance and network traffic for the best selling [15] RTS game, *Warcraft III*.⁴ We quantify the effect of latency on user performance in *Warcraft III* by analyzing the results of controlled research experiments designed to measure quality of service over a range of induced latencies. As a foundation for RTS research, we divide *Warcraft III* up into fundamental game components of building, exploring and combat. We then develop multiple criteria for measuring user performance in *Warcraft III* and use these criteria in very carefully designed experiments to determine user performance over a range of latency conditions.

We find that latencies up to several seconds have little effect on the final outcomes of building, exploring, and most combat. Although, the effectiveness of certain strategies that involve precise timing of events are influenced by the amount of latency, very few such strategies prevail in typical *Warcraft III* games. Overall, strategy plays a much larger role in determining the outcome of the game than does latency. We conclude that *Warcraft III* should be placed in a different QoS class than applications with stringent latency constraints, such as FPS games or audio-conferences, since *Warcraft III* has latency requirements more similar to those of Web browsing.

The rest of this paper is organized as follows: Section 2 presents background information on *Warcraft III*; Section 3 describes our approach to measure the effects of latency on *Warcraft III*; Section 4 analyzes the application, network and user results from our experiments; Section 5 summarizes our conclusions; and Section 6 presents possible future work.

¹Top Ten Industry Facts, IDSA, <http://www.idsa.com/-pressroom.html>

²<http://www.idsoftware.com/games/doom/>

³<http://www.dune2k.com/duniverse/dune2/>

⁴<http://www.blizzard.com/war3/>

2 Background

Warcraft III is an RTS game in which players construct buildings and fighting units, and issue commands that cause the units to move, engage enemy units in battle, and build structures. Games are played on one of many possible maps, which are either provided with the game or custom built by players. During the game, Warcraft III uses a centralized server in a client-server architecture, either over the Internet with possibly thousands of participants or in a LAN scenario with a few players. For most Internet games, the server is via Battle.net,⁵ a free service that allows Blizzard’s Starcraft, Diablo and Warcraft players to initiate multi-player games over the Internet. For a LAN game, users can have a client’s machine act as a server, too, by choosing a map and then letting other clients join the game.

At the beginning of a game, players can choose to be a member of one of four “races” (Humans, Orcs, Undead and Night Elves). Our research focuses on the Humans, but since Blizzard put great effort into making the races equivalent in power, our results should generalize to the other races. There are a number of ways in which players can be competitively grouped. In a free-for-all game, all players vie to have the last remaining army on the map. Players can also team up against each other and/or against artificially intelligent computer-controlled players in myriad ways.

Figure 1 shows a screenshot of a Human town under attack from an Undead army. The Undead are in the upper left area of the screen. Human peasants can also be seen carrying lumber to the town hall and other activities.

Town control and unit control are two major aspects of the game. Town control consists of selecting what buildings are to be built or upgraded, what units are to be produced and what technologies are to be researched. In order to accomplish these tasks gold and lumber are required which are gathered by peasants. The Human buildings range from farms, which provide food, to town halls, which are the central building in any Human town. Buildings such as barracks produce standard army units such as footmen and knights while buildings



Figure 1: Screenshot of Undead Army Attacking a Town.

such as the arcane sanctum produce units such as priests and sorceresses. Town control involves a lot of strategy in knowing when and where to build, upgrade, and research.

Unit control can be broken up into three sub-categories: building, exploration and combat. Building overlaps with town control as it is the management of peasants in harvesting gold and lumber as well as building and repairing buildings. Exploration allows players to determine geography, find creeps (neutral enemies on which to gain experience for their heroes), enemy towns or units, and neutral shops at which their heroes can buy various items. Combat allows units to kill other units, to defend towns, secure territory or gain experience points for Heroes. There are various battle strategies that can be deployed, but at a minimum the player can let the computer’s artificial intelligence handle the units. Warcraft III has special units called heroes which are more powerful than standard units, and have additional abilities.

3 Approach

In order to empirically measure the effects of latency on Warcraft III, we employed the following methodology:

- Categorize user interactions in typical RTS

⁵<http://www.battle.net/>




Figure 2: RTS Component Maps: Build (left), Explore (middle), Combat (right).

games and construct Warcraft III campaigns that exercise each category (see Section 3.1).

- Determine criteria to quantitatively measure Warcraft III performance (see Section 3.2).
- Construct an environment for measuring the effects of latency on Warcraft III (see Section 3.3).
- Conduct pilot studies (see Section 3.4) and then numerous user studies for each RTS category over a range of latencies, recording the performance measurements.
- Analyze the results (see Section 4).

3.1 Categories of RTS Interaction

Through pilot studies, we determined there are three main user interaction components of an RTS game: *building* when players gather resources, construct defenses and recruit units; *exploring* when players send units out to determine geographic layout and location of other players' units; and *combat* when players engage their units with other units in battle. Since all components require user interaction, we hypothesized that under each component, user performance would degrade as latency increased. We built multi-player maps⁶ that isolated each component so that we could use experiments to measure the effects of latency on that component.

For the building map (Figure 2 (left)), we divided the map into four quarters using mountain ranges that units could not cross. Each player started with a town hall and four peasants, had unlimited gold and lumber available, and had to research, build, and upgrade the complete Human technology tree as fast as possible. We added triggers to the map

⁶The maps can be downloaded at <http://perform.wpi.edu/downloads/#war3>

that disabled players' ability to build more than one building in order to provide consistency and reduce confusion, as well as a trigger to display the total time since the beginning of the game.

For the exploration map (Figure 2 (middle)), we designed a raised path that kept units on a general exploration course. The player had to guide a unit along the winding path and step on numerous waypoints. Map triggers kept track of the player's time to complete the map.

For the combat map (Figure 2 (right)), we designed a small player versus player arena in which each player controlled a small army consisting of a level 6 Hero (a Mountain King), two Knights, four Footmen, two Riflemen, a Sorceress, and two Priests.

3.2 Warcraft III Performance Criteria

For both the building and exploration maps we recorded the game length as a measure of performance. For the combat maps, in addition to the game length, we recorded each player's unit score and which player won. The number of units a player starts with plus the number of units killed determines the unit score. The breakdown of points for the individual Human units used in our combat map are listed in Table 1.

Unit	Score
Footman	160
Priest	170
Sorceress	200
Rifleman	270
Knight	350
Level 6 Hero	600

Table 1: Unit Score Values

3.3 Experimental Setup

Figure 3 depicts our experimental testbed setup, which consisted of PCs connected on a private network subnet. Computer A was a dual-processor Pentium-2 300 MHz machine running Mandrake Linux that routed packets with 100 Mbps connections to the Warcraft computers B, C and D. Computers B and C were Pentium-2 350 MHz machines with 256 MB of RAM, and a 64 MB Geforce 2 Ti

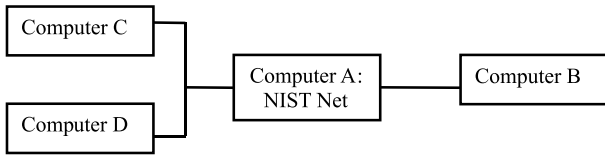


Figure 3: Experimental Testbed Setup.

graphics card running Windows 98. Computer D was a Pentium-4 1.3 GHz with 256 MB of memory and a 64 MB Geforce2 graphics card running Windows XP.

The recommended specifications for Warcraft III are a 400 MHz Pentium-2 or equivalent, 128 MB of RAM, and an 8 MB 3D video card (TNT, i810, Voodoo 3, Rage 128 equivalent or better) with DirectX[®] 8.1 support. Although machines C and D were only 350 MHz machines, the graphics cards and extra memory that they contained made up for this slight deficiency, and all machines were capable of rendering 30 frames per second⁷ even during battle. We used Warcraft III version 1.04 for all user tests and version 1.05 for the network traces due to the Battle.net requirements.

We installed NIST Net⁸ on computer A. NIST Net allows emulation of a wide variety of network conditions by giving control at the IP level, including fine tuning of latency and variation in latency (jitter). We used NIST Net to induce latency (and jitter) for one of the machines in a game, while the others played with no induced latency. Also, in order to analyze the network footprint of Warcraft III, we ran Ethereal⁹ to capture packet traces for network analysis.

3.4 Pilot Studies

First, we conducted pilot studies to help determine the range of viable latencies on which to focus. Our first pilot studies consisted of two-player games in which one player was subjected an increasing amount of latency and the other player experienced none. Initially, each player had a town hall and a gold mine placed a fixed distance away from the town hall; and second, each player had

two identical units that did one point of damage per hit. We setup triggers in the maps so games could be run automatically and ran repeated tests with one player (the lagged player) having increasingly greater latency. We found both players did equally well, gaining gold and inflicting damage at exactly the same rate. In addition, both players saw exactly the same events on each screen, except the player with added latency saw events later than the player without added latency.

From these pilot studies, we made two important observations about latency compensation in Warcraft III:

First, the game does not use handicapping in the game to equalize latencies across all players. Both lagged and non-lagged players see events happen at the real-time rate, regardless of the latency of the other player. The lagged player has events executed later by an amount equal to the induced latency.

Second, the game does not have inconsistent game states, which implies no dead reckoning [8]. The actions that occur on each machine are identical; there is no prediction of user actions and then correction upon some later time if the predictions are inconsistent with the actual game state.

Thus, clients must communicate any user actions to the server before executing them. After that, the commands themselves are executed identically on all machines in the game.

For the real experimental runs, the maps were not automated and we pitted one player against another player. The first player was the server with no induced latency. The second player was the client that was subjected to induced latencies ranging from 0 to 3500 ms, a range even broader than typically found in dialup modems [12]. Hence, we concentrated our data points on ranges of more typical latencies [1] which are less than 1000 ms.

From traces collected during our pilot studies, we determined that clients communicate only with the server but not directly with other clients. Servers combine data from multiple clients before distributing data. Each machine maintains a complete copy of the game state, and to an extent, all outcomes are predetermined upon initiation of the action. Command data is only transferred upon the issuance of a command, and never again during

⁷Tested with *fraps*, <http://www.fraps.com/>

⁸<http://snad.ncsl.nist.gov/itg/nistnet/>

⁹<http://www.ethereal.com/>

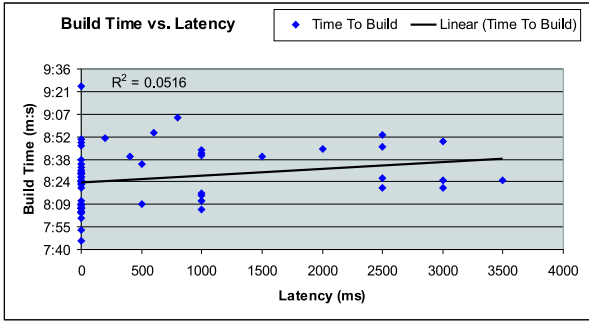


Figure 4: Build Time versus Latency.

the life of the event. For instance, the commands to initiate a large-scale battle are propagated to all clients once, resulting in an increase in the packet payload size, but the battle itself has no effect on traffic unless further commands are issued as the battle is carried out.

4 Analysis

We analyzed our experimental data at three levels: Section 4.1 contains our analysis of the application level data we collected from our Warcraft III user studies; Section 4.2 analyzes network level traffic for a LAN game and two Battle.net Internet games as well as network level traffic for combat games with three levels of induced latency; and Section 4.3 summarizes the observation data we collected during the user studies.

4.1 Application Level Analysis

This section analyzes the results from each of our test maps, starting with building (Section 4.1.1), then exploring (Section 4.1.2) and lastly combat (Section 4.1.3).

4.1.1 Building

Figure 4 illustrates the effect of latency on the total time required to construct every building and research every upgrade (the technology tree) for the Human race from our test map. The graph shows the build time versus latency for all runs, as well as a best-fit line for the data. Under conditions with no induced latency, building the technology tree

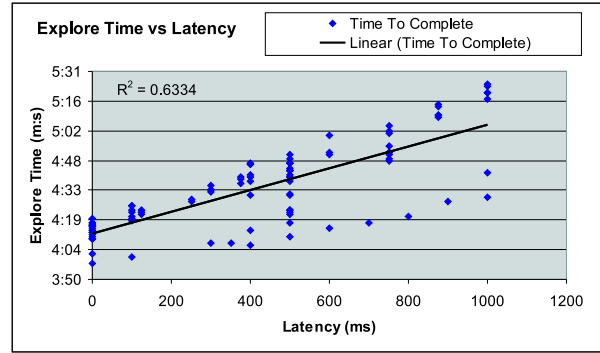


Figure 5: Explore Time versus Latency.

takes about 8 minutes. Latency values of up to 3.5 seconds increase total build time by at most 14 seconds, which is less than 1% of the total time for this short game. The coefficient of determination¹⁰ is 0.05, indicating there is very little statistical correlation between latency and building. In addition, the statistical correlation observed in a real game environment is likely to be even lower. This is due in part to a longer game time and that different numbers of buildings would be produced (such as more than one farm). In addition, in a real game, players may build their towns in a strategic layout instead of in a random pattern. Finally, time is often spent in a real game attending to other matters so that the speed of building the base is not of utmost importance. Our conclusion is that any effect latency may have on building would have no significant impact on the outcome of typical Warcraft III games.

4.1.2 Exploring

Figure 5 illustrates the effect of latency on the exploration of our test map. The graph shows the exploration time versus latency for all runs, as well as a best-fit line for the data. The overall correlation between explore time and latency is modest (0.63), but can be high (0.95) for individual users. The first 8-10 games of a test typically showed a vertical

¹⁰The coefficient of determination (R^2) represents the fraction of variability in y that can be explained by the variability in x . In the simple linear regression case, R^2 is simply the square of the correlation coefficient. An R^2 of 1 represents perfect correlation.

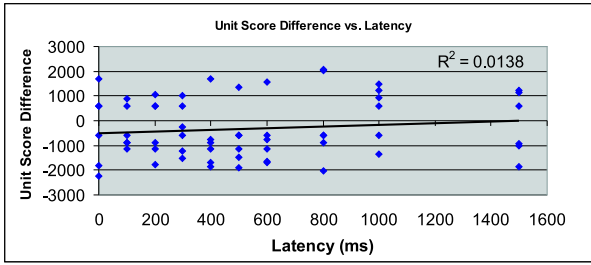


Figure 6: Unit Score Difference versus Latency.

component where exploration times decreased. We attribute this to the player learning the map, gaining from the knowledge in subsequent games. Once the map is known, all data show a nearly perfect linear relationships between latency and time to explore. Overall, while there is a statistical correlation for explore time versus build time, the effect of an additional 6 seconds of exploration time for every 100 ms of latency would be insignificant in a real game. In addition, it is likely that high latency players in a real game may try to adapt in various ways during exploration. For instance, high latency players may discover that they achieve better results by spending less time actively controlling their units during exploration and thus decide to send them for long distances with each move command instead of micro-managing them for shorter distances.

4.1.3 Combat

Figure 6 shows the unit score difference versus latency for all runs, as well as a best-fit line for the data. The unit score difference is the non-lagged player’s unit score minus the lagged player’s unit score. The maximum difference (if one player loses all units and the other loses none) is +/-3020. From Figure 6, there is a slight upward trend in that the score difference increases as latency increases, but the coefficient of determination is an extremely low 0.01. Moreover, the difference in points from no induced latency to one second of induced latency is only about one unit, an insignificant amount in the large battles that are typical in Warcraft III. Thus, we conclude that latency has little effect on the individual units in combat.

Figure 7 illustrates the effect of latency on com-

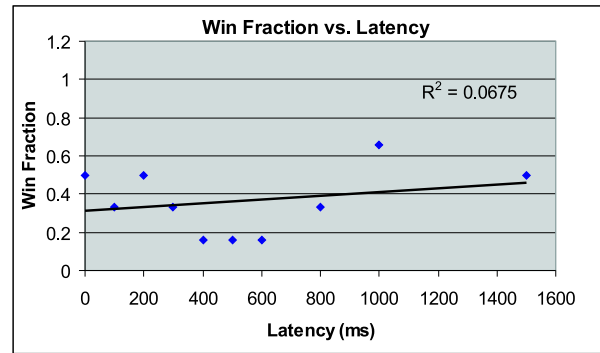


Figure 7: Combat Games Won by Host (non-lagged) versus Latency (on the Client).

bat from our test map. The graph shows the percentage of games won by the non-lagged host versus the latency of the lagged client. Even though there is a slight upward trend in the data, the coefficient of determination is an extremely low 0.07, indicating there is little statistical significance. Thus, we conclude that latency has little effect on the overall outcome of combat.

While the previous studies measured the effects of fixed latency on user performance, we also examined the effects of variable latency. For these tests, we set NIST Net to induce an average latency of zero¹¹ with a range of standard deviations.

Figure 8 shows games for 2 pairs of users with a standard deviation of latencies from a normal distribution with mean zero. Figure 8 (top) shows player 1 winning two games, one at 100 ms, and then again at 750 ms, while losing the games in between, all by similar margins of 1 or 2 units. Figure 8 (bottom) shows Player 3 consistently beating his opponent in every game, but by varying margins. Neither graph shows a significant statistical relationship between the variable latency and success in battle, similar to the results with constant latency.

Overall, both from a direct conclusion from our data and with extrapolation into a complete game, we found that the effect of latency on the outcome of a Warcraft III game is negligible over a range of practical latencies.

¹¹Our testbed had 1 ms of base latency from client to server.

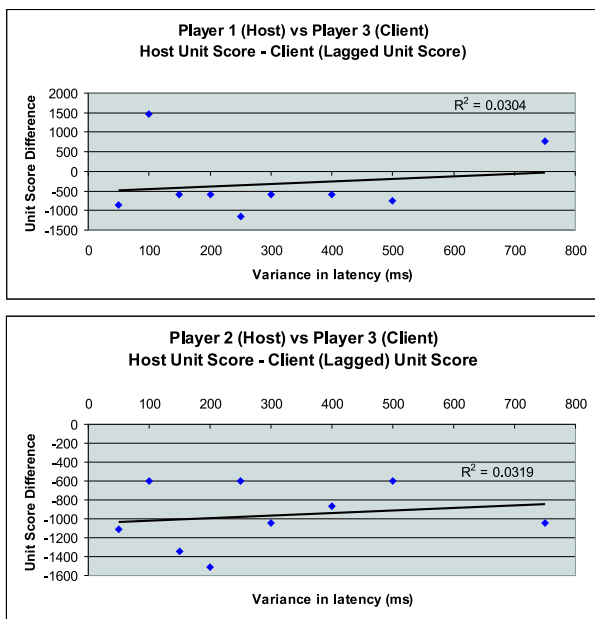


Figure 8: Unit Score Difference versus Variable Latency: Player 1 versus Player 3 (top), Player 2 versus Player 3 (bottom).

4.2 Network Level Analysis

Among other things, a better understanding of network game traffic can help design networks and architectures that more effectively accommodate their traffic footprints. Furthermore, careful empirical measurements of network games can provide the data required for accurate simulations, a typical tool for evaluating network research.

4.2.1 Traffic for Full Games

We packet traced three full (20-30 minute) games, two played over Battle.net and one played over a LAN.¹² The LAN game was 1 player versus 1 player (1v1), and the Battle.net games had a 1 player versus 1 player game and a 2 player team versus another 2 player team (2v2) game. Unlike other popular networked games [7], Warcraft III always uses TCP as the transport protocol with port 6112 for the server. All IP traces were performed on the client machines. For reference, the average round-trip times for the Battle.net games was about 100

¹²All network traces can be downloaded at <http://perform.wpi.edu/downloads/#war3>

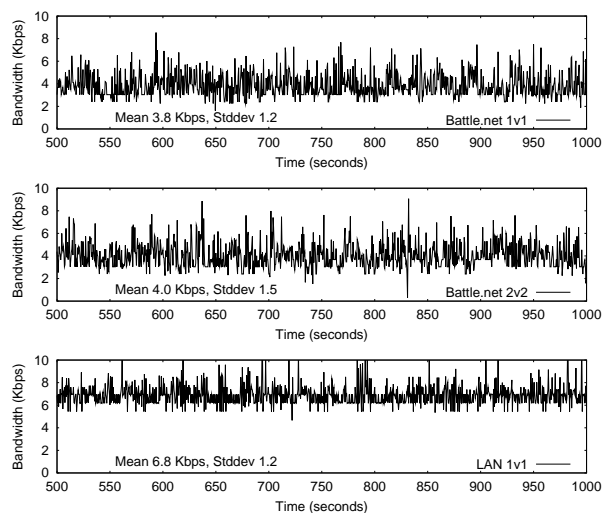


Figure 9: Bandwidth versus Time.

ms and there was less than 0.1% data loss.

Figure 9 depicts the bandwidth (including IP headers) taken in 500 ms intervals for the three packet traces. Only the intervals 500-1000 seconds are shown to illustrate more detail, but the bandwidth pattern throughout each game is similar to the shown interval. Overall, the variance in network bandwidth for all three is similar, with the average data rate for the LAN being slightly higher (6.8 Kbps) than the Battle.net traces (3.8 Kbps and 4.0 Kbps). All three traces have very low data rates that can easily be achieved with a modem. In comparison, Starcraft,¹³ the previous generation RTS game from Blizzard, has a data rate of about 5 Kbps for a 2 player game [7], similar to that of Warcraft III.

Figure 10 depicts the cumulative density functions (CDFs) of the payload sizes for all packet traces (incoming and outgoing). The median payload sizes are all very small, only 9 bytes. The two most common payload sizes are 6 and 9 bytes. Less than 1% of the payloads for any game are over 40 bytes with the Battle.net games having slightly more larger packets. The 2v2 player Battle.net game has a distribution with slightly larger payloads, most likely because of command aggregation across users at the Battle.net server.

Overall, Warcraft III sends considerably smaller

¹³<http://www.blizzard.com/worlds-starcraft.shtml>

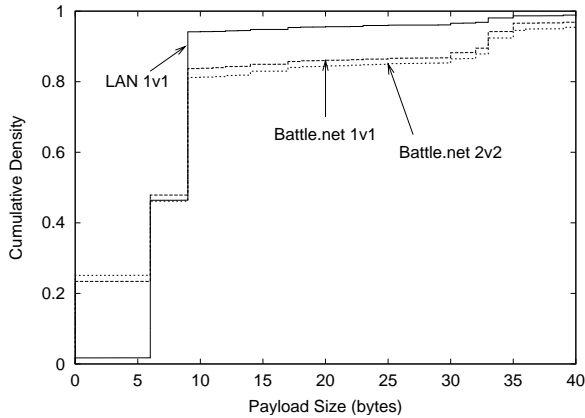


Figure 10: Payload Distributions.

packets than the typical Internet traffic packet size of over 400 bytes [13]. The number of players does not have a significant effect on the packet sizes, either. Warcraft III packet sizes are consistent throughout the game and are not significantly influenced by the action in the game. Since current Internet routers are designed for large transfers with large packets, there may be opportunities to improve network architectures to better manage and support game traffic. Starcraft has typical packet sizes of 122 and 132 bytes [7], while Warcraft III packets are most commonly 46 or 49 bytes in size (including headers).

Warcraft III sends out packets at regular intervals. Table 2 shows the inter-packet times that we observed for incoming and outgoing packets during the games we traced. In our local area network game, Warcraft maintained a very steady inter-packet rate of approximately one packet every 1/10th of a second both incoming and outgoing. With our Battle.net games the timing interval was lower, down to one packet every 200 ms incoming and every 160 ms outgoing.

	1v1 LAN	1v1 B.net	2v2 B.net
In Mean	104	201	201
In Std Dev	18.6	79.1	78.1
Out Mean	104	165	159
Out Std Dev	19.4	87.4	88.2

Table 2: Inter-packet Summary Statistics (ms).

Figure 11 depicts the CDF for inter-packet times

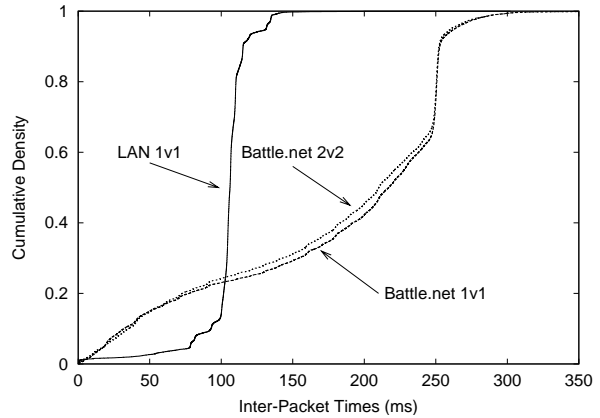


Figure 11: Inter-packet Distributions.

(incoming and outgoing). The LAN game has a much more consistent packet rate while the Battle.net Internet game varies considerably more. The median times for the Battle.net games are around 225 ms compared with around 100 ms for the LAN game. The 1v1 player Battle.net game exhibits about the same inter-packet times as does the 2v2 player Battle.net game.

4.2.2 Combat Traffic and Latency

From Section 4.2.1, the differences between the Battle.net game traces which had latencies around 100 ms and the LAN game traces which had latencies around 1 ms, suggest Warcraft III network traffic patterns change at least slightly with changes in latency. In this section, we analyze traces over a range of latencies in an attempt to quantitatively determine how Warcraft III network traffic differs with different latencies.

We packet traced games with our combat map at latencies of 0 ms, 500 ms, and 1000 ms with three games at each latency. All games took similar amounts of time (around 2 minutes each). The first phase (about 30 seconds long) of the combat games mostly involved the two armies moving towards each other, so there were few user commands and little network traffic. Thus, we removed the first 30 seconds of data from each trace for all subsequent analysis.

Table 3 shows the mean number of packets sent and the standard deviation across the three runs for each latency. Also shown is the mean bandwidth

Latency	Packet Count		Bandwidth (Kbps)	
	Mean	Std Dev	Mean	Std Dev
0 ms	1886	230	7.2	0.6
500 ms	550	292	2.3	0.6
1000 ms	255	123	2.1	0.5

Table 3: Packets and Bandwidth.

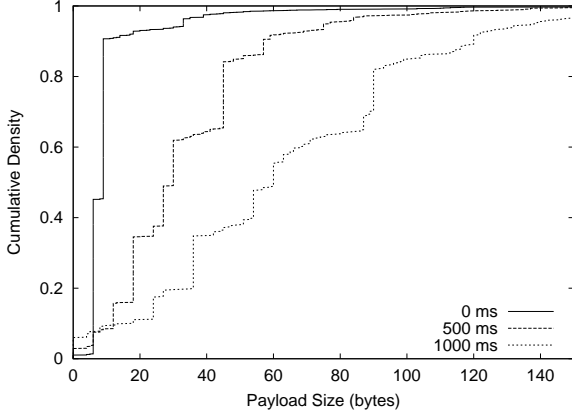


Figure 12: Combat Payload Distributions.

(including IP headers) over 500 ms intervals as well as the standard deviation. The number of packets (incoming and outgoing) decreases as the latency increases, with the combat games with 500 ms and 1000 ms of latency sending only about 1/3rd and 1/8th as many packets, respectively, as the game with no added latency.

The 0 ms latency combat game produces about the same bandwidth as does the full LAN game, shown in Figure 9. The 500 ms latency and the 1000 ms latency combat games have about 1/4th the bandwidth as the 0 ms latency game and both the 500 ms latency and the 1000 ms latency games produce less bandwidth than do the Battle.net games. This data suggests that Warcraft III bandwidth decreases with an increase in latency up to 500 ms, but remains constant for latencies beyond 500 ms.

Figure 12 depicts the CDFs of the payload sizes for all packet traces (incoming and outgoing), grouped into the three latencies. The median payload sizes increase from 9 bytes at 0 ms, to 30 bytes at 500 ms and to 60 bytes at 1000 ms. Less than 10% of the packets for any game are empty ac-

Latency	Commands Payload
0 ms	45.2 Kbytes
500 ms	46.3 Kbytes
1000 ms	45.0 Kbytes

Table 4: Commands Payloads.

knowledgments (payload size of 0). Overall, the distributions vary considerably with latency with higher latencies having larger packets. This suggests that at higher latencies, there is command aggregation at either the TCP or application level, meaning more Warcraft III commands are placed into each IP packet.

Based on Warcraft III traffic analysis during our pilot studies, we assume that there is an application overhead of 6 bytes for each packet issued, possibly used by Warcraft to indicate command sequence numbers or timing information. If we remove this overhead from the traces by subtracting 6 bytes from each packet, we can assume the “left-over” payloads are the result of user commands. Table 4 shows the sum of the command payloads over all the traces for each latency. The sum of the command payloads is very similar for each latency, which suggests that the commands issued by users are very similar, regardless of the network latency.

4.3 User Level Analysis

While we did not provide a way to quantify player perceptions, we did note game player comments and observed trends during and after our user studies.

Players observed that it was relatively easy to adjust their strategy to compensate for latencies between 0 ms and 500 ms. The game still appeared to run smoothly, and although the delays in executing commands were perceptible, it was relatively easy to estimate this delay and react accordingly. For latencies above 800 ms, the game appeared erratic which made for a degraded game experience. Without a short response time when executing commands, gamers found it difficult to implement particular strategies.

The exact point at which a player perceived a degraded game experience was between 500 ms and 800 ms but varied from person to person based

on strategy and skill level. A strategy that relied heavily on micro-management of units was more sensitive to latency than a strategy that was less focused on individual unit control. What game aspects that a player chooses to micro-manage also had an effect on how perceived latency affected the gaming experience. A player that micro-managed the town rather than the combat units was much less likely to be aggravated by latency than a player that micro-managed units in combat. Also, a mistake during combat that appeared to be the result of high latency was viewed, rightly or wrongly, as more serious than a mistake during building.

Thus, while latency does not necessarily affect the outcome of a Warcraft III game, if high enough, perceived latency does affect a user's gaming experience.

5 Conclusions

Understanding the effects of latency on application performance is important in order to design networks that meet application requirements. The growth in interactive network games demands better understanding the effects of latency on user performance in network games.

In this work, we investigated the effects of latency on user performance in Warcraft III, a popular Real Time Strategy (RTS) game. We divided RTS games into their fundamental components of build, explore and combat and designed experiments to isolate and measure the effects of latency on each component.

We find that overall user performance is not significantly affected by Internet latencies ranging from hundreds of milliseconds to several seconds. There is some statistical correlation between latency and the exploration game component, but the overall impact is minimal and there is even less correlation between latency and building and between latency and combat.

While these results are, at first glance, somewhat surprising they can be explained by the fact that the nature of RTS game play emphasizes strategy more than the interactive aspects. While Warcraft III is played in real-time, reaction time plays a small role compared to understanding the game, knowing a campaign map, and having a good strat-

egy. Since RTS user strategies take seconds or even minutes to carry out, the effects of typical network latencies (less than a second) do not impact the overall outcome. This relative insensitivity to latency is further illustrated by Warcraft III's use of TCP as the underlying transport protocol. TCP retransmits lost packets, with the retransmissions increasing latency on the order of a round-trip time, at best, and several seconds (upon timeout) at worst. Overloading at the game server is another factor which potentially adds to game latency. The fact that Warcraft III's game play over the Internet via Battle.net is very effective further underscores the lack of significant impact of latency on game outcome.

Overall, in terms of general classification of traffic, RTS games do not have the very strict latency requirements (on the order of hundreds of milliseconds) of audio-conferencing or First Person Shooter network games, but instead have latency requirements most similar to that of Web browsing (on the order of seconds).

At the network level, Warcraft III basically produces small, regularly-spaced packets and modest aggregate bandwidths which make it suitable for play over a low speed modem. At higher latencies, Warcraft III aggregates multiple commands in each packet, resulting in fewer, but larger packets. By placing multiple commands in each packet, Warcraft III amortizes the overhead of each IP header cost, thus reducing network bandwidth slightly. Our network analysis suggests that the aggregate of user commands sent are comparable over a range of latencies.

6 Future Work

The component-based studies presented here do not allow users to choose long-term strategies. Evaluating the effects of latency on how users choose what components to micro-manage, how they select and form long-range, even full-game strategies may provide insights beyond the results presented here.

While the results in this project have focused on Warcraft III, we assume that they generalize

to other RTS games (such as Age of Empires¹⁴ or Command and Conquer¹⁵) as well, since most RTS games have the same fundamental components (build, explore, combat) and similar user-interaction models. Studies to confirm this through select user-studies, possibly less extensive than those presented here, would be useful to verify our assumption and help generalize our results.

The effects of latency on user performance in other game genres, such as First Person Shooters or Massively Multi-player Online Role Playing Games, is also still an open issue. However, it is clear that several network games consist of distinct phases which vary greatly in their interaction model and hence network behavior. The component-centric methodology presented here, which entails categorization of the game play and running of controlled users studies in each category, can perhaps be applied to these games as well, in order to increase overall understanding of network games.

References

- [1] Mark Allman. A Web Server's View of the Transport Layer. *ACM Computer Communication Review*, 30(4), October 2000.
- [2] Grenville Armitage. Lag Over 150 Milliseconds is Unacceptable, May 2001. [Online] <http://gja.space4me.com/things/quake3-latency-051701.html>.
- [3] Paul Bettner and Mark Terrano. 1500 Archers on a 28.8: Network Programming in Age of Empires and Beyond. *Gamasutra*, March 2001. [Online] http://www.gamasutra.com/features/20010322/terrano_02.htm.
- [4] Nina Bhatti, Anna Bouch, and Allan Kuchinsky. Integrating User-Perceived Quality into Web Server Design. In *Proceedings of the 9th International World Wide Web Conference*, Amsterdam, Netherlands, May 2000.
- [5] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An Architecture for Differentiated Services. *IETF Request for Comments (RFC) 2475*, December 1998.
- [6] U.S. Census Bureau. Home computers and internet use in the united states, September 2001. [Online] <http://www.census.gov/prod/2001pubs/p23-207.pdf>.
- [7] Mark Claypool, David LaPoint, and Josh Winslow. Network Analysis of Counter-strike and Starcraft. In *Proceedings of the 22nd IEEE International Performance, Computing, and Communications Conference (IPCCC)*, April 2003.
- [8] Chris Haag. Targeting, a Variation of Dead Veckoning, May 2001. [Online] <http://www.gamedev.net/reference/articles/article1370.asp>.
- [9] Tristan Henderson. Latency and User Behaviour on a Multiplayer Game Server. In Jon Crowcroft and Marcus Hofmann, editors, *Proceedings of the Third International COST Workshop (NGC 2001)*, number 2233 in LNCS, pages 1–13, London, UK, November 2001. Springer-Verlag.
- [10] International Digital Software Association. Economic Impacts of the Demand for Playing Interactive Entertainment Software, 2001. [Online] <http://www.idsa.com/releases/EIS2001.pdf>.
- [11] Peter Lincroft. The Internet Sucks: Or, What I Learned Coding X-Wing vs. Tie Fighter. *Gamasutra*, September 1999. [Online] http://www.gamasutra.com/features/19990903/lincroft_01.htm.
- [12] Dmitri Loguinov and Hayder Radha. Measurement Study of Low-bitrate Internet Video Streaming. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, November 2001.
- [13] S. McCreary and k claffy. Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange. In *Proceedings of ITC Specialist Seminar on Measurement and Modeling of IP Traffic*, pages 1 – 11, September 2000. [Online] <http://www.caida.org/outreach/papers/2000/AIX0005/AIX0005.html>.
- [14] Yu-Shen Ng. Designing Fast-Action Games for the Internet. *Gamasutra*, September 1997. [Online] http://www.gamasutra.com/features/19970905/ng_01.htm.
- [15] Blizzard Entertainment Press Release. Warcraft III - Shatters Sales Records Worldwide with over 1 Million Copies Sold, October 2002. [Online] <http://www.blizzard.com/press/020722war3.shtml>.

¹⁴<http://www.microsoft.com/games/empires/>

¹⁵<http://westwood.ea.com/>