

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188		
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 14/11/2017		2. REPORT TYPE Technical		3. DATES COVERED (From - To) 01/2011 – 04/2011	
4. TITLE AND SUBTITLE The Effect of MTU Size on Internet-Controlled Unmanned Ground Vehicles			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
			5d. PROJECT NUMBER		
6. AUTHOR(S) Ty Valascho, Greg Czerniak			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
			8. PERFORMING ORGANIZATION REPORT		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army, Tank Automotive Research Development and Engineering Command (TARDEC) Warren, MI 48397			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Army, Tank Automotive Research Development and Engineering Command (TARDEC) Warren, MI 48397					
12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution Statement A. Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions, and/or findings contained in this report are those of the authors and should not be construed as an Official Department of the Army position, policy, or decision, unless so designated by other documents.					
14. ABSTRACT Two methods were tested to reduce latency within an Unmanned Ground Vehicle system controlled by passing messages over the internet. The first finding of this work was to have the system refrain from sending video messages when a communications path did not exist, found to reduce end to end system latency from an average of 409 milliseconds to 185 milliseconds. The second finding was that using an MTU size of 400 for UGV messages provided the highest percentage of acceptable delay.					
15. SUBJECT TERMS UGV, latency, MTU, robotics					
16. SECURITY CLASSIFICATION OF: Distr. A			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON Ty Valascho
a. REPORT Distr. A	b. ABSTRACT Distr. A	c. THIS PAGE Distr. A			19b. TELEPHONE NUMBER (include area code) (586) 459-7485

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

US Army - TARDEC

The Effect of MTU Size on Internet-Controlled Unmanned Ground Vehicles

Ty Valascho & Greg Czerniak
4-22-2011



Contents

Introduction	3
Background	3
Dealing with Delays.....	4
System Architecture / Test Setup	5
Overall System Architecture	5
Measuring Latency	6
Performance Test Procedure	6
System Configuration 1: Video Reduction by not sending messages on connection loss.....	7
System Configuration 2: MTU Size Modifications.....	8
Test Results and Analysis	9
Pareto Analysis.....	9
Test Results	10
Conclusions and Recommendations	12
References	13
Appendices.....	13
Abbreviations and Acronyms	13

Introduction

Tele-operational control of UGVs (Unmanned Ground Vehicles) over the internet has many challenges. One important aspect of performing this type of real-time control is the variable latency that occurs when messages are passed over the internet, due to different routing paths and other factors out of control of the sending and receiving entities.

In an attempt to optimize this type of system, one aspect that was considered was the effect of MTU (Maximum Transmission Unit) of the messages employed for this real time control. This is the maximum size of the packets that are sent over the internet. The MTU size can be controlled by the sending software for internet transmitted messages. Messages and information can span over multiple packets and be reassembled at the receiver, allowing MTU size to be independent of the content that is sent.

Background

In a typical single UGV system, there exists the following three components:

- An OCU (Operator Control Unit) – a computer, handset, or other device which the operator utilizes to “drive” the UGV being controlled
- The UGV being controlled
- A communications medium of some kind to allow information to pass between the two devices

Figure 1 shows the relationship between these three elements with more detail about the types of information being sent over the communications medium.




Figure 1 - Idealized Control of a UGV

UGV feedback signals are the sensor returns from the UGV such as accelerometers, gyroscopes, wheel speed encoders, etc. The Video Signals are also a type of feedback signal but are separated in this discussion because the amount of information – and therefore the size of the messages – is usually an order of magnitude or more larger in size and may be treated differently than the rest of the signals passing over the communications medium.

For an operator to effectively control the UGV, an idealized system is envisioned where the communications medium allows instantaneous passing of information in both directions. The operator commands to the UGV and feedback from the UGV are sent and received with zero delay. The operator can maneuver the UGV in real time, avoiding obstacles and performing tasks.

In practice, however, there are delays at every stage of the system due to a variety of factors. Processing delays, switching delays, routing differences, message parsing, and time slicing of shared resources are some examples of factors that add to the total system latency.

Many of these delays are imperceptible to the operator, but not all. In addition, the cumulative effect of many small, unnoticed delays can add up to a delay that is perceived by the operator and can make control difficult or impossible.

In Figure 2, the addition of signal delay is shown in the video link between OCU and UGV.




Figure 2 - UGV Control with Significant Video Latency

The work of this document is mainly concerned with the transmission of real-time video packets. In general, the size of messages has an effect on how quickly the message is passed over the internet. Smaller messages can be shown to pass quicker than larger ones. Larger messages may be broken up into smaller messages and re-assembled at the receiver, which adds delay for this processing and also waiting for all the messages to arrive.

The control messages sent to a robot and the status and feedback messages back are generally small, when compared to other types of messages. However, video messages contain a lot of data and are much larger than these. Therefore, controlling a robot in real-time using video as the main source of driving feedback can be a challenge for two reasons: latency and variable latency.

Latency is the temporal delay between an event happening and the feedback response of the operator to the event. In the case of this work, the latency between the video signal being recorded at the source – the UGV – and the operator seeing the video signal is the latent time of concern. Since the operator's control commands are being received with lower latency than the user's video, it makes control of the vehicle difficult. The video latency determined at the beginning of this testing was 409 milliseconds on average.

Dealing with Delays

Human operators can still operate UGVs with delayed feedback signals by a couple methods.

One method is to make small movements and wait for the feedback. This slows down operations and delays responses to environmental factors.

Another method is to anticipate the movements. Anticipation occurs when the operator gets a “feel” for the latency and accounts for it while operating. For example, if the operator is aware that latency is about 1 second, they can make movements and control actions about 1 second before they are actually required, knowing the actual response will roughly correspond to reality “on the ground.”

These difficulties are made worse for teleoperation of a robotic vehicle when the latency is variable. With video signals as a feedback to the operator or driver, variable latency adds another impediment to control because the strategy of anticipation no longer works.

This document records two methods to reduce the latency of video signals when teleoperating a vehicle over the internet.

System Architecture / Test Setup

Overall System Architecture

The system that was first used for this work consisted of a UGV communicating over a dedicated Verizon 3G cellular network air card to the internet and then to a server called the Switchboard Server. The OCU was also connected by a 3G cellular network to the internet, in this case a shared AT&T air card.

A diagram of this system configuration is provided in Figure 3.




Figure 3 - System Configuration 1

When initializing the system, all UGVs being used will connect to the switchboard after initial power on. The OCU is also powered on and the user is given a choice to take control of any robot that is available at the switchboard. Only one robot at a time can be controlled with this system.

Once a connection is made to the switchboard, it sends the IP (Internet Protocol) addresses of the UGV and OCU to both endpoints so that communications can be made directly from OCU to UGV and vice versa. Connections between both endpoints and the server are still maintained, however, in case reconnections are needed and also to measure the status and latency of the communication links.

Measuring Latency

Latency measurements were made by the use of small messages with time stamps in them called “ping” messages which were responded to with a “pong” message that contained the original time stamp and the time stamp of when the ping message was received. The original sending party would then have the following information:

- The time when the message was sent
- The time the other party received it
- The time when the response was received.

From this information the time the message took to reach each point could be calculated. Note that these times can be used to show trends and relative message delays but are not accurate to more than a few milliseconds due to clock inaccuracies and drift, processing times, and other factors.

The ping / pong messaging scheme was utilized from the OCU to the Switchboard Server and from the UGV to the Switchboard Server. The data for both links was recorded in raw data files on the Switchboard Server.

Performance Test Procedure

The testing was performed by configuring the system in various ways and gathering data of several hundred or several thousand packets. The UGV system was connected to - but not driven - during this testing.

The initial baseline testing that caused this investigation to be initiated showed that the system was operating with average latency of 409 milliseconds. Also, during 90 seconds of testing, 17 increases in latency were observed that were 1000 – 5700 milliseconds. The spikes were occurring every 40th message on average. These increases were short in duration – less than 1000 milliseconds – but caused noticeable lags and pauses in the video feedback to the operator. Figure 4 shows a plot of this phenomenon based on measurements made by the Switchboard Server.




Figure 4 - Message Latency over Time

The issue was theorized to be caused by temporary losses of connectivity over one or both cellular links. Under this condition, it was thought that the large video packets would stay at the UGV, then be resent when connectivity was restored. This may cause messages to queue while sending and receiving. In addition, the video information would be “stale” and not useful to the operator any more, having been replaced by more recent information which rendered the older data useless and wasteful.

Based on this theory, it was decided to not send video messages when connectivity could not be confirmed by the system.

[System Configuration 1: Video Reduction by not sending messages on connection loss](#)

The software was modified on the robot to only send video packets when the connection to the Switchboard server was present.

Results from testing this change are shown in Figure 5.




Figure 5 - Message Latency Using Video Reduction Technique

This change appeared to reduce the frequency and severity of the latency spikes, but did not remove them completely.

The average latency dropped from 409 to 185 milliseconds. 11 latency spikes were observed in 1090 messages which reduced the average spike occurring from every 40th message to every 67th message. The severity of the spikes also decreased from a range of 1000 – 5700 milliseconds in the baseline data to a range of 400 – 1900 milliseconds using the modified software.

System Configuration 2: MTU Size Modifications

Only sending the video packets when the connection is stable decreased the overall latency as well as the frequency and severity of the intermittent spikes in latency. However, the behavior of the system was still unacceptable and the UGV was difficult to control.

The spikes in latency that were now experienced were most likely the result of the key frames being sent, which were large packets with the complete information for one video frame. Most of the time small packets are sent with only the changed information from one video frame to the next, but occasionally the video system will send a key frame to “synchronize” the video stream.

The next part of the investigation was measuring the effect of MTU size on the time the messages took to be received. It was hoped a “best MTU size” could be determined from testing that would facilitate the fastest transmission of messages.

For this part of the testing, a change to the network configuration was implemented, as depicted in Figure 6.




Figure 6 – System Configuration 2

In the new configuration, the connection between OCU and Switchboard Server was through copper and optical fiber - a high speed connection. This allowed the testing to concentrate on the communications between Switchboard Server and the UGV, which was over a 3G cellular wireless link - a less stable and slower connection.

Test Results and Analysis

Testing was performed by setting the MTU size of messages from Switchboard Server to UGV to 400 bytes then varying the MTU size of messages from UGV to Switchboard Server. All data passing between all the parts of the system were recorded along with time of sending and time of receiving. The measurement of delay for each message was analyzed afterward.

The test was then repeated with the MTU setting at 500 bytes for the Switchboard Server to UGV messages.

Pareto Analysis

To provide an interaction experience to an UGV operator, it was desired that data should be no more than 250 milliseconds delayed. The delay for messages passed over the internet is variable and non-deterministic for the reasons presented in the introduction. The data confirmed this.

Pareto analysis was performed on the measured delay between message sending and receipt to determine the percentage of messages that fell into several different "delay bins." This analysis was used to gauge the relative quality of the interactive experience achieved for each MTU size setting. The more packets in the smaller delay bins, the better the operator control experience.

The bins for Pareto analysis were the percentage of packets received in:

- Less than 125 milliseconds
- Less than 250 milliseconds
- Less than 375 milliseconds
- Less than 500 milliseconds

Note that messages could appear in more than one bin. That is, a message that arrived in 100 milliseconds would be in all 4 bins, but a message that arrived in 450 milliseconds would only be in the “Less than 500 milliseconds” bin.

These categories were chosen because previous informal testing had determined that these delays were rough limits for different categories of user experience when operating small tele-operated vehicles in real time that did not move faster than 1 meter / second.

- Message delays of less than 125 milliseconds were not discernable
- Message delays between 126 and 250 milliseconds were discernable but didn’t impact control¹
- Message delays between 251 and 375 milliseconds were acceptable but made control challenging
- Message delays between 376 and 500 milliseconds were acceptable but made control difficult
- Message delays greater than 500 milliseconds were unacceptable for control

The Pareto analysis was performed on the communication links between

- UGV and Switchboard Server only
- OCU and Switchboard Server only
- UGV and OCU - “End to End”

Test Results

The results of the testing are provided in Table 1 for a UGV MTU setting of 400 and Table 2 for a UGV MTU setting of 500.

¹ According to https://infogalactic.com/info/Input_lag , “Testing has found that overall “input lag” (from controller input to display response) times of approximately 200 ms are distracting to the user”

Table 1 - Pareto Analysis of Message Delays, UGV MTU of 400

UGV Only		Switchboard Server MTU Size									
		150	200	250	300	350	400	450	500	550	600
	Packets <125 ms	80.4%	75.7%	80.0%	81.9%	75.7%	85.7%	77.5%	83.7%	83.1%	86.7%
	Packets <250 ms	89.4%	86.8%	91.6%	90.8%	88.6%	95.5%	93.8%	92.3%	92.1%	95.2%
	Packets <375 ms	92.3%	90.6%	94.9%	93.8%	93.0%	97.5%	96.6%	95.0%	95.0%	96.9%
	Packets <500 ms	94.0%	92.8%	96.5%	95.7%	95.3%	98.3%	97.6%	96.6%	96.8%	97.8%
	Total Packets	34364	21503	48034	59433	44515	123604	53011	97228	95570	79358
OCU Only	Packets <125 ms	99.7%	98.5%	99.6%	98.9%	96.3%	99.3%	96.5%	99.6%	99.6%	99.0%
	Packets <250 ms	99.9%	99.5%	99.8%	99.4%	98.5%	99.6%	97.9%	99.8%	99.8%	99.6%
	Packets <375 ms	100.0%	99.7%	99.9%	99.7%	99.4%	99.7%	98.5%	99.8%	99.9%	99.8%
	Packets <500 ms	100.0%	99.7%	99.9%	99.8%	99.8%	99.8%	98.8%	99.9%	99.9%	99.8%
	Total Packets	24185	19642	35571	40949	42132	96813	36012	81260	69500	33792
OCU to UGV	Packets <125 ms	41.6%	34.5%	26.2%	38.3%	28.4%	31.9%	25.1%	56.8%	48.1%	34.9%
	Packets <250 ms	83.0%	83.1%	85.9%	84.3%	82.9%	92.3%	86.6%	87.2%	87.4%	87.8%
	Packets <375 ms	88.1%	88.6%	92.4%	89.8%	90.4%	96.0%	93.0%	92.1%	92.3%	92.3%
	Packets <500 ms	90.9%	91.4%	94.7%	93.0%	93.9%	97.4%	95.2%	94.7%	95.0%	94.4%
	Total Packets	24185	19640	35570	40947	42130	96791	35891	81231	69488	33759

Table 2 - Pareto Analysis of Message Delays, UGV MTU of 500

		Switchboard Server MTU Size											
	150	200	250	300	350	400	450	500	550	600	650	700	
UGV Only	Packets <125 ms	86.9%	66.7%	95.5%	98.4%	73.7%	92.5%	68.7%	65.6%	53.5%	84.7%	55.5%	53.2%
	Packets <250 ms	97.1%	95.2%	99.4%	99.6%	98.8%	99.3%	98.3%	98.9%	97.4%	98.2%	94.3%	82.3%
	Packets <375 ms	98.3%	98.2%	99.7%	99.8%	99.5%	99.7%	99.3%	99.7%	99.2%	99.3%	98.4%	88.4%
	Packets <500 ms	98.8%	98.8%	99.8%	99.9%	99.7%	99.8%	99.6%	99.9%	99.6%	99.6%	99.4%	91.6%
	Total Packets	100165	73652	123157	120124	99084	241644	105982	94440	95322	121566	97266	47603
OCU Only	Packets <125 ms	99.3%	89.0%	42.2%	95.6%	83.3%	91.1%	72.5%	69.9%	87.8%	84.7%	81.6%	99.1%
	Packets <250 ms	99.7%	98.3%	93.5%	99.3%	93.9%	96.1%	93.0%	93.6%	97.4%	98.2%	92.8%	99.5%
	Packets <375 ms	99.9%	99.4%	98.4%	99.3%	95.5%	96.9%	95.3%	96.8%	98.7%	99.3%	94.1%	99.7%
	Packets <500 ms	100.0%	99.8%	99.0%	99.3%	95.5%	98.4%	97.3%	98.1%	99.3%	99.6%	95.4%	99.8%
	Total Packets	4355	1270	2694	405	245	257	258	156	304	121566	152	39819
OCU to UGV	Packets <125 ms	24.4%	7.9%	3.0%	49.1%	0.4%	19.5%	0.0%	0.6%	0.7%	3.0%	0.0%	0.6%
	Packets <250 ms	68.8%	33.8%	51.2%	92.1%	65.2%	84.0%	58.9%	60.3%	76.6%	74.5%	62.3%	76.1%
	Packets <375 ms	75.4%	43.1%	88.0%	97.5%	81.6%	93.8%	82.2%	88.5%	92.4%	82.8%	80.8%	84.8%
	Packets <500 ms	79.6%	47.9%	95.1%	98.3%	86.1%	95.7%	89.5%	94.2%	95.4%	86.4%	84.1%	89.1%
	Total Packets	4346	1259	2694	405	244	256	258	156	304	2764	151	39794

The cells with the highest percentage of OCU to UGV messages that arrive within 250 milliseconds are highlighted in green in both tables. From the user perspective of driving an UGV, this is the most important aspect of the information.

In most cases, orders of magnitude more packets were gathered for the UGV MTU size of 400 than the UGV MTU size of 500.

Conclusions and Recommendations

The first finding of this work was to have the system refrain from sending video messages when a communications path did not exist. This change was found to reduce end to end system latency from an average of 409 milliseconds to 185 milliseconds.

Average latency alone does not provide a good measure of the user interactive experience. The metric was subsequently changed to be the percentage of messages transmitted within different time bins.

MTU sizes were adjusted at the Switchboard Server and UGV. The effect on the percentages of packets transmitted within the time bins was measured. The goal was to get the best user experience by finding the combination of MTU sizes that maximized the quantity of messages in the smallest delay bins.

Testing showed that using an MTU size of 400 for UGV messages and MTU of 400 for Switchboard Server messages provided the highest percentage of acceptable delay when controlling a UGV over the internet. With these settings in place it was determined that:

- 31.9% of messages have no discernable delay to the operator
- 92.3% of the messages have a delay of 250 milliseconds or less
- 2.6% of the messages have a delay of greater than 500 milliseconds

It is recommended that further study be performed by gathering more data with the UGV MTU setting at 500 and additional data be gathered with the UGV MTU set to 300 to verify this is the best possible outcome.

The testing was performed without driving the UGV or otherwise causing the camera view to change. If the camera view were to change, it would increase the amount of data sent over the stream which could have an impact on the latency of the messages. Additional testing should be performed while driving the UGV to measure this effect.

References

“Latency and the Quest for Interactivity” Stuart Cheshire November 1996

<https://www.humanbenchmark.com/tests/reactiontime/statistics>

Appendices

Abbreviations and Acronyms

UGV: Unmanned Ground Vehicle

MTU: Maximum Transmission Unit

OCU: Operator Control Unit

AT&T: American Telephone and Telegraph (Company)

3G: 3rd Generation (mobile telecommunications technology)

IP address: Internet Protocol address

ms: milliseconds