January 2010

Investigating Effectiveness of Small-scale Lunar Excavators

CMU SCS Senior Thesis

Author: Matthew F. Delaney
Advisor: Dr. David Wettergreen

Abstract

This work presents a small-scale (< 100kg) lunar excavation rover “Lysander”, an investigation into its effectiveness for lunar excavation tasks, and an analysis of the sensitivity of system productivity to changes in design and operational parameters of a small-scale lunar excavator performing point-to-point excavation.

A 2008 Astrobotic study determined that robots with mass of 300kg or less are a viable option for lunar excavation tasks including building protective berms for future lunar outpost landing pads. For that work, a regolith construction simulator dubbed REMOTE was developed to provide a model of the time required to complete lunar excavation tasks for a given platform in addition to a numerical analysis of the sensitivity of that platform’s productivity to changes in design and operational parameters.

With the lessons learned from the Astrobotic study, the Lysander rover was developed in 2009 by a Carnegie Mellon University team (including myself) for lunar excavation research and entrance into the 2009 NASA Regolith Excavation Challenge. The Lysander rover’s scraper style excavator design was inspired by the CRATOS rover; it employs a centrally located scraper bucket for regolith excavation and transport.

Using REMOTE, this work models the high and low sensitivity operational and design parameters of the Lysander rover and provides experimental validation of the actual sensitivities. This analysis of the Lysander rover shows that the productivity of small-scale lunar excavators is highly sensitive to changes in the payload ratio and transport drive speed, but not the number of wheels. The design process of future small-scale excavators can benefit from these findings when aiming to optimize the productivity of the system.

Acknowledgements

- I would like to acknowledge my fellow teammates of the Moon Diggers team for making the creation of Lysander possible.
- I would like to thank Kryzsztof Skonieczny for his ongoing support and dedication to this work.
- I would like to thank Dr. David Wettergreen for his guidance and supervision of my senior thesis work.

Table of Contents

1 Introduction

The Moon has experienced a revival of political and public interest in recent years due to such events as the evidence of lunar ice by the LCROSS mission and the announcement of the Google Lunar X Prize (NASA) (GLXP). The creation of permanent lunar outposts is a next step in exploration that will enable the discovery of lunar resources, provide a better understanding of working safely in a harsh environment, and act as a stepping-stone to further exploration. Lunar site survey and preparation are the vital first steps in the construction of such lunar outposts. Especially amidst the current economic climate, it is vital to consider cheaper and faster ways of accomplishing these next steps in exploration. Small-scale (< 100kg) lunar excavators provide an attractive solution for the immediate need of surveying future lunar outpost sites and for initial site preparation such as leveling terrain and berm building. Additionally, they provide an effective long-term solution for point-to-point excavation as is needed for resource extraction and other in-situ resource utilization tasks.

1.1 Motivation

The use of small-scale lunar excavators for lunar site survey and preparation and in-situ resource utilization has many benefits. The reduced cost and time of development and deployment affords the option of sending a team of small-scale lunar excavators for initial site survey and preparation while subsequent missions are planned. This would allow for a low cost initial mission in which many assumptions about operating autonomous or remote controlled vehicles for lunar excavation could be tested. In addition, such a platform can be repurposed as an effective long-term excavation system if proven effective once on the lunar surface. Most importantly, with no need for local human operators and minimal to no infrastructure requirements, they allow for a low-cost quick initial mission that risks zero human lives.

1.2 Lunar Excavation Background Work

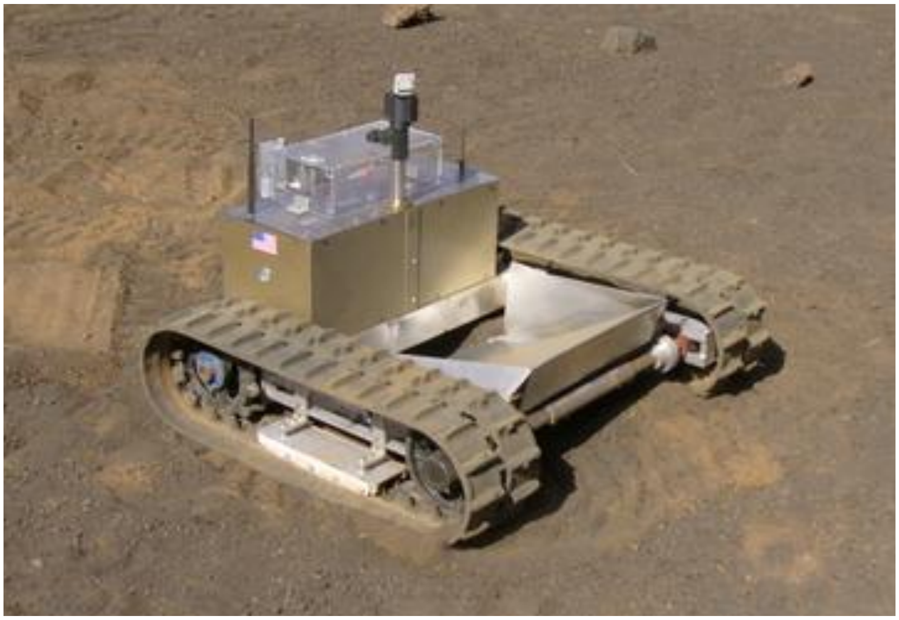

Figure 1 - Cratos Rover

Caruso (John J. Caruso, 2006) demonstrated the capability of a small low-power tracked rover named “Cratos” to pick up, carry, and dump sand. This allowed the rover to accomplish delivering regolith simulant to bury a simulated inflatable habitat, to supply an oxygen production plant with excavated regolith, and to construct a ramp. Cratos weighed in at 80kg and consumed 100W with the ability to operate for 16h out of 24h off of two 12v, 18Ahr Sealed Lead Acid batteries charged at intervals (John J. Caruso, 2006).

A study by Astrobotic Technology, Inc. determined that small-scale lunar excavators were capable of constructing a protective berm to surround a landing pad at a polar output in less than 6 months [2]. The study identifies many critical aspects of lunar excavation using small-scale excavation robots to suggest the most important factors affecting the productivity of a lunar excavation mission. A regolith construction simulator named REMOTE (Regolith Excavation, MObility and Tooling Environment) was developed in order to determine these findings by modeling an inputted rover platform’s operation in excavation scenarios.

1.3 Scope

This work investigates the effectiveness of small-scale lunar excavators through the further development and analysis of the Lysander rover. In particular, the main focus of this work is devoted to analyzing the sensitivity of the Lysander rover's productivity to changes in key operational and design parameters. The intent is to identify which parameters may have the largest effect on overall productivity of this class of lunar excavators to aid future excavator design processes and operations development. Additionally, a discussion is presented on the effectiveness of the Lysander rover and the control of rovers under medium latency (4-10 sec round-trip) and limited bandwidth.

This work does not investigate or attempt to measure the trade-offs between different designs or features of small-scale lunar excavators, though many of the design decisions of the Lysander rover are retold. This work focuses chiefly on analyzing the productivity of a capable lunar excavator rover to better understand the effect of varying its design and operational parameters. This work in no way argues the necessity of certain design or operational parameters' settings over another, but instead only considers their effect on productivity. As such, no argument is made over the necessity of choosing 4 or 6 wheels. All operational parameters' settings that are used are considered to be within the reasonable range of operation and all design parameters' settings are considered typical choices for vehicles of this nature. As such, this work wouldn't consider analyzing the effect of unreasonably slow drive speeds or the use of an unreasonable amount of wheels.

2 Development of Lysander Rover

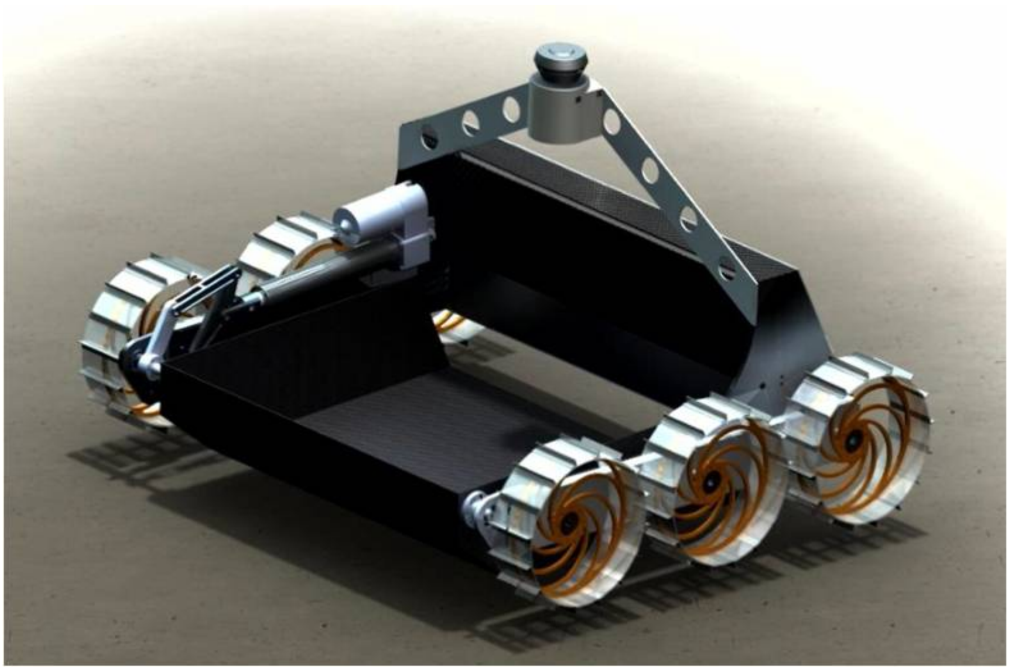

Figure 2 – Lysander CAD Model

2.1 Design Motivations

The mechanical design of the Lysander rover was greatly influenced by the CRATOS project. Like CRATOS, Lysander employs a scraper style excavation system in the form of a centrally located bucket as seen in Figure 2.The central bucket is surrounded by a six wheel skid steer system. Lysander uses an above mounted laser scanner and field fiduciary poles to localize.

Lysander has a low center of gravity and wide base in order to maintain stability in driving and pinpoint turning as well as to maximize regolith collection during scraping. The wide base allows Lysander to employ a bucket with a scraping mouth measuring approximately three quarter meters wide. The bucket can hold upwards of 90kg of regolith simulant, but a nominal load ranges safely between 20kg and 60kg in the current design due to actuator limitations; small changes and fixes could allow Lysander to achieve full bucket loads. The large bucket is meant to double as a collection bin where the majority of the mass of the collected regolith is shuffled to the back for ease of bucket actuation. The rover is meant to take multiple small 10-15kg bites of regolith that are shuffled to the back of the bucket to minimize the amount of load put on the lip of the bucket at any given time and clear the lip for the next bite. After excavation, the rover dumps the bucket's load by actuating the bucket upwards until it flips upside down to drop the load directly behind the rover.

The six active wheel skid steer system was chosen for simplicity and was thought to be a good pairing with the scraper style excavation tooling. Six active wheels and a low wide base help to direct all the forward driving force into the lip of the bucket when scraping through more challenging soil or lower scraping depths. Six active wheels also help the rover traverse the inconsistent terrain of a lunar surface. Given that it is possible for the rover to double its effective mass when carrying a large load, the low-slung body of the rover is susceptible to bottoming out and having difficulty climbing large slopes. Thus, six wide wheels – all active – were chosen to increase traction while also adding more floatation and ground coverage to get the rover out of these tougher terrain conditions. Increased floatation from more wheels means less chance of the wheels burying themselves in the sand especially during turning or movement in looser soils. Six wheels provide larger ground coverage to minimize the chance of bottoming out or getting a wheel stuck.

Under the assumption that most equatorial regions of the moon are relatively flat and incredibly long distances aren't necessary in lunar excavation tasks, Lysander uses an above mounted laser scanner and fixed field fiduciary poles for localization. A SICK LMS-111 laser scanner was used in conjunction with 1 inch PVC poles wrapped in highly reflective red tape. This laser scanner has a 270-degree field of view with half angle resolution and is capable of capturing distance and reflectivity values very accurately of up to at least tens of meters. The absolute positions of the poles and their reflectively are known a priori and used to distinguish the poles in the field to then localize based off of them.

2.2 System Details

Parameter	Value
Mass	56.38 kg
# Of Wheels	6
Wheel Diameter	30 cm
Wheel Width	12 cm
Drive Speed	80 cm/s
Bucket Width	.75 m
Max Bucket Payload	94 kg
Nominal Payload Range	20-60kg
Rover Width	1.25 m
Rover Length	1.2 m
Rover Height	.6 m

Figure 3 - Lysander stat breakdown.

In Earth gravity, the Lysander rover is able to achieve greater than 1m/s drive speeds with pneumatic wheels on a flat concrete surface, and 80cm/s on compacted lunar regolith simulant with specially designed grouser-studded non-pneumatic wheels (as seen in Figure 2). Due to a lack of a suspension on the Lysander rover, the effects of reduced gravity and tough terrain on the Moon would lessen these speeds in the actual lunar environment (Heiken, Vaniman, & French, p. 528).

Without a regolith payload, Lysander is within the 50-60kg range depending upon the number of wheels and internal electronics setup used. Despite observed mechanical limitations in the bucket's ability to actuate large loads exceeding 90kg, Lysander has proven itself capable of successfully excavating, shuttling, and dumping regolith payloads as massive as itself.

With a low and wide base, Lysander keeps quite a low center of mass well below its highest point of .6m atop the head of its above-mounted Lidar unit. This low and wide base allows for a wide bucket mouth of .75m creating a large scraping surface that is ideal for collecting the powdery top layer of lunar regolith on the Moon.

2.3 Control Software Design

2.3.1 Earth to Moon Communication Limitations

Considering that the most likely use case for this class of rover is site survey and preparation work prior to humans revisiting the Moon, they must be controlled from Earth. Reasonable estimates for the latency and bandwidth that would be encountered via satellite communication from Earth to Moon would be up to 6 s round-trip latency and 1Mbps bandwidth limitations (Sheridan, 1993). This amount of latency makes direct remote control a psychologically tiring task for any expert operator that can greatly hamper the productivity of even the most capable machines (Skaar, 1994). While a well-implemented fully autonomous system can achieve a perfect operational efficiency and possible productivities that no human operator can achieve, it poses many significant risks given the current lack of knowledge and experience in lunar excavation. Thus, a control approach incorporating the best of both autonomy and direct human control would be most suitable for initial site survey and excavation on the lunar surface.

2.3.2 Remaining Effective and Safe under Limitations

There are many sensible control schemes for a small-scale lunar excavation system to remain safe and productive, each having its own tradeoffs. Supervised autonomy poses a solution with the highest potential for productivity. Given that rover state and video feeds are at least 2 seconds latent, the soonest that a supervisor could alert or stop the rover from entering a hazardous state (i.e. stuck in a crater) is overall at least 4 seconds after the fact. Therefore, any supervised autonomous system must have a high degree of safety built into its operations.

A purely tele-operated rover is not ideal for many reasons. The strain on the driver caused by the latency causes poor operational efficiency that results in poor overall productivity (Skaar, 1994). Tele-operation approaches that employ predictive models to give the driver a sense of where the rover is currently – by estimating off the current yet latent rover status – present a promising alternative that may dramatically increase productivity but still retain the inherent safety issues.

The Lysander rover uses a hybrid approach by having a sliding range of autonomy from full autonomous operation, to semi-autonomous routines, down to basic tele-operation commands. A key design feature was to present the driver with as much rover system state as possible without inundating the driver while also having a more detailed overview of system state displayed in a separate display for a “wing-man” driver. This focus on visualizing system state is to help the driver realize and diagnose any crucial system problems in as little time as possible. Since the latent environment already exaggerates safety hazards, minimizing any additional latency due to the operator’s interactions and human element is vital.

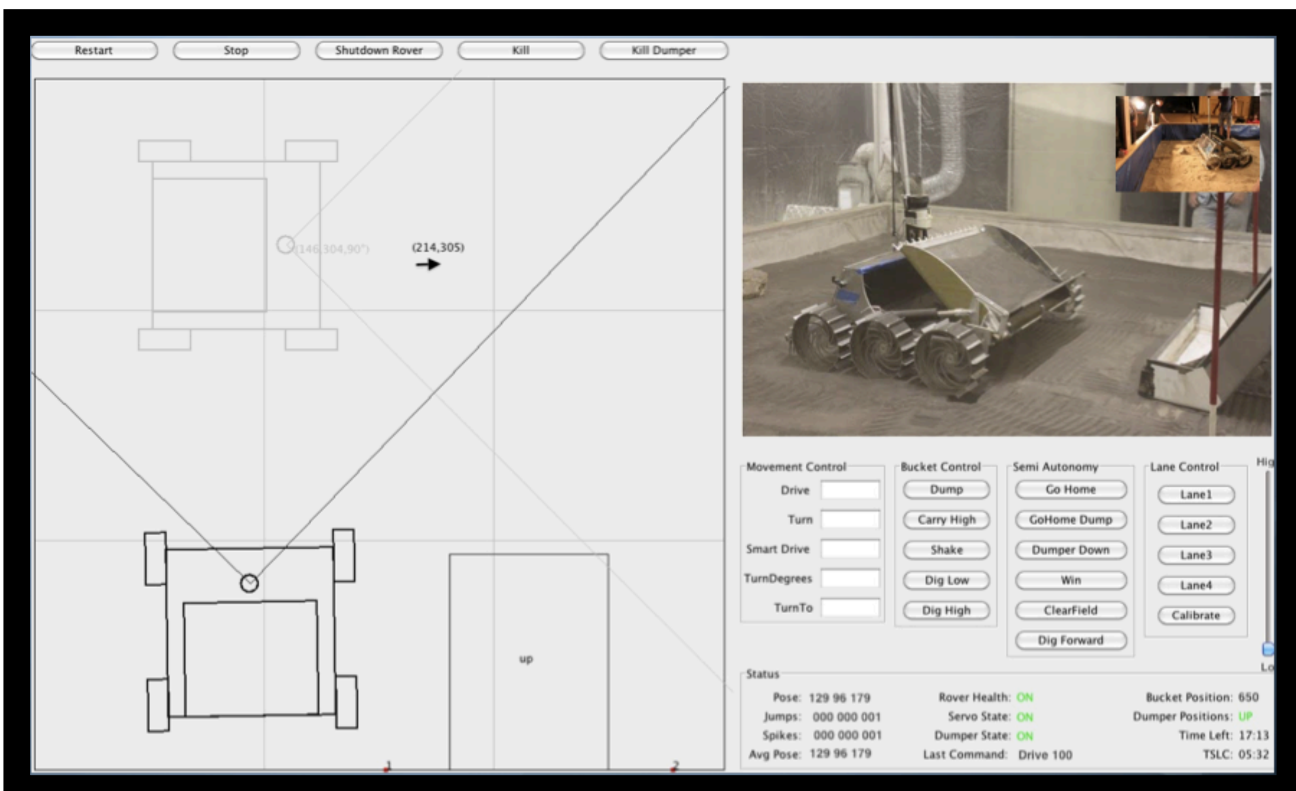

Figure 4 – Screenshot of client-side control GUI.

3 Operational and Design Parameter Sensitivity Analysis

3.1 Motivation: Importance of Analysis and Validation

The design process for any robotic platform can be a very daunting task where many assumptions need to be made to overcome the uncertainties inherent in any untested design. For example, the decision between four or six wheels in a wheeled rover platform can spur a debate over the amount of floatation, traction, and stability needed. Beyond the design parameters of a rover, it can also be a tough task understanding the effect on productivity of varying operational parameters such as drive speed or the payload ratio achieved by an excavator during general site tasks. One might ask whether or not the speed has more effect on the productivity than does the payload ratio when deciding how to best operate the rover.

In order to aid the design process of future lunar excavators, this work isolates and analyzes particular design and operational parameters of small-scale lunar excavators. The intent is to uncover and then experimentally validate the sensitivity of the overall productivity of the excavator in lunar excavation tasks to changes in key design and operational parameters. As such, this work uses REMOTE (Regolith Excavation, MObility & Tooling Environment), a task-level site work simulator, to characterize the Lysander rover's high and low sensitivity design and operational parameters and then validates these sensitivities through experimentation. With a better understanding of how sensitive the productivity of a small-scale lunar excavator system is to the varying of certain parameters, future design processes can optimize their system's productivity by knowing where to focus their design efforts.

3.2 Background Work: REMOTE

REMOTE (Regolith Excavation, MObility & Tooling Environment), a task-level site work simulator, analyzes the total time required to complete a regolith-moving task by combining traction and excavation models with transport shuttling, power draw, and recharge tracking (Astrobotic Technology Inc., February 2009). Most importantly to this work, REMOTE is also used to determine parametric sensitivity of system designs and parameters with respect to task productivity.

In other words, REMOTE provides a numerical prediction of the sensitivities of system parameters (e.g. drive speed, dig angle, number of wheels) with respect to the overall productivity of the system during particular excavation tasks. Thus, for a given general excavation task and particular excavation platform it can predict how much of an effect on productivity varying a specific parameter, such as drive speed, can have. Since these predictions are based solely off of a numerical analysis leveraging traditional traction and excavation models, it is not clear how accurate or correct they are in the real world for actual excavation systems.

3.2.1 Description of Observed Parameters

This section provides definitions of the major design and operational parameters considered in this work – whether or not they are considered in the sensitivity analysis.

- **Cut Angle** – The angle the excavation tool's (bucket) cut edge creates against the soil surface during excavation.
- **Cut Depth** – The depth of the excavation tool's cut edge from the surface of the soil during excavation.
- **Cut Speed** – The speed of the rover during excavation.
- **Driving Speed** – Defined to be the average driving speed of the rover during transport to and from the digging area for a point-to-point excavation task.

- **Number of Wheels** – This is simply the number of wheels used for the rover locomotion. The Lysander rover is capable of running 4 or 6 wheels.
- **Operational Efficiency** – Defined to be the percentage of time the rover is doing work, where work is essentially whenever a motor is being actuated. Conceptually this measurement hopes to capture the percentage of time the operator is making forward progress.
- **Payload Ratio** – Defined to be the ratio of the mass of the regolith payload to the mass of the empty rover. In other words, this is how many of itself in regolith the rover can carry with respect to mass.
- **Wheels Radius** – The radius of the wheels used for rover locomotion measured from the center of the axel to the point of ground contact when the rover has a zero moment resting on flat soil. In this work, it is assumed that all wheels have the same radius.
- **Wheel Width** – The width of the wheels used for rover locomotion measured at the contact surface. In this work, it is assumed that all wheels have the same width.

3.3 Simulator Results

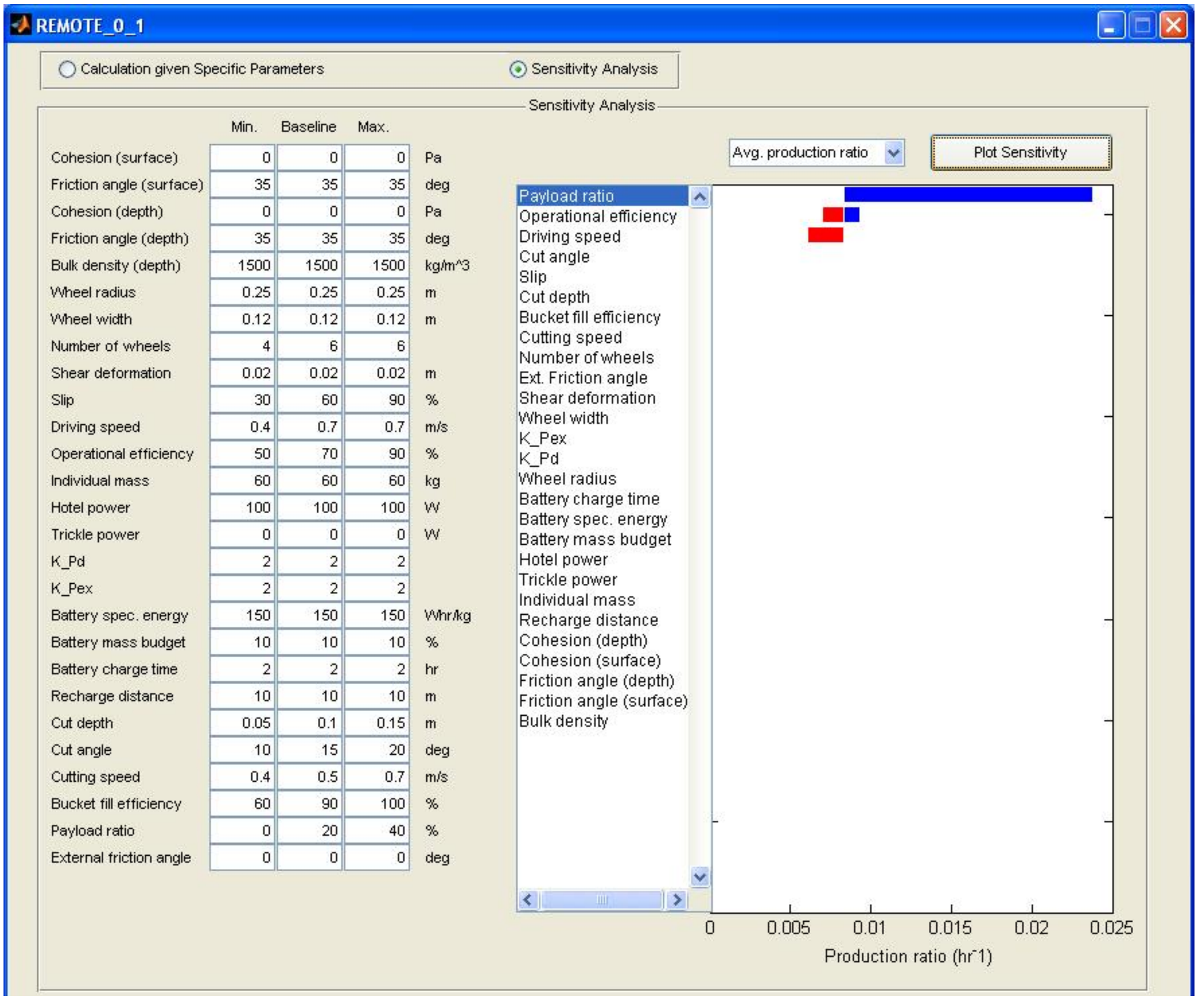

Figure 5 – A screenshot of the result page of REMOTE's parametric sensitivity analysis for Lysander.

REMOTE reported Lysander's productivity to be most sensitive to changes in drive speed, payload ratio, and operational efficiency while parameters such as the number of wheels and cut angle showed very low sensitivity. This is shown in the parametric sensitivity plot in Figure 5, where for each parameter there is an accompanying red-blue horizontal bar that plots the effect on overall productivity that will be experienced by exaggerating that specific parameter from its nominal setting. The production ratio on the x-axis serves as a measure of the payload ratio achieved per hour, hence the units of hr^{-1}. The nominal setting of all the rover's parameters creates the nominal production ratio mapped by the line created between the red and blue bars.

On the left hand side of the image, we can see each system parameter with three accompanying values: Min., Baseline, and Max. These are the minimum, nominal, and maximum settings for the system, respectively. As these are the major inputs to the REMOTE simulation, they are used to generate the sensitivity plot seen to the right. The baseline – or nominal – settings dictate the nominal productivity of the system. The minimum and maximum settings' effect on productivity is characterized by the end of the red and blue bars from the nominal productivity line, respectively. We can see that increasing the payload ratio parameter from its nominal setting of 25% to its maximum of 50% can have a dramatic effect on the overall productivity according to REMOTE's analysis.

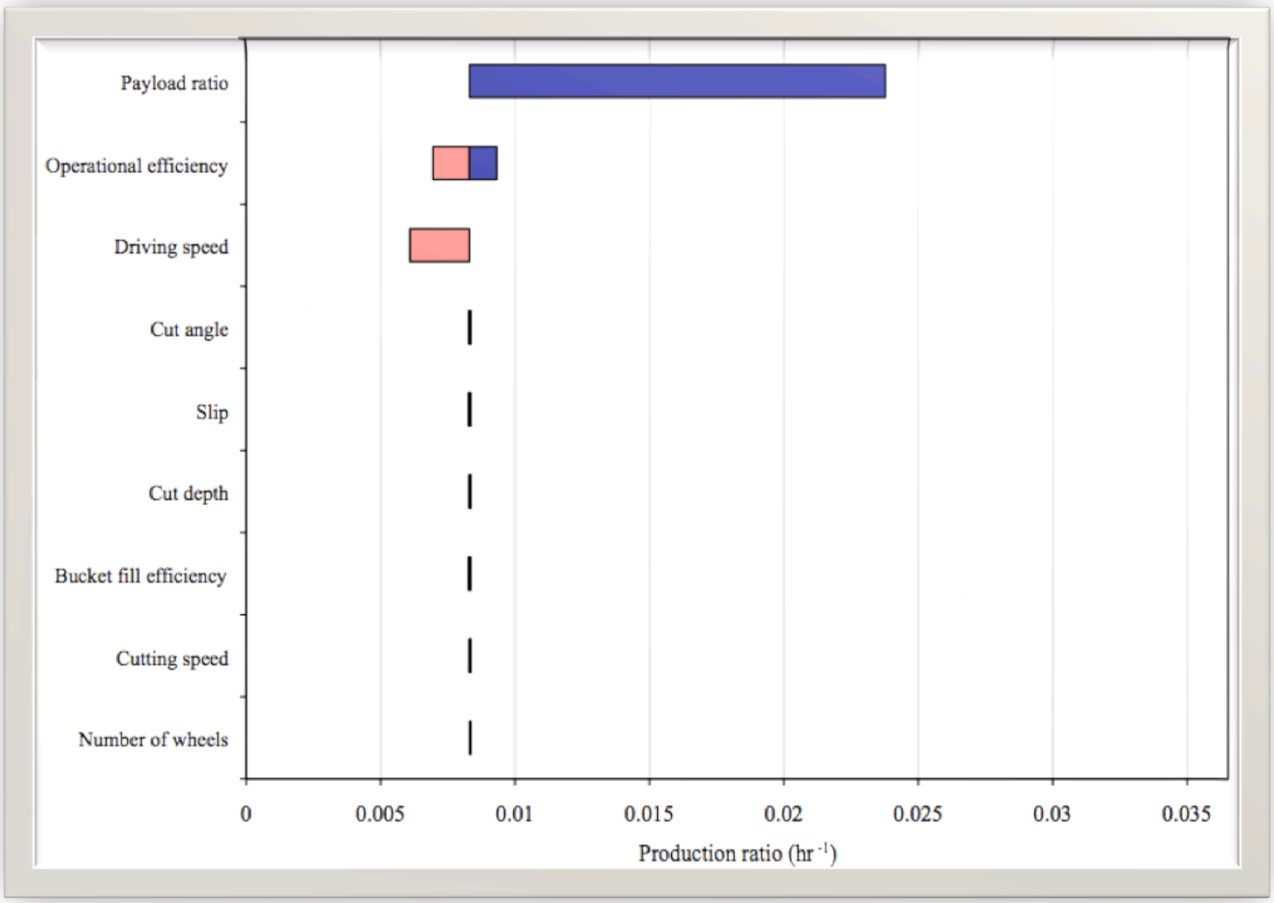

Figure 6 – Top nine most sensitive parameters of Lysander rover with respect to productivity.

3.4 Experimental Validation

Now that the predicted sensitivities of major system parameters are known, we must measure the actual sensitivity of these parameters. For the purposes of this work, we narrow our focus to consider only the drive speed, payload ratio, and number of wheels as variable parameters of interest. Note that the plot in Figure 6 shows that the number of wheels is the ninth most sensitive variable and shows little to no visibility on the plot where as payload ratio and drive speed show a large sensitivity bar. These three parameters are interesting because they are commonly discussed system parameters and because two are predicted to have a large effect on the productivity while the third should not, thus including parameters from both side of the sensitivity spectrum.

Validation of the sensitivity of these parameters is confirmed through careful experimentation. First, we choose a general excavation task cycle of point-to-point excavation over which to measure each parameter's sensitivity. The setup for this task must include an excavation area and a dump location separated by a non-trivial traversal. The task cycle consists of traveling from the dumpsite to the excavation area, collecting a pre-determined payload size from within the excavation area, transporting it to the dump location, dumping the load, and then repeating.

In order to measure the sensitivity of the system's productivity to changes in a particular design or operational parameter, we must measure the effect that varying that specific parameter has on overall productivity while holding all other parameters constant. We define productivity as the payload ratio of regolith simulant collected in a task cycle divided by the amount of time taken to complete the cycle. This can be thought of as how many times the rover collected its own self in mass per hour. Thus, if a 100kg rover collects 150kg in 1 hour, then it has a productivity of 1.5 hr^{-1}.

Next, we must define the nominal settings for each parameter to characterize the normal operating conditions and setup. For each parameter that we wish to measure the system's sensitivity to, we define either a high or low attainable value that is within a reasonable range but far from nominal to exaggerate the effect of the controlled variable.

In Figure 7, the nominal and exaggerated values for the Lysander rover are displayed. While experimenting on the drive speed, payload ratio, and number of wheels, we shoot to keep the operational efficiency constant at 75%, the cut angle at 5 degrees, and the cutting speed at 20cm/s.

The first step in measuring the sensitivities of parameters is to characterize the baseline productivity of the system. Setting all parameters to their nominal value and then measuring the productivity of the resulting task cycle with these parameters serves as the baseline. Then, to measure the sensitivity with respect to a given parameter, we perform the task cycle with that parameter in its exaggerated value while holding all other parameters to their nominal value. We then compare the productivity of this task cycle against that of the baseline.

Parameter	Low Setting	Nominal Setting	High Setting
Drive Speed	40 cm/s	65 cm/s	—
Payload Ratio	—	25%	50%
Operational Efficiency	—	75%	—
Number of Wheels	4	6	—
Cut Angle	—	5°	—
Cutting Speed	—	20 cm/s	—

Figure 7 – Low, Nominal, and High Setting Values for Lysander Rover.

3.4.1 Experimental Setup

A sandbox was constructed in the Roundhouse at Robot City in Hazelwood, PA measuring 3.8m wide by 3.8m long containing a 0.75m deep mixture of commercially available sand and silica as seen in Figure 8. The dump location was chosen to be in the front left corner and the dig location in the front right corner with a short row of rocks between them forcing a U-shaped traversal of approximately 7m as seen in Figure 8.

The field fiduciary poles wrapped in highly reflective red tape can be seen in the foreground. The field is protected from the leaky ceilings of the Roundhouse by a large suspended green tarp partially observable in the background. For these experiments, the rover was powered via a tethered yellow extension cord to eliminate the otherwise frequent and time-consuming need to recharge the rover's batteries.

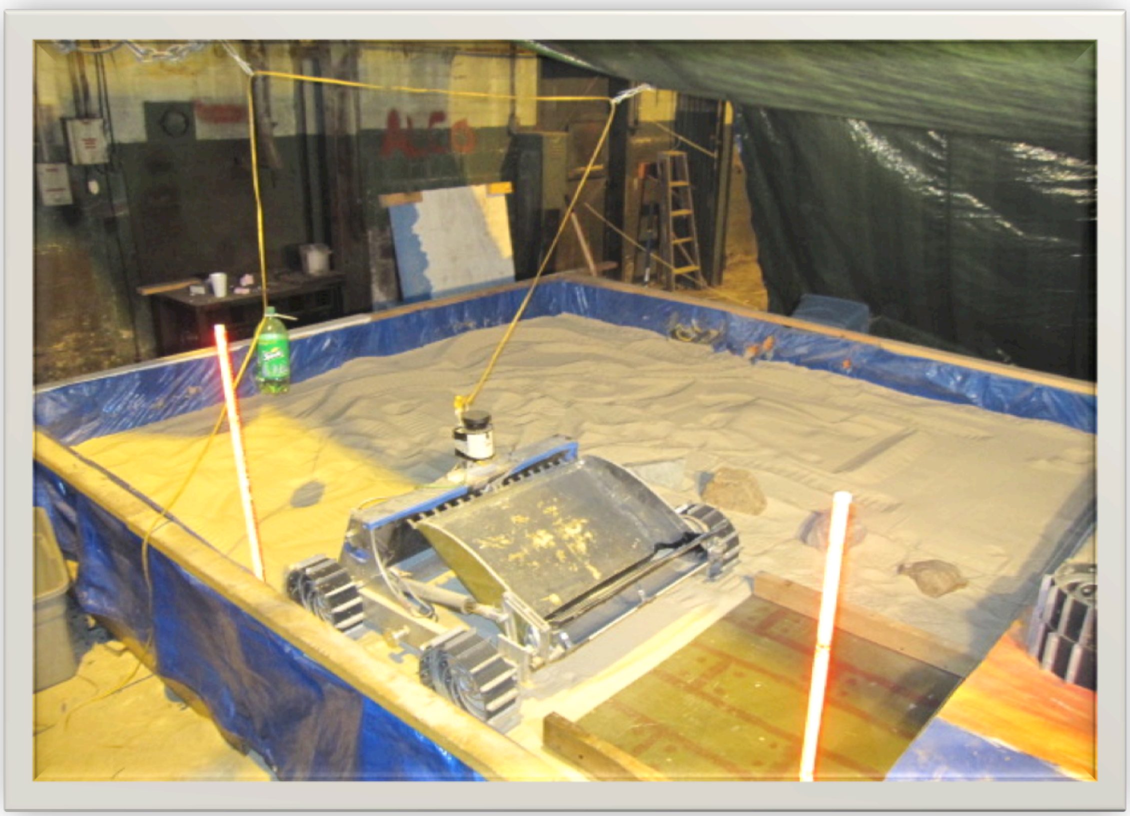

Figure 8 – View of testing sandbox at Robot City.

The point-to-point excavation task illustrated in Figure 9 was chosen as the most representative single task amongst general excavation site tasks since it is a major element, if not the most integral element, of most excavation tasks such as berm-building, trench-digging, and regolith collection for resource extraction.

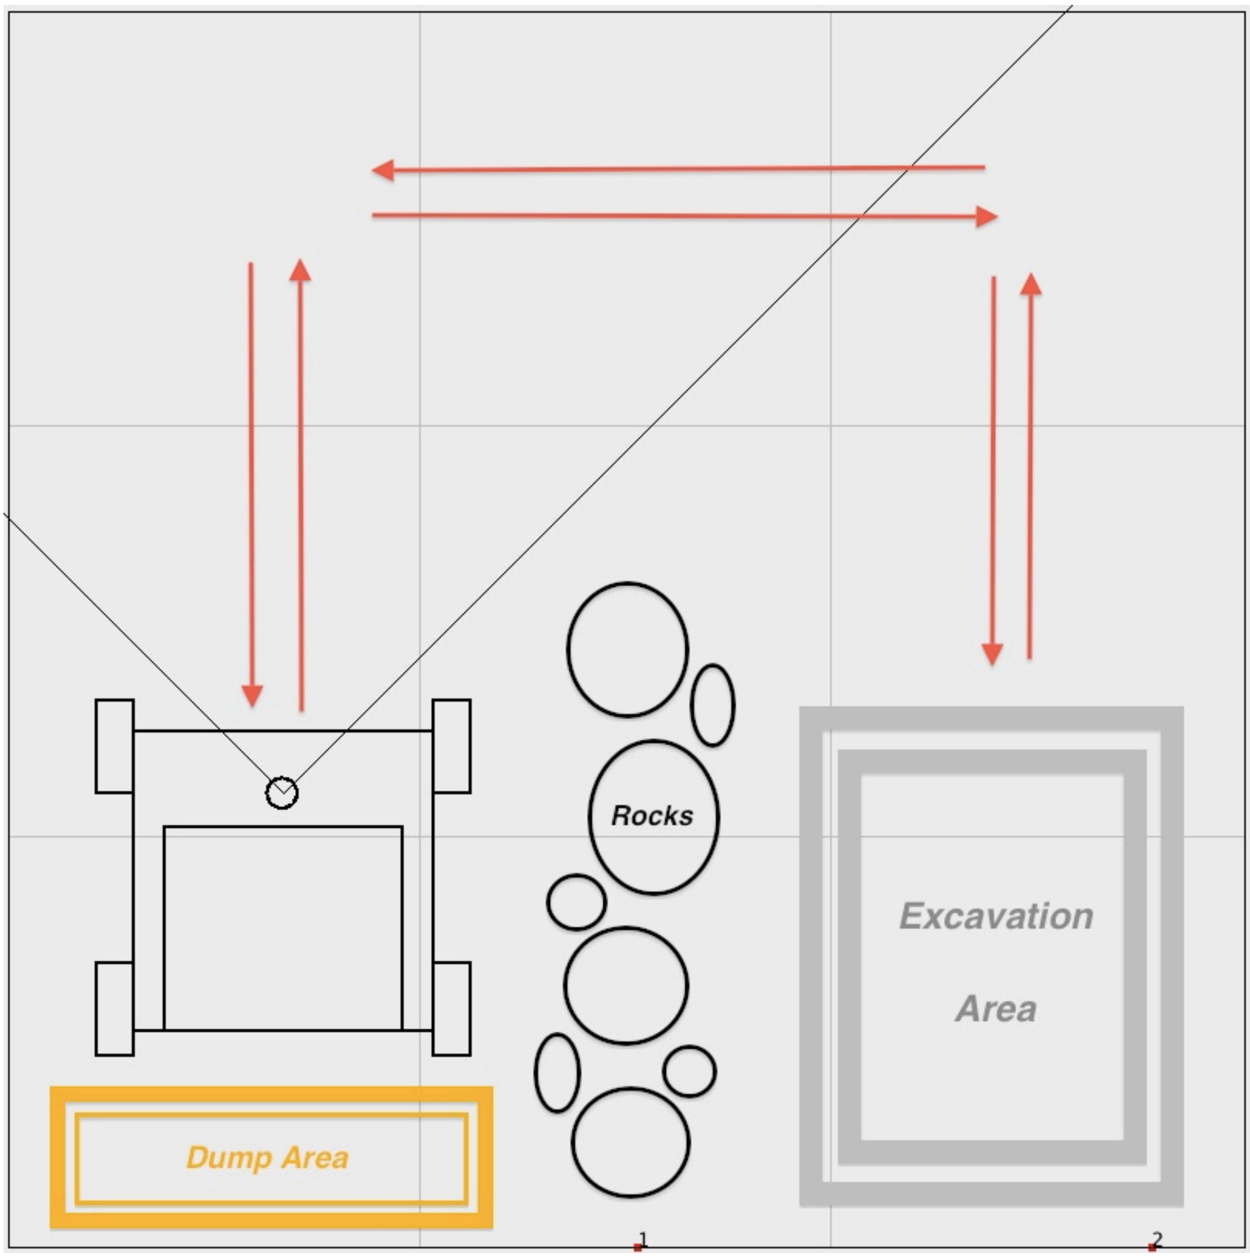

Figure 9 – Point-to-point excavation task layout.

During a test run, the rover starts adjacent to the dump area with an empty load and with the bucket in a normal carry position. The test clock starts when the rover drives forward. The stone barrier between the dump area and excavation area force the rover to follow a 7m U-shaped corridor to arrive at the excavation area. The excavation area is a roughly 1.5 meter by 2 meter carefully prepared patch of soil. After achieving the desired payload ratio from within the excavation area, the rover travels back along the same path back to the dump area and dumps its collected regolith into the dump area for measurement. The test clock stops when the bucket begins actuation for dumping.

Between each test run, the field is reset to keep consistency between runs. The most important preparation is of the soil in the excavation area to create similar digging situations for each run. A 3-step approach is used to prep the soil. First, the soil is churned with a shovel to break up any chunks. Next, the soil is compacted with a hand compactor. Lastly, the soil is flattened out with a flat rake to create a flat uniform surface.

3.4.2 Issues with Initial Experiments

During this initial batch of tests, keeping all system parameters consistent with their desired values proved to be difficult for many reasons. In fact, many test runs were aborted or neglected due to these problems. In the end, three successful test runs were conducted for each parameter being examined as well as three to establish the baseline productivity.

Initial review of the best three test runs for each parameter shows that the data agrees with the predictions of REMOTE. However, upon further analysis it was found that the issues with maintaining consistency in the system parameters rendered the experiments inconclusive of showing the trends REMOTE predicts.

After this initial round of experiments, many measures were taken to correct all the previous issues before beginning the second round of experiments. These corrections will be described in section 3.4.5 Correcting Previous Issues. Here follows a description of the major problems experienced in the initial round of experiments.

<u>Issue #1:</u> Difficulties Achieving Consistent Payload Ratios

The payload ratio tests - to measure the effect on productivity of varying the payload ratio from 25% to 50% - were run many times without much success due to difficulties in reproducing a consistent digging behavior to achieve the desired payload ratios. Since the current system had no way measuring the fullness of the bucket on the fly, the digging behavior of the rover and excavation soil were carefully monitored and instrumented to recreate the same pattern each time and thus hopefully the same payload ratio. However, due to mechanical slip in the bucket's actuation stack the bucket controller wasn't able to accurately achieve the same bucket dig angle each time.

<u>Issue #2:</u> Pose Dropouts

Due to the geometry of the field and the U-shaped traversal of the rover during the excavation cycle, placing two fiduciary posts in the field so that the rover's Lidar could always detect them and thus localize proved to be difficult. Since measuring the rover's drive speed and distance traveled was vital in order to ensure these parameters stayed close to their intended values, any losses of pose during a test run could easily ruin the data. Unfortunately, it was not detected until after these initial tests that the pose cut out briefly at a particular area of the field to muddle both the measurements for the distance traveled and drive speed.

<u>Issue #3:</u> Operational Efficiency Measurement Faults

Another major problem resulted from the method in which the operational efficiency was measured. During each run, the rover will log regular status updates at a predetermined frequency to a log file. These status updates are time-stamped and include information such as the vehicle's pose, bucket position, and the health of subsystems. Initially, the operational efficiency was measured as the percentage of status updates in which the rover was "operating" with the frequency of updates being 100 hertz. The rover was considered to be "operating" during a given status update if its speed was above a certain threshold or if the bucket position value was in the digging range and the rover was moving forward. At first glance, this way to measure the operational efficiency seems to agree with the definition of operational efficiency in section 3.2.1 Description of Observed Parameters. In actuality, it has many problems.

First, considering the problem explained in the previous paragraph concerning the occasional loss of pose, this measurement missed all of the data points in which the rover was actually operating but had a reported speed of zero due to pose loss. Second, it doesn't make sense to have a speed threshold because even if the rover is experiencing slip or driving slowly it should still be considered operating and not awaiting a command. Similarly, if the rover is attempting to actuate the bucket, this should also be considering operating. Thirdly, even when the rover wasn't moving, the raw Lidar data nonetheless jittered which caused the pose to subsequently bounce around ever so slightly. This bounce could be perceived as movement and thus if quick enough would be considered a speed above the threshold. This jitter in the data was likely due to mechanical forces and vibrations acting on the Lidar unit.

3.4.3 Analysis and Results of Initial Experiments

Despite the faults detected in the data of the initial experiments as described above, the results of the experiments are nonetheless interesting.

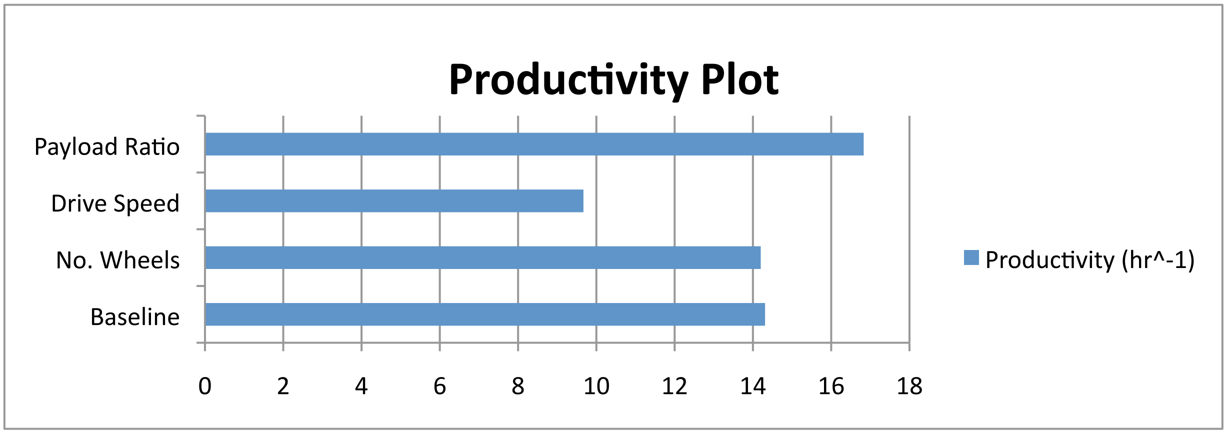

Figure 10 – Productivity plot for each varied parameter and the baseline.

In Figure 10 we can see how the average productivity of each of the sets of test runs varied from that of the baseline set of tests. In Figure 11, the effect of exaggerating these parameters is more apparent. The drive speed makes the largest percentage difference in productivity with over a 32% change while the number of wheels hardly makes more than 1% difference.

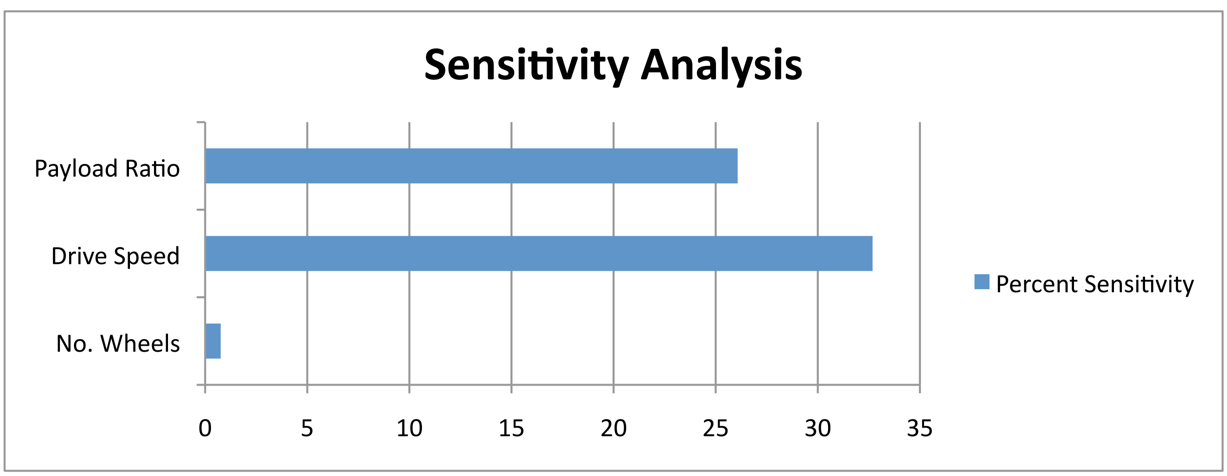

Figure 11 – Plot of the percent sensitivity of parameters against the baseline.

Although these results are interesting, the consistency issues described in section 3.4.2 Issues with Initial Experiments cause the data to have too large of an error to argue any correlation between varying the instrumented system parameters and the observed change in productivity. In Figure 12, it is easy to observe the major spread of payload ratios achieved in the best three exaggerated payload ratio tests. In addition to the issues discovered in the method by which operational efficiency was calculated, the values measured were themselves also highly variable.

Test Name	Distance (m)	Speed (cm/s)	Operational Effic. (%)	Payload (kg)	Payload Ratio (%)	Prod. (hr^{-1})
Baseline Test 1	6.2	64.5	46	15.6	22.9	11.5
Baseline Test 2	6.1	62.1	51	16.9	25.2	13.4
Baseline Test 3	5.5	65.2	50	19.2	29.3	17.8
Payload Ratio Test 1	6.3	62.1	52	29.8	48.0	17.0
Payload Ratio Test 2	6.6	64.5	45	23.9	37.6	15.5
Payload Ratio Test 3	6.4	62	52	26.2	41.7	18.0
Drive Speed Test 1	6.5	41.3	44	16.8	25.0	9.2
Drive Speed Test 2	6.3	45.3	53	15.3	22.3	9.0
Drive Speed Test 3	6.2	43.5	49	17	25.4	10.8
4 Wheel Test 1	6.1	60.5	40	18	27.1	13.4
4 Wheel Test 2	6.2	62.7	52	18	27.1	15.1
4 Wheel Test 3	5.9	60.1	49	17.2	25.7	14.1

Figure 12 – Digest results from each test of 1st round of experiments.

The mean baseline productivity measured 14.2 hr^{-1} with a large standard error of 3.24, thus having a confidence interval ranging from 10.5 hr^{-1} to 17.9 hr^{-1}. The payload ratio tests had a mean productivity of 16.8 hr^{-1} with a standard error of 1.4, thus having a confidence interval ranging from 15.4 hr^{-1} to 18.2 hr^{-1} nearly overlapping the upper half of the confidence interval of the baseline tests. Thus, we cannot say with confidence that the system's productivity is sensitive to increasing the payload ratio.

Similarly, the drive speed tests' data showed the same problem but to a smaller degree. The drive speed tests had a mean productivity of 9.7 hr^{-1} with a standard error of 1.1, thus having a confidence interval ranging from 8.6 hr^{-1} to 10.8 hr^{-1} and again overlapping slightly with the baseline interval. Therefore, we cannot say with confidence that the system's productivity is sensitive to changing the transport drive speed.

3.4.4 Summary of Initial Experiments

Although there was an apparent difference in the mean productivity experienced by changing the payload ratio and drive speed and not in the number of wheels as REMOTE had predicted, the problems with data collection and holding system parameters consistent between test runs forces us to take these results with a grain of salt. These problems necessitate further experimentation using better practices and preparation to successfully validate or debunk the predictions of REMOTE.

Nevertheless, the effect that varying the payload ratio and drive speed have on the overall productivity as compared to varying the number of wheels is noticeable. The productivity observed for the baseline test runs and the 4 wheel tests are close compared to the difference between the baseline and either of the other two observed parameters. Thus, one could hypothesize from these experiments that the system productivity is less sensitive to a switch between 4 or 6 similar wheels than to changing the drive speed or increasing the payload ratio. It is important to note that this work makes no claim as to the importance or necessity of 6 over 4 wheels nor the effect either setup may have on the ability of the rover to traverse tough terrain. Instead, this experiment assumes the rover to be excavating easily traversable terrain and only attempts to discover whether or not switching between 4 or 6 wheels has any noticeable effect on the amount of the soil the rover may excavate per unit time.

3.4.5 Correcting Previous Issues

The initial experiments' major issues included trouble maintaining accurate actuation of the rover's bucket, inconsistent soil preparation, pose dropouts, and a flawed method of computing the operational efficiency. A description of how these problems were corrected before the second round of experimentation is described below.

<u>Issue #1:</u> Difficulties Achieving Consistent Payload Ratios

It was found that there was a significant amount of mechanical slip in the actuation stack of the bucket, which made it difficult for the system to have an accurate reading of the bucket's position. Without an accurate measurement of the bucket's position coupled with the fact that the difference between two discrete position values in the range of values that corresponded to digging angles equaled a significant difference in digging angles, it proved incredibly difficult for the current system to maintain a consistent dig angle during excavation. Furthermore, any inconsistencies in the soil pattern between runs would result in even more varied digging behavior. Consider if a bucket scrapes flat soil as opposed to upward sloping soil. The upload sloping soil will be more difficult to excavate since the dig angle is effectively steeper than the rover had intended it to be had it been flat soil.

These digging problems were corrected first by attempting to reduce mechanical slip and then retune the bucket controller in software to account for the slip and difficulty in maintaining consistent bucket position values. In addition, more care was taken in soil preparation to ensure a flat excavation area. This was done by laying down a 2x4 wood plank on each side of the excavation area, carefully leveling them with respect to each other, and then dragging another 2x4 along the other two to flatten the excavation area as a last step in the soil preparation. Although this flat excavation area cannot be guaranteed on the lunar surface, realistic soil patterns were not nearly as vital to these experiments as consistent soil patterns are so as to best compare the performance of the vehicle between test runs.

Lastly, the new bucket controller was tested by commanding the rover to set the bucket to a large set of values in the digging range and measuring the angle of the bucket to the frame of the rover with a set of levels. This time around, the bucket angles were very consistent. To ensure consistency of the bucket angle between runs, the method was periodically used to ensure that bucket position values matched up with actual bucket angles.

Issue #2: Pose Dropouts

The pose dropouts were corrected by finding a better location for the fiduciary posts so that the only drop out location would be in the corner going into the excavation area on the rover's U-shaped traversal of the field. Rather than taking the corner with a rigid right angle turn, there existed a path around this corner consisting of a more rounded traversal that ensured no loss of pose. With a little operator practice in maneuvering the rover around the corner in this manner, the pose loss issue was mitigated.

After the pose issue was corrected and the speed estimates became much more accurate, it was discovered during the beginning of the second round of experiments that the rover's speed during the increased payload ratio tests was consistently smaller than during the baseline tests. Since the rover does incorporate its speed estimates from the pose in its control for the rover's driving, this issue was likely a product of the fact that these payload ratio tests increased the mass of the rover and thus slowed it down a little. Since the drive speed needed to be held consistent to its nominal value for all tests except the ones examining a decreased drive speed, small modifications to the rover's driving pulse width modulation needed to be made.

Issue #3: Operational Efficiency Measurement Faults

Lastly, the operational efficiency had a much easier solution than was previously thought. Instead of computing when the rover was operating through examining the data after the test, it was much easier for the rover to report in software when it was attempting to operate through a boolean flag kept in the logs. This method produced very repeatable measurements of the operational efficiency so that in subsequent tests it could be examined to ensure consistency between test runs. Of course, this also took some training for the operator because the operator needed to have consistent driving and digging behavior to ensure identical operational efficiency between runs.

3.4.6 Analysis of Second Round Experiments

This second round of experiments produced much more consistent and valuable results after incorporating all the fixes described above and checking the consistency of the data between each run. Of course, many runs had to be repeated for many reasons including operator difficulties; yet at least three successful runs were performed for each parameter being examined as well as the baseline.

In Figure 13 the results of each set of experiments are shown. Recall that parameters that weren't being examined, such as distance traveled and operational efficiency, were to remain as constant to their desired nominal values as possible during all tests. Thus, it is important to verify that these parameters did in fact remain constant enough so that any slight changes in their values between tests can be considered negligible toward any change in the productivity of these tests. Applying a t-test to the distances and operational efficiencies measured in each round of tests against those of the baseline shows no significant variations in these parameters.

Test Name	Distance (m)	Speed (cm/s)	Operational Effic. (%)	Payload (kg)	Payload Ratio (%)	Prod. (hr^{-1})
Baseline Test 1	6.3	64.9	73	14.1	25.0	17.9
Baseline Test 2	6.8	65.3	66	13.2	23.4	14.3
Baseline Test 3	6.8	61.5	72	10.9	19.3	12.4
Payload Ratio Test 1	6.6	66.4	76	28.6	50.7	28.7
Payload Ratio Test 2	6.8	63.9	72	29.3	52.0	29.3
Payload Ratio Test 3	7.2	62.0	72	29.3	52.0	26.6
Drive Speed Test 1	7.4	41.9	75	12.6	22.3	8.7
Drive Speed Test 2	7.3	43.7	77	14.3	25.4	10.3
Drive Speed Test 3	6.9	43.1	80	14.6	25.9	11.0
4 Wheel Test 1	6.8	60.5	77	14.3	26.2	14.0
4 Wheel Test 2	6.9	66.1	79	14.9	27.3	15.2
4 Wheel Test 3	6.5	65.3	74	13.5	24.8	13.1

Figure 13 - Digest results from each test of 2nd round of experiments.

Shown in Figure 14 are the confidence intervals of the productivity for each set of tests. Each red bar represents the confidence interval of the productivity of the corresponding set of tests listed on the y-axis. Interval endpoint values are displayed just next to the ends of the bar.

We can see that the mean baseline productivity measured 14.9 hr^{-1} with a standard error of 3.2, thus having quite a large confidence interval of 11.7 hr^{-1} to 18.0 hr^{-1}. The number of wheels had an interval that fits nicely in the middle of that of the baseline. This shows that the actual productivity of the system with 4 wheels is not statistically different from that of using 6 wheels with high confidence.

The interval for drive speed sits just below that of the baseline. Although the endpoints of these intervals are close, we know with 95% confidence that the productivity of the system with a low drive speed is statistically different from that of the baseline. The interval of the payload ratio tests on the other hand is seen to be far removed from that of the baseline. Thus, we know also with high confidence that the actual system productivity in the exaggerated payload ratio tests is statistically different than that of the baseline.

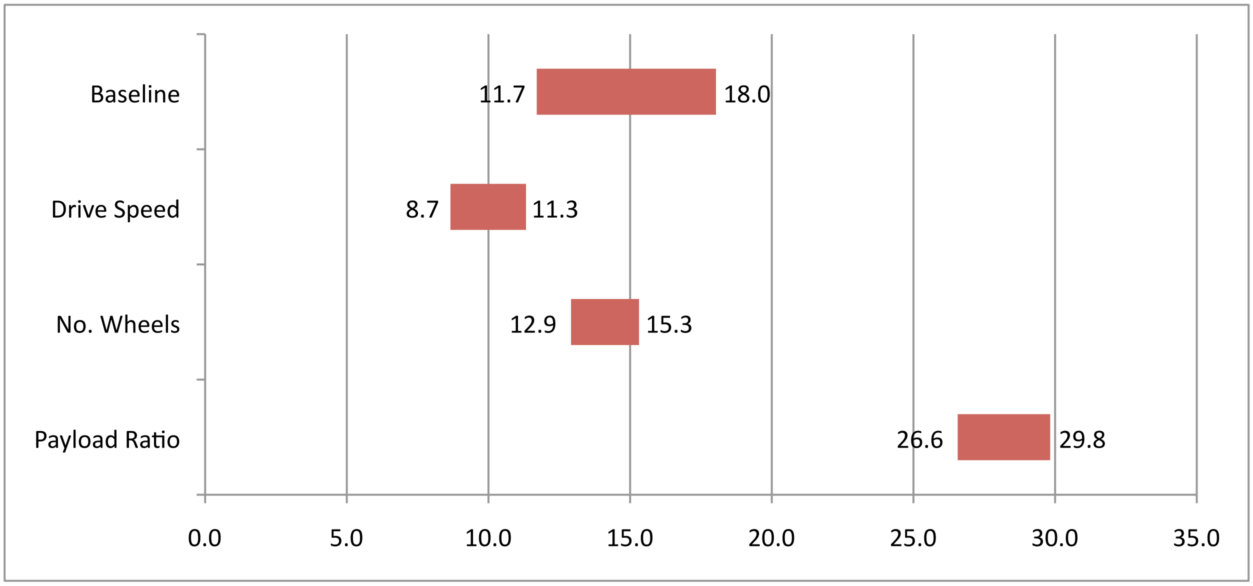

Figure 14 – Confidence intervals of productivity.

3.4.6 Results of Second Round Experiments

The results of the second round of experiments agree with the results of the first round of experiments in showing the drive speed and payload ratio are highly sensitive parameters while the number of wheels is comparatively much less sensitive. However, there are some major differences.

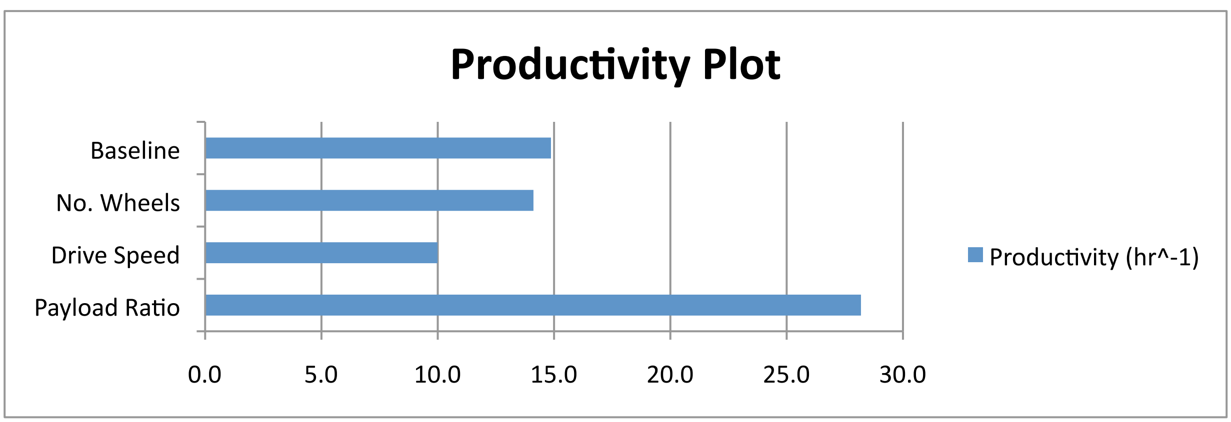

Figure 15 - Plot for the mean productivity of each set of tests.

As seen in Figure 15, the productivity of Lysander was greatly increased by doubling the payload ratio from 25% to 50% despite the relatively small traversal of about 7 meters total for the excavation cycle. The drive speed had a large effect on the productivity and the number of wheels had less of an effect, as was the prediction of REMOTE, however the magnitude of their effect is interesting.

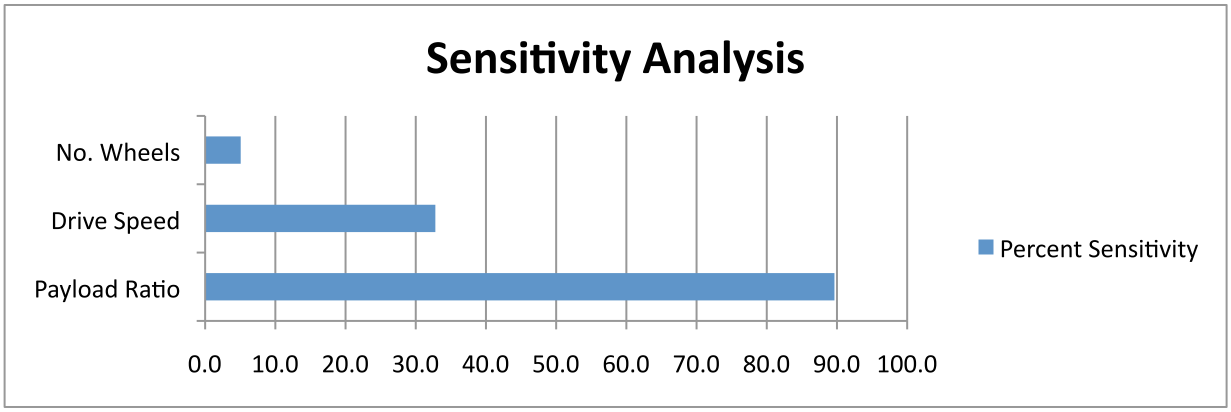

Figure 16 - Plot of the sensitivity of each set of tests compared to the baseline.

In Figure 16 we can see that varying the payload ratio had a much bigger effect on the productivity than the drive speed did. Of course, the sensitivities of these parameters are dependent upon the predetermined values for the nominal and exaggerated settings. Thus, it is not clear if the payload ratio is somehow more sensitive in general for any choice of nominal and exaggerated values as compared to drive speed. In fact, this is surely not the case if we had chosen the exaggerated payload ratio to be very close to the nominal setting. Thus, if we had perhaps doubled the drive speed just as we had doubled the payload ratio from its nominal setting, the difference in drive speed could have had a larger effect on productivity than the difference in payload ratios. Nevertheless, finding these relative differences is not the focus of this work. Instead, highly exaggerated settings of parameters were chosen in an attempt to discover which parameters are most sensitive with respect to the productivity.

3.4.7 Summary of Second Round Experiments

The results of the second round of testing fit much closer to the predictions of REMOTE in showing that payload ratio is very highly sensitive for this range of values, drive speed is highly sensitive as well, and changing the number of wheels does not in fact affect the productivity much. It is interesting to note that REMOTE predicted that changing the number of wheels should produce less than 1% difference in productivity while the experiments showed it having about a 5% difference. This discrepancy could be explained by a higher variance in controlled parameters such as speed or payload ratio and a product of having only 3 test runs worth of data to compute this difference. Alternatively, it is also possible that REMOTE might underestimate the difference between operating with 4 and 6 wheels.

4 Future Work

There is still much work to be done to fully understand the effectiveness of small-scale lunar excavators and how best to design and operate them. Although this work analyzes the Lysander rover as the benchmark for understanding small-scale lunar excavators, it would be very beneficial to perform the identical design and operational parameter sensitivity analysis on other small-scale lunar excavators to see how the results match up. Additionally, another insightful undertaking would be to design the next iteration of the Lysander rover based off the results of this work and then characterize the productivity of this new vehicle.

Beyond the transport drive speed, payload ratio, and number of wheels, there are many other important design and operational parameters whose effect on productivity could be better understood through experimentation. For example, REMOTE identified the operational efficiency of Lysander to be a parameter for which the productivity of the rover was highly sensitive. It's not clear how to instrument this parameter to measure its effect on productivity without just trivially adding time to the cycle period to decrease its value or finding a more efficient operator or even fully automating the system to increase its value. Both of these approaches bring up a philosophical dilemma as to the meaning and subsequent importance of the operational efficiency. Given that a fully autonomous system should theoretically have little to no downtime and thus nearly 100% operational efficiency, can we compare this to a 70% efficient human operator and argue based of their results how sensitive the system is to changes in operational efficiency without considering *how* both the computer and human perform the task?

In this work, the operator trained to perform the excavation task in the same manner each cycle while using methods of digging that were found to be effective. Unfortunately, it is currently not well understood how to best dig any area of soil for a scraper style rover such as Lysander especially if we're unsure of the composition of its particles and compactness at each depth. On top of that, major differences and uncertainties in lunar regolith composition from that of terrestrial soils make the task even harder. It is important also to realize that the best method of digging for a scraper style rover will likely be far different than that of a bucket wheel or any other excavator type.

Beyond finding the best method in which to excavate a patch of regolith for a given platform, it is equally unclear how to best plan an excavator's digging cycles for a given excavation task such as removing a predetermined volume of regolith from a site. Furthermore, regardless of the planning a scraper style rover will have troubles excavating out a perfect square volume of regolith smaller than itself or much larger than itself due to physical constraints of the workspace of the bucket. A task such as berm-building can get even more complicated when considering planning around the need to scale the berm mid-task to dump atop the berm to make it taller as well as choosing where to get excavate that soil from. Poor planning here can lead to an unstable berm and the dangers of landslides.

Ultimately, much about the effectiveness of small-scale lunar excavators will not be known until their mettle is put to the test on the Moon. The actual effects of gravity on the productivity of a system or the safeness and efficiency of a given control mode will not be fully understood until they are experienced by actual rovers on the lunar surface.

5 Conclusion

Small-scale lunar excavators currently provide an effective and enticing solution to initial lunar site survey and preparation with the potential of being a highly productive long-term solution if understood how best to design and operate them. This work has shown through simulation and experimental validation that small-scale excavators' productivity during general excavation tasks is highly sensitive to changes in transport drive speed and payload ratio achieved, but not the number of wheels. The design process of future small-scale excavators can benefit from this when aiming to maximize their system's productivity by understanding which design and operational parameters to focus their efforts on and which to not focus their efforts upon.

Bibliography

Astrobotic Technology Inc. (February 2009). "Configuring Innovate Regolith Moving Techniques for Lunar Outposts". *Final Study Report Submitted to NASA Exploration Systems Mission Directorate, Contract NNJ08TA79C Lunar Surface Systems Concepts Study* .

Foundation, X. (n.d.). *GLXP*. Retrieved from Google Lunar XPrize: http://www.googlelunarxprize.org/

Heiken, g., Vaniman, d., & French, b. (1991). *Lunar Sourcebook: A User's Guide to the Moon.* Cambridge University Press.

John J. Caruso, D. C. (2006). *Excavation on the Moon: Regolith Collection for Oxygen Production and Outpost Site Preparation.* Cleveland, Ohio: NASA Gleen Research Center.

NASA. (n.d.). *LCROSS*. Retrieved from arc.nasa.gov: http://lcross.arc.nasa.gov/

Sheridan, T. B. (1993). Space Teleoperation Through Time Delay: Review and Prognosis. *IEEE Transactions on Robotics and Automation* , 592-606.

Skaar, S. B. (1994). *Teleoperation and Robotics in Space.* Notre Dame, IN: AIAA.