2016 CU Boulder Systems Engineering Paper

School:

Colorado University at Boulder

Team members:

Rod Jafari, Alexander Silverman, Brandon Romero, Sam Oliver, Ryan McMullen, Raza Qazi, Bo Cao, Burak Karaoglu, Bharat NC, Luo Jingwen, Yanyu Xiong, Ryan McCormick, Nikolai Tangdit, Alex Mulvaney, Robert Belter, Mathew Hanley, Erik Pronk, Jake Bomark, Ahmed Shahid

Faculty Advisor:

Nikolaus Correll

Introduction:

Space programs around the world, public and private alike, are beginning to solidify plans of visiting and eventually creating settlements on the Moon and Mars. Although the cost of space travel is decreasing dramatically with the advent of re-usable systems like SpaceX's Falcon 9, it is currently, and will most likely continue to be, extremely inefficient to bring heavy building materials into space. In order to further reduce the cost of settling the Moon and or Mars, the resources available on location must be utilized. This concept of ISRU, or In Situ Resource Utilization, is what drives NASA's Robotic Mining Competition. If a team establishing a presence on Mars can bring a versatile robot capable of harnessing the planet's resources, then it can use the payload space saved to bring more valuable equipment or more redundant life support to increase the mission's effectiveness and chances of success.

Our goal is to, through competition, refine the priorities of such a robot and optimize its performance in the harsh conditions of another world. One day what our team learns by striving to win this competition may help NASA take on the challenges of ISRU on another planet, make mining more efficient and safer here on Earth, or produce engineers with the experience to do so.

The competition's rules and scoring attempt to mirror the challenges that would be faced on Mars. Each team must produce a robot capable of navigating through rough, obstacle ridden terrain to excavate simulated Martian regolith and pierce through the layers to reach "ice" simulated by gravel, then returning its payload to a collection bin as many times as possible in a 10 minute match. In completing this task the robot must withstand the harsh terrain and abrasive regolith while being designed to minimize power consumption and weight and maintaining as much autonomy as possible. The robot is to be scored on its performance in all of these categories.

Our aim is to place among the veteran teams in the competition by creating a robot that is effective in the short run and effortless to use and maintain in the long run, making it ideal for a settlement on mars where every hour of manpower must be efficiently used. We believe that, through our systems engineering approach, we have produced a robot that is capable of demonstrating a hands-off level of autonomy, an efficient use of power and data resources, an effective and weight conscious design, and a simple, easy to maintain construction.

Table of Contents:

Background and Systems Engineering Approach:

As a new RMC team, our budget, experience, and membership were all limited. The team grew from a small handful of members working on the project to more than a dozen contributing members over the course of the school year. The ramp up was slow but by the 2016 spring semester a solid membership was formed and productive meetings began to be held regularly.

Because of the team's limited time, and budget, less time and money could be spent on prototyping. This meant that the team had to research the competition and the concepts involved in a more in-depth fashion. We looked to real world problems that were similar to the challenges faced in the competition for inspiration and looked to previous competitions to check if our theories were correct. For example, our excavation tool and collection device were derived from the popular bucket wheel excavator, used in industrial mining, and altered to match our needs. We were also able to analyze videos from previous years to find common problems and avoid them in our design without having to test and run into them ourselves. Problem's included wheel traction, tool jamming, and accounting for the robot's shifting center of gravity. Although problems will still arise, we believe that learning from past competitions is an invaluable tool.

Each subsystem of the robot was drafted into an inclusive design to maintain compatibility, but because of time and budget constraints the prototyping phase had to be kept concise. Prototypes would often be built to simply test concepts, and could rarely be mounted onto the robot before finalizing. Because of the lack of test data in integrated systems, many parts of the robot were designed with arguably overkill factors of safety to the detriment of weight in order to improve confidence and prevent failure. In upcoming years we will be able to draw on this year's design and data to produce a more efficient and perhaps better integrated robot.

The basic design process for a system on the robot was to first brainstorm solutions to the problem faced by the system. After compiling some ideas they were checked against real world applications. After altering our plans we would verify and further alter the concepts and components involved through research and testing in order to improve confidence in our design. Once a base level of confidence was reached the team would convene to discuss the compatibility of different systems with each other and modifications would be made in order to guarantee that the systems would integrate properly to create a working robot. Some sacrifices in the process had to be made in the interest of time due to our late start, but the process seemed to produce a clean result.

Priorities in Competition:

To streamline the design process and improve our chances of success, a set of priorities had to be decided upon in the competition. The priorities chosen would determine our design direction and our allocation of team resources. The priorities were influenced by the team's limiting factors, like budget, time, and member expertise.

After analyzing the rules, viewing previous competitions, and examining why some robots succeeded and others failed, we had a basic understanding of what factors were involved in success. These factors were the excavation speed, regolith depositing speed, robot speed, autonomy, weight, power consumption, and last but not least: reliability. We found that winning robots were defined by their ability to quickly excavate regolith and then travel to the collection bin and deposit said regolith quickly. The importance of robot speed was relative to the robot's capacity and speed in excavating and depositing regolith. The ability to mine deep enough to collect "ice" also had a large impact on score. However, the most important factor in success appeared to be reliability. We saw several robots that we deemed extremely effective in other categories fail almost completely when a key component of the robot failed.

The team determined that, with reliability to be considered heavily in all of our designs, we would prioritize the ability to slowly but surely excavate "ice" and then autonomously return it to our collection bin. Weight, speed, and efficiency would be sacrificed for the goal of reliably and autonomously delivering "ice" to the collection bin.

This decision was influenced by the fact that about half of the team studies computer science or has more interest in the autonomous aspects of the competition, and by NASA's decision to score autonomous function so highly. The decision to attempt to mine the "ice" was made because of the large disparity in value between regolith and "ice", if "ice" was successfully mined its score would outweigh any sacrifices made to get to it, and if we failed to excavate the "ice", our tool could still succeed in excavating regolith. Reliability (and resistance to the elements) was chosen as a major priority because of the number of failures we witnessed in our research. The prevalence of mechanical failure paired with the team's newcomer status and inexperience in the competition made mechanical failure a threat to be considered. Speed was considered to be a low priority because our calculations showed that excavation speed and capacity had a much larger impact on our final score than speed would. Our low budget and lack of time, in addition to the scoring weights of other criteria, made weight and power consumption a low priority.

It's also important to note that pursuing autonomy as a priority yielded few sacrifices or tradeoffs for the club as many of our members joined only to work on that aspect. If the robot failed to be autonomous we would simply revert to a manual control system without altering the robot.

Chassis and Drive Train:

The chassis of the robot is extremely important in guaranteeing that each individual system could be well integrated into the robot, as they are all mounted onto the chassis. After considering our preliminary excavator and collection bin designs, we found that a chassis made up of a strong box-like design would yield an easy to work with robot that could accept a variety of different systems without much modification. The following picture shows the design chosen for the robot, it was chosen due to its simplicity and ease of mounting.

The material used is a tubular aluminum stock that provides excellent resistance to bending and a high factor of safety so that the frame could be trusted while under full load. The end bars were extended so that the wheels could be wider while still mounting on the smaller output shafts of the gearboxes.

An “I” shaped chassis with suspension was considered but ultimately discarded due to the apparent lack of need for suspension and the ease of mounting systems to a rectangular frame.

The next component in the drivetrain is the motors and wheels. The motors were chosen due to their small size, light weight, and ease of integration with pre-made, sealed, and encoder integrated gearboxes. The motors were also chosen due to their ability to remain in stall for extended periods of time without failure, as we would not be able to easily afford extra motors. Our other options were lighter weight but were air-cooled and provided less resistance to heat failure under stall.

After plenty of research and testing, we determined that wheels needed to be large in diameter and high in traction to be able to drive over smaller obstacles and keep the robot from sinking into the regolith during turns. It was found that traction in regolith was best created by adding flaps or tabs to the wheels so that the robot would have excellent forwards and backwards traction while maintaining low traction laterally for turning maneuvers. Because of the regolith's tendency to slip upon itself, the material of the wheel appeared to have little bearing on traction while the addition of tabs had a large effect on traction due to their ability to push on a larger amount of regolith than a flat wheel could.

In the interest of cost we constructed wheels from premade plastic lawnmower wheels. While seemingly unprofessional, the wheels were easy to work with and after removing the tread and replacing it with rolled sheet aluminum they were light weight and effective. The following image shows a wheel with its tabs removed for testing indoors.

A track-style system was considered but not pursued due to the increased complexity and weight involved with a very small increase in performance.

Mining and Excavation Tool:

The team considered three designs for the mining tool. The first was a bulldozer style design that would simply scoop a large amount of regolith and bring it back to the collection bin. Analysis showed that this design, while good in theory, was difficult to execute and resulted in a much smaller regolith capacity while removing the possibility of excavating "ice". The second design was a drill type design that would slowly bring material up with a corkscrew design held within a tube. This design was found to be prone to high friction and low efficiency but perhaps being able to reach the "ice".

The third design was a bucket wheel excavator. This seemed to be the most efficient and straight forward design but lacked the ability to reach the "ice" due to its circular design. This is an image of an industrial bucket wheel for reference.

The bucket wheel is unable to be fully plunged into a surface due to how it collects the material it removes. We determined that in order to adapt the design to our smaller application and need of excavation depth, we would mount the buckets on a chain instead of a wheel.

Our "bucket-chain" design then had to be mounted on a linear motion device in order to reach the desired depth during competition. Once the design allowed for the desired depth, it was found to be too tall to fit on the robot, which lead to it being mounted on a rotational device so that it could be deployed when needed and stashed for the rest of the competition.

These complications were deemed worth the trouble after analyzing previous competitions and weighing the costs of extra systems vs their benefits.

The design of the individual buckets was heavily inspired by one team members experience with shovels. The buckets were designed to be sharp, pointed buckets that would move around a pulley without contacting the walls of the hole being dug anywhere but at the tip of the bucket. The buckets were constructed out of bent 0.032” 6061 aluminum sheeting as it was found to be strong and light enough for the application. Thinner aluminum was tested but had a tendency to fold over its edge when digging through gravel.

Storage and Deposition System:

The robot’s storage needed to be able to handle a large amount of material while still maintaining the ability to deposit said material into the collection bin. Two main designs were analyzed.

The first design examined by the team was the dump-truck style system. Seen with varying levels of success in competition it was a preferred design due to its versatility and speed in delivering material to the collection bin. We found that the use of other systems in similar applications was rare.

The second design was the tabbed conveyer belt, this type of system is commonly used in very large scale systems like the bucket wheel excavator, and similar applications like in factories where materials of similar densities are moved from place to place. The downside of this design is a slower speed and a possible tendency to be prone to regolith friction. However, this design could use a lower current system to slowly deposit a larger amount than the dump truck system could.

Ultimately the tabbed conveyer belt system was chosen due to its lack of large moving parts and its lower cost. Although the dump truck’s speed was highly valued, the costly linear actuators and high forces necessary to make the system a reality outweighed its benefits. The conveyer belt also gave the team more freedom in mining tool design and placement due to the lack of a large clear area for the dump truck to move through.

Automation:

The National Aeronautics and Space Administration's Robotic Mining Competition (NASA RMC) provides unique challenges in creating software that can autonomously navigate and mine a simulated environment. Because the competition simulates conditions similar to those on Mars conventional simultaneous localization and mapping methods will not work. In this paper we outline our software framework to address this challenge.

Challenges

Localization and Odometry

Given a global map, conventional localization methods use its spatial relations to predetermined features in the environment and odometry to determine where it is located in that global map. However, the constraints given by RMC and the environment in which the robot is driving through requires us to create our own localization methods. In terms of odometry, the arena floor is made out of a basaltic regolith simulant called Black Point1. Because it is a granular material the robot will be driving over a dynamic surface. Thus, shifts in the surface and wheel slip will make the measurements given by the wheel odometry inaccurate and unusable. In terms of localizing given predetermined features in the environment, the competition prohibits the use of the arena walls to localize. Therefore, we cannot add predetermined features in the global map. All we can do is place a banner on a bin that we can use to localize. With that banner we can localize, however, this means we have to rely on complex computer vision algorithms or additional hardware to use this.

Navigation

Once we can localize ourselves with a reasonable amount of certainty, we can then navigate the arena to mine and deposit our materials. However, from a software perspective this provides a greater challenge. As we move through the area we must determine where randomly placed obstacles are, and then maneuver around those obstacles if the pose a large enough risk to the

Tools and Sensors

To facilitate the software development process we are using two opensource Software packages, Open Computer Vision (OpenCV) and Robot Operating System (ROS). OpenCV will enable us to quickly develop and prototype computer vision algorithms, that we can optimize later. Meanwhile, ROS will provide us with a meta operating system that facilitates the integration of new sensors. In terms of sensors, we plan on using an RGBcamera such as the Asus® Xtion5, a nine degrees of freedom inertial measurement unit, wheel encoders, and a web camera. The RGBcamera will serve as our obstacle detector, while the rest of our sensors will be used to localize our robot.

System Design

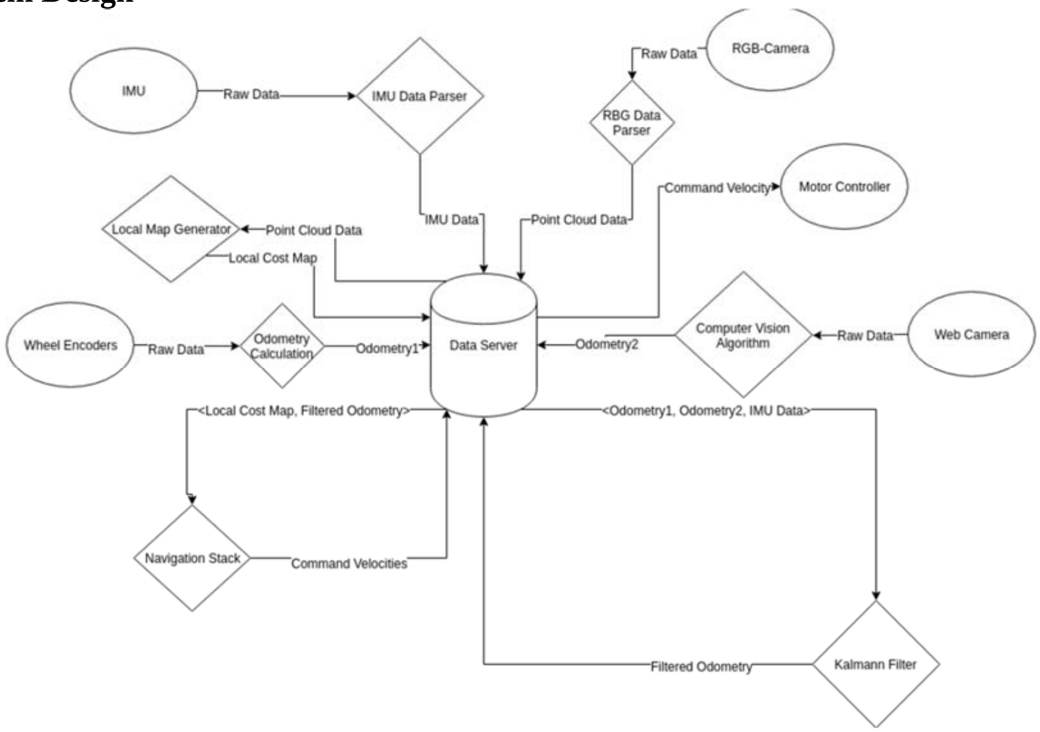

Figure 1. Interactions between nodes and server software

Overview

Using ROS we will follow the publish subscribe design pattern in which we have nodes send messages of a particular topic to a server, then the server will send the appropriate message to a node requesting a specific topic. In particular we will have a series of nodes (designated as diamonds in **Figure 1**) that receives and/or sends topics. For example the Navigation Stack node receives messages from the topic of local cost map and filtered odometry from the server, and then publishes messages to the server of the command velocity topic. Nodes can have various task such as read raw data from a sensor, parse it, then publish it. While some servers do purely computational task like the Kalman Filter. In the following sections we will look at some of the nontrivial nodes. In the following sections we will give an overview of the computational nodes (the ones that do not parse raw data) and the computer vision algorithm.

Computer Vision Algorithm

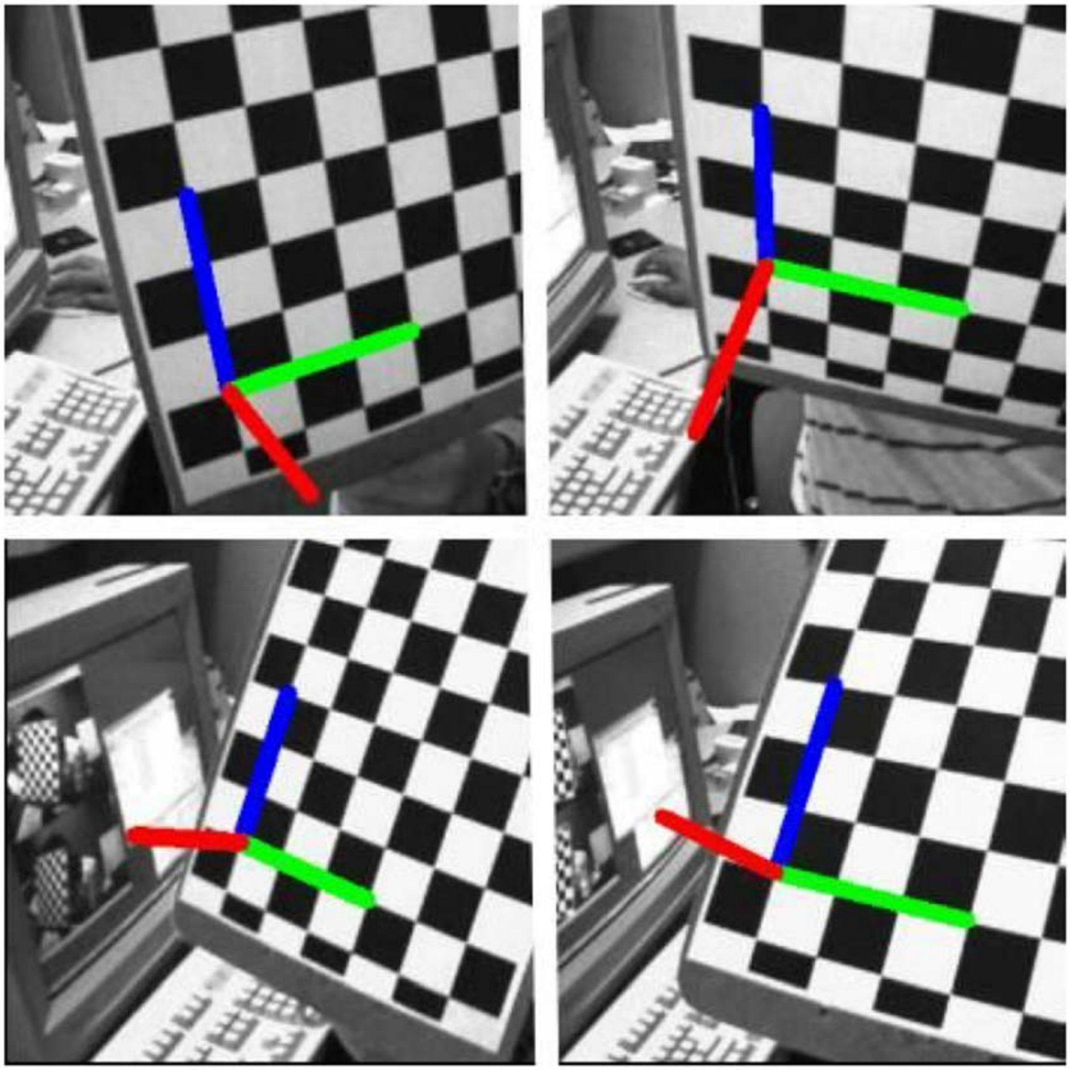

Figure 2. Pose Estimation using a textured surface

With a textured surface we can estimate the orientation and distance of that object using a standard web camera. Therefore, if we know the initial location of the textured object (which we will since this will be our marker) we can do an inverse transformation to get the location and orientation of our robot in the arena.

Kalman Filter

All nodes that publishes the odometry will have some noise or error associated with it, for example the odometry calculation node has error due to wheel slippage. However, even if there is noise or error within the system it can reveal some information. Therefore, we use multiple sources of odometry to get a potentially better estimate. To merge these measurements we use a Kalman Filter.

Local Map Generator

As the robot moves throughout the arena, it will need to know if there are obstacles within its vicinity. We keep track of these obstacles with a local map. To determine if there is an obstacle we use point cloud data. When looking at a flat surface at an angle with an RGBcamera, it generates point cloud data that has a linear gradient in terms of depth. Since the arena will be somewhat level, except with the obstacles, we can determine whether or not our linear gradient has been disrupted. However, since the surface is not completely level, we have to determine a tolerance appropriate enough to determine whether not a gradient change is significant enough to be considered an obstacle.

Navigation Stack

Once we have a local map and have a good enough source of odometry to localize with we can navigate throughout the arena. Given the robot's location and its destination we can generate a path to that destination using the A* algorithm. Here we discretize the arena, so we can tell the robot which grid location it is at and which grid location it has to go to. Now if no obstacle are detected in the vicinity it moves towards the next grid determined by A*. If an obstacle is detected in the next grid cell, then A* will generate a path around and resume its navigation to the final goal.

Mining Strategy

To effectively mine the area for 10 minutes without letting the robot fall into one of its own holes, the robot will generate an obstacle in whichever location it has mined in to prevent contact with the location again. The robot will start at the rear right side of the mining area and slowly move to the left and then forward and to the right once again, much like a typewriter’s motion. This will allow the robot to avoid becoming trapped and provide fresh mining area for the entire 10 minutes. If possible, after testing in Florida, we will reuse holes we have dug in order to mine more of the “ice” with less effort.

Future Work

Because the surface in which the robot is traversing is dynamic, we may find ourselves stuck in a hole or other object that was not detected. Thus we have to design escape maneuvers to get out of those situations.

Budgeting:

The team was funded mostly by the CU Boulder's Engineering Excellency Fund and its student organization allocations committee. Our budget came out to approximately $4000 dollars for the robot with a 10% error for any issues we would encounter.

Because we created a detailed bill of materials for the robot our estimates were very close to accurate. We ended up spending about $4500, but were also given minor sponsorships by SparkFun and Advanced Circuits, two local companies. SparkFun provided the team with $150 dollars of merchandise, which was spent on an IMU, while Advanced Circuits provided the team with one custom PCB that was used to take in encoder data and output motor control pwm signals.

The team was also partially "funded" by our advisor's lab in borrowing some key electronics, namely the nvidia tk1 test board, a Kinect, camera, and a robot test platform. These contributions saved the club over $1000 which allowed us to compete.

Conclusion:

We believe that a systems engineering approach helped us streamline our design process and will make it easier to adapt our design in the coming years and help our successors develop better designs by using our in-house documentation as guidance.

Although we were limited in resources in preparation for the 2016 NASA Robotic Mining Competition, We're confident that our design decisions and our choices to prioritize certain factors like reliability, deep excavation, autonomy and simplicity will steepen competition in Florida and place us near the veterans of the competition.