ILR10 - Progress Review 11

Justin Morris
Teammates: Awadhut Thube, Alex Withers
Team G: The Pit Crew

November 13, 2020

1 Teamwork

As the Fall Validation Demonstration approaches, our efforts have been focused on integrating our subsystems and testing the integrated system. After the previous progress review, we felt that we had the major components of our system (brinkmanship, planning, and camera operation) in working order. Brinkmanship and planning had been integrated in the simulator, allowing the rover to execute a simulated mission in which it navigated around the pit and approached the pit edge at multiple waypoints. Having tested these features independently on the surrogate rover, it was time to integrate them in real life as well. We spent several days at our testing site, Gascola, operating our rover in a field environment. These field tests have provided valuable information about how our brinkmanship and navigation algorithms operate in the real world.

On the first day of field testing, 10/30, we were unable to complete a dry run of the Terrestrial Pit Edge Validation Test. Each attempt to execute this code resulted in new errors, some of which we were able to fix in the field but some of which required additional development. We had also not integrated the camera operation program with the rest of the rover's code at this point. This day of testing was frustrating, but provided insight about how to improve and adjust our code, as well as which supplies we needed to bring with us in order to test effectively.

We spent several days making the necessary adjustments to the code from home. Awadhut and I worked on improving the brinkmanship code to more robustly detect sheer edges and unsafe slopes. We also integrated the image capturing code into the system. Meanwhile, Alex adjusted the waypoint navigation subsystem to correct some errors we had observed as the rover backed away from the brink. The next time the team gathered for a field test, we first confirmed the functionality of all subsystems in the area around my house before we returned to Gascola. Only once we were confident that every portion of the system was working properly did we proceed to the field test site.

When we arrived, we found that our previous testing site was directly in the middle of a gathering of ATV riders and dirt bikers. We were able to find a second location with a suitable cliff to test with. Once again, our first few tests exposed issues with our software. In particular, the brinkmanship stopping conditions were triggering as soon as the rover activated its brinkmanship subroutine. We had learned from our previous test, and purchased a portable monitor that could be plugged into the rover computer. With this, we were able to visualize the point clouds and meshes that the rover was generating. We determined that there was a layer of particles in the point cloud that were above the terrain surface, possibly due to glare from the sun. We implemented a filter for the point cloud in the software, and this resolved the issue.

After fixing the bugs that we found, we were able to achieve our first complete execution of the Terrestrial Pit Edge Validation Test. The rover drove to three vantage points, used the brinkmanship subroutine to navigate close to the edge (as seen in Figure 1), then captured three panoramas at different tilt angles at each location. The first time this test was performed, the rover still stopped short of the cliff edge, beyond what we would consider

Figure 1: The surrogate rover, Blue, approaches a cliff edge at the Gascola site. Photo credit: Awadhut Thube

appropriate for the goals of the project. We ran several successive tests, adjusting the brinkmanship parameters each time to improve performance. We encountered some cases where the rover behaved inconsistently, such as when the rover's trajectory was not perpendicular to the edge of the cliff. W mounted several cameras around the test site to capture a variety of perspectives of the rover, and also filmed with a handheld smartphone camera. In the last test of the day, we also collected a rosbag file that contained the point clouds published by the Realsense camera. We will analyze this rosbag in order to further optimize our brinkmanship and navigation software, in order to achieve the best possible performance in future demonstrations of the system.

Having successfully completed a dry run of the real life portion of our Fall Validation Demonstration, we are now confident that our system will be fully operational for the real thing next week. We are currently working on bringing our updated version of the brinkmanship algorithm back into the simulator. Once this is done, we will be able to run an entire mission in the simulator, although this process takes several hours to complete. Alex and Awadhut have been running the simulated mission over and over, fixing small issues along the way. Once the rover is able to execute the entire mission without failing, we will have accomplished both our goals for the Fall Validation Demonstration.

2 Challenges

Traveling to Gascola to test in the field was a significant undertaking, and it brought with it a variety of challenges. During the first trip, we were reliant on Alex's car to power a monitor so that we could make edits to code and view the results of tests, which slowed us down significantly. We had to carry the rover back and forth from the car to the cliff side every time we wanted to make significant changes. One particular issue that we encountered was that the rover used a different coordinate frame than the simulated rover. This meant that several sections of our code had to be adjusted to match the rover's coordinate system. In the field, we were forced to use trial and

error to determine the appropriate coordinate frame, which took up a significant amount of time. Furthermore, because this was not the only bug in our code, it was difficult to separate issues caused by the coordinate frame from issues caused by other bugs.

When we returned for our second test, we were unable to access our original testing site due to the dirt bikers. We were unsure we would be able to test at all, but we knew that our opportunities were limited and so we had to find a way. We were able to find another suitable cliff edge nearby. This time, we were more prepared due to the lessons we had learned during our first field test, so the challenges we had previously faced were not as much of an issue. Our portable monitor was powered by the rover computer, so it was much easier to connect to the rover and examine the results of our tests. This was how we discovered that the point cloud was producing a significant number of outlier points that were causing the mesh generation to be inconsistent. At first we were not sure how to resolve that issue, but thankfully the outliers were easily identifiable as being far from the median Z position of the point cloud, and could therefore be filtered out without too much trouble.

3 Future Plans

We are now in the final stretch of the project before the Fall Validation Demonstration next week. We feel that we are prepared for the FVD, and we are confident that all our planned demonstrations will succeed. We will be performing one more test at Gascola this weekend, so that we can fine-tune the surrogate rover's performance and make sure that it gets appropriately close to the edge in all cases. During this test, we will also confirm that we can stream video at a decent frame rate from Gascola, so that we can provide a live demonstration of our rover in action.

Simultaneously, we will be continuing to work on the simulator until we reach a point where the simulated rover can execute an entire mission successfully. This will include implementing image capturing within the simulation, to mimic the image capture routine that we run on the surrogate rover. Because the simulated rover does not have a pan/tilt mount, the simulation will simply capture one high-resolution image, instead of stitching together a panorama from several images. This will allow us to collect an appropriate amount of image data over the course of the mission, even if the method of doing so is somewhat different from the hardware implementation.

For the FVD itself, we will prepare a sped-up video showing the entire mission execution in simulation. We will also endeavor to have a live feed of the simulation as the rover executes the mission, so that the reviewers can compare to the premade video. Meanwhile, our team will be at Gascola, where we will prepare the rover for the Terrestrial Pit Edge Validation Test. We will show this test being executed live through a phone camera that will be connected to the Zoom call. This method should maximize the amount of functionality that we can demonstrate live, instead of through pre-recorded video.