Individual Lab Report

Erik Sjoberg

Team C – Column Robotics

Rohan Thakker, Job Bedford, Cole Gulino

IRL 11

April 12, 2016

Individual Progress

Debugged lighting-induced optical flow failures

I spent time this week getting to the root cause of some unexpected behavior. Both of our quadcopters began to exhibit very significant drift during hovering, although nothing significant had changed in our code. At first we thought the erratic behavior was due to bugs in our movement code, however after testing it became clear that the quadcopters were no longer capable of stabling holding their position, although this had been working quite well recently.

Once we had narrowed the issue to something which was known to be working previously, I began to examine the environment for changes which could possibly explain the behavior. Since we knew optical flow is dependent on the texture of the floor surface we thought this might be related to the problem, however nothing obvious had changed about the layout of the wooden panels we had been using to improve the visual texture of the floor as can be seen in Figure 1 below.

Eventually, I noticed that one of the overhead lights in the corner was actually not turned on as it usually was, and the team spent 10 minutes or so searching for the light switch to turn it back on. In fact, it looks like the switch was already on but for whatever reason this particular light had not actually turned on as expected. Flipping the switch off and back on (then waiting 10 more minutes for the lights to turn back on) did the trick. With this light on again, all of our problems were suddenly gone. It turns out the multiple lights had been mitigating the effect of the significant flicker of each light, which the optical flow system could not handle.

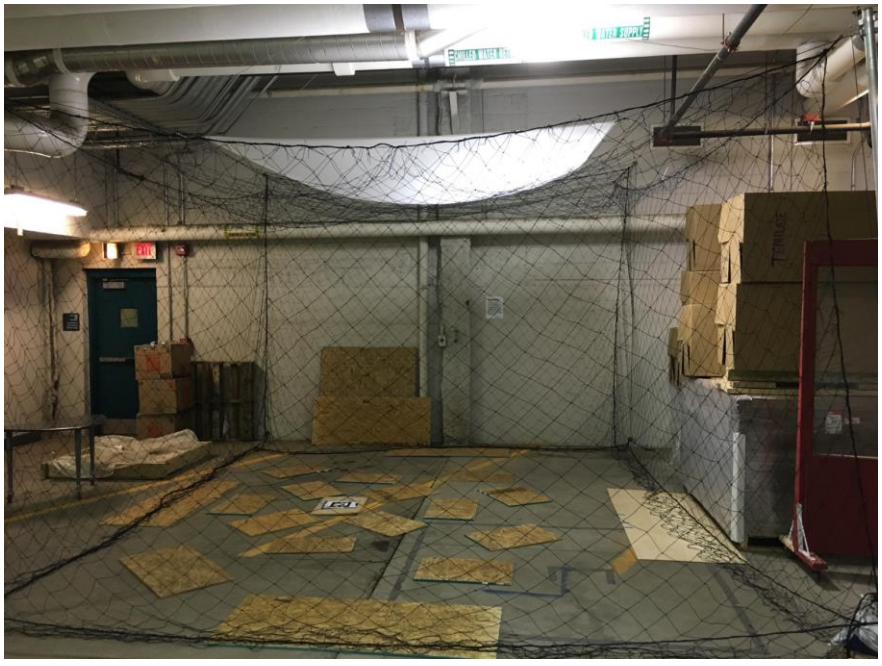

Figure 1: B-level test area with arc lighting

Integrated all subsystems into a working demonstration

I completely integrated all of our work to date resulting in a complete and working system.

The first task was to normalize the various coordinate frames used for the camera and the quadcopter's internal pose estimates. I used some simple transforms to provide the higher-level "global" planning code with a uniform interface which treats movement to the right as positive X, and movement forwards as positive Y. Movement up was left as positive Z. This enabled for more simple reasoning about the desired behavior of the system at a high level.

I also merged together the two large feature branches for cone search and april-tag hovering. Since these branches had been worked on individually, this required quite a bit of re-plumbing and tweaking of message types etc.

In order to improve the speed of development, I also implemented ROS parameters for various tunable C++ code features. This meant that we could tune the behavior of our april-tag filtering RANSAC algorithm without requiring a complete re-compilation of the code. Using these knobs, I tested and tuned the complete system to exhibit stable flight and reliably identify and land on the april-tag after finding it.

Challenges

Optical flow sensitivity to lighting

High frequency flickering of the lights in the B-level caused a significant amount of trouble this week in the form of difficult-to-diagnose poor performance of optical-flow based odometry. This was eventually traced back to the fact that one of the overhead lights was actually turned off, which presumably meant that there was fewer sources of light to overcome the flicker and as a result the optical flow algorithms failed. This has been a valuable lesson in the importance of lighting for vision-based systems which operate at high frequencies.

Teamwork

Teamwork went smoothly this week; I performed most of my work early in the week to integrate our two disparate code branches, then Cole Rohan and Job built upon this work by developing the yaw control functionality and debugging the final demo. This split of work was effective as it allowed me to focus on the integration which required a single person to tackle it, while the other teammates were able to build upon this. Starting earlier in the week also meant that we didn't suffer the same integration failure due to a last-minute branch merge which we experienced last week.

Plans for Upcoming Work

Over the next two weeks I will be working to further polish our demo to enable an expanded search area. I will also be focusing on our higher level goals of global pose estimation via RGBD SLAM as well as obstacle avoidance.

Meanwhile, Job will be working on tuning our precision docking and Rohan and Cole will be focused on finalizing our yaw control and then providing the infrastructure to integrate global pose estimates into our Kalman filter.