

# Technical Report: Test 3

SkyNet

## 1. Methods

Our approach to Test 3 can be broadly broken into three main components: localization, path planning, and control. Our localization is designed to take the IR beacon locations, stereo camera inputs and the inertial measurement units (IMU) to determine the current location of the drone in the simulation relative to the position of the gates. Path planning seeks to determine the optimal path through multiple gates based on the localization output. Finally, our control algorithms translate the path planning output into a specific action plan, then iteratively executes on these.

For localization, we chose to make use of the visio2 library built in to ROS, which takes stereo camera image as input, tracks harris corners and computes depth information, and using the resultant visual odometry to return the position within the simulated warehouse in three dimensions. Since the visio2 library cannot always be relied upon to give the exact position of the drone (particularly in instances where not enough features can be tracked for accurate visual odometry), we intended to also apply approximate gate locations (as measurement update) based EKF, along with IMU based state estimations. However, unfortunately due to lack of time, our gate based distance estimation was not fully developed into a satisfactory performance level, and IMU could not be incorporated. So for the submission we solely relied on the visio2 package. See future works section for the ideas planned to be implemented. Nominal gate location based EKF would have been critical as it is the only source that is robust against drift over time (which are inevitable for the visual odometry and IMU based state estimations).

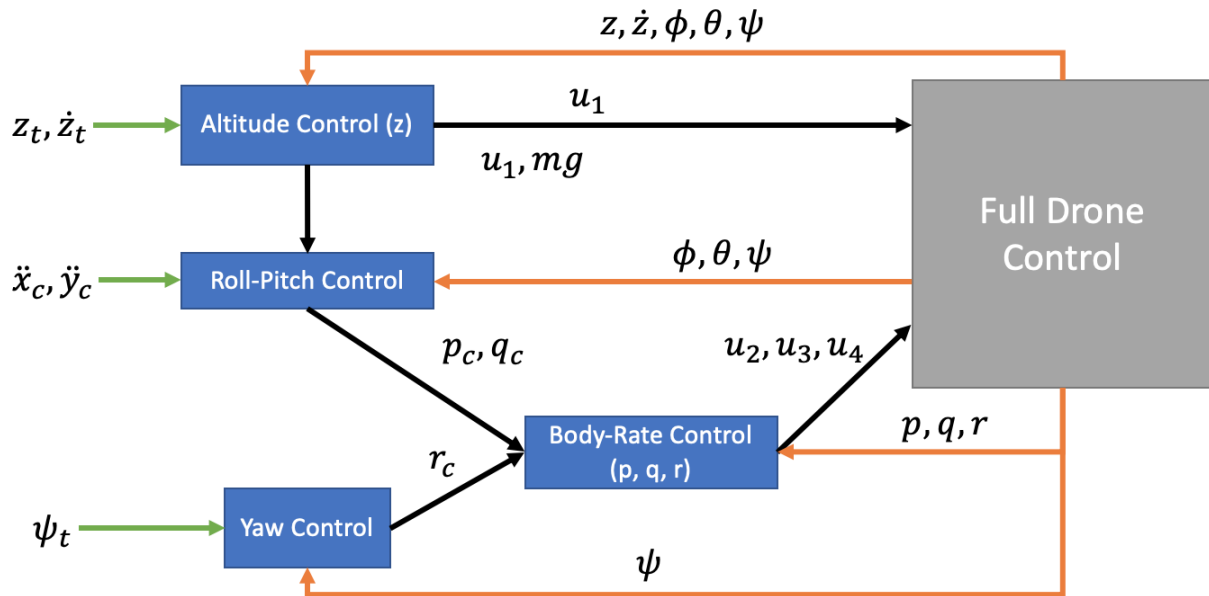
Path planning takes into account the positions of the following four gates as well as the coordinates of the ego-drone from the localization algorithm, then finds a three-dimensional spline that passes through them. The function then returns a series of “waypoints” which are intended to divide the extended flight path into smaller, more manageable chunks. For each of the waypoints, the 3d position and velocity vector are calculated. As soon as a “waypoint” has been passed by the drone, the path is re-calculated. In the event that fewer than four gates are left to completion, the path planner adds an interpolated gate position that generates an identical number of waypoints no matter how many gates remain to be passed. The velocity is derived as a linear relation to the approaching curvature of the line from one waypoint to the next. Here we define the curvature as a value bounded between 0 and 1, which is the ratio of the straight-line distance to the line integral when looking a short distance ahead. In other words, the straighter the approaching segment of the spline is, the faster the drone will proceed. The more curved the approaching segment is, the slower the drone will proceed. We found this to be a helpful setting so the drone does not crash while making close-cutting turns. Also the path planning provides 3 directional vectors that is tangent to the waypoint location on the spline, which then later controller uses to compute the yaw angle of the desired direction at the next waypoint.

Our process of drone control can be summed up in Figure 1. There are three separate controllers that are all working in concert to give directives to the drone. These include an altitude controller, a yaw controller, and a roll-pitch controller. The altitude controller takes the desired altitude and the z-direction velocity and position at a given time  $t$  as input from the localization algorithm and returns the desired thrust. The computed thrust incorporates the required force to counter gravity as well as the force to ascend. The yaw controller acts to return the angular rotation along the z-axis based on the desired output from the path planning. The roll-pitch controller takes the desired next waypoint, the current drone position and orientation and computes pitch rate and roll rate. The pitch rate is proportional to the desired pitch which is proportional to the optimal velocity computed by the planner. The roll, on the other hand, calculates the translational error from the current location and the planned path and corrects to that. The

body-rate controller takes each of these values and controls the drone accordingly. Together with the thrust calculation from the altitude controller, we have all the control parameters necessary to fully control all axes of motion that the drone can utilize.

A number of difficulties are encountered when trying to optimally control the drone. Each of the controllers is a proportional-integral-derivative (PID) controller, and therefore must be tuned at all three levels simultaneously for proper control to be achieved. There are still a number of issues that can arise with the controllers being ill-tuned, which can cause the drone to over-correct and crash. We ran out of time to apply PID tuning such as twiddle and for more complicated scenarios, but we did what we could to tune them as best as possible.

While there are certainly a lot of improvements that should be made, we believe we have the architecture in place to reliably adjust and significantly enhance our performance in the near future. **However, going forward with non-simulated environment, using PIDs and tuning them would become prohibitive to suit all unknown environment noise conditions such as wind. See Future Work section below for details on how we plan to tackle them.**



**Figure 1:** Drone control diagram

## 2. Future Work

For improved localization, we hope to also incorporate an advanced monocular/stereo visual odometry that takes into account pixelwise depth estimation as well which directly regresses the translational and orientational offsets per timestep. This combined with IMU based estimation and gate-location based estimation in the framework of EKF, we hope to achieve the ideal localization accuracy we need.

For path planning, we will either employ A\* or similar path planning algorithm to go around obstacles in the scene effectively with the lowest cost path. For heuristic measures, we will incorporate the curvature information so that the least curved path can be selected.

Finally, for controls, we are looking to apply reinforcement learning. Since the action space is infinite, rather than discretizing the action space and using Q learning variants, we will seek policy gradient methods such as PPO. With such algorithm, the controls will be trained to be robust against external forces such as wind drafts and perturbations or the presence of other drones. We will use the simulation environment for maximum training, but will mix the memories gathered from real environments as much as possible to train the policy optimally.