Summary

Imagine an unmanned, autonomous vehicle (UAV) deployed into a remote area. This UAV has a robot arm used to collect soil samples. The robot arm is positioned via a sequence of unit step inputs. A proportional-integration-derivative (PID) controller is used to produce smooth robot arm movements. The PID controller gains K_p, K_I and K_d were tuned prior to deployment. The UAV communicates with a rear area base station over a radio link. Suddenly a fault has occurred in the robotic arm. The exact nature of this fault is unknown, but the PID controller gains no longer produce acceptable arm movement. Fault recovery must be initiated to restore, as much as possible, the previous behavior. No redundant hardware is available so the only viable fault recovery method is to readjust the PID controller gains. Unfortunately, the UAV does not have the computational resources to compute new gain values. It does, however, have the capability of setting the gains to any value received from the base station over the radio link.

If a duplicate UAV was located at the base station then under ideal circumstances the fault could be duplicated at the base station and the new controller gains could be determined using say Zeigler-Nichols rules. Unfortunately the exact nature of the fault is unknown so this approach won't work. However, the base station can send the UAV candidate gain values and then direct the UAV to apply a unit step input and report response times, steady-state errors, and so forth. In other words, the base station works with the UAV to get observation data for trial PID gain values.

It is difficult to know exactly how to reconfigure a faulty system if the type of fault is unknown. In this research effort *Bayesian optimization* is used to reconfigure a PID controller. The Bayesian optimizer software will propose new candidate K_p, K_I and K_d gain values which will be implemented in the faulty UAV. A step input will produce a new step response which will be recorded and then reported via the radio link. The Bayesian optimizer will then propose new gain values. This iterative process continues until acceptable behavior in the physical system is restored.

To the best of our knowledge Bayesian approaches have been used for fault detection, but not fault recovery, which is the focus of this research effort. A commercial off-the-shelf ball & beam system will be used to validate our approach.

The research team will consist of the principal investigator and an under-graduate research assistant. Emphasis will be given to a student from a historically underrepresented minority group. All results, data collected and source code will be uploaded to GitHub making it available to the general public. A conference paper detailing our findings will be presented at an appropriate conference. The under-graduate student will be encouraged (and tutored) to present the paper.