Stealthy Attacker Complexity in Cyber-Physical Systems: Milestone 7

https://c-er.github.io/15400/

Uday Shankar <us@cmu.edu>

Major Changes: There seems to have been a slight adjustment to the goal of the project. Instead of focusing on finding and proving lower bounds for attacker complexity, we're more focused on the related problem of trying to determine if attackers with restricted capabilities can carry out successful attacks. Discussion of attacker complexity may still take place, but that is more of a stretch goal.

What I Have Accomplished So Far: Thus far, I've completed my initial background knowledge acquisition, which included the first three chapters of this book[1] as well as two papers[2;3]. I have continued by brainstorming some potential measures of stealthy attacker complexity. Some measures seem to be addressed by the framework from the original paper[2], but others seem to require some additional machinery. Currently, I am focused on formally defining the simpler notions of complexity and proving that they are handled by the framework from the original paper.

For one of the simpler notions of complexity that isn't handled by the original paper[2], I have a formally described potential solution. I have been working on the proof of correctness for the past few weeks, and it is nearly complete.

My advisor and I have also discussed methods of solving the problem we are looking at in another framework. In short, the method we are taking right now constructs a graph (different from the graph proposed in the original paper) that acts as a search space for attackers of limited capability. The alternative method uses the graph constructed in the original paper and alters the search process so it will only find attackers of limited capability. This is more in line with the original goal of measuring attacker complexity, but no concrete work has been done yet.

Meeting My Milestone: I think I am almost on track as far as milestones go. I am supposed to be part way between a formal write-up of any proofs and a final clean-up of all the results, and that sounds accurate.

Surprises: It is surprising to me how difficult it is to work on research in short periods of intense effort. I only seem to be able to make meaningful progress when I'm totally relaxed over a period of a few days.

Revisions to my 15-400 Milestones: As mentioned in previous reports, my first few milestones were modified in the sense that I started directly by brainstorming measures of complexity. Additionally, instead of looking for complexity lower bounds, I'm focusing on the related problem of determining whether an attacker of complexity at most C exists (under various different notions of complexity). My advisor and I are discussing the possibility of addressing the original complexity problem after the current construction is proved correct, but that will likely extend past the end of the semester.

Resources Needed: I've obtained software from the authors of the original paper[2] that I can extend to implement my constructions as needed.

References

[1] Cassandras, C. G., and Lafortune, S. *Introduction to Discrete Event Systems*, 2nd ed. Springer Publishing Company, Incorporated, 2010.

[2] Góes, R. M., Kang, E., Kwong, R., and Lafortune, S. Stealthy deception attacks for cyber-physical systems. *2017 IEEE 56th Annual Conference on Decision and Control (CDC)* (2017), 4224–4230.

[3] Yin, X., and Lafortune, S. A uniform approach for synthesizing property-enforcing supervisors for partially-observed discrete-event systems. *IEEE Transactions on Automatic Control 61* (2016), 2140–2154.