Intrusion Detection for Stochastic Task Allocation in Robot Swarms

Florian Maushart

Professors : Alcherio Martinoli (EPFL), Vijay Kumar (UPenn)
Assistants : Amanda Prorok (UPenn), Zeynab Talebpour (EPFL)

Swarm Robotic Systems are recently becoming a reality, after they had mostly been used in theoretical frameworks and simulations due to the high production cost of real robots in the past. Networks of autonomous taxis, quadcopters for autonomous aerial inspection of farms and industrial structures or large scale mobile sensor networks are only a few examples where Swarm Robotic Systems theory might be applied in the near future. While this seems very promising, it is important to understand how well we will be able to control and supervise the behavior of each individual robot when the swarm becomes very large. If the swarm can be disturbed by maliciously changing the behavior of some agents and we are not able to detect this intrusion early enough, the consequences could potentially be detrimental to the performance and security of the system. Interesting candidates for the analysis of such large-scale systems could be autonomous farms, as depicted in Fig. 1 below.

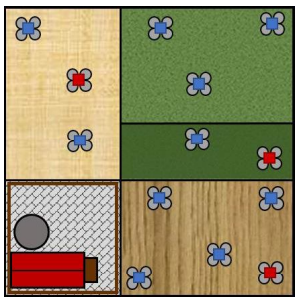

Fig. 1: A farming scenario in which malicious agents (red) disturb the optimal distribution of robots across all fields as they are not following our control policy.

In this thesis, we therefore present a novel framework for the integrity analysis of Swarm Robotic Systems using an approximation of the symmetric Kullback-Leibler Divergence. The objective is to understand a robot swarm's vulnerability to malicious intrusion and to develop the necessary computational tools that would detect the presence of malicious agents within the swarm. Using ensemble approaches for modeling and analyzing stochastic task allocation, we measure the performance of the proposed strategy subject to different system parameters, and show how different design choices can facilitate early intrusion detection.

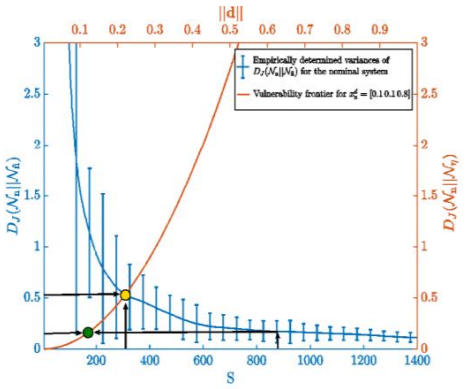

Fig 2: Microscopic detection model: For a sample size of $S \approx 300$ an attacker can achieve a disturbance $||d|| \approx 0.22$ before crossing the detection threshold (yellow point), while for $S \approx 900$ the limit is $||d|| \approx 0.08$ (green point).

We further evaluate the performance of our method in realistic scenarios through stochastic simulations for different team sizes and at different abstraction levels. The main contribution is an analysis framework whose output can be used to avoid system-inherent design flaws and to decrease the damage that can be inflicted by an undetected attacker. Key results of our analysis are that our ability to detect an attack increases quadratically with the fraction of malicious agents within the swarm and linearly with the total number of agents, while the uncertainty of the detection decreases linearly with the number of observations of the complete distribution of agents. Using our framework, we can also quantify the disturbance which malicious agents are able to inflict without revealing their presence (see Fig. 2). We further present a sub-microscopic application for our framework in the Webots robotic simulation environment. We can show that our framework generates useful results which correspond well to our predictions, even at the sub-microscopic level. This simulation could therefore serve as an interesting tool for further developments and extensions regarding this project.