Breaking Blue:Automated Red Teaming Using Evolvable Simulations

Stephen C. Upton[1], Sarah K. Johnson[2], and Mary L. McDonald[2]

[1] The MITRE Corporation, 2907 W Bay to Bay Blvd., Suite 303, Tampa, FL 33629
upton@mitre.org
[2] The MITRE Corporation, 2750 Killarney Drive, Suite 100, Woodbridge, VA 22192
{skjohn,mmcdonald}@mitre.org

Abstract. We continue to develop a complementary approach to manual Red Teaming we call AutoRedTeaming, which uses a combination of evolutionary algorithms and agent-based simulations to generate novel solutions that 'break Blue'. Information obtained during this automated process can then be used to either enhance or assist the manual effort. During the past year we have conducted the first test of the AutoRedTeaming concept using a scenario that concerns the defense of a fixed structure. We will discuss the overarching concept, results of our first experiments, the analysis of the data obtained, and describe our follow-on work to incorporate the concept of Evolvable Simulations.

1 Background

Red teaming is a technique that has been used successfully for some time in the military defense community to uncover system vulnerabilities or to find exploitable gaps in operational concepts. This technique has also recently been proposed as a tool for homeland security in evaluating the protective measures of key facilities (e.g. nuclear power plants, ports, harbors), events (e.g., Super Bowl), and national infrastructure (e.g., electric power grid, rail lines)[1]. Red teaming is currently a manually intensive technique that typically brings together experts relevant to the system under consideration and who are then charged with identifying weaknesses.

We have begun, and are continuing to develop, a supplemental approach we call AutoRedTeaming, whereby we automate this vulnerability discovery process using a combination of evolutionary algorithms and agent-based simulations. In essence, the approach is to use agent-based simulations to simulate proposed security procedures and then allow the Red (threat) agents to evolve capabilities, using a variety of evolutionary algorithms, over many millions of simulation runs with the goal of discovering means to thwart, evade, or otherwise exploit gaps in Blue's security procedures, hence 'breaking Blue'. Our approach has first focused on using simple, reactive agents whose behavior is controlled by a set of parameters. Their behavior, and the set of capabilities they possess, e.g., sensor and weapon systems, are evolved using an evolutionary algorithm. Information and analysis obtained during this automated process could then be used to enhance and assist the manual effort, with the overall goal of reducing surprise.

2 Initial Experiments

During the past year we have conducted the first tests of the AutoRedTeaming concept. We developed an agent-based scenario that concerns the defense of a fixed structure, ran one of the AutoRedTeaming algorithms against the scenario in order to quickly and effectively find ways to 'break Blue', and acquired a significant volume of data. Specifically, we used an Evolutionary Programming algorithm with a simple mutation operator and tournament selection (see, e.g., [2]). We used a population size of 80, 50 replications per individual in the population, using the mean across those replications as the final fitness, and ran the EP algorithm for 25 generations, resulting in 100,000 runs of our agent-based simulation. Each individual in a population represents a particular setting of a Red agent's parameters, with the Blue agent's parameters kept fixed. This might correlate to a specified number of Blue guards, a sensor configuration, or a security procedure, but something that is within Blue's ability to control. In the presentation, we will discuss additional details of the scenario, the algorithm settings, and some results from preliminary analysis of the data. We will also discuss other challenges with using this technique, such as generating solution diversity and the notion of simulation expressiveness.

3 Future Work

One of the challenges discovered during our work was the limited flexibility in generating solutions using only parameter settings of the agents. In order to generate better and more diverse solutions, we developed a concept we call Evolvable Simulations. This takes the AutoRedTeaming idea one step further by searching not just for simulation parameter settings that will break Blue, but also for changes in the simulation agent's structure, similar to the idea of Genetic Programming but applied to agent-based simulations. We believe this extension will allow for much more robust and powerful AutoRedTeaming. In addition to the Evolvable Simulation concept, we will discuss the potential use of other algorithms, such as artificial immune algorithms, as the generating engine.

4 Conclusion

This extensive simulation of security procedures and automated 'what-if' generation of potential threat tactics could be used as a first step of an enhanced process that includes high-resolution simulations, wargaming, and 'threat drills'. We envision applications in homeland defense, force protection, systems acquisition analysis, and course of action assessment.

References

1. Branscomb, L.M., Klausner, R.D.: Making the nation safer: the role of science and technology in countering terrorism. Committee on Science and Technology for Countering Terrorism, National Research Council (2002)

2. Chellapilla, K., Fogel, D.: Two new mutation operators for enhanced search and optimization in evolutionary programming. In Bosacchi, B., Bezdek, J., Fogel, D., eds.: Applications of Soft Computing. Volume 3165., Proc. SPIE (1997) 260–269