# Is This a Good Time? Deciding When to Launch Attacks on Process Control Systems

Marina Krotofil
Hamburg University of Technology
21079 Hamburg, Germany

Álvaro A. Cárdenas
University of Texas at Dallas
Richardson, TX75080, USA

## ABSTRACT

We introduce a new problem formulation for understanding the threats and vulnerabilities of process control systems. In particular, we consider an adversary that has compromised sensor or actuator signals of a control system and needs to identify the best time to launch an attack. We have previously shown that attackers might not be able to reach if the timing of their Denial-of-Service (DoS) attacks is not chosen strategically: Therefore, if the timing of a DoS attack is not chosen correctly, the attack can have limited impact; however, if the timing of the attack is chosen carefully, the attack has higher chances of succeeding. We formulate the problem of selecting a good time to launch DoS attacks as an optimal stopping problem that the adversary has to solve in real-time. In particular, we use the theory for the Best-Choice Problem to identify an optimal stopping criteria and then use a low pass filter to identify when the time series of a process variable has reached its peak. We identify some of the complexities associated with solving the problem and outline directions for future work.

## 1. INTRODUCTION

One of the growing research areas for securing cyber-physical systems is the work on threat models that considers an adversary which can manipulate sensor or actuator signals in order to drive the physical process under control to an undesired state. While most of the work in this area explores the implications of manipulating signals, there has been little work on understanding the complexity of launching successful attacks, and in particular, on finding a "good time" to launch an attack. In our study we consider an attacker that can read sensor signals for a given process variable, and then has to select a time to launch a DoS attack on the communication channel, which in turn will freeze the process value in the controller's memory [3]; and therefore the controller will select control commands based on a value that is not being updated.

We formulate the problem as an optimal stopping time problem for the attacker. In particular, we formulate the problem as a Best-Choice Problem (also known as the "Secretary Problem"), in which the adversary is presented with a time series of system states (obtained by sensors) and has to select the optimal time to attack based on these measurements.

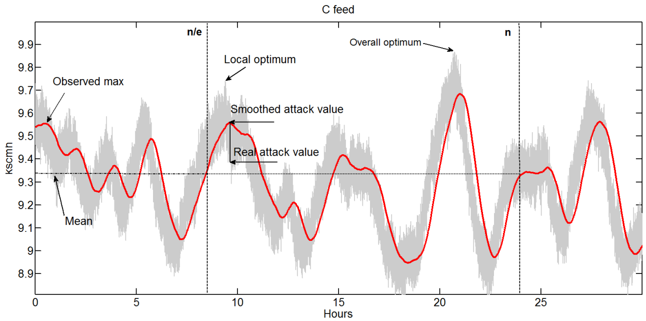

Figure 1: The learning phase $\frac{n}{e}$ of the Best-Choice problem identifies a peak that is used as a reference point to find a better peak between $\frac{n}{e}$ and $n$. Because the signal is noisy it is difficult to identify a peak in real-time–a downward trend might be a noisy artifact masking a general upward trend–therefore we use a low-pass filter so we can identify a peak as soon as the signal has a downward trend. If we do not find a peak higher than the learning period before reaching $n$ we call this a "non-selection."

Optimal stopping problems are studied in statistics and they are generally concerned with the problem of choosing a time to take a particular action in order to maximize an expected reward. Recall that the attacker faces the following problem: given a time-series of $N$ sensor signal samples that will exhibit a sequence of peaks, how should the attacker select one of peaks to launch her attack? This particular formulation in the context of the Secretary Problem was solved by Freeman [2], which showed that one can select the best signal sample with maximum probability $(1/e)$ by using the following rule: do not select any sample among the first $(N/e)$ candidates (the learning period), and after that, select the first sample whose value exceeds the value of all candidates seen so far, or the last sample in the series.

There are multiple variations or formulations to the secretary problem. In this paper we consider the classical solution, and a recent result that assumes the order in which the candidates arrive is not completely random, but has a probability distribution satisfying a hazard rate condition [4]. This assumption is commonly used in standard engineering applications. Gaussian, uniform, and exponential distributions satisfy this property. Under this assumption it can be shown that the learning period can be cut down to $N/log(N)$; which is significantly smaller than $N/e$.

Because the Best-Choice Problem formulation assumes non-correlated time measurements, the attacker has to add an additional stopping criteria to identify when the sensor signal has reached its peak; this is a non-trivial task in many

*HiCoNS'14,* April 15–17, 2014, Berlin, Germany.
ACM 978-1-4503-2652-0/14/04.

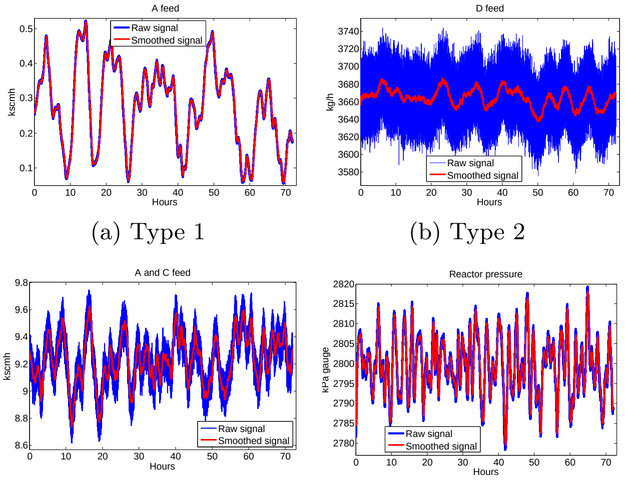

(a) Type 1 (b) Type 2

(c) Type 3 (d) Type 4

Figure 2: Different types of sensor signals with their correspondent smoothed signals, $m$=250

practical situations as sensor measurements tend to be noisy and have sudden fluctuations. To this end we explore two methods: signal smoothing with low pass filters (completed work to appear in the 2014 International Conference on Critical Infrastructure Protection) and quickest change detection theory (future work). A summary of our approach can be seen in Fig. 1.

## 2. EXPERIMENTAL RESULTS

For our experimental analysis we use the C and Matlab model of the full Tennessee Eastman (TE) challenge process developed by Ricker [5]. The TE process is a modified model of a real plant-wide industrial process [1] which was created to equip the process control community with a realistic and accurate model for testing process control technologies. We have extended this line of work to test and evaluate the performance of the system under adversarial attacks [3].

Although the TE model provides a solid platform for process and control engineering experimentation, we had to extend the system to incorporate two new requirements. First, we modified the original process code to introduce randomness into the sensor and actuator signals noise, without disturbing the underlying dynamic behavior of the process. This enabled statistical evaluation of the experiments. Second, due to low sampling frequency of the final simulated data available in Matlab (100 samples per hour), we implemented in C a way to export the exact process simulation data into the Matlab workspace, which are a sampled at a rate 2000 samples per hour and thus are much closer to real-world scenarios.

As mentioned before, in addition to using the Best-Choice theory to identify a peak, we use a low pass filter to remove short-term fluctuations or "noise" and highlight long-term trends of the data. The simplest form of smoothing is the "moving average" which is the mean of the previous $m$ samples. In TE process, sensor signals can be roughly divided into 4 groups (Fig. 2). Type 1 is characterized by few distinct peaks and low noise level. Type 2 is distinguished by the slow signal amplitude change with high frequency noise. Type 3 can be described as a very noisy variation of Type 1 signal. Type 4 signal distinguishes itself with the multiple noisy signal peaks. Our experimental results show that to succeed, the attacker has to take into account type of the signal she is dealing with.

We evaluate the performance of our approach based on three metrics: (1) fractional error in identifying the peak, in %; (2) fractional error in selecting the highest possible peak in the series, in %; and (3) number of non-selections (when no peak is higher than the one observed in the learning phase, and last sample in the series is taken). Our results confirm that the learning period can be indeed cut down to $N/log(N)$ while achieving results comparable to the the $N/e$ strategy; in particular, because if the learning phase is short, the number of non-selections reduces substantially (almost to zero). For the same reason the fractional error in selecting the highest possible peak increases as the attacker has less time to achieve sufficient aspiration level. Since for the classic solution the number of non-selections reaches on average 25%, it can be a decisive factor to favor $N/log(N)$ strategy.

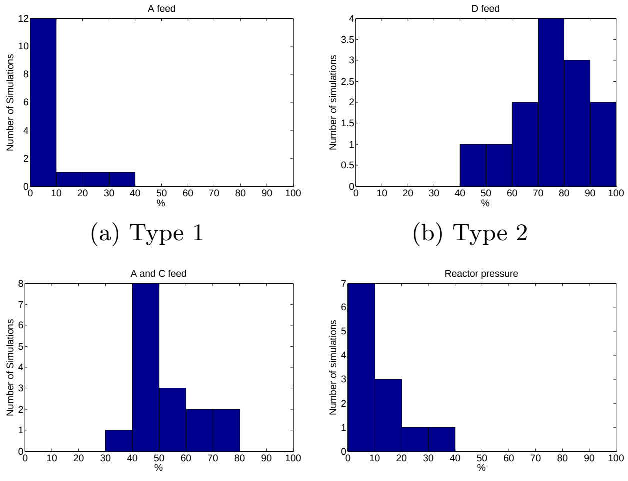

(a) Type 1 (b) Type 2

(c) Type 3 (d) Type 4

Figure 3: Distribution of the fractional error in selecting the highest possible value, $m$ = 250

Because of space considerations, we illustrate the performance of our approach in a form of the histogram of the distribution of the fractional error in selecting the highest possible value in the time series (Fig. 3). As can be seen, the best results can be achieved for sensor signals of Type 1 and 4. In contrast the methodology proposed in this paper is not well suited for conducting attacks on the sensor signals of Type 2 and 3 because of their noise. While applying a low-pass filter delivers good results for the attacks on the low-noise signals, an alternative solution is required for dealing with the noisy process variables.

## 3. REFERENCES

[1] J. J. Downs and E. F. Vogel. A plant-wide industrial process control problem. *Computers & Chemical Engineering*, 17(3):245–255, 1993.

[2] P. Freeman. The secretary problem and its extensions: A review. *International Statistical Review/Revue Internationale de Statistique*, pages 189–206, 1983.

[3] M. Krotofil and A. A. Cárdenas. Resilience of Process Control Systems to Ccyber-Pysical Attacks. In *NordSec'14, Secure IT-Systems*, volume 8208 of *LNCS*, pages 166–182, 2013.

[4] M. Mahdian, R. P. McAfee, and D. Pennock. The secretary problem with a hazard rate condition. In *Internet and Network Economics*, pages 708–715. 2008.

[5] N. L. Ricker. Tennessee Eastman Challenge Archive. `http://depts.washington.edu/control/LARRY/TE/download.html`. retrieved: May, 2013.