
IDENTIFYING VULNERABILITIES OF INDUSTRIAL CONTROL SYSTEMS USING EVOLUTIONARY MULTI-OBJECTIVE OPTIMISATION

A PREPRINT

Nilufer Tuptuk

Department of Computer Science
University College London
London, United Kingdom
n.tuptuk@ucl.ac.uk

Stephen Hailes

Department of Computer Science
University College London
London, United Kingdom
s.hailes@ucl.ac.uk

ABSTRACT

In this paper we propose a novel methodology to assist in identifying vulnerabilities in a real-world complex heterogeneous industrial control systems (ICS) using two evolutionary multi-objective optimisation (EMO) algorithms, NSGA-II and SPEA2. Our approach is evaluated on a well known benchmark chemical plant simulator, the Tennessee Eastman (TE) process model. We identified vulnerabilities in individual components of the TE model and then made use of these to generate combinatorial attacks to damage the safety of the system, and to cause economic loss. Results were compared against random attacks, and the performance of the EMO algorithms were evaluated using hypervolume, spread and inverted generational distance (IGD) metrics. A defence against these attacks in the form of a novel intrusion detection system was developed, using a number of machine learning algorithms. Designed approach was further tested against the developed detection methods. Results demonstrate that EMO algorithms are a promising tool in the identification of the most vulnerable components of ICS, and weaknesses of any existing detection systems in place to protect the system. The proposed approach can be used by control and security engineers to design security aware control, and test the effectiveness of security mechanisms, both during design, and later during system operation.

Keywords Evolutionary multi-objective optimization · Vulnerability assessment · Security management · Industrial control systems · Cyber-physical systems

1 Introduction

Industrial Control Systems (ICS) are command and control systems that are found at the core of the national critical infrastructure services such as gas; electricity; oil; water supply; telecommunication; transportation; process manufacturing (chemicals, pharmaceuticals, paper, food, beverages and other batched-based manufacturers); and discrete manufacturing (automobiles, ships, computers and many other durable goods). The security of ICS is of critical importance in industrialised economies: they are so pervasive that national security, public health and safety, and economic growth all rely on their correct operation. In the past, security of ICS was achieved simply through isolation and control of physical access. However, ICS are making increasing use of network technologies, commercial-off-the-shelf (COTS) components, and wireless systems because of factors such as low-cost, increased sensing and communication capacity and convenience.

With these technological advances, factory and plant networks are becoming highly connected over multiple layers, and signals sent between control components (i.e. sensors, controllers, and actuators), are sent through a common network, known as networked control system [1]. The number of motivated and highly skilled adversaries carrying out complex attacks against networked control system is on the increase. Some of the past attacks include the attack against the operational systems of Evraz Steel in North America [2]; attack on Ukraine's power grid [3] that targeted the electric transmission system in Kiev; attack against a German steel mill [4] that caused unspecified but "massive" physical damage; malware attacks such as Duqu [5] and Havex [6] that targeted ICS for industrial espionage; and Stuxnet [7] that targeted Iran's Natanz nuclear plant, and destroyed centrifuges installed at the time of the attack. As the evidence from these attacks show the potential outcome of a successful attack on a critical service ranges from injuries and

fatalities, through serious damage to the environment, to catastrophic nation-wide economic loss due to production losses or degradation of products and services. Shutting down or preventing access to these systems, even for a short time, may cause significant harm to people and may impact public confidence, leading to a general feeling of insecurity.

Despite technological advances in ICS, understanding the vulnerabilities of ICS at the process level and the potential physical consequences when these vulnerabilities are exploited remains an important aspect of ICS security that is currently lacking research. It is critical to national economic resilience that we explore better ways of searching for vulnerabilities in the industrial processes and defence mechanisms, and understand the impact of these vulnerabilities would be if they were to be exploited. Considerable research [8, 9, 10, 11, 12, 13, 14] has focused on developing threat models and analysing a variety of attacks, aimed at modifying process measurements (sensor measurements) or manipulated variables (values going from controller to actuators), or manipulating the control algorithm (e.g. set points), however there is a lack of research work on investigating impact of combinatorial attacks, and automating attack generation.

In this paper, we seek to identify the most vulnerable components in the ICS by searching for process level attacks that cause the greatest damage in terms of plant safety and economics, using the least effort, as well as identifying weaknesses in detection methods, at the lowest risk of being caught. To solve this, we establish the problem as a multiobjective optimisation problem, and investigated effectiveness of EMO algorithms as a tool for generating attacks.

Multiobjective optimisation has been widely applied to real world complex scientific and engineering problems, but, we are not aware of any studies using EMO algorithms to identify the most vulnerable components of cyber-physical systems, including ICS. However, evolutionary algorithms have been extensively used to improve the security mechanisms used especially within IT networks in a wide variety of ways. Others have used genetic algorithm (GA) [15, 16, 17, 18, 19, 20, 21], and genetic programming (GP) [22, 23] to evolve new rules to detect new forms of network intrusion. Our previous work [24] [25] used GP to attack a wireless sensor network (WSN) protected by an artificial immune intrusion detection system. Kayacik et al. [26] used GP to evolved variants of buffer overflow attacks against an open-source signature-based IDS. John et al. [27] applied GA to the improvement of moving target defence, in which the defence changes the system's attack surface to disrupt the intelligence gathered by the attacker. Dewri et al. [28] used GA with multiobjective optimisation to investigate optimal security measures for a system. Garcie et al. [29], Hemberg et al. [30], Rush et al. [31] used co-evolution to model attacker and defence dynamics for network security. Co-evolutionary concepts have also been investigated to prevent faults and cascading blackouts in electric power transmission systems [32, 33]; automate red teaming for military scenarios [34]; and improve the performance of malware detectors [35].

Our work complements the existing studies in applying evolutionary multiobjective optimisation to identify ICS vulnerabilities by generating attacks. We investigate both the worst-case condition where ICS has no security protection, and also where there are some measures against attacks, in the form of a novel intrusion detection system.

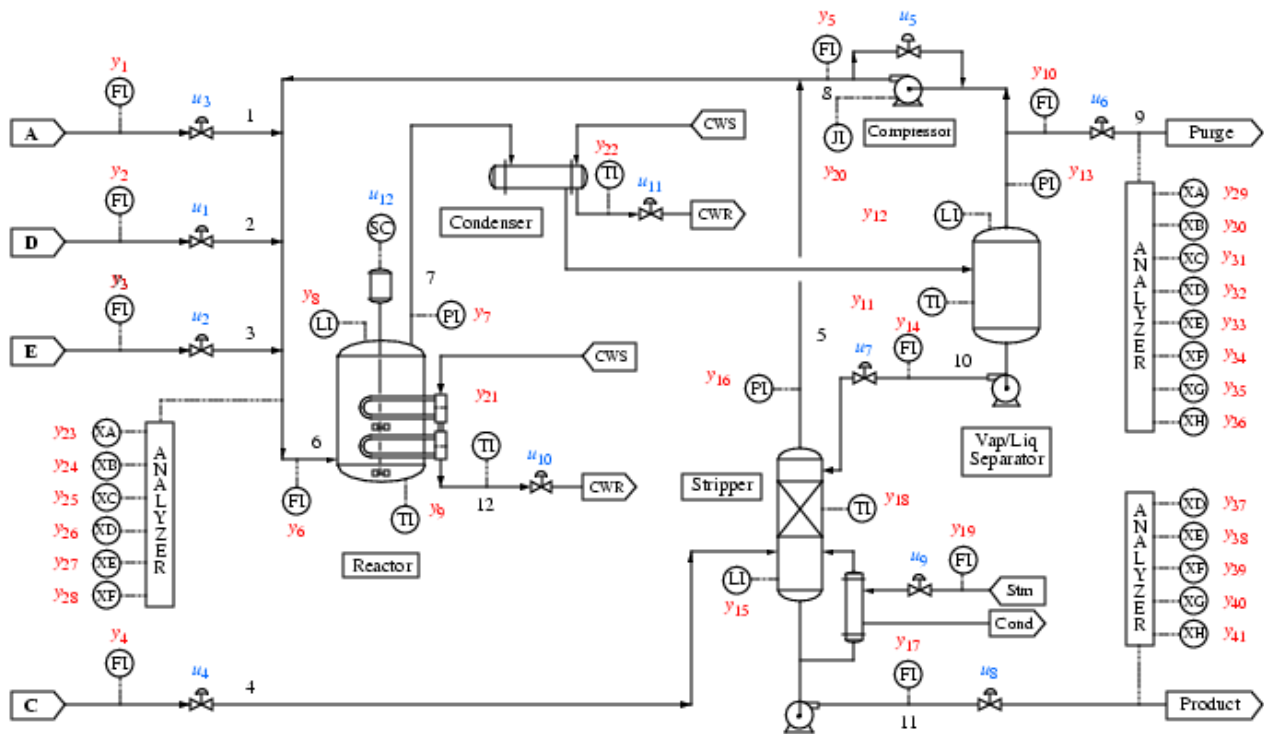
The remainder of the paper is organised as follows. Section 2 presents the background material related to our approach. This includes the characteristics of the the case study, the Tennessee Eastman (TE) process; description of the multiobjective optimisation problems; and attack model used for generating attacks. Section 3 covers the baselines to compare EMO approach; details of the EMO algorithms; the detection methods used to evolve attacks against detection; the performance metrics used for comparing EMO algorithms; and describes the random approach used to compare EMO approach. Section 4 presents experimental results. In section 5, application of results are discussed. Conclusion and future work are presented in Section 6.

2 Background

2.1 Case study: The Tennessee Eastman (TE) process


To develop and explore the effectiveness of evolutionary multiobjective optimisation as an possible candidate for identifying vulnerabilities in ICS, a well-known chemical plant, the Tennessee Eastman (TE) process control model [36] was selected.

Our reasons for selecting this process are: i) it is a well-known model that has been widely studied; ii) it is a complex, highly non-linear system with a number of components that reflects a real process; iii) safety and economic viability can be quantified; iv) the code and model is available, and have been revised and validated over the years; v) it continues to be a relevant model for both the control, and more recently, the security communities. We are not aware of any other open source model that has these properties. The TE model is based on a real chemical process; however, the identity of the reactants and products are hidden to maintain commercial confidentiality. The process has eight components: four gaseous reactants (A, C, D, E), two products (G, H), an inert component (B) and a by-product (F). These reactions are [36]:



PI: Proportional-integral, **FI:** Flow Indicator, **TI:** Temperature Indicator, **LI:** Level Indicator

Figure 1: Tennessee Eastman challenge model [36, 37]



The process is illustrated in Figure 1, and consists of five major components [36]: a reactor, a product condenser, a vapour-liquid separator, a recycle compressor and a product stripper.

There are 41 process measurement variables known as XMEASs (sensors, denoted as y signals in red) and 12 manipulated variables known as XMV (valves/actuators, denoted as u signals in blue), illustrated in Figure 1, that are involved in controlling and monitoring the plant. The main control objectives of the plant are [36]: to maintain the process variables at the desired values; to ensure that the operational conditions are within the equipment constraints; and to minimise variability of the product rate and product quality during disturbances. To protect the safety of the process, the TE model has a set of operating constraints that are captured as normal operating limits and shutdown operating limits. If the process reaches the shutdown limits, it automatically shuts the plant down. These constraints (low/high limits of reactor pressure, reactor level, reactor temperature, product separator level and stripper base level) [36] are put in place to protect the personnel, equipment, production, and meet compliance requirements. The operating costs of the TE process are calculated according to the following equation [36]:

$$\text{total costs} = (\text{purge costs})(\text{purge rate}) + (\text{product stream costs})(\text{product rate}) + (\text{compressor costs})(\text{compressor work}) + (\text{steam costs})(\text{steam rate})$$

The TE process problem makes no recommendation as to what needs to be controlled, and leaves the selection of controlled variables and control strategies to the control engineers. Most proposed solutions do not control all the variables. The control strategy used in this paper is that described by Larsson et al. in [37] using 16 process measurements and 9 manipulated variables. The process model used in this paper is developed by Ricker, available

from his home page [38]. The code is implemented in C, with a MATLAB/Simulink interface via an S-function implementation. Isakov and Krotofil [39] extended the original Simulink model by enhancing it with Simulink blocks that enable integrity and denial-of-service (DoS) attacks to be carried out on the sensors and manipulated variables. We extended their model with replay attacks, and made further small changes needed to carry out the work in this paper.

We first analyse the vulnerabilities of the system by considering two types of adversary: i) an adversary targeting the safety of the plant by attempting to shut it down using the least effort; ii) an adversary targeting the operating cost of the plant to increase economic loss using the least effort. The effort of the attack is calculated as the number of sensors and actuators that must be compromised.

We, then, use machine learning techniques to detect attacks, and evolve new attacks against the detection methods. We assume an adversary who has control over her attack vector and who knows the feature space used by the detection system. However, she does not know the details of the underlying detection methods. So, essentially, the detection is a blackbox that can be queried. An adversary queries the detection with the attack vector, and obtains a detection probability. Her goal is to find attacks that cause damage whilst evading detection, and using the least effort. Thus, the problem is formulated as a search problem with objectives.

2.2 Multiobjective optimisation

In many real-world problems, decisions need to be made on the basis of multiple competing or conflicting objectives and constraints. This is often the case when making decisions related to cyber security investment or security hardening, attempting to balance risks of attack against a limited budget to buy defensive countermeasures. In such situations, formulating the problem as a multiobjective optimisation (MOO) with multiple choices can help to determine the trade-offs among the objectives in a more effective manner [40, 41, 28]. These approaches search for the set of *non-dominated* or *Pareto-optimal* solutions. A solution is defined non-dominated if there are no other solutions that would improve any objective without degrading one or more of the other objectives. Once the set of Pareto-optimal solutions, has been identified, a decision maker can make a decision by examining the tradeoffs represented by individual solutions within the set. MOO takes a problem with multiple objectives and simultaneously seeks to optimise all objectives, providing solutions in, or close to, the true Pareto-optimal set. More often than not, this is an estimate because determining the true Pareto-optimal set for real world problems is hard, either because the search space is too large or because obtaining solutions is costly in time and computation.

In this piece of work, we are concerned with finding optimised solutions for the following multiobjective problems:

1. **Attack safety of the plant (shutting down the plant):** Minimise plant operating time of the plant, minimise effort required to carry out attacks.
2. **Increase operating cost (economic loss attack):** Maximise operating cost of the plant, minimise effort required to carry out attacks.
3. **Increase economic loss and avoid detection:** Maximise operating cost of the plant, minimise detection (alarm) probability, and minimise effort required to carry out attacks.

The generation of optimal attacks against real systems with large number of components involves selection of many parameters: attack targets (controllers, sensors, actuators); attack types; and attack parameters for these attacks (e.g. mode of attacks, attack start times and attack duration). As explained earlier, in practice identifying the true Pareto-optimal set to such problems may not be feasible. Evolutionary multiobjective optimisation is a promising approach to identify an estimate of best trade-off attacks in such complex systems at reasonable computational cost.

2.3 Attack modelling

Figure 2 shows our underlying threat model that is based on common attacks against networked systems. The adversary is capable of intercepting communication from the sensor to controller (process measurements), and controller to actuator (manipulated variables). The attacks we consider are categorised as DoS, integrity (man-in-the-middle) and replay attacks. We will investigate the impact of these attacks in terms of measuring the impact on safety and operating cost of the plant. In the following section, we briefly discuss what this means, and explain how the attacks are modelled.

Past studies have investigated the safety and economic impact of DoS and man-in-the-middle attacks on the TE model [9, 8, 10]. We are following their attack model; however, our focus is to generate optimised combinatorial attacks. We extend their analysis by undertaking a more comprehensive search and examining the possibility of forming combinations of attacks. The attack parameters are as follows:

Attack Targets: The control strategy selected for our investigation, Larsson et al. in [37], uses 16 XMEAS variables, and 9 XMV variables. An adversary may attempt to manipulate signals that are sent from XMEASs to controllers




Figure 2: ICS attack model against the networked control system

(process measurements), and/or from controller to the XMVs (manipulated variables). The process run time used in this study is 72 hours. An attack begins at time t_s and ends at t_e , it can start any time, between start of the plant and the end, $t \in [0 - 72]$. The manipulated control (XMV) signal $y_i^a(t)$ and process measurement (XMEAS) $u_i^a(t)$ are as follows:

$$\begin{aligned} & y_i^{(t)}, \text{ for } t \notin I_a \\ & \hat{y}_i^{(t)}, \text{ for } t \in I_a \end{aligned} \quad (1)$$

$$u_i^a(t) = \begin{cases} u_i^{(t)}, \text{ for } t \notin I_a \\ \hat{u}_i^{(t)}, \text{ for } t \in I_a \end{cases} \quad (2)$$

where y_i are the $y_i^a(t)$ and $\hat{u}_i^{(t)}$ are the modified values the adversary sends.

To investigate the impact of attacks, we considered three types of attacks that could be used: *DoS*, *integrity* and *replay* attacks.

A **DoS attack**, is an interruption attack in which a signal is not received by its intended destination. For example, let y_i be the output of sensor i at time t , and let u_i be the output from the controller to actuator signal i at time t . When the attack occurs, the response strategy for the controller or the actuator is to use the last received value as the current reading:

$$\begin{aligned} \hat{y}_i^{(t)} &= y_i(t_{s-1}) \\ \hat{u}_i^{(t)} &= u_i(t_{s-1}) \end{aligned} \quad (3)$$

where $\hat{y}_i^{(t)}$ and $\hat{u}_i^{(t)}$ are the modified values the adversary sends.

An **integrity attack** involves an attacker manipulating signals by changing their values. A naïve attacker may listen to the transmitted values and modify them so that they are still within the ranges of possible plant values, since this has the potential to cause some damage. One way to achieve this is to try to modify the sensor measurements (y_i) and manipulated variables (u_i) using observed upper minimum ($integrity_{min}$) and lower maximum ($integrity_{max}$) values:

An $integrity_{min}$ is where the actual output of the sensor or controller signal i at time t is replaced with a minimum value:

$$\begin{aligned} \hat{y}_i^{(t)} &= \min_{t \in T} (y_i(t)) \\ \hat{u}_i^{(t)} &= \min_{t \in T} (u_i(t)) \end{aligned} \quad (4)$$

An $integrity_{max}$ is where the actual output of the sensor or controller signal i at time t is replaced with a maximum value:

$$\begin{aligned}\hat{y}_i^{(t)} &= \max_{t \in T} (y_i(t)) \\ \hat{u}_i^{(t)} &= \max_{t \in T} (u_i(t))\end{aligned}\quad (5)$$

A **replay attack** involves forging sensor measurements or manipulated variables as in the integrity attack but, this time, it repeatedly replays legitimate data it observed earlier:

$$\begin{aligned}y_i^r &= [y_i^{(r_{start})}, \dots, y_i^{(r_{end})}] \\ u_i^r &= [u_i^{(r_{start})}, \dots, u_i^{(r_{end})}] \\ \hat{y}_i^{(t)} &= \hat{y}_i^r [t \bmod \text{len } y_i^r] \\ \hat{u}_i^{(t)} &= \hat{u}_i^r [t \bmod \text{len } u_i^r]\end{aligned}\quad (6)$$

where y_i^r and u_i^r are the signals recorded by the adversary from the replay period, r_{start} to r_{end} .

3 Methodology

3.1 Baselines

To provide a baseline cost for normal operation of TE plant, 1000 independent runs were executed without disturbances. For this, and all subsequent runs, the plant was operated in Mode 1, the most commonly used configuration in the literature [42]. The following plant operating costs were obtained:

Plant operation	Max (\$)	Mean (\$)
Normal (without disturbances)	8,218	8,208

To inform later $integrity_{min}$ and $integrity_{max}$ attacks, the minimum and maximum values of the XMEAS signals and XMV signals observed under normal operating conditions were recorded.

3.1.1 Single random attacks

To provide a baseline against which to compare the impact of evolved attacks, 500 random attacks were launched on each of the XMEAS and XMV signals per attack type. Attacks were started at hour 2 with the intention of running them for the remainder of the simulation time (i.e. 70 hours). Table 1 and Table 2 shows the impact of each attack in terms of the operating cost of the plant and safety (fastest shutdown time). The variable name and range column describes the name of the sensor or actuator, and the observed minimum and maximum values, which are also used to devise the integrity attacks. The *shutdowns* column denotes the total number of times the plant shut down as a consequence of the attack out of 500 runs. The *shutdown range* column indicates the time at which the plant shuts down after the attack was started. Our experiments show that the fastest attack that can bring the plant down are an $integrity_{min}$ attack on A and C feed flow (XMEAS 4), requiring an attack to last for a duration of 0.52-0.65 hours (31.2-39 minutes), and an $integrity_{max}$ attack on reactor temperature (XMEAS 9), resulting in a shut down at between 0.59-0.63 hours (35.4-37.8 minutes). In these cases, the controller receives fake values from XMEAS 4 and XMEAS 9, which are lower for XMEAS 4 and higher for XMEAS 9 than the legitimate values, and calculates the control values that are sent to the actuators using these fake values. This behaviour causes a significant increase in the reactor pressure, and the plant shuts down as a result.

As reported in Table 1, the experiments carried out showed the following single attacks against process measurements (XMEAS) signals increased the operating cost of the plant significantly: i) $integrity_{max}$ attack on reactor pressure (XMEAS 7) increased the operating cost to \$24,507; ii) $integrity_{max}$ attack on the sensor measuring component C in purge (XMEAS 31) increased the operating cost to \$20,515; iii) DoS attack on reactor pressure (XMEAS 7) increased the operating cost to \$24,299; iv) replay attack on recycle flow (XMEAS 5) increased the operating cost to \$22,429; and v) DoS attack on C in purge (XMEAS 31) increased the operating cost to \$17,205.

Table 2 shows the impact of attacking the manipulated variables issued by the controller to actuators (X MVs). The attack that caused the fastest damage is the $integrity_{min}$ attack on the condenser cooling water flow (X MV 11), resulting in a shutdown time of 0.64 hours (38.4 minutes) due to low separator liquid level. Carrying out a $integrity_{max}$ attack on reactor cooling water flow (X MV 10) is able to shut down the plant in 1.65 hours (99 minutes) due to high reactor pressure. Integrity attacks on A and C feed flow (X MV 4), purge valve (X MV 6) and reactor cooling water flow valve

Variable Number	Variable Name and Range	Attack	Max Cost	Mean Cost	Shutdowns	Shutdown Range(hrs)
XMEAS 1	A-Feed (stream 1) 0.25-0.27 kscmh	Max	8449	8407	0	-
		Min	8364	8260	0	-
		DoS	8447	8331	0	-
		Replay	8429	8316	0	-
XMEAS 2	D Feed (stream 2) 3579.20-3744.46 kg ^h ⁻¹	Max	1158	1152	500	3.43-3.48
		Min	915	902	500	4.31-4.4
		DoS	3149	1761	500	6.9-21.86
		Replay	3897	2765	500	13.64-30.42
XMEAS 3	E Feed (stream 3) 4339.06-4536.27 kg ^h ⁻¹	Max	739	730	500	2.66-2.72
		Min	1320	1312	500	4.17-4.22
		DoS	2480	1494	500	3.57-18.16
		Replay	3350	2330	500	11.65-22.9
XMEAS 4	A and C Feed (stream 4) 8.98 9.48 9.24 kscmh	Max	384	378	500	1.25-1.34
		Min	392	361	500	0.52-0.65
		DoS	1318	743	500	1.38-6.48
		Replay	1227	825	500	2.32-5.99
XMEAS 5	Recycle flow (stream 8) 31.32-33.13 kscmh	Max	2099	2040	500	8.15-8.42
		Min	3456	3415	500	10.58-10.78
		DoS	21942	10675	322	10.4-69.7
		Replay	22429	9169	3	64.67-64.67
XMEAS 7	Reactor pressure 2793.54-2806.12 kPa	Max	24507	24468	0	-
		Min	671	650	500	8.37-8.78
		DoS	24299	10430	254	9.01-60.3
		Replay	23889	10894	233	9.73-66.15
XMEAS 8	Reactor level 62.77-67.24 %	Max	381	375	500	2.81-2.87
		Min	1017	1008	500	2.83-2.89
		DoS	7435	1989	494	3.77-35.79
		Replay	11302	5058	418	15.67-67.34
XMEAS 9	Reactor temperature 122.85-122.95 °C	Max	364	356	500	0.59-0.63
		Min	377	366	500	1.23-1.28
		DoS	10464	1554	489	0.92-61.64
		Replay	10774	3681	467	2.35-67.04
XMEAS 10	Purge rate (stream 9) 0.1545-0.2689 kscmh	Max	8228	8195	0	-
		Min	8246	8218	0	-
		DoS	8235	8201	0	-
		Replay	8236	8201	0	-
XMEAS 11	Product separator temperature 91.46-92.12 °C	Max	10427	10364	0	-
		Min	10444	10358	0	-
		DoS	10386	10270	0	-
		Replay	10393	10264	0	-
XMEAS 12	Product separator level 45.28-54.74 mol %	Max	896	888	500	5.85-5.95
		Min	1256	1245	500	8.57-8.68
		DoS	8210	3816	463	8.38-66.43
		Replay	8217	7802	143	39.66-69.96
XMEAS 14	Product separator underflow 51.64-55.89 %	Max	1370	1357	500	9.09-9.79
		Min	1449	1385	500	9.71-10.41
		DoS	4038	2274	500	11.07-32.98
		Replay	4431	3004	500	19.17-36.51
XMEAS 15	Stripper level 45.29-54.57 %	Max	1756	1740	500	13.48-13.71
		Min	1804	1793	500	13.31-13.54
		DoS	8234	5368	413	19.77-69.61
		Replay	8234	8181	16	57.03-68.13
XMEAS 17	Stripper underflow (stream 11) 22.37-23.41 m ³ h ⁻¹	Max	312	305	500	1.02-1.03
		Min	387	379	500	1.01-1.02
		DoS	4512	2298	500	6.68-37.75
		Replay	5716	4312	500	28.45-48.24
XMEAS 31	C in Purge 11.87-14.22 mol%	Max	20515	20438	500	65.65-66.1
		Min	13493	13455	500	50.02-50.4
		DoS	17205	9841	0	-
		Replay	10951	8818	0	-
XMEAS 40	G in product 51.64-55.89 mol%	Max	8312	8298	0	-
		Min	8117	8106	0	-
		DoS	8267	8208	0	-
		Replay	8268	8212	0	-

Table 1: Impact of random attacks on XMEAS variables (500 runs for each attack)

Variable Number	Variable Name and Range	Attack	Max Cost	Mean Cost	Shutdowns	Shutdown Range(hrs)
XMV 1	D feed flow (stream 2) 62.89-63.12 kgh ⁻¹	Max	8209	8198	0	-
		Min	8226	8217	0	-
		DoS	8216	8204	0	-
		Replay	9347	8223	0	-
XMV 2	E Feed (stream 3) 52.99-53.24 kgh ⁻¹	Max	8234	8221	0	-
		Min	8204	8194	0	-
		DoS	8215	8204	0	-
		Replay	8216	8203	0	-
XMV 3	A Feed (stream 1) 25.12-27.045 kscmh	Max	8352	8342	0	-
		Min	8259	8248	0	-
		DoS	8277	8216	0	-
		Replay	8247	8210	0	-
XMV 4	A and C Feed (stream 4) 59.93-61.32	Max	10101	10085	0	-
		Min	9339	9290	500	38.49-39.08
		DoS	14972	8433	2	68.09-68.09
		Replay	8873	8248	0	-
XMV 6	Purge valve (stream 9) 19.39-32.76 %	Max	9552	9542	0	-
		Min	7223	7215	0	-
		DoS	9064	8221	0	-
		Replay	8876	8234	0	-
XMV 7	Separator pot liquid flow (stream 10) 37.21-37.46 m ³ h ⁻¹	Max	2382	2313	500	17.65-18.92
		Min	3355	3273	500	25.83-27.35
		DoS	8217	7578	131	26.01-69.42
		Replay	8217	7975	82	44.05-68.92
XMV 8	Stripper liquid product flow (stream 11) 46.36-46.55 m ³ h ⁻¹	Max	2014	1962	500	15.18-15.99
		Min	2097	2060	500	15.38-16.11
		DoS	8231	5698	366	22.32-68.72
		Replay	8235	6387	356	24.54-69.52
XMV 10	Reactor cooling water flow 35.46-36.33 m ³ h ⁻¹	Max	786	701	500	1.65-2.01
		Min	11703	6651	489	18.88-67.49
		DoS	6093	1976	500	6.18-34.73
		Replay	5909	2283	500	6.4-34.61
XMV 11	Condenser cooling water flow 5.20-19.69 m ³ h ⁻¹	Max	325	319	500	1.59-1.6
		Min	414	408	500	0.64-0.65
		DoS	10803	2268	477	1.61-54.62
		Replay	11596	7782	108	38.32-69.49

Table 2: Impact of random attacks on XMV variables (500 runs for each attack)

(XMV 10) have the potential to increase the operating cost; however, the effect (\$9,064-11,596) is smaller than the attacks on XMEAS, as shown in Table 2.

3.2 Evolutionary multiobjective optimisation approach

Due to the computationally intensive nature of this work, and the slowness of the selected TE model simulator in MATLAB, we selected only two of the well-known Pareto-based EMO algorithms, the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) [43, 44] and Strength Pareto Evolutionary Algorithm 2 (SPEA2) [45]. Both algorithms were implemented using the Distributed Evolutionary Algorithms in Python (DEAP) [46] library. Using the selected algorithms we investigate the following three optimisation problems:

1. **Shutdown attacks** is defined as a two objective optimisation problem: minimise time required to shutdown the plant (f_1) and minimise the effort required to carry out the attack (f_2). Time required to shutdown the plant is the length of the plant operating after attack started. Plant shuts down as a result of exceeding plant operating constraints. We define effort as the total number of sensors and actuators being attacked.
2. **Operating cost attacks** defined as a two objective optimisation problem: maximise the total operating cost of the plant (f_1), and minimise the effort required to bring the plant down (f_2).
3. **Attacks against detection** defined as two/three objective optimisation: maximise economic loss (f_1), minimise detection (alarm) probability (f_2), and minimise detection probability (f_3).

3.2.1 Detection methods

The TE model generates a total of 51 data variables, 41 process measurements (XMEAS) and 12 manipulated variables (XMV). This data was used to train the detection methods. To detect attacks, three supervised learning methods –

decision tree (CART, tree depth=50), AdaBoost (with CART, number of estimators=100), random forest (with CART, number of estimators=25) – and one unsupervised learning method – one-class SVM (kernel=RBF, $\nu = 0.00346$ and $\gamma = 0.018$) – were used. The training data for supervised learning were selected from $integrity_{min}$ and $integrity_{max}$ attacks on measured variables (XMEAS). Attack samples were generated by carrying out integrity attacks on XMEAS signals with durations of between 20 minutes and 3 hours. The dataset used for training the unsupervised learning method consists of normal operational data without any attacks. The performance of the detection methods were measured using F1 score and false positive rate:

$$F_1 = 2 \frac{precision \times recall}{precision + recall} \quad (7)$$

$$\text{False Positive Rate (FPR)} = \frac{FP}{FP + TN} \quad (8)$$

where FP is false positives and TN is false negatives.

The performance of the detection algorithms on test data, unseen cases of integrity attacks on both XMEAS and XMV are illustrated in Table 3.

Each execution of the plant produces a data matrix of size 51 x 36000 data-points (500 points per hour for 72 hours). Systems like the TE plant are prone to natural noise due to behaviour of the physical components of the systems (e.g. actuator and sensors degrading over time, components of the plant wearing, or other kinds of natural noise in the environment). Detection should be robust against natural noise, and distinguish between the normal plant disturbances and attack conditions. Figure 3 shows random forest classifying data-points for a normal execution of the plant, under no attack, for 72 hours. False positives are the lines pointing at 1.

Algorithm	F1	FPR
Decision Tree	87.74	0.19
Random Forest	90.54	0.15
AdaBoost	72.40	0.016
One-Class SVM	86.62	0.40

Table 3: Performance of detection mechanisms




Figure 3: Anomaly detection using random forest (under normal operating condition)

Declaring that an attack is taking place at present requires the detection to be robust to false positives. To cater for this, we use a sliding window of size 100 to declare that an attack is present only if the percentage of anomalous data points in a window exceeds a threshold to ensure false positives do not overwhelm the operator.

We executed the TE model 1000 times under normal conditions, without any attacks, using a different random seed for each replicate to ensure that randomness was achieved. Then, using each detection method, we investigated the




Figure 4: Percentage of false positives for detection methods in a window size of 100 (under normal operating conditions)

highest false positive percentage seen for each run. Results obtained are shown in Figure 4. Based on these results, plant operator will need to define a threshold for raising an alarm or declaring an attack is taking place. For this study, we select the 99% percentile as the threshold for false alarm: 99% for decision tree; 99% random forest; 98% One-Class SVM; and 47% for AdaBoost. The objective of the attacker is not to raise any alarms while causing some damage by keeping the probability of attack as low as possible by reducing the number of attack data points in the sliding window below the defined threshold to ensure no alarms are raised.

3.2.2 Representation

Individuals (chromosomes) are represented as a list, consisting of 25 positions, each position (gene) denoting a sensor or an actuator. The types of genes are represented as integers, denoting the attack and its starting hour. Initial population of individuals are generated randomly.

3.2.3 Fitness function

The fitness of individuals is evaluated by converting individuals into MATLAB scripts that is executed on the TE plant. The performance of the shutdown attacks is evaluated based on plant run time (f1) and number of variables attacked (f2). The performance of the economic loss attack is evaluated based on the operating cost (f1) of the plant and number of variables attacked (f2). The performance of the attacks against detection is evaluated based on the detection probability (f1), increased operating cost of the plant (f2), and number of variables attacked (f3, for 3-objective optimisation). Data collected from TE model is tested against the detection model to determine the detection probability.

Parameters	Value
Chromosome size	25
Representation of genes	Integers
Number of generations	500-1000
Parent population (μ)	100-400
Crossover	Two-point crossover
Mutation	Uniform mutation
Selection	NSGA-II, SPEA2
Crossover probability (cxp)	0.8-0.90
Mutation probability ($mutpb$)	0.05-0.15
Probability of mutating a gene	0.05-0.08

Table 4: Operators and parameters for evolutionary multiobjective optimisation

3.2.4 Evolution

Table 4 shows the genetic parameters and the operators used in our experiments. Once the fitness of the initial population₀ has been evaluated, the evolutionary loop begins to generate the next generation (gen=1) population, as illustrated in Algorithm 1.

```

Function NSGA-II( $\mu$ ,  $mutp$ ,  $cspb$ ):
    ParetoFront=[];
     $\mu$ =pop size,  $pop$ =generateRandomPop( $\mu$ );
     $pop$ =evaluateFitness( $pop$ );
     $gen$ =1;
    while  $gen \leq ngens$  do
         $offspring$ =selTournament( $pop, \mu$ );
         $offspring$ =crossover( $offspring, cspb$ );
         $offspring$ =mutate( $offspring, mutp$ );
         $offspring$ =evaluateFitness( $offspring$ );
         $pop$ =selectNSGA2( $offspring+pop, \mu$ );
        ParetoFront.update( $pop$ );
         $gen$ = $gen + 1$ ;
    end
    return ParetoFront;

Function SPEA2( $\mu$ ,  $mutp$ ,  $cspb$ ):
    ParetoFront=[];
     $\mu$ =pop size,  $pop$ =generateRandomPop( $\mu$ );
     $pop$ =evaluateFitness( $pop$ );
     $gen$ =1;
    while  $gen \leq ngens$  do
         $offspring$ =vary( $pop, \mu, mutpb, cspb$ );
         $offspring$ =evaluateFitness( $offspring$ );
         $pop$ =selectSPEA2( $offspring+pop, \mu$ );
        ParetoFront.update( $pop$ );
         $gen$ = $gen + 1$ ;
    end
    return ParetoFront;

Function vary( $pop, \mu, mutp, cspb$ ):
     $offspring$ =[];
    for  $i = 0$  to  $\mu$  do
         $random$ =randomGenerator(0,1);
        if  $random < cspb$  then
             $ind1, ind2$ =selectTwoParents( $pop$ );
             $child1, child2$ =crossOver( $ind1, ind2$ );
             $offspring.add(child1)$ ;
        else if  $random < cspb+mutp$  then
             $ind$ =selectOneParent( $pop$ );
             $child$ =mutate( $ind$ );
             $offspring.add(child)$ ;
        else
             $ind$  = selectOneParent( $pop$ );
             $offspring.add(ind)$ ;
        end
    end
    return  $offspring$ ;

```

Algorithm 1: Evolutionary multiobjective optimisation algorithm for generating attacks

First, the individuals in population₀ are subject to the genetic variation operators to generate the next population of offspring. For NSGA-II we used the same genetic variation operators as used in the original algorithm [44]: tournament selection, two-point crossover, and uniform mutation. Crossover and mutation rates were manually tuned based on

values that are usually chosen within the literature: cross-over probabilities used are between 0.8-0.9 and mutation probability between 0.05-0.15. For SPEA2, we used the *vary* function shown in Algorithm 1, where, on each of the μ iterations, randomly picked individuals are subject to one of the three operations: two-point crossover, uniform mutation, or reproduction. Our initial experiments showed these operators performed better than the evolutionary operators that were used in standard SPEA2 [45]: binary tournament selection, single-point crossover, and bit-flip mutation. Functions *selectNSGA2* and *selectSPEA2* are the selection operators of NSGA-II and SPEA2, readily available in DEAP library. Both selection operators use the (parent+offspring) population to select the next generation: in other words the next generation of the population is produced from both the generated offspring and current parent population₀.

NSGA-II creates a Pareto rank of individuals from (parent+offspring) population using non-dominated sorting to select the next generation of individuals. If individuals have the same ranking score, then the crowding distance assignment method based on density estimation is used to select the individuals that are in the least crowded regions within the rank. SPEA2 calculates the strength of the individuals based on domination and density information, to select the new population for the next generation.

The obtained Pareto front set is updated with the new population, and we use the elements in this set to calculate the hypervolume at each generation of the evolution to monitor the convergence speed of the EMO algorithms. All experiments were carried out using DEAP library.

3.2.5 Performance metrics for evolutionary multiobjective optimisation

Multiobjective evolutionary algorithm have three important goals [47]: i) to minimise the distance between the obtained non-dominated solution set and the true Pareto-optimal set ii) to obtain a good, in most cases uniform distribution of solutions; iii) to maximise the extent of the obtained non-dominated solutions. Therefore, a wide variety of performance metrics [40] have been proposed to measure and compare the performance of EMO algorithms. In this study, we selected three of most frequently used EMO performance metrics, namely, hypervolume, spread, and inverted generational distance (IGD) with which to compare the performance of EMO algorithms.

The **hypervolume** [48], also known as the *S-metric* or *Lebesgue measure* is used for comparing convergence and diversity of a Pareto front. It measures the size of volume of the region between the estimated Pareto-optimal front, P and, a reference point r . We followed common practice by calculating r as a point defined by taking the worst known values for each of the objectives and shifted it slightly towards some unattainable values to ensure it is placed in a way that will be dominated by everything else. The hypervolume is defined as follows:

$$HV(P, r) = \mathcal{L} \left(\bigcup_{i=1}^{|P|} [r, i] \right) \quad (9)$$

It is the union of Lebesgue measure L for all points in P with respect to the reference point r . A Pareto front with a higher hypervolume is considered to indicate a better performing EMO algorithm. The **spread** metric Δ [44], also known as the diversity metric, measures the diversity of the solution by calculating Euclidean distance, d_i , between consecutive solutions. d_f and d_l are the Euclidean distances between the *extreme* solutions of the true Pareto front and the boundary solution of the obtained Pareto front [44]. Assuming there are N solutions in the obtained Pareto front, $i=1, i=2, \dots, N-1$, \bar{d} is the average of all distances d_i . Smaller value of Δ is desired, indicating better spread of solutions.

$$\Delta = \frac{d_f + d_l + \sum_{i=1}^{N-1} |d_i - \bar{d}|}{d_f + d_l + (N - 1)\bar{d}} \quad (10)$$

The **inverted generational distance (IGD)** [49] is another widely used metric to measure both convergence and diversity. IGD is defined as:

$$IGD = \frac{\sqrt{\sum_{i=1}^{P^*} d(i, P)}}{|P^*|} \quad (11)$$

where P^* denote the number of solutions in the optimal Pareto front, and d_i is the Euclidean distance between solution i and the nearest member in the obtained Pareto front, P [49]. A value, near $IGD = 0$ indicates better coverage of Pareto front and near true Pareto front. Due to size of the problem it was not possible to calculate the true Pareto front, and instead, a reference true Pareto front was computed for each problem by aggregating the obtained non-dominated

solutions from all runs to obtain a single front. These values were used to estimate the extreme values for the spread metric, and the reference Pareto front for the IGD metric.

To compare the performance of the EMO algorithms, the Kruskal-Wallis test was used for data that do not fit a normal distribution, and a one-way ANOVA test was used for data with a normal distribution. The confidence interval for all experiments is 95%.

3.3 Random generation of combined attacks

To determine the effectiveness of using EMO algorithms for attack generation some initial simple experiments were carried using a random generator and compared with EMO approach. As explained previously, there are in total 25 sensors and actuators to attack, and four possible attacks (*DoS*, *integrity_{min}*, *integrity_{max}* and *replay*). For the comparison study, we decided to keep the search simple by investigating only shutdown attacks. 10 sets of 50,000 randomly generated attack strategies were generated, limiting number of attacks in each strategy to maximum of 7, out of 25. For each set, a new seed was used for the random number generator of the TE model to ensure randomness in the simulation of the plant. Attack type and targets (sensors and actuators to attack) were randomly selected. Attacks were started at hour 2, and they were left to run until the end of the simulation time, (i.e. remaining 70 hours).


To compare the performance of the EMO against the randomly generated combinatorial attacks, 10 sets of experiments were carried out using SPEA2 algorithm and the genetic operators in Table 4 with a cross-over probability of 0.9, mutation probability of 0.05, and the independent probability of mutating a gene was 0.05. All experiments started from a random initial population of 100 individuals, and EMO was run for 500 generations. The same seeds used for the random number generator of the TE plant in randomly generated attacks were used here to prevent variability.

4 Experiments and results

4.1 Comparison of random generation and EMO approach

In this section, we report the results obtained from experiments to compare the random generation of shutdown attacks with EMO approach.

After removing duplicates, random generation produced a total of 442,125 unique attacks. Just over 18.5% of these attacks were able to bring the plant down in less than 1 hour. Figure 5a illustrates the distribution of these attacks. The best attack strategy random search generated was the attack that shut down the plant in 0.158 hours by attacking 6 sensors and actuators (XMEAS4, XMEAS8, XMEAS10, XMEAS11, XMEAS17, XMV1) using *integrity_{max}* attack.



(a) Distribution of shutdown attacks generated using random generation

(b) Distribution of shutdown attacks generated using EMO approach

Figure 5: Comparison of random and EMO approaches

Figure 5b show the results obtained using EMO approach. In total, EMO generated 35,658 unique attacks, and 86% of these attacks were under 1 hour, as indicated by the distribution skewed towards right in Figure 5b. These attacks performed far better than those generated by random generation, generated a high number of attacks with shutdown time of 0.138-0.156 hours. These results show the EMO approach is more effective at generating attacks than random chance, and justified our work to use the approach to generate further attacks. The results obtained for generating attacks using both EMO algorithms is presented in the next subsections.

4.2 Attacking the safety of the plant: shutdown attacks

In this section we report the performance of EMO algorithms that targetted the safety of the plant by searching for attacks that shut down the plant by attacking the least number of sensors and actuators. To compare the performance of NSGA-II and SPEA2 algorithm, results were collected over thirty runs for each EMO algorithm using a cross-over probability of 0.9, a mutation probability of 0.05, and the independent probability of mutating a gene was 0.05. For each of the thirty runs of evolution, a new seed was used to produce a different initial random population of size 100. The same sets of seeds were used for NSGA-II and SPEA2 to ensure both algorithms started with the same initial population. Similarly, the seed used for the TE Plant was kept the same for both algorithm. Each evolution was run for 500 generations. All experiments were carried out on a HPC platform facility at the University College London.

Table 5 shows the average and standard deviation of hypervolume, spread and IGD metrics. SPEA2 achieved better hypervolume and IGD with statistical confidence, yielding a better Pareto front and converged faster than NSGA-II. Meanwhile, NSGA-II obtains lower hypervolume and IGD than SPEA2 despite producing significantly better spread. The solutions in the obtained Pareto front were not widely spread across the front, only a small number of signals are required to cause the plant to shut-down. SPEA2 was able search this area better, and converged to a better Pareto front faster than NSGA-II.

Performance Measure	NSGA-II	SPEA2	p-value
Shutdown Attacks			
Hypervolume	0.8782 (0.0316)	0.8963 (0.0064)	0.000486
Spread	0.6878 (0.2346)	0.8786 (0.1473)	0.000054
IGD	0.0598 (0.025)	0.0368 (0.0154)	0.000245
Fastest Shutdown (hrs)	0.138	0.138	
Opcost Attacks			
Hypervolume	0.8235 (0.055)	0.8877 (0.074)	0.003
Spread	0.7117 (0.0853)	0.7356 (0.095)	0.234
IGD	0.1861 (0.0554)	0.1110 (0.0494)	0.000260
Max. Operating Cost (\$)	46,814	56,090	


Table 5: Performance of shutdown and operating cost attacks averaged over all runs

One of the best run obtained from shutdown attacks using NSGA-II and SPEA2 is shown in Figure 6a and Figure 6b. The elements of the final Pareto front obtained at the end of the evolution are plotted in red dots, and some of the members of the Pareto set are shown in Table 6. For this particular run, SPEA2 outperformed NSGA-II in that it found several different attacks with equal fitness functions and effort values in the Pareto set, and it had better Pareto front cardinality; in this case, a total of 12 elements, against the 9 found by NSGA-II. In the plot of the Pareto front some of the points on the diagram refer to multiple attacks with equal fitness; for example, there were 3 attacks using 6 effort and shutdown time of 0.140 hours; 2 attacks with 4 effort and shutdown time of 0.1560 hours.


Attack Strategy	Shut-down (hrs)	Effort
$XMEAS4_{IntegrityMin}, XMEAS5_{IntegrityMin}, XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}, XMEAS17_{IntegrityMin}, XMEAS31_{IntegrityMin}, XMV6_{Replay}$	0.138	7
$XMEAS4_{IntegrityMin}, XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}, XMEAS17_{IntegrityMin}, XMEAS31_{IntegrityMin}, XMV6_{Replay}$	0.140	6
$XMEAS4_{IntegrityMin}, XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}, XMEAS17_{IntegrityMin}, XMEAS31_{IntegrityMin}$	0.1460	5
$XMEAS4_{IntegrityMin}, XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}, XMEAS17_{IntegrityMin}$	0.1560	4
$XMEAS4_{IntegrityMin}, XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}$	0.1760	3
$XMEAS8_{IntegrityMin}, XMEAS11_{IntegrityMin}$	0.2579	2
$XMEAS4_{IntegrityMin}$	0.5640	1
Do Not Attack	72	0

Table 6: Some elements of the Pareto front for shutdown attack


Figure 6c shows the hypervolume for each of the 500 generations, averaged over all 30 runs. At the end of each generation, the hypervolume was computed according to the Pareto front achieved at that generation to compare the speed of the convergence. For convenience, plotted hypervolume results are normalised to the interval [0-1] according to the best hypervolume value possible, estimated based on the maximum measurement obtained. SPEA2 converges faster




(a) Pareto front for shutdown attacks < 1 hour (NSGA-II)




(b) Pareto front for shutdown attacks < 1 hour (SPEA2)



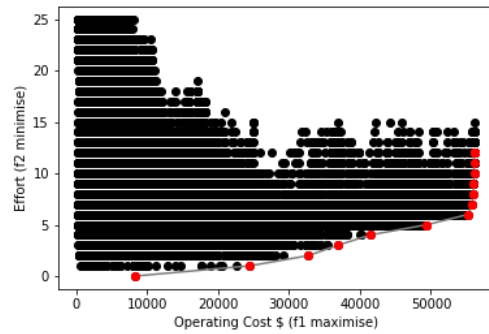
(c) Hypervolume for shutdown attacks




(d) Measurement metrics for shutdown attacks




(e) Pareto front for operating cost attacks (NSGA-II)



(f) Pareto front for operating cost (SPEA2)



(g) Hypervolume for operating cost attacks



(h) Measurement metrics for operating cost attacks

Figure 6: Results of shutdown and operating cost attacks

Attack Strategy	Opcost (\$)	Effort
XMEAS2 _{DoS(2,70)} ,XMEAS7 _{IntegrityMax(2,70)} ,XMEAS8 _{DoS(10,20)} ,XMEAS11 _{DOS(30,42)} , XMEAS12 _{Replay(50,12)} ,XMEAS15 _{IntegrityMax(50,12)} ,XMEAS31 _{IntegrityMin(30,42)} , XMV1 _{DOS(10,62)} ,XMV2 _{IntegrityMin(2,70)} ,XMV3 _{DOS(2,70)} ,XMV4 _{DOS(50,12)} ,XMV11 _{DOS(50,12)}	56,090	12
XMEAS2 _{DoS(2,70)} ,XMEAS7 _{IntegrityMax(2,70)} , XMEAS31 _{IntegrityMin(30,42)} ,XMV1 _{DoS(10,62)} , XMV2 _{IntegrityMin(2,70)} ,XMV3 _{DoS(2,70)}	55,275	6
XMEAS2 _{DoS(2,70)} ,XMEAS7 _{IntegrityMax(2,70)} , XMEAS31 _{IntegrityMin(30,42)} , XMV1 _{DoS(10,62)} , XMV3 _{IntegrityMin(2,70)}	49,449	5
XMEAS2 _{DoS(2,70)} ,XMEAS7 _{IntegrityMax(2,70)} ,XMV1 _{DoS(10,62)} ,XMV3 _{IntegrityMin(2,70)}	41,430	4
XMEAS7 _{IntegrityMax(2,70)} ,XMEAS31 _{DoS(10,62)} ,XMV3 _{IntegrityMin(2,70)}	36,866	3
XMEAS7 _{IntegrityMax(2,70)} ,XMEAS31 _{IntegrityMin(30,42)}	32,761	2
XMEAS7 _{IntegrityMax(2,70)}	24,479	1
Do Not Attack	8,210	0

Table 7: Some elements of the Pareto front for operating cost attacks using SPEA2

to a better Pareto front whereas NSGA-II requires more time to reach a slightly worse Pareto set. This is supported by IGD metric, as shown on the boxplot in Figure 6d SPEA2 scores a better IGD score.

Results obtained show combined simultaneous attack generated using EMO approach performed better than single and combined attacks generated randomly against the XMEAS and XMV variables. As discussed in Section 3.1, it took 0.52 hours to bring down the plant, whereas, EMO found a range of attacks that could bring down the plant faster, in 0.138-0.5640 hours. Random combined attacks (Section 4.1) at the very best produced an attack that could bring down the plant in 0.158 hours.

4.3 Causing economic loss: Operating cost attacks


In this section we report the performance of EMO algorithms that targetted the operating cost of the plant to cause economic loss by attacking the least number of sensor and actuator signals. Due to the slowness of the fitness function, only a small set of attack start times (2, 10, 20, 30, 50 hours) and attack duration (10, 12, 20, 42, 50, 52, 62, 70 hours) were used to generate attacks. As before, attack types were DoS, replay, *integrity_{max}* and *integrity_{min}*. Individuals were represented as chromosomes encoded as a list of 25 integers (genes), each position denoting a sensor or an actuator. In total each gene could have a value between 0-37 (0 representing not to attack, remaining 36 representing 9 different attacks per attack type with different start time and attack duration). This is a search space of size 37^{25} . Twenty runs were carried out for each EMO algorithm to analyse the impact of attacks on the operating cost of the plant. As before, attacks were generated using the same genetic operators, shown in Table 4, with a cross-over probability of 0.85, mutation probability of 0.10, and independent probability of mutating a gene was 0.08. Each evolution started with a random population of 400 individuals, and ran for 1000 generations.

One of best runs of operating cost attacks using NSGA-II and SPEA2 is shown in Figure 6e and Figure 6f. As with shutdown attacks, SPEA2 performed better than NSGA-II both in the quality of obtained Pareto front set, and the time it took to converge. As indicated in Table 5, the average over all the runs showed that SPEA2 has a hypervolume average of 0.8877, against NSGA-II scoring a hypervolume of 0.8235. Figure 6h shows the boxplots for measurement metrics. SPEA2 produced significantly higher hypervolume and IGD, but no significant difference was found between the two algorithms for spread. Figure 6g shows the comparison of hypervolume between NSGA-II and SPEA2, averaged over all runs. These results show that SPEA2, as before, is able to search the space faster and produce a better Pareto front with a higher cardinality. SPEA2 (Table 7) increased the cost of the operating the plant from an average of \$8,208 to \$56,090, whereas NSGA-II increased the operating cost to \$46,814.


Table 7 shows some of the elements of the obtained Pareto front set for SPEA2. The numbers in the bracket denotes the start time and duration of the attack, for example XMEAS8_{DoS(10,20)} means a DoS attack was carried out against XMEAS 8 starting at hour 10, for a duration of 20 hours. Overall, these results show that EMO approach can be used successfully to generate attacks that could cause economic loss, by identifying the components that increase the operating cost of the plant.

4.4 Generating attacks against detection methods


As before individuals are represented as a list, consisting of 25 positions, each position denoting a sensor or an actuator. One key obstacle we had with our experiments was time, the execution of the TE model in MATLAB could take up to several minutes. This performance issue influenced the way we encoded our individuals, the size of the population, and




(a) Pareto front for AdaBoost



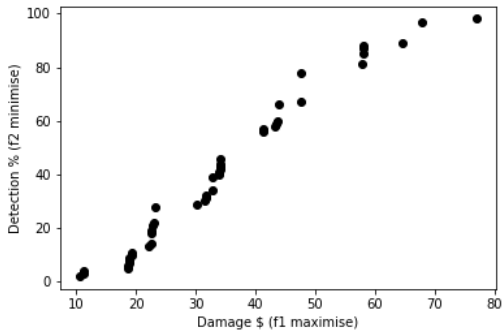
(b) Hypervolume for AdaBoost




(c) Pareto front for decision tree




(d) Hypervolume for decision tree




(e) Pareto front for random forest



(f) Hypervolume for random forest




(g) Pareto front for one-class SVM




(h) Hypervolume for one-class SVM


Figure 7: Pareto front and hypervolume for detection methods, averaged over all runs




(a) Metrics for AdaBoost



(b) Metrics for decision tree



(c) Metrics for random forest



(d) Metrics for one-class SVM

Figure 8: Performance metrics for classifiers, averaged over all runs

the number of generations. To make the problem computationally tractable, we limited the types of genes to a pool of 140 in which half denoted DoS attacks and the remaining half denoted replay attacks. Genes were represented as integers, each number denoting the start time of the attacks (between hour 2-70). The duration of the attacks against detection methods were kept constant for all attacks, 2 hours. This is a combinatorial search problem of size 140^{25} . Results were collected over 10 runs for each EMO algorithm against each detection method (AdaBoost, decision tree, random forest and one-class SVM) using a cross-over probability of 0.8, mutation probability of 0.15, and the independent probability of mutating a gene was 0.08.

Each evolution started from a random population of 200 individuals, and ran for 1000 generations. Attacks that lead to plant to shutdown were penalised, with a score -1, as our aim was to generate attacks that kept the plant running while causing some damage and evading detection.

Despite using a limited number of attack parameters (start time, duration), and using a small number of population and generation size, the obtained result capture some of the weaknesses of the detection methods.

Figure 7 shows the results obtained. Figure 7a, Figure 7c, Figure 7e and Figure 7g show one of the best Pareto fronts obtained against four classifiers. As expected, EMO algorithms were able to exploit low detection rate of AdaBoost, and cause more damage (\$406-6) while keeping the attack detection (alarm) probability at a lower rate. EMO algorithms performed less damaging attacks against decision tree (\$190-10), random forest (\$77-10), and one-class SVM (\$21-13). As before, hypervolume was computed at each generation to analyse the convergence. Figure 7b, 7d, 7f, and 7h shows a steady increase of hypervolume suggesting that more generations will yield a better search and convergence. Figure 8 shows the hypervolume, spread and IGD boxplots for four classifiers. As indicated in Table 8 statistical test shows there was no significant difference between NSGA-II and SPEA2 against decision tree, random forest and one-class SVM. However, for AdaBoost, NSGA-II did significantly better, both in terms of obtaining a better Pareto front (as

Performance Measure	NSGA-II	SPEA2	p-value
AdaBoost			
Hypervolume	0.6020 (0.2257)	0.3910 (0.2525)	0.032663
Spread	0.7952 (0.0910)	0.7165 (0.0841)	0.000054
IGD	0.2273 (0.1024)	0.3345 (0.1203)	0.0376
Damage Range (\$)	15-490.63	12.5-424.99	
Decision Tree			
Hypervolume	0.1442 (0.0464)	0.1452 (0.0416)	0.8798
Spread	0.6825 (0.0564)	0.6450 (0.0441)	0.1340
IGD	0.1505 (0.0506)	0.1580 (0.0382)	0.7284
Damage Range (\$)	9-190.63	13-182.26	
Random Forest			
Hypervolume	0.1170 (0.0290)	0.1240 (0.0304)	0.2895
Spread	0.5846 (0.0384)	0.5943 (0.0361)	0.5877
IGD	0.1375 (0.0110)	0.1382 (0.01328)	0.7622
Damage Range (\$)	4-76.92	1.56-56.10	
One-Class SVM			
Hypervolume	0.0785 (0.0222)	0.0788 (0.0225)	0.9768
Spread	0.5887 (0.0363)	0.5909 (0.03789)	0.8205
IGD	0.1108 (0.0389)	0.1181 (0.04144)	0.7070
Damage Range (\$)	1.5-21.28	2.0-20.75	

Table 8: Performance of detection attacks against classifiers, averaged over all runs

indicated by hypervolume and IGD) metrics and converged a lot of faster as indicated by hypervolume plot (Figure 7b). Although, this time SPEA2 had a better spread compared to NSGA-II, it did not yield a better Pareto front.


As the damage increases, denoted by the increased cost of operating the plant, the probability of detection also increases. EMO algorithms failed to evolve highly damaging attacks against the one-class SVM as it was able to detect DoS and replay attacks better than other detection methods. However, we were able to cause some economic damage with low detection probabilities for other detection methods.

4.4.1 Seeding initial population, and generating attacks using 3-objective optimisation


In this section, we report some preliminary experiments that require further investigation, but show promising results. Generating attacks against a strong classifier can be a very time consuming task; this is especially true for cases like our plant model, for which the evaluation of individuals (fitness function) is slow. The slowness of our fitness function also influenced the attack parameters, the size of the population, and the number of generations. To test if we could generate better attacks against decision tree and random forest classifiers, we started some experiments using a seeded initial population that included some good individuals that were obtained previously from experiments carried out against the decision tree classifier (i.e., Figure 7c). Seeding is a common practice in single-objective evolutionary algorithms where prior knowledge obtained from previous experiments or expert knowledge is included in the initial population as a good initial estimate, however advantages and disadvantages of using seeding in EMO algorithms, especially for solving real-world combinatorial optimisation problems require further studies [50].

A comprehensive study is left as future work, and here we report some initial results. We carried out an experiment where we seeded the initial population with 10 of the best individuals obtained from the previous runs of decision tree experiments, and started the EMO from this modified population against the decision tree and random forest classifiers. Figure 9a shows the performance of decision tree after seeding the initial population (compare with no seeding in Figure 7c), and Figure 9b shows the performance of random forest after seeding the initial population (compare with no seeding in Figure 7e), showing attacks with higher damage and low detection probability. These results indicate seeding could significantly reduce the duration of experiments, and more rapidly identify those attacks that are likely to evade detection (raise alarms) and, at the same time cause some economic loss. However, more experiments are required to understand the full benefits and weaknesses of the variety of strategies for this approach (e.g. such as the number of seeds used), as seeding can reduce the diversity of the population, and so fail to explore other feasible region in the attack space.

One of the weaknesses of the two objective optimisation against the detection methods was that the effort was not optimised, and we were not able to tell if attacks could be generated using less effort. To investigate this, if attackers could cause same damage using less effort, we carried out a 3-objective optimisation (maximise operating cost of the



(a) Attack generated against decision tree




(b) Attack generated against random forest

Figure 9: Seeding population with knowledge gained from prior experiments


plant, minimise detection, minimise effort). Given the computational cost, only one detection method, AdaBoost was selected for further investigation. Table 9 shows the performance of the NSGA-II and SPEA2 averaged over two runs, for 800 generations.

	NSGA-II	SPEA2
Damage Range (\$)	0-1633	0-1114
Hypervolume	0.7140	0.5423

Table 9: Performance of 3-objective attack generation against AdaBoost



(a) NSGA-II



(b) SPEA2

Figure 10: Obtained Pareto front for 3-objective attack generation against AdaBoost

Both runs required more time to converge, but NSGA-II appears to search a wider region using 3-objectives compared to SPEA2. NSGA-II showed a better spread of attacks over the search space, and a better value for hypervolume. The best attacks are those with lower detection probability, and these tend to lie in the lower regions of the detection which cause damage less than \$200. As shown in Figure 10, NSGA-II finds attacks with higher damage, damage over \$1500, but, SPEA2 tends to generate more attacks in the lower regions where there is a better probability of avoiding detection.

A more reliable comparison requires further experiments, which we leave as part of future work; however, overall, both algorithms found attacks that avoid detection and cause some harm. Assuming the detection probability threshold is less than 5%, the highest cost attack NSGA-II found is \$266.92, attacking 12 sensors and actuators, with a detection probability of 3%, whereas SPEA2 found a slightly worse attack, \$179.72 with effort 16 and detection probability of 4%. If the intention is to use the smallest effort and cause maximum damage, then the optimal attack strategy produced using NSGA-II is an attack that costs \$179.72 using effort 6. SPEA2 found a similar attack using effort 6, at a cost of

\$170.69. The attacks generated using less effort were either detected or caused very little economic damage, that is $\leq \$50.00$.

5 Application of results

The EMO approach developed in this paper can be employed to test the vulnerabilities of cyber-physical systems with a broad attack surface, where attackers often target process measurements and manipulated variables, as these values can impact the process directly, and cause physical damage.

Results obtained provide a rich set of attacks that can be analysed and help plant operators identify the most vulnerable combinations of sensors and actuators. Table 10 shows a small subset of vulnerable combinations of the sensors and actuators that could shut down the TE plant in under 17 minutes. For example, carrying out a single attack on XMEAS 8 (reactor level) was able to shut down the plant in over 2.8 hours, and attacking XMEAS 11 (separator temperature) avoided shutting down the plant for all types of attacks, as shown in Table 1. Results show that attacking both of these sensors simultaneously could shut down the plant in 14.8 minutes. These results also showed that the plant is less resilient to attacks on process measurements (sensors), and, if an adversary wants to bring down the plant in a very short period of time such as less than 10 minutes, attacking sensors is more likely to cause this to happen than attacking manipulated variables. This is, possibly because the selected attack parameters for actuator signals (based on the observed signals), were not as effective as attacking process measurements. Similarly, DoS attacks were slower and less successful against manipulated signals. Future research will investigate this further, focusing more on the attack parameters and timing of the attacks.

Vulnerable Combinations	SDT (min)
XMEAS4,XMEAS8,XMEAS11, XMEAS17,XMV6	8.9
XMEAS4,XMEAS8,XMEAS11, XMEAS31	9.6
XMEAS3,XMEAS4,XMEAS8, XMEAS11	10.2
XMEAS8,XMEAS9,XMEAS11	10.8
XMEAS8,XMEAS11,XMV6	11.4
XMEAS8,XMEAS11,XMV10	12.0
XMEAS5,XMEAS8,XMEAS11	13.8
XMEAS9,XMV7,XMV11	14.4
XMV7,XMV10,XMV11	14.5
XMEAS4,XMEAS7,XMEAS17	14.7
XMEAS8,XMEAS11	14.8
XMV9, XMV10	16.2
XMEAS9,XMV11	16.8

Table 10: A selection of vulnerable combinations that could bring the plant down under 17 minutes

Similarly, we found a wide variety of combined attacks that are capable of increasing the operating cost. Most of these attacks involved carrying out an *integrity_{max}* attack on XMEAS 7 (reactor pressure), which means sending higher values than expected. When the controller receives these values, it attempts to reduce the pressure by opening the purge valve, and this causes raw materials to be lost in the purge stream. This in turn increases the operating cost. Carrying out a single attack on XMEAS 7 increases the operating cost to \$24,479 but, as the results indicate, more damage can be inflicted using the right combinations of sensors and actuators.

Obtained results show EMO algorithms can be used to generate attacks to identify vulnerable parts of a system, and this knowledge can be used to design systems that are more secure and resilient to cyber attacks. One way to achieve this is to consider the vulnerable combinations when designing network segmentation. The *zone and conduit model* is a framework for network segmentation to manage security threats for industrial automation and control systems, recommended as part of the standard such as ISA/IEC 62443 [51]. Zones are defined as a group of logical or physical assets sharing common security requirements, and conduits are the paths of communication between the zones. Using the knowledge from EMO, vulnerable combinations of sensors and actuators that cause the plant to shut down or increase the operating cost of the plant can be aggregated in different zones to build a more secure network.

As we were able to generate a large number of attacks that were not detected by all detection methods, this approach can be also used as a tool to test and develop better detection methods. However, generating attacks that evade detection,

and at the same time cause some significant economic damage on a system like TE process is a challenging task, since it requires attacks to be long in duration. Our results show that carrying out a successful attack requires knowledge of the system to ensure that the correct combination of sensors and actuators are attacked, and to avoid detection or the triggering of the safety system. The significant attacks generated against AdaBoost, decision tree and random forest classifiers involved attacking multiple components in the system to evade detection while causing economic damage. A naïve attacker that randomly attacks multiple targets is unlikely to achieve similar damage, and is highly likely to be detected. The focus of our future work will involve investigating and designing more complex attacks that could learn the behaviour of the plant and the detection system, and generate attacks using this knowledge.

Results obtained show that evolutionary multiobjective optimisation can be used successfully as an effective tool to model the behaviour of the adversary against a defence method, and illustrate the conflicts and associated trade-offs among common security objectives: attack impact, detection and effort required. Using the results obtained from such an analysis, security engineers can take measures to understand and eliminate system vulnerabilities before they are exploited by malicious actors.

6 Conclusions and future work

This paper demonstrates a novel and important application of evolutionary multiobjective optimisation for security of industrial control systems, and more general cyber-physical systems. Using a simulation of a complex and realistic plant, of the sort used routinely in factories and plants, it is possible to discover (and remove) vulnerabilities through an automatic process. The threat to such systems is both realistic and of critical importance, and there is no comparable mechanism that is robust, efficient and that allows the identification of vulnerable combinations of sensors and actuators. The knowledge gained from our work can be used in number of ways. The first is in determining the criticality of security decisions such as patching, carrying out risk assessment and designing resilient network segments. Secondly, control engineers can use the insight gained from this work to analyse the implications of security on process control, and design security-aware control algorithms. The attacks generated against the detection methods show our approach can also be used as a tool to find the vulnerabilities in the detection before these are exploited by the adversaries.

The existing implementation of the TE model in MATLAB was too slow to be ideal for stochastic population-based optimisation algorithms like evolutionary algorithms and, therefore our experiments were limited by this. This was particularly a problem when generating attacks against the detection as evolution required more time to converge properly. Future work will involve optimising the code to improve the performance, and enable us to carry out attacks with better primitives. Planned future work will involve using the approach to design more advanced attacks against detection methods. TE model was used to illustrate the approach, however the approach is agnostic, and future work will focus on demonstrating this on other cyber-physical systems using more advanced attacks. The unavailability of accessible benchmark methods is a major obstacle for ICS security research. We are currently working on building cyber-physical testbeds, and we hope to test our approach on a different type of industrial process with physical and network components. Finally, will also investigate how other EMO algorithms, in particularly those based on aggregation-based and indicator-based algorithms to see if the obtained results can be further improved.

7 Acknowledgements

This work was conducted partly under UK EPSRC Grant No: EP/G037264/1 as part of University College London's Security Science Doctoral Training Centre, and partly under UK EPSRC Grant No: EP/N023234/1 as part of PETRAS Internet of Things Research Hub.

References

- [1] R.J. Patton, C. Kambhampati, A. Casavola, P. Zhang, S. Ding, and D. Sauter. A Generic Strategy for Fault-Tolerance in Control Systems Distributed Over a Network. *European Journal of Control*, 13(2):280 – 296, 2007.
- [2] Cyber attack shuts down Evraz IT systems across North America, but company says no data compromised. <https://www.cbc.ca/news/canada/saskatchewan/evraz-regina-shut-down-ransomware-attack-1.5487017>. Last Accessed: 11/03/2020.
- [3] CRASHOVERRIDE Analyzing the Threat to Electric Grid Operations Version 2.20170613. <https://dragos.com/wp-content/uploads/CrashOverride-01.pdf>, June 2017. Last Accessed: 04/05/2020.

- [4] Die Lage der IT-Sicherheit in Deutschland 2014. Technical report, Bundesamt für Sicherheit in der Informationstechnik (BSI), Nov 2014.
- [5] W32.Duqu: The Precursor to the Next Stuxnet (Version 1.4). Technical report, Symantec Security Response, Nov 2011.
- [6] Havex Hunts For ICS/SCADA Systems. <https://www.f-secure.com/weblog/archives/00002718.html>. Accessed: 02/05/2020.
- [7] W32.Stuxnet Dossier (Version 1.4). Technical report, Symantec Security Response, Feb 2011.
- [8] Alvaro A. Cárdenas, Saurabh Amin, Zong-Syun Lin, Yu-Lun Huang, Chi-Yen Huang, and Shankar Sastry. Attacks against Process Control Systems: Risk Assessment, Detection, and Response. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '11, page 355–366, New York, NY, USA, 2011. Association for Computing Machinery.
- [9] Yu-Lun Huang, Alvaro A. Cárdenas, Saurabh Amin, Zong-Syun Lin, Hsin-Yi Tsai, and Shankar Sastry. Understanding the physical and economic consequences of attacks on control systems. *International Journal of Critical Infrastructure Protection*, 2(3):73 – 83, 2009.
- [10] Marina Krotofil and Alvaro A. Cárdenas. *Resilience of Process Control Systems to Cyber-Physical Attacks*, pages 166–182. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [11] B. Genge, C. Siaterlis, M. Hohenadel, Béla Genge, Christos Siaterlis, and Marc Hohenadel. Impact of network infrastructure parameters to the effectiveness of cyber attacks against industrial control systems. *INTERNATIONAL JOURNAL OF COMPUTERS, COMMUNICATIONS & CONTROL*, page 673, 2012.
- [12] Wei Wang, Antonio Cammi, Francesco [Di Maio], Stefano Lorenzi, and Enrico Zio. A monte carlo-based exploration framework for identifying components vulnerable to cyber threats in nuclear power plants. *Reliability Engineering & System Safety*, 175:24 – 37, 2018.
- [13] Antonio Di Pietro, Chiara Foglietta, Simone Palmieri, and Stefano Panzieri. Assessing the Impact of Cyber Attacks on Interdependent Physical Systems. In Jonathan Butts and Sujeet Sheno, editors, *Critical Infrastructure Protection VII*, pages 215–227, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [14] K. Huang, C. Zhou, Y. Tian, S. Yang, and Y. Qin. Assessing the Physical Impact of Cyberattacks on Industrial Cyber-Physical Systems. *IEEE Transactions on Industrial Electronics*, 65(10):8153–8162, 2018.
- [15] Wei Li. Using genetic algorithm for network intrusion detection. In *In Proceedings of the United States Department of Energy Cyber Security Group 2004 Training Conference*, pages 24–27, 2004.
- [16] T. Xia, G. Qu, S. Hariri, and M. Yousif. An efficient network intrusion detection method based on information theory and genetic algorithm. In *PCCC 2005. 24th IEEE International Performance, Computing, and Communications Conference, 2005.*, pages 11–17, April 2005.
- [17] T. Vollmer, J. Alves-Foss, and M. Manic. Autonomous rule creation for intrusion detection. In *2011 IEEE Symposium on Computational Intelligence in Cyber Security (CICS)*, April 2011.
- [18] AA Ojugo, AO Eboka, O Okonta, RE Yoro, and FO Aghware. Genetic algorithm rule-based intrusion detection system (GAIDS). *Journal of Emerging Trends in Computing and Information Sciences*, 2012.
- [19] Mohammad Sazzadul Hoque, Md. Abdul Mukit, and Md. Abu Naser Bikas. An Implementation of Intrusion Detection System Using Genetic Algorithm. *CoRR*, abs/1204.1336, 2012.
- [20] Pedro A. Diaz-gomez, Ingenieria De Sistemas, and Dean F. Hougen. Improved off-line intrusion detection using a genetic algorithm. In *in Proceedings of the Seventh International Conference on Enterprise Information Systems*, 2005.
- [21] Anup Goyal and Chetan Kumar. GA-NIDS: A Genetic Algorithm based Network Intrusion Detection System. In *ElectricalEngineering & Computer Science, North West University, Technical Report*, 2008.
- [22] Wei Lu and Issa Traore. Detecting New Forms of Network Intrusion Using Genetic Programming. *Computational Intelligence*, 20(3):475–494, 2004.
- [23] S. Pastrana, A. Orfila, and A. Ribagorda. A Functional Framework to Evade Network IDS. In *2011 44th Hawaii International Conference on System Sciences*, pages 1–10, Jan 2011.
- [24] Kinga Mrugala, Nilufer Tuptuk, and Stephen Hailes. Evolving Attackers against Wireless Sensor Networks. In *GECCO 2016 Companion Volume*, page pp306, Denver, USA, 20-24 July 2016. ACM.
- [25] K. Mrugala, N. Tuptuk, and S. Hailes. Evolving attackers against wireless sensor networks using genetic programming. *IET Wireless Sensor Systems*, 7(4):113–122, 2017.

- [26] Hilmi Güneş Kayacik, Malcolm Heywood, and Nur Zincir-Heywood. On Evolving Buffer Overflow Attacks Using Genetic Programming. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation, GECCO '06*, pages 1667–1674, New York, NY, USA, 2006. ACM.
- [27] David J. John, Robert W. Smith, William H. Turkett, Daniel A. Cañas, and Errin W. Fulp. Evolutionary Based Moving Target Cyber Defense. In *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation, GECCO Comp '14*, pages 1261–1268, New York, NY, USA, 2014. ACM.
- [28] Rinku Dewri, Nayot Poolsappasit, Indrajit Ray, and Darrell Whitley. Optimal Security Hardening Using Multi-objective Optimization on Attack Tree Models of Networks. In *Proceedings of the 14th ACM Conference on Computer and Communications Security, CCS '07*, pages 204–213, New York, NY, USA, 2007. ACM.
- [29] Dennis Garcia, Anthony Erb Lugo, Erik Hemberg, and Una-May O'Reilly. Investigating Coevolutionary Archive Based Genetic Algorithms on Cyber Defense Networks. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '17*, pages 1455–1462, New York, NY, USA, 2017. ACM.
- [30] Erik Hemberg, Joseph R. Zipkin, Richard W. Skowrya, Neal Wagner, and Una-May O'Reilly. Adversarial Co-Evolution of Attack and Defense in a Segmented Computer Network Environment. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '18*, page 1648–1655, New York, NY, USA, 2018. Association for Computing Machinery.
- [31] George Rush, Daniel R. Tauritz, and Alexander D. Kent. Coevolutionary Agent-based Network Defense Lightweight Event System (CANDLES). In *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation, GECCO Companion '15*, pages 859–866, New York, NY, USA, 2015. ACM.
- [32] T. Service, D. Tauritz, and W. Siever. Infrastructure Hardening: A Competitive Coevolutionary Methodology Inspired by Neo-Darwinian Arms Races. In *Computer Software and Applications Conference, 2007. COMPSAC 2007. 31st Annual International*, volume 1, pages 101–104, July 2007.
- [33] Travis Service and Daniel Tauritz. Increasing Infrastructure Resilience through Competitive Coevolution. *New Mathematics and Natural Computation*, 05(02):441–457, 2009.
- [34] J. Decraene, M. Chandramohan, M. Y. H. Low, and C. S. Choo. Evolvable simulations applied to Automated Red Teaming: A preliminary study. In *Simulation Conference (WSC), Proceedings of the 2010 Winter*, pages 1444–1455, Dec 2010.
- [35] R. Bronfman-Nadas, N. Zincir-Heywood, and J. T. Jacobs. An Artificial Arms Race: Could it Improve Mobile Malware Detectors? In *2018 Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–8, June 2018.
- [36] J.J. Downs and E.F. Vogel. A plant-wide industrial process control problem. *Computers & Chemical Engineering*, 17(3):245 – 255, 1993.
- [37] Truls Larsson, Kristin Hestetun, Espen Hovland, and Sigurd Skogestad. Self-Optimizing Control of a Large-Scale Plant: The Tennessee Eastman Process. *Industrial & Engineering Chemistry Research*, 40(22):4889–4901, 2001.
- [38] N. Lawrence Ricker. Tennessee Eastman Challenge Archive. <http://depts.washington.edu/control/LARRY/TE/download.html#Multiloop>, Dec 1998.
- [39] A. Isakov and M. Krotofil. Damn Vulnerable Chemical Process - Tennessee Eastman. <https://github.com/satejnik/DVCP-TE>, 2015.
- [40] N. Riquelme, C. Von Lüken, and B. Baran. Performance metrics in multi-objective optimization. In *2015 Latin American Computing Conference (CLEI)*, pages 1–11, Oct 2015.
- [41] Andrew Fielder, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. Decision support approaches for cyber security investment. *Decision Support Systems*, 86:13 – 23, 2016.
- [42] N.L. Ricker. Optimal steady-state operation of the Tennessee Eastman challenge process. *Computers & Chemical Engineering*, 19(9):949 – 959, 1995.
- [43] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. *A Fast Elitist Non-dominated Sorting Genetic Algorithm for Multi-objective Optimization: NSGA-II*, pages 849–858. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.
- [44] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, April 2002.
- [45] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. Technical report, Swiss Federal Institute of Technology (ETH), 2001.

- [46] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner Gardner, Marc Parizeau, and Christian Gagné. DEAP: Evolutionary Algorithms Made Easy. *J. Mach. Learn. Res.*, 13(1):2171–2175, July 2012.
- [47] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evol. Comput.*, 8(2):173–195, June 2000.
- [48] Eckart Zitzler and Lothar Thiele. Multiobjective optimization using evolutionary algorithms - A comparative case study. In Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature - PPSN V*, pages 292–301, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.
- [49] D. A. Van Veldhuizen and G. B. Lamont. On measuring multiobjective evolutionary algorithm performance. In *Proceedings of the 2000 Congress on Evolutionary Computation. CEC00 (Cat. No.00TH8512)*, volume 1, pages 204–211 vol.1, July 2000.
- [50] Tobias Friedrich and Markus Wagner. Seeding the initial population of multi-objective evolutionary algorithms: A computational study. *Applied Soft Computing*, 33:223 – 230, 2015.
- [51] ISA99: Developing the Vital ISA/IEC 62443 Series of Standards on Industrial Automation and Control Systems (IACS) Security, 2015.