Quality Control Tools for Cyber-Physical Security of Production Systems

Ahmed Essam Elhabashy

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Industrial and Systems Engineering

Jaime A. Camelio, Chair
Lee J. Wells, Co-Chair
Zhenyu (James) Kong
William H. Woodall

November 9, 2018
Blacksburg, VA

Keywords: Advanced Production Systems, Cyber-Physical Attacks, Cyber-Physical Security, Quality Control, Statistical Process Control

Quality Control Tools for Cyber-Physical Security of Production Systems

Ahmed Essam Elhabashy

ABSTRACT

With recent advancements in computer and network technologies, cyber-physical systems have become more susceptible to cyber-attacks; and production systems are no exception. Unlike traditional Information Technology (IT) systems, cyber-physical systems are not limited to attacks aimed at Intellectual Property (IP) theft, but also include attacks that maliciously affect the physical world. In manufacturing, such cyber-physical attacks can destroy equipment, force dimensional product changes, alter a product's mechanical characteristics, or endanger human lives.

The manufacturing industry often relies on modern Quality Control (QC) tools to protect against quality losses, such as those that can occur from an attack. However, cyber-physical attacks can still be designed to avoid detection by traditional QC methods, which suggests a strong need for new and more robust QC tools. Such new tools should be able to prevent, or at least minimize, the effects of cyber-physical attacks on production systems. Unfortunately, little to no research has been done on using QC tools for cyber-physical security of production systems.

Hence, the overarching goal of this work is to allow QC systems to be designed and used effectively as a second line of defense, when traditional cyber-security techniques fail and the production system is already breached. To this end, this work focuses on: 1) understanding the role of QC systems in cyber-physical attacks within manufacturing through developing a taxonomy encompassing the different layers involved; 2) identifying existing weaknesses in QC tools and exploring the effects of exploiting them by cyber-physical attacks; and 3) proposing more effective QC tools that can overcome existing weaknesses by introducing randomness to the tools, for better security against cyber-physical attacks in manufacturing.

Quality Control Tools for Cyber-Physical Security of Production Systems

Ahmed Essam Elhabashy

GENERAL AUDIENCE ABSTRACT

The recent technological developments in computers and networking have made systems, such as production systems, more vulnerable to attacks having both cyber and physical components; i.e., to cyber-physical attacks. In manufacturing, such attacks are not only capable of stealing valuable information, but can also destroy equipment, force physical product changes, alter product's mechanical characteristics, or endanger human lives.

Typically, the manufacturing industry have relied on various Quality Control (QC) tools, such as product inspection, to detect the effects caused by these attacks. However, these attacks could be still designed in a way to avoid detection by traditional QC methods, which suggests a need for new and more effective QC tools. Such new tools should be able to prevent, or at least minimize, the effects of these attacks in manufacturing. Unfortunately, almost no research has been done on using QC tools for securing production systems against these malicious attacks.

Hence, the overarching goal of this work is to allow QC systems to be designed in a more effective manner to act as a second line of defense, when traditional cyber-security measures and attackers have already accessed the production system. To this end, this work focuses on: 1) understanding the role of QC systems during the attack; 2) identifying existing weaknesses in QC tools and determining the effects of exploiting them by the attack; and 3) proposing more effective QC tools, for better protection against these types of cyber-physical attacks in manufacturing.

ACKNOWLEDGMENTS

The research work presented here was funded by the National Science Foundation (NSF) Civil, Mechanical and Manufacturing Innovation (CMMI) grant number 1436365 – *GOALI: Robust Quality Control Tools for Cyber-Physical Manufacturing Systems: Assessing and Eliminating Cyber-Attack Vulnerabilities* and supported by Virginia Tech's Cyber-Physical Security Systems Manufacturing (CPSSMFG) Group. However, any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

In addition, I wanted to acknowledge the efforts of all the co-authors I have collaborated with on publications outside my dissertation work while being enrolled in the PhD program at the Grado Department of Industrial and Systems Engineering: Ms. Romina Dastoorian, Dr. Sherif Abdelhamid (Abdelhamid *et al.*, 2015; Elhabashy *et al.*, 2015a), Dr. Zach DeSmit (DeSmit *et al.*, 2016; 2017), Dr. Yao Pan (Pan *et al.*, 2017), Dr. Mathew Keefee (Keefe *et al.*, 2017), and Dr. Wenmeng Tian (Dastoorian *et al.*, 2018).

I would also like to thank my advisor, for offering me the opportunity to earn a PhD degree from Virginia Tech; my co-advisor for his persistent guidance and continuous constructive feedback; my remaining committee members, for their helpful advice and valuable comments; and all my colleagues in the Center of Innovation-based Manufacturing (CIbM). Finally, I would like to thank my family members, who have supported me throughout this lengthy journey; I couldn't have done it without any of them.

Contents

ABSTRACT ii

GENERAL AUDIENCE ABSTRACT iii

ACKNOWLEDGMENTS iv

Contents v

List of Figures x

List of Tables xii

List of Abbrevations xvi

1 Introduction 1

1.1 Significance 2

1.2 Motivation 3

1.3 Research Objectives 5

1.4 Outline 6

2 Related Work 8

2.1 Difference from Traditional IT Security 8

2.1.1 Cyber-Physical Systems 8

2.1.2 Cyber-Physical Production Systems 9

2.2 Cyber-Physical Security Efforts in Production Systems 10

2.2.1 Attack Case-studies 11

2.2.2 Novel Attack Defense Techniques 13

2.2.3 Frameworks and Classifications 17

2.2.4 Other Related Efforts 19

2.2.5 Closely Related Personal Efforts 19

2.3 Quality Control for Cyber-Security 22

2.4 Existing Gaps and Proposed Research Areas 23

3 A Cyber-Physical Attack Taxonomy for Production Systems: A Quality Control Perspective 25

3.1 Background and Motivation 25

3.1.1 Cyber-physical Attacks in Manufacturing Environments 25

3.1.2 The Need for New QC Tools 27

3.2 Cyber-Physical QC Attack Taxonomy for Manufacturing 30

3.2.1 Attack Objectives 31

3.2.2 Affected Production Sub-system 34

3.2.3 Attack Method and Attack Location 36

3.2.4 Taxonomy Overview 38

3.3 Taxonomy Implementation Example 39

3.3.1 Attack Objective: Disrupting Manufacturing Operations 40

3.3.2 Attack Objective: Altering Product Design Intent 42

3.3.3 Attack Objective: Reducing Outgoing Quality 44

3.4 System-Specific Attack Design Considerations 47

3.4.1 Component Accessibility 47

3.4.2 Attack Pervasiveness 48

3.4.3 Attack Impact 48

3.4.4 Attack Coordination 48

3.5 Summary and Discussion 49

4 Cyber-Physical Attack Vulnerabilities in Manufacturing Quality Control Tools 51

4.1 Introduction 51

4.1.1 Background 51

4.1.2 Motivation 52

4.2 Overview of the Exploitation Classes 54

4.3 Exploiting Bad Decisions 56

4.3.1 Misusing QC Tools Categories 56

4.3.2 The Effects of Cyber-Physical Attacks on the Categories of QC Tools Misuse 60

4.4 Making Decisions Bad 70

4.4.1 Making Decisions Bad Levels 70

4.4.2 The Effects of Cyber-Physical Attacks on Class 2 Exploitation Levels........ 71

4.5 Potential Mitigation Guidelines 74

4.5.1 Applying QC Best Practices 74

4.5.2 Selecting Alternate KQCs........ 74

4.5.3 Developing More Resilient QC Tools 75

4.6 Summary 75

5 Introducing Randomness into Control Chart Sampling Plans for Better Security against Cyber-Physical Attacks in Production Systems – Univariate Case 77

5.1 Background and Motivation 77

5.2 Methodology 79

5.2.1 Sampling Plans 81

5.2.2 Attack Scenarios 82

5.3 Assessment........ 82

5.3.1 The Used Performance Metrics 83

5.3.2 Performance Metrics Evaluation 84

5.4 Results and Discussion 86

5.4.1 For Each Attack Scenario 86

5.4.2 Across the Attack Scenarios 91

5.4.3 Conclusions........ 92

5.5 Summary 93

6 Introducing Randomness into Control Chart Sampling Plans for Better Security against Cyber-Physical Attacks in Production Systems – Multivariate Case 94

6.1 Background and Motivation 94

6.2 Developed Approach 96

6.2.1 Differences from other Approaches 96

6.2.2 Sampling Plans 97

6.2.3 Attack Details 99

6.3 Approach Evaluation 99

6.3.1 Evaluation Criteria 99

6.3.2 Naïve Variable Selection Approach 104

6.4 Results and Discussion 106

6.4.1 Sampling Plan 1 - Fixed-sized Subsets 106

6.4.2 Sampling Plan 2 - Variable-sized Subsets 107

6.4.3 Discussion 108

6.5 Summary 108

7 Contributions and Future Work 110

7.1 Contributions 110

7.1.1 Overall Contributions 110

7.1.2 Cyber-physical Attack Taxonomy from a QC Perspective 111

7.1.3 Cyber-Physical Attack Vulnerabilities in QC Tools 111

7.1.4 Introducing Randomness to Control Charts 112

7.2 Future Work 113

7.2.1 Introducing Randomness to Univariate Control Charts 113

7.2.2 Introducing Randomness to Multivariate Control Charts 114

7.2.3 Improving QC Tools to be More Resilient 114

References 115

Appendix A *Journal of Intelligent Manufacturing* Copyright Permission 128

Appendix B Theoretical Evaluation of Some of the Performance Metrics – Individual Control Charts …… 131

Average Run Length (ARL) 131

Using the Probabilities of Getting an Out-of-Control Signal within a Group of *N* products 131

Using Discrete Time Markov Chains (DTMCs) 136

Average Time to Signal (ATS) 141

True Positive and False Positive Probabilities 141

Number of Successfully Attacked Products until Signal Metrics 142

Percentage of Attacked Products 143

Detection Probability 143

Out-of-Control Metrics 143

Appendix C Comparison of the Metrics Theoretical and Simulated Values – Individual Control Charts 145

Appendix D Details of the Economic Cost Model Proposed by Lorenzen and Vance (1986) 148

Model Assumptions 148

Quality Cycle Time 148

Quality Costs 149

Related Error Probabilities 150

Appendix E ARL Values Obtained from Simulation for Both of the Proposed Random Sampling Plans – Individual T^2 Control Charts 152

Appendix F Cost Values Obtained from the Simplified Cost Model – Individual T^2 Control Charts 156

List of Figures

Figure 1-1: Inspection outline for a) intended product and b) resulting attacked product. 4

Figure 3-1: Cyber-physical attacks objectives in manufacturing. 33

Figure 3-2: Overview of the production system components (adapted from Groover (2010)). ... 34

Figure 3-3: Summary of affected systems within each attack type on production systems. 36

Figure 3-4: Overview of the proposed cyber-physical QC attack taxonomy, along with a summary of the details of each layer. 38

Figure 3-5: Implementation example part.. 39

Figure 3-6: Production system example flow diagram 39

Figure 4-1: The resulting X-bar chart without (a) and with (b) a cyber-physical attack, as an example of the effects of improperly using QC tools. Samples affected by the cyber-physical attack are designated with a magenta asterisk. 63

Figure 4-2: A corresponding R chart that would have been able to detect the effects of the cyber-physical attack in this case. Out-of-control points are designated with a cyan circle. 64

Figure 4-3: Individual (I) control charts for product's KQC during Phase I for a) independent samples without a cyber-physical attack and b) negatively correlated samples as a result of a cyber-physical attack; as an example of an attack taking advantage of the second misuse category. Red dashed lines represent the original control chart limits without the attack, red solid lines represent the control limits with the attack, whereas green dotted dashed lines represent the charts center lines. 66

Figure 4-4: The I-chart during Phase II resulting from the cyber-physical attack for the autocorrelated samples, where the attacked samples outside the original limits are designated with cyan circles. Red dashed lines represent the original control chart limits when there was no attack, red solid lines represent the control limits with the attack, whereas green dotted dashed lines represent the new chart center lines. 67

Figure 4-5: The resulting X-bar control chart with both limits for misuse example 3. Red solid lines represent the product's specification limits, red dashed lines represent the charts control limits, whereas green dotted dashed lines represent the chart's center line........................... 68

Figure 4-6: X-bar control charts with a cyber-physical attack using specification (a) and control (b) limits. Red solid lines represent the product's specification limits, red dashed lines represent the charts control limits, whereas green dotted dashed lines represent the charts center lines. Samples outside the control limits are designated with cyan circles. 69

Figure 4-7: A manufactured product without an attack (a) and with a Passive cyber-physical attack (b), as an example of a Class 2 - Level 1 exploitation. 72

Figure 4-8: The manufactured bracket without an attack (a) and with a Passive cyber-physical attack adding an extra feature (b), as an example of a Class 2 - Level 2 exploitation........... 73

Figure 4-9: An example of a traditional data inspection decision of collecting data every fixed number (3) of products. The hollow products are the ones being inspected, whereas the solid products (also marked with an “x”) are not inspected. 73

Figure 5-1: An example of a traditional inspection plan, where sampling occurs every 5 products. The products that are sampled for inspection are colored in green and designated with an “I”, whereas the red products are the ones that are not inspected and could be affected by an intelligently designed Product-Oriented C2P attack. 78

Figure 5-2: An example of sampling Plan 1, where the selected products for inspection, colored in green and designated with an “I”, are randomly chosen within each group of N = 5 products. 81

Figure 6-1: Examples of fixed-sized random subsets used in the sampling plan of multivariate T^2 control charts, instead of monitoring all the variables as shown on the far left. 98

Figure 6-2: Examples of variable-sized random subsets used in the sampling plan of multivariate T^2 control charts, instead of monitoring all the variables as shown on the far left................ 98

List of Tables

Table 3-1: Examples of affected production systems. 35

Table 3-2 Summary of taxonomy details for the implementation example. 46

Table 4-1: Summary of the exploitation classes characteristics discussed in this chapter. 56

Table 4-2: A numerical example to show the difference between physically altered KQC values for one sample during an attack and the values they would have taken in case of no attack implemented. 62

Table 4-3: Summary of bad decisions within the different misuse category examples. 69

Table 4-4: Class 2 exploitation levels summary. 71

Table 5-1: Performance metrics values for the 1st attack scenario across the different sampling plans obtained from the simulation procedure with 1,000,000 iterations. 87

Table 5-2: Performance metrics values for the 2nd attack scenario across the different sampling plans obtained from the simulation procedure with 1,000,000 iterations. 89

Table 5-3: ARL metric values for both attack scenario across the proposed random sampling plans. 91

Table 5-4: PAOC (%) metric values for both attack scenario across the proposed random sampling plans. 92

Table 5-5: DP (%) metric values for both attack scenario across the proposed random sampling plans. 92

Table 6-1: Recommended Phase I sample size (m) for T^2 control charts with individual observations, according to Chen and Pan (2011). 97

Table 6-2: Cost savings (%) resulting from both approaches, when p = 10 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 1. 106

Table 6-3: Cost savings (%) resulting from both approaches, when p = 100 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 1. 107

Table 6-4: Cost savings (%) resulting from both approaches, when p = 10 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 2. 107

Table 6-5: Cost savings (%) resulting from both approaches, when p = 100 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 2. 108

Table B-1: Probabilities of getting an out-of-control signal and true alarms, along with the ARL values, corresponding to different transient shift sizes caused by a cyber-physical attack for the fixed attack scenario – Plan 1. 133

Table B-2: Possibilities for the next part to be inspected after a specific part has been currently chosen for inspection within sampling Plan 0b. An “I” indicates that the part could be inspected within groups of size 10. 137

Table B-3: Description of the states within the used DTMCs. 138

Table B-4: ARL and ATS values, along with TPP and FPP values (as percentages), corresponding to different transient shift sizes caused by an attack, for the fixed attack scenario – Plan 1. 142

Table C-1: Comparison of the ARL metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the current sampling plans. 145

Table C-2: Comparison of the ARL metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans. 145

Table C-3: Comparison of the ATS metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans. 146

Table C-4: Comparison of the TPP metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans. 146

Table C-5: Comparison of the FPP metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans. 147

Table E-1: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values. 152

Table E-2: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values. 153

Table E-3: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values. 153

Table E-4: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values. 153

Table E-5: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values. 154

Table E-6: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values. 154

Table E-7: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values. 155

Table E-8: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values. 155

Table F-1: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 156

Table F-2: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 156

Table F-3: Cost values obtained (in $) for both approaches, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 157

Table F-4: Cost values obtained (in $) for both approaches, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 157

Table F-5: Cost values obtained (in $) for both approaches, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 157

Table F-6: Cost values obtained (in $) for both approaches, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 158

Table F-7: Cost values obtained (in $) for both approaches, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 158

Table F-8: Cost values obtained (in $) for both approaches, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values, using sampling Plan 1. 158

Table F-9: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 159

Table F-10: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 159

Table F-11: Cost values (in $) obtained for both approaches, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 159

Table F-12: Cost values (in $) obtained for both approaches, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 160

Table F-13: Cost values (in $) obtained for both approaches, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 160

Table F-14: Cost values (in $) obtained for both approaches, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 160

Table F-15: Cost values (in $) obtained for both approaches, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 161

Table F-16: Cost values (in $) obtained for both approaches, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values, using sampling Plan 2. 161

List of Abbrevations

ARL	Average Time to Signal
ATS	Average Run Length
C2P	Cyber-to-Physical
CAD	Computer-Aided Design
CAE	Computer-Aided Engineering
CAM	Computer-Aided Manufacturing
CMM	Coordinate Measuring Machine
CNC	Computer Numeric Control
CPS	Cyber-Physical System
CUSUM	Cumulative Sum
DM	Direct Manufacturing
DP	Detection Probability
DQ	Direct Quality
EWMA	Exponentially Weighted Moving Average
FDM	Fused Deposition Modeling
FEA	Failure Element Analysis
FMEA	Failure Mode and Effects Analysis
FPP	False Positive Probability
GD&T	Geometric Dimensioning and Tolerancing
IC	Integrated Circuit
ICS	Industrial Control System
ICS-CERT	Industrial Control Systems Cyber Emergency Response Team
IoT	Internet of Things
IIoT	Industrial Internet of Things
IP	Intellectual Property
IT	Information Technology
KPC	Key Performance Characteristic
KQC	Key Quality Characteristic
KSC	Key Security Characteristics

LCL	Lower Control Limit
LSL	Lower Specification Limit
NDIA	National Defense Industrial Association
NSAS	Number of Successful Attacks until Signal
OEM	Original Equipment Manufacturer
PA	Percentage Attacked
PAOC	Percentage Attacked Out-of-Control
POC	Percentage Out-of-Control
PSAS	Percentage of Successful Attacks until Signal
PZTs	Piezoelectric Transducers
QC	Quality Control
QDMS	Quality Data Management System
RCA	Root Cause Analysis
RL	Run Length
SaaS	Software as a Service
SCADA	Supervisory Control and Data Acquisition
SPC	Statistical Process Control
TPP	True Positive Probability
UCL	Upper Control Limit
USL	Upper Specification Limit

1 Introduction

With the latest advancements in networking and internet technology, cyber-attacks have been an increasing phenomenon. Recent examples of cyber-attacks included, but were not limited to, the data breaches of large retailers, such as Target in December of 2013 (Target, 2014) and Home Depot in mid-2014 (Kurtzman Carson Consultants, 2016), the hacking of the corporate network of Sony Pictures Entertainment in November of 2014 (Lee, 2014), acquiring individuals' private information from Anthem Health Insurance starting from early December of 2014 (Anthem, 2016), the multiple Yahoo data breaches reported in 2016 (Fahey and Wells, 2016; Castillo, 2017), and the "WannaCry" ransomware virus that affected thousands of computers worldwide in May of 2017 (Kessem, 2017; Larson, 2017).

In addition to these instances of traditional cyber-attacks, other incidents also involved cyber-attacks on physical systems, i.e., cyber-physical attacks. As the name implies, cyber-physical attacks are those whose influencing elements are in the cyber domain, but the victim elements are in the physical domain (Yampolskiy *et al.*, 2013). The key difference between cyber-*physical* attacks and traditional cyber-attacks is that the latter can be viewed as attacks where *both* the influencing and victim elements are in the cyber domain (Yampolskiy *et al.*, 2013).

Cyber-physical attacks are not necessarily new but are typically less publicized. For instance, a "logic bomb" was reportedly inserted in the Trans-Siberian pipeline's control software to abnormally change the pumps and valves setting, causing a massive explosion in 1982 (Rost and Glass, 2011). Perhaps the most famous cyber-physical attack was the "Stuxnet" virus between late 2009 and early 2010. This virus was allegedly responsible for destroying about 1,000 Iranian high-speed centrifuges used for Uranium enrichment, through periodically changing the rotational speeds of the centrifuges, significantly shortening their life spans (Albright *et al.*, 2010). The attack's success was only due to it being able to exploit the system and display misleading equipment readings to operators (Cherry, 2011).

A more recent instance of a cyber-physical attack was a computer virus attacking a turbine control system at a U.S. power company during the fall of 2012, causing the plant to go offline for three weeks (Finkle, 2013). A spear-phishing attack on a German steel mill in 2014 was another case of a cyber-physical attack, where attackers were able to eventually gain access into the plant's network, causing multiple system components failure and massive physical damage (Lee *et al.*, 2014). Another example also involved a computer virus that affected a German nuclear plant in 2016, but without posing a threat (Steitz, 2016). Furthermore in 2016, there was an attack on a power grid in Ukraine, depriving over 100,000 people of power (Tuptuk and Hailes, 2016). Finally, in May of 2017, the “WannaCry” ransomware virus managed to cause production to stop in several automotive factories in Europe as well (Liptak, 2017). Given that these examples expand across a variety of fields, they demonstrate that no system is beyond reach of cyber-physical attacks, with production systems being no exception.

1.1 Significance

John Chambers, Cisco’s executive chairman and former CEO, stated that “I am convinced that there are only two types of companies: those who *have been hacked*, and those who *don’t yet know they have been hacked*” (Chambers, 2015). Such a statement should be alarming for any industry because given enough time, computational resources, and luck, cyber-attackers will find their way into almost any system. Hence, despite the best network security possible, cyber-physical production systems should not be considered 100% secured from any type of cyber-attacks.

In fact, manufacturing was the most targeted sector in 2014 by spear-phishing attacks, accounting for 20% of the total attacks with an attack odds of 1 in 3 (Symantec, 2015). Furthermore, the critical manufacturing sector alone accounted for the majority (33%) of all security incidents reported to the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) during the fiscal year 2015 (ICS-CERT, 2016). Cyber-attacks of this nature on production systems are usually after gaining unauthorized access to users’ information or valuable trade secrets (Deloitte, 2014; Symantec, 2014; 2015; 2016; 2017). Gaining such information is typically accomplished through spear-phishing attacks, which are targeted e-mail scams to access sensitive data, steal valuable information, or install malware on compromised computers (Kaspersky, 2015).

Cyber-attacks in manufacturing could also go beyond accessing the system for the sake of intellectual property theft and could alternatively affect the physical component itself within the production environment. An example of such cyber-physical attacks is the Aurora project conducted at the Department of Energy's Idaho lab in 2007, where a diesel generator’s circuit breakers were remotely manipulated with, causing eventual equipment breakdown (Meserve, 2007; Brenner, 2013; Tucker, 2014). Examples of other potential cyber-physical attacks in manufacturing include altering product designs or modifying manufacturing processes (Wells *et* *al.*, 2014). One way an attack altering product design could be accomplished is through altering Computer Aided Design (CAD) files, whereas attacks affecting manufacturing processes could be implemented through changing the used tools and equipment or the manufacturing parameters.

If cyber-physical attacks within manufacturing go undetected, their effects could be devastating. The undetected attacks could negatively affect a product's design intent, performance, quality, or perceived quality, posing a risk to human safety as operators and consumers could be using faulty products (Wells *et al.*, 2014). Such effects could cause products to fail prematurely during use, which would be catastrophic for critical product components such as car brakes, jet engines, or aircraft turbine blades. The financial consequences would then be overwhelming, potentially causing delayed products' launch, ruined equipment, increased warranty costs, avoidable injury settlements, or eventually loss of customer trust (Wells *et al.*, 2014). Hence, such cyber-physical attack types in manufacturing are the focus of this work.

1.2 Motivation

One way to ensure that the effects of cyber-physical attacks are detected is to rely on traditional inspection and monitoring techniques which are part of the Quality Control (QC) system within manufacturing. However, although there exists a variety of QC tools available for use, they could be ineffective in detecting the effects of cyber-physical attacks against production systems quickly enough or at all. This inefficiency is mainly because current QC methods were not designed to detect these types of malicious attacks (Wells *et al.*, 2014; Vincent *et al.*, 2015). Specifically, QC tools are used to detect quality losses and ensure manufacturing processes stability, but are based on specific assumptions that may no longer be valid during cyber-physical attacks (Wells *et al.*, 2014; Vincent *et al.*, 2015). Cyber-attackers can take advantage of such specific assumptions about QC tools to design their attacks to not be easily detectable, preventing proper product quality assessment (Wells *et al.*, 2014).

For instance, consider the bracket to be manufactured in Figure 1-1a, which has 10 mounting holes. The locations and sizes of these holes are inspected using a Coordinate Measuring Machine (CMM). An attack could be as simple as adding an 11th hole to the product, as in Figure 1-1b. Since this extra hole never should have existed, there is no way for the CMM to notice it, making the attack undetectable. This QC system has been compromised because the assumptions that the CMM is the best tool to check for the holes' sizes and location and that such an assessment is enough to confirm the product's overall quality have been invalidated by the cyber-physical attack.

Although one could argue that such an attack is rather obvious, this may not necessarily be always true, especially if the product is mass-produced in an automated environment, the product's complexity is increased, or the product's size is significantly decreased.

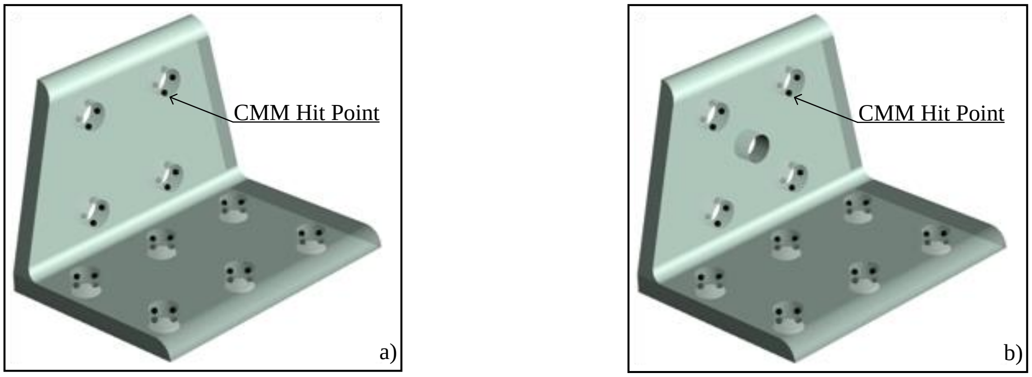

Figure 1-1: Inspection outline for a) intended product and b) resulting attacked product.

In addition to cyber-physical attacks being potentially difficult to detect and being able to compromise QC tools to reduce detectability even more, a considerable period of time might elapse between the start of those attacks and the time its existence is realized (if at all). It has been reported that the median number of days between the onset of a cyber-attack and its detection in organizations was over 200 days in 2013 and 2014, with the longest presence being 2,982 days (Mandiant, 2015). Unlike traditional cyber-attacks resulting in data breaches, where attackers would collect all the required data within a relatively short period of time (probably just in the first few days), not being able to detect the attack for that long would be equivalent to at least 200 days (a little over half a year) of unknowingly producing defective products in a manufacturing setting. What makes matters worse is that 69% of the attacks in 2014 were not discovered by the victims themselves, but by third parties such as law enforcement (Mandiant, 2015). Meaning that, for the manufacturing industry, the effects of a substantial portion of the attacks would not be detected during testing within the production facilities but could be discovered when products are in-use by customers.

Since cyber-physical attacks in manufacturing will mostly lead to tangible physical changes in the appearance and structure of the products, there are opportunities for detecting cyber-physical attacks efficiently (Cárdenas *et al.*, 2009; 2011), i.e., faster and with higher probability. With QC systems being heavily relied on in manufacturing, inspection and monitoring techniques should be designed to detect the physical changes resulting from cyber-physical attacks more frequently and in a timely fashion. In other words, current QC systems needs to be improved for better cyber-physical security in manufacturing environments.

Little to no research has been done on using QC tools for cyber-physical security of production systems[1]. On one hand, research work in the field of cyber-physical security in manufacturing systems have not provided any QC-related solutions nor ones that can significantly reduce the occurrence of cyber-physical attacks in the first place. Still being in its infancy, the research efforts in the field of cyber-physical security in manufacturing is limited to discussing case-studies demonstrating cyber-physical attacks for varying manufacturing processes, suggesting alternative potential defense approaches, or proposing frameworks and classifications. On the other hand, research work in the field of QC has not tackled the problem of cyber-physical security and only addressed some issues within cyber-security in general. Most of the work that involved the use of QC tools is typically related to network intrusion detection such as in anomalous network behavior detection or other network-related applications.

1.3 Research Objectives

From the discussion in the previous section, it has become clear that there are opportunities for incorporating QC-related approaches within the field of cyber-physical security for production systems. Given the fact that QC systems can already detect the physical changes within products' physical attributes caused by cyber-physical attacks, more effective QC tools are needed for improved cyber-physical security. QC systems would then have a better chance against the usually hidden malicious cyber-physical attacks on production systems. Hence, the overarching aim of this work is to allow QC systems to be designed and used effectively as a second line of defense, when traditional cyber-security techniques fail and the production system is already breached. Such an overall aim will be accomplished through fulfilling three specific objectives:

1) **Understanding the role of QC systems in cyber-physical attacks within manufacturing:** This objective is concerned with understanding the relationships between production systems, QC systems, and cyber-physical attacks in the context of malicious process changes. Doing so is the first step towards an enhanced production system, as system defenders would gain critical insights on the potential attack details (attack goals, types, and implementation methods) and weaknesses in their systems.

[1] Further literature review details supporting this statement will be discussed in more length throughout Chapter 2.

2) **Exploring the effects of exploiting QC tools weaknesses by cyber-physical attacks within manufacturing:** The next step is concerned with investigating the effects of exploiting QC-related cyber-physical security weaknesses existing within the production systems. The focus here is going to be more on identifying these weaknesses, which occur during the use of specific QC tools. Through illustrating the different possible QC tools exploitation classes, it would be possible to assess the resulting consequences on the production system under cyber-physical attacks.

3) **Proposing more effective QC tools:** By extending traditional cyber-security approaches from the Information Technology (IT) domain, the final objective is to propose improvements to existing QC tools. This will be achieved in a manner to prevent attackers from exploiting the identified QC-related cyber-physical security weaknesses in production systems, without significantly reducing the tools' performances or increasing associated costs.

1.4 Outline

This document is composed of seven chapters, which are organized as follows. A comprehensive review of the related work in the fields of cyber-physical security in manufacturing and quality control for cyber-security is discussed in the next chapter, along with the existing gaps that needs to be addressed. In particular, sub-section 2.2.5 within Chapter 2 contains a summary of a few efforts in the field of cyber-physical security for production system, to which I have contributed, but are beyond the scope here. While the efforts within the scope of this work to address each of the aforementioned research objectives are discussed in much more details within the following four chapters, respectively.

In Chapter 3, a taxonomy to better understand the relationships between QC systems, manufacturing systems, and cyber-physical attacks is presented. Content from this chapter has already been published as a research article in the *Journal of Intelligent Manufacturing*. The existing weaknesses in current QC tools, such as the tools misuse, and the effects of exploiting them by cyber-physical attacks within manufacturing are examined in Chapter 4. The outcome of the work in this chapter will be submitted to the *Quality and Reliability Engineering International*. The details of more effective QC tools, through including randomness, are discussed next across Chapters 5 and 6. The fifth chapter will be prepared for submission to the *Journal of Quality*

Technology. Finally, Chapter 7 contains a summary of the expected contributions of this work and ideas for extending it further.

2 Related Work

For several decades, the CIA-triad of confidentiality, integrity, and availability has been serving as the main conceptual model for computer and Information Technology (IT) security (Whitman and Mattord, 2012). The CIA-triad presents security concerns more related to data, and although it can be extended to manufacturing (Pan *et al.*, 2017), it is important to recognize that simply applying IT-based security measures to cyber-physical production systems is not enough (Cárdenas *et al.*, 2009; 2011).

Hence, this section starts with an introduction about how exactly a Cyber-Physical System (CPS) is different from an IT system and how production systems have some additional unique characteristics themselves, making it not suitable to solely apply traditional IT security measures. Then, related cyber-physical security efforts more specific to production systems are discussed next. Lastly, current efforts using Quality Control (QC) tools for cyber-security are presented, along with opportunities for extending the use of those tools to provide enhanced cyber-physical security in manufacturing.

2.1 Difference from Traditional IT Security

2.1.1 Cyber-Physical Systems

The main differences between IT and cyber-physical systems security in general are highlighted by Cárdenas *et al.* (2009); Cárdenas *et al.* (2011) and the Manufacturing and Cyber divisions of the National Defense Industrial Association (NDIA, 2014). Although the authors use Industrial Control Systems (ICSs) as representatives of cyber-physical systems, these differences, which are discussed next, are still true for almost any other cyber-physical system, such as production systems.

The first difference is that patching and continuous updates, which are common practices in IT systems security, are not well suited for CPS (Cárdenas *et al.*, 2009; 2011; NDIA, 2014). Due to high equipment and downtime costs, frequent updates may not be possible in CPSs, since updates would require getting certain equipment offline and not being able to perform its function for a certain amount of time. Real-time availability is of utmost importance in CPSs (Cárdenas *et al.*, 2009; 2011; NDIA, 2014; Wu *et al.*, 2017). Furthermore, it may take months just to plan for an upgrade and could be difficult to financially justify on regular basis (Cárdenas *et al.*, 2009; 2011; NDIA, 2014; Pan *et al.*, 2017).

Another difference is the presence of a large number of legacy equipment in CPSs when compared to IT systems (Cárdenas *et al.*, 2009; 2011). The issue with legacy equipment is that it is, as a result of age, less secure than the more advanced recent equipment. Due to being old and not able to necessarily get frequent updates, this equipment would contain more weaknesses, which could be exploited more easily during cyber-physical attacks. In addition, much of this equipment have a diverse range of operating systems (NDIA, 2014). Such a wide range of operating systems not only makes it hard for different equipment to communicate together, but it also makes the equipment more difficult to update and consequently more vulnerable.

As previously discussed, perhaps the most obvious and significant difference here is the "physical" aspect of CPSs, i.e., their interaction with the physical world (Cárdenas *et al.*, 2009; 2011). With the focus of computer and IT security on information protection, it is not enough to rely on computer and IT systems security measures as they do not consider how cyber-attacks would affect the physical world; i.e., IT security is necessary but not sufficient for CPS security (Cárdenas *et al.*, 2009; 2011).

With such a physical aspect, there are now human safety concerns associated with the security of CPSs (NDIA, 2014), including operators, personnel, and customers' safety. This physical aspect also means that there are additional intrusion detection options in CPS. Traditional firewalls may prevent intrusions from outside the system, but not attacks from inside the system, such as those by disgruntled employees (Slay and Miller, 2008; Kravets, 2009; Poulsen, 2009; Brenner, 2013).

2.1.2 Cyber-Physical Production Systems

The nature of production systems also plays a role in the inefficiency of just applying regular IT security measures to achieve cyber-physical security. Current manufacturing enterprises consist of a large number of heterogeneous components with a wider range of security requirements such as controllers, machining equipment, assembly equipment, post-processing equipment, and inspection equipment. Furthermore, cyber-physical production systems have a more complex IT structure. Those diverse technological components communicate together through a wide variety of architectures, protocols, and network technologies that are not necessarily available in regular IT systems. For instance, MTConnect is a new emerging communication protocol between manufacturing equipment (MTConnect Institute, 2016).

Additionally, manufacturing has currently entered the Industry 4.0 era with more emphasis on easier automation and data exchange. Along with the Internet of Things (IoT), more manufacturing equipment (other than computers) is now connected to the Internet such as sensors, cameras, etc. Industry 4.0 has also resulted in even more connectivity to easily extract and share data with cloud computing and new communication protocols. This increased connectivity has led to a growing reliance on creating almost everything within a digital environment using Computer-Aided Engineering (CAE) support tools, such as Computer-Aided Design (CAD) and Manufacturing (CAM) software, resulting in more security concerns.

All of these aspects show that cyber-physical production systems have additional security requirements. Therefore, for securing cyber-physical production systems, more tools are needed along with the available traditional IT cyber-security tools, ones accounting for the physical aspect in particular. A review of the current research efforts in the field of cyber-physical security in production systems is presented next.

2.2 Cyber-Physical Security Efforts in Production Systems

The focus of current cyber-physical security research efforts for production systems has been traditionally on security issues regarding ICSs and Supervisory Control and Data Acquisition (SCADA) networks (Cárdenas *et al.*, 2009; Huang *et al.*, 2009; Cárdenas *et al.*, 2011; Huang *et al.*, 2015). As a result, the security of production systems has been grouped into the generic area of critical infrastructure (Stamp, 2003; Giraldo *et al.*, 2017) such as electric power grids, water and waste management systems, and transportation systems (Wells *et al.*, 2014). However, production systems are becoming more sophisticated and, in addition, control systems are just one component within production systems and a cyber-physical attack can occur anywhere within the manufacturing enterprise and the corresponding supply chains. In fact, ICSs and manufacturing could be viewed, despite potentially overlapping, as two separate domains (Giraldo *et al.*, 2017), having different objectives and implementation techniques (Yampolskiy *et al.*, 2017). Hence, although the security of ICSs could encompass manufacturing processes (Giraldo *et al.*, 2017), the review here is more concerned with the security of manufacturing processes alone, rather than focusing on the control systems included.

As previously mentioned, research in the field of cyber-physical security of production systems is relatively new, but has been expanding quickly. The majority of the current relevant research efforts in this specific field can be broken down into three main categories: 1) cyber-physical attack case-studies, 2) novel cyber-physical attack defense techniques, and 3) cyber-physical attack frameworks and classifications.

2.2.1 Attack Case-studies

The first area within the field of cyber-physical security for manufacturing includes research efforts showcasing different small-scale cyber-physical attack examples and their effects on the manufactured products. Researchers in this area typically aim to also raise awareness about the issue of cyber-physical security in manufacturing through demonstrating the relative ease of applying a cyber-physical attack on a specific manufacturing process.

As an example, Wells *et al.* (2014) discussed some of the cyber-security related weaknesses existing in production systems before presenting a case-study for a subtractive manufacturing process. In the case-study, the manufacturing of a tensile test specimen on a Computer Numerical Control (CNC) milling machine was attacked through altering the tool path files as part of an undergraduate student project. The purpose of the case-study was not only to demonstrate the attack feasibility, but also to assess the diagnostic abilities of future engineers (Wells *et al.*, 2014). None of the student groups involved were able to identify the existence of an attack, with three out of seven students groups not even bothering to measure the final product since it "looked correct" (Wells *et al.*, 2014). Turner *et al.* (2015) went into more detail about this case-study and echoed the concerns by Wells *et al.* (2014) about the lack of enough awareness regarding cyber-physical security. In addition to discussing this case study, Turner *et al.* (2015) also analyzed some potential attack surfaces[2] within manufacturing such as the design tool chain, control, and direct equimpent attack surfaces, among others.

In another publication, Zeltmann *et al.* (2016) presented an example of a cyber-physical attack on an additive manufacturing process. The authors first started by providing an overview of potential risks existing in the field of additive manufacturing before presenting their attack case-study. Also using tensile test specimens, they investigated the effects of two types of changes in additive manufacturing; namely, embedding internal defects and altering the printing orientation (Zeltmann *et al.*, 2016). Two types of non-destructive testing were then used to demonstrate the decreased performance and evaluate attack detectability. Specifically, ultrasonic inspection and Finite Element Analysis (FEA) techniques were used for the cyber-physical attacks detection and to evaluate their effects, respectively (Zeltmann *et al.*, 2016). The results showed that both types

[2] A system's attack surface is a term that represents all possible means through which an attacker can access the system and potentially cause harm (Manadhata and Wing, 2011).

of attacks were not easily detectable and would have a negative impact on the material's performance.

In the same setting of additive manufacturing, Belikovetsky *et al.* (2016) presented an interesting case-study with the aim of sabotaging a manufactured functional part. In doing so, the authors also proposed an approach to identify attack opportunities[3], analyzed the attack's full chain, and developed a methodology to assess attack difficulty in additive manufacturing (Belikovetsky *et al.*, 2016). In the demonstrated attack, a 3D printed propeller of the quadcopter was compromised remotely through maliciously changing its design file, causing the quadcopter to collapse while flying (Belikovetsky *et al.*, 2016). To make the attack successful, the authors had to design it in such a way that the change was not noticeable and, at the same time, the product does not fail right away during operation, requiring them to experiment with different design iterations first (Belikovetsky *et al.*, 2016). In addition to having to do some trial-and-error experimentation, this attack also required at least some type of basic knowledge with the product and process.

Moore *et al.* (2017b) illustrated another possibility for a cyber-physical attack in an additive manufacturing process through tampering with the 3D printer's firmware and showing the ensuing negative effects. The authors took advantage of the firmware having an open source code and were able to send two maliciously attacked versions of the code to the printer via USB communication ports (Moore *et al.*, 2017b). One version affected the printer's control flow, resulting in a different product; whereas the other increased the extruder's feed rates by increments of 10%, up to 40% (Moore *et al.*, 2017b). Although both types of those cyber-physical attacks were successful, they were rather easy to detect (even visually) since the first attack changed the entire shape of the final product, while the other changed its size considerably.

Sturm *et al.* (2014, 2017) pointed out existing cyber-related weaknesses in the additive manufacturing process and discussed a variety of potential cyber-physical attacks on a 3D printed tensile test specimen. Then, they showed how an attack of adding a void could be accomplished through changing the associated design files and its effects on the material's properties such as strength and loading conditions (Sturm *et al.*, 2014; 2017). In a manner similar to the work of Wells *et al.* (2014), five students groups unknowingly participated in this attack experiment and

[3] This was based on the cross-domain attack classification for CPS that one of the authors has previously created (Yampolskiy *et al.*, 2013) and further developed in sub-sequent publications.

none of them were able to detect the occurrence of a cyber-physical attack, despite destructively testing the manufactured part (Sturm *et al.*, 2017). Finally, Sturm *et al.* (2014); (2017) concluded their work with some recommendations for protection against similar attacks. It should be noted that the successful implementation of the case-studies by Wells *et al.* (2014) and Sturm *et al.* (2017) and the fact that the students were not able to identify them demonstrates the relative ease of attacking the cyber-physical aspect of a manufacturing environment without getting visually noticed (Vincent *et al.*, 2015).

2.2.2 Novel Attack Defense Techniques

Another research area in this field is developing novel cyber-physical attack defense techniques. Adopting the different defense mechanisms mentioned by Giraldo *et al.* (2017), the defense techniques discussed in this sub-section will be grouped into either: 1) preventing the attack from occurring in the first place; 2) detecting attacks that were able to bypass the used prevention techniques; or 3) responding, perhaps even automatically, to attacks in order to mitigate their effects. It should be just noted that the detection approaches will be presented first, due to the considerably larger number of efforts done in their research area, followed by the other two defense mechanisms presented together.

2.2.2.1 Detection Approaches

One approach that has been quite popular is the use of side-channels to detect product alterations resulting from cyber-physical attacks in a fashion analogous to using them to detect Trojans in Integrated Circuits (ICs)[4] (Vincent *et al.*, 2015). In IC literature, side-channels are used as non-destructive Trojan detection approaches where measurements of external characteristics of the ICs are taken to determine if alternations have been made (Chakraborty *et al.*, 2009; Vincent *et al.*, 2015).

Side-channels could include, but are not limited to, impedance signals, vibration signals, force signals, weight, acoustic emissions, and temperature measurements[5]. Side-channels typically contain useful information that can easily be related back to the product characteristics. In fact, it has been demonstrated that it is possible to use some sensor data, such as acoustic emissions,

[4] Trojans are inserted, as extra code, into ICs to force them to carry out a certain task (or more), that is not related to its original intent, when a certain trigger or situation occurs, unleashing its harmful effects (Vincent *et al.*, 2015).

[5] As an example, Chhetri and Al Faruque (2017) analyzed different potential side-channels in FDM-based additive manufacturing.

thermal images, and magnetometer measurements, from 3D printers and CNC mills as side-channel information to re-construct the manufactured product (Al Faruque *et al.*, 2016a; 2016b; Hojjati *et al.*, 2016; Song *et al.*, 2016). Therefore, it would be advantageous to use side-channels as a cyber-physical attack detection technique since attackers wouldn't know which ones (if any) are being used, rather than having to test a product visually or destructively.

As an example of such research efforts, Vincent *et al.* (2015) proposed using side-channels to enable real-time detection of cyber-physical attacks in production systems. The authors proposed applying an impedance based structural health monitoring technique, using Piezoelectric Transducers (PZTs) attached on a "removable antenna" within the manufactured part (Vincent *et al.*, 2015). The idea is to compare resulting impedance signatures from part excitation to baseline measurements when no part variation was present. Any deviation from the baseline measurements in the collected signatures would then indicate that the part's characteristics have changed (Vincent *et al.*, 2015). However, the design of the part itself had to be modified to accommodate the PZTs used for capturing the impedance signatures during part excitation.

Similarly, Albakri *et al.* (2015) explored the use of the impedance signatures, captured by PZTs, as side-channels for defect detection in additive manufacturing processes. Although the authors did not focus on specific defect sources, their method could still be applied to those resulting from cyber-physical attacks. Specifically, this effort was aimed at detecting generalized defects common to all additive manufacturing technologies, namely, dimensional inaccuracies, positional inaccuracies, and internal porosities (Albakri *et al.*, 2015). For this assessment, Albakri *et al.* (2015) used 2 control samples and 6 test samples each containing unique defect characteristics. All of the 8 samples were fabricated by material jetting and had the PZT sensors attached onto them using superglue (Albakri *et al.*, 2015). Utilizing appropriate metrics, the majority of defect types could be easily detected, with the porosity defects being a little less easily detectable (Albakri *et al.*, 2015). As a follow-up to this effort, the authors reported similar results in a later publication where they included one additional part type, more defect combinations, and more test specimen sets with various fabrication methods (Albakri *et al.*, 2017).

The same authors also further examined the use of impedance signatures captured by PZTs for the purpose of in-situ defect monitoring (Sturm *et al.*, 2016) in additive manufacturing processes. Sturm *et al.* (2016) experimented with embedding the PZTs within the manufactured part during material jetting using Polyjet 3D printers. Controlled specimens along with the defective ones were used to assess if the signatures would be: 1) the same at any given layer within a manufactured part; 2) able to detect the change from layer to layer within the same part; and 3) able to detect the existence of defects during the manufacturing process (Sturm *et al.*, 2016). The results have shown that the collected signatures from the PZTs could be successfully used for those 3 assessments. However, the results also showed that this approach is less effective for the third assessment since smaller defects were not detected quickly enough (Sturm *et al.*, 2016). In addition to this issue, other shortcomings are that, despite being an in-situ approach, the production process has to be stopped periodically to obtain the measured signatures and the approach also requires changing the part design a little to embed the sensors.

Additionally in the context of additive manufacturing, Chhetri *et al.* (2016) also proposed using side channels as a cyber-physical attack detection technique. More specifically, they suggested to use analog emissions, such as acoustic emissions, as side channels for attack detection in Fused Deposition Modeling (FDM) based additive manufacturing (Chhetri *et al.*, 2016). Their proposed method essentially compares the observed resulting analog emissions with statistically estimated emissions (inferred from initial control parameters) to check if they are different, implying the existence of an attack. Several steps are required to implement their method including emissions pre-processing and feature extraction, control signals extraction, and applying a detection algorithm (Chhetri *et al.*, 2016). Through an experimental setup, the authors showed that their method can detect a simple attack which changes the velocity, displacement, and movement of the printer's axes when printing a base plate for a quadcopter with an accuracy of 77.45% (Chhetri *et al.*, 2016).

Also in the additive manufacturing setting, Moore *et al.* (2017a) presented another method for detecting cyber-physical attacks, such as sabotage, by using the electric current traces produced from all actuators during the manufacturing process. This proposed approach has the extra benefit of being air-gapped from other computerized components in the system and employs an absolute-deviation-based threshold for comparing the different current traces (Moore *et al.*, 2017a). In an experimental study using FDM technology, Moore *et al.* (2017a) monitored the current traces from four motors, the 3-axes motors and filament extrusion motor, to assess their approach's ability to detect 4 types of G-codes alterations. More specifically, the alterations included inserting, deleting, reordering, and replacing G-code commands and the approach was able to detect the 1st three alterations with 100% precision using the traces from the X and Y axes traces (Moore *et al.*, 2017a). Lastly, Moore *et al.* (2017a) discussed their approach's limitations and its applicability in metal additive manufacturing.

In a more broad setting, Wu *et al.* (2017) used side channels to detect cyber-physical attacks in "CyberManufacturing" systems during the manufacturing process (real-time). Although the authors of this work do not refer to it as a side-channel detection approach, they presented two examples where physical data was extracted from manufacturing processes for evaluating the existence of cyber-physical attacks. More specifically, they used images and acoustic signals as the extracted physical data sources in additive and subtractive manufacturing processes, respectively (Wu *et al.*, 2017). Then, through the application of machine learning methods (such as k-nearest neighbor classifier, random forest multi-way classifier, and an anomaly detection algorithm) the authors were able to detect a range of cyber-physical attacks with varying accuracies (Wu *et al.*, 2017).

Similar to the work of Chhetri *et al.* (2016), Belikovetsky *et al.* (2018) also presented a method for detecting cyber-physical attacks in additive manufacturing based on using audio emissions as side-channels. Through recording and digitally signing the audio signals generated by the 3D printer's axes movement the authors were able to compare the resulting signals to detect any potential deviations in real-time (Belikovetsky *et al.*, 2018). Both Fast Fourier Transformation and Principal Component Analysis were applied to the original audio recording, which was taken using free mobile device applications, to obtain a comparable signal (Belikovetsky *et al.*, 2018). Then, the authors validated their proposed approach's ability to detect different manipulations affecting the tool path. Specifically, Belikovetsky *et al.* (2018) tested for detecting adding or removing G-code commands, modifying length parameters and extruder's speed, and reordering G-code commands. The proposed approach was able to detect all these modifications along with identifying detection limits for some of these modifications (Belikovetsky *et al.*, 2018). Finally, Belikovetsky *et al.* (2018) discussed additional aspects such as the effects of some recording-related factors on the accuracy of the obtained signal and the limitations of their approach.

2.2.2.2 Prevention and Response Approaches

The authors of such efforts mainly discussed innovative approaches for securing production systems against cyber-physical attacks from different perspectives. As one example, Wegner *et al.* (2017) presented a direct-to-machine security approach that ensures authentication and authorization by using two newly proposed devices, namely, comptrollers and manufacturing security enforcement devices. This approach would ensure that each step taken within the whole manufacturing enterprise has proper authentication and authorization (Wegner *et al.*, 2017). This work could be considered as a method to protect the production system from cyber-physical attacks occurring in the first place, i.e., more of a prevention mechanism.

As an example for a response approach, Bayanifar and Kühnle (2017) proposed an agent-based structure to achieve a CPS's dependability and security objectives. The proposed structure allows supervising and controlling the system to meet these objectives both autonomously and in real-time (Bayanifar and Kühnle, 2017). The authors specifically targeted cyber-physical production systems and intend for this structure to be a part of the system's inherent properties. The structure is composed of a core model and a control loop, both containing multiple agents that are responsible for data filtering, monitoring, analyzing, and eventually making the most appropriate decision (Bayanifar and Kühnle, 2017). The authors also present an example of the autonomous actions taken when a machine is no longer available in a job-shop unit. This effort could be thought of as a reaction plan to protect the system in case a cyber-physical attack does occur, without having to sacrifice performance.

2.2.3 Frameworks and Classifications

A final area currently being researched is developing cyber-physical security frameworks and classifications within production systems. The focus of scholars in this research area has been on providing systematic methods to quantify different issues related to cyber-physical security in manufacturing. Examples of such issues are identifying and/or assessing corresponding risks, vulnerabilities[6], and classifications for various types of cyber-physical attacks, among others. The results of these efforts allow researchers to gain better understanding of the nature of cyber-physical attacks on production systems and serves as another step towards better cyber-physical security of manufacturing enterprises.

2.2.3.1 Frameworks

As an example of research efforts discussing a framework, Hutchins *et al.* (2015) outlined a framework for identifying cyber-security risks in manufacturing. The focus of the proposed framework was more on data transfer within the manufacturing and supply chain environments

[6] In IT security literature, a *risk* is formally defined as "the probability that something unwanted will happen," whereas a *vulnerability* is "a weakness or fault in a system or protection mechanism that opens it to attack or damage" (Whitman and Mattord, 2012).

(Hutchins *et al.*, 2015). More specifically, activities in a manufacturing enterprise were grouped in this framework into five activity levels with the main objective of quantifying the following: 1) the data structures at each level; 2) the transfer of data between the levels; and 3) the data transfer across the supply chain (Hutchins *et al.,* 2015). With the proposed framework, several mechanisms for identifying generic and manufacturing-specific vulnerabilities would be possible within such data flows (Hutchins *et al.*, 2015). However, the authors did not consider cyber-physical security in their framework.

As another example, Chhetri *et al.* (2017) presented a cross-domain security analysis framework for CPSs, applicable to cyber-physical production systems. In their framework, the relation between the cyber domain model and physical domain information flows is abstracted using data-driven model estimation (Chhetri *et al.*, 2017). This model estimation requires the generation of training data and a variety of learning models in order to be able to select the most appropriate model (Chhetri *et al.*, 2017). Although the authors show two manufacturing specific case studies implementing their framework, there don't seem to be enough details about how exactly the implementation is done. Also, it is not that clear if these two case-studies, which are two of their previous publications (Al Faruque *et al.*, 2016a; Chhetri *et al.*, 2016), really fit into their suggested framework.

2.2.3.2 Classifications

As an example for research efforts discussing a classification, Yampolskiy *et al.* (2016) discussed how 3D printers could be used as cyber-physical weapons within an additive manufacturing environment and provided an assessment of the factors that would be involved within this environment in the form of different taxonomies. More specifically, the authors presented different classifications[7] for this situation such as classifications for: 1) compromised elements; 2) manipulations (or the influenced elements) and how those relate to the compromised elements; and 3) the effects of using 3D printers as a weapon and their relationship with the manipulations (Yampolskiy *et al.*, 2016). Finally, the authors discussed the characteristics of using 3D printers as a weapon and how much these characteristics could influence the effects of such usage (Yampolskiy *et al.*, 2016). In a subsequent publication, the authors further extend their proposed

[7] Again, these classifications are based on the cross-domain attack classification for CPS by Yampolskiy *et al.* (2013) and this publication is one of its extensions.

taxonomy to the subtractive manufacturing domain and compared the different elements within each taxonomy from both domains (Yampolskiy *et al.*, 2017).

As another example, Wu and Moon (2017) used some of the concepts discussed by Yampolskiy *et al.* (2013) and Wu *et al.* (2017) to build their own taxonomy for cyber-physical attacks in "CyberManufacturing" systems. Four attack dimensions were proposed for this taxonomy; namely, attack vector, attack impact, attack target, and attack consequence, with each dimension having components in both the cyber and physical domains (Wu and Moon, 2017). Although the authors define the impact as the "direct consequence" of the attack, they did not really explain what the two remaining dimensions exactly mean and how they are different. In addition, there seems to be some similarities between this attack taxonomy and the cyber-physical attacks one proposed by Pan *et al.* (2017).

2.2.4 Other Related Efforts

In addition to the research efforts discussed within the previous sub-sections, it is also worthwhile to mention other effort(s) that do not necessarily fit within any of these 3 groups. For example, through using multi-agents systems along with Petri nets, Yu *et al.* (2017) proposed a formal model for validating the design of cyber-physical manufacturing systems to improve their trustworthiness. More specifically, different components within the manufacturing systems are treated as agents and, then, the whole system is represented by object-oriented Petri nets (Yu *et al.*, 2017). In addition, Yu *et al.* (2017) also suggested a model for what they refer to as the "malicious software spreading" and a control method for the system. However, it is not clear how exactly this control method is applicable in practical situations and whether it is just a theoretical concept that still needs to be validated.

2.2.5 Closely Related Personal Efforts

The purpose of this sub-section is to discuss in a little more detail a couple of closely related research efforts involving Virginia Tech to which I contributed. These efforts can still be considered within the grouping outlined in the first three sub-sections, with the first being a vulnerability assessment framework and the second a classification. However, they are discussed separately because of my direct participation in these efforts and due to the close relation they have with the work presented in this dissertation.

2.2.5.1 A Cyber-Physical Vulnerabilities Assessment Framework

In the work by DeSmit *et al.* (2017), a framework for assessing cyber-physical vulnerabilities within manufacturing was proposed which actually included both the cyber and physical aspects of production systems. Such a framework was composed of a sequence of intersection mapping and decision tree analysis to identify potential vulnerabilities and analyze their impact, respectively (DeSmit *et al.*, 2017). The cyber-physical vulnerability impact analysis was performed at each intersection for the different processes under consideration within a facility. This analysis was designed in a manner to provide manufacturers with a stoplight-like scale between low, medium, and high for five different metrics, as an indication of potential vulnerabilities (DeSmit *et al.*, 2017). The output of such a framework was an overview of the cyber-physical vulnerabilities within a facility, which could help manufacturers to prepare appropriate mitigation strategies (DeSmit *et al.*, 2017).

The main advantages of the proposed framework are: 1) proposing a vulnerability assessment approach that takes both cyber and physical aspects of manufacturing into account, as opposed to other approaches that either do not consider the physical aspect or are not well-suited for manufacturing; 2) requiring minimum amount of input from manufacturers since only the production process maps are needed to implement the framework; 3) being a systematic approach that can be easily applied and repeated as needed; and 4) providing manufacturers with an overview of the areas that need attention to better enhance their cyber-physical security defense approaches. Yet, several improvements could be made to this framework such as including more capabilities within (e.g. threat identification, threat likelihood, and risk analysis) and incorporating all those capabilities into a more comprehensive cyber-physical vulnerabilities tool (DeSmit *et al.*, 2017).

In addition to discussing the proposed framework, DeSmit *et al.* (2017) also presented a cyber-physical attack example carried out at a partner facility, the Commonwealth Center for Advanced Manufacturing. The attack, which was part of a practical case study for implementing the framework, included altering the CAD file of a product to change three specific features within (DeSmit *et al.*, 2017). Although the product had to pass through a quality inspection phase during its manufacturing, the attack was still successful because the compromised features were not considered in the product's Geometric Dimensioning and Tolerancing (GD&T) information used by the CMM nor could the changed features be easily detected visually (DeSmit *et al.*, 2017). The success of this attack demonstrates that cyber-physical attacks in manufacturing could indeed be designed in a way to avoid detection by regular QC approaches, which reinforces the need for devising better QC tools or, at least, improving the current ones.

For this cyber-physical vulnerability assessment framework, my main contribution was suggesting the use of decision tree analysis for evaluating the impact of the different metrics at each intersection. The benefit of using decision trees for such an analysis is ensuring that the results could be easily repeated and are objective enough, without being influenced by any individuals' subjective opinions. Furthermore, I was part of the team that implemented the cyber-physical attack case study at our partner facility. This is in addition to participating in developing the used metrics, their decision rules, and assessing the validity of implementing them in practical situations.

2.2.5.2 Cyber-Physical Security Taxonomies

Pan *et al.* (2017) discussed two taxonomies for IoT-based production systems to better understand the potential dangers and provide better cyber-physical security in manufacturing. The first one is for cyber-physical attacks against manufacturing processes where elements of an attack flow can be classified into either vulnerabilities, attack vectors, attack targets, or attack impact (Pan *et al.*, 2017). Accordingly, a vulnerability is typically a software weakness that can be exploited by the attacker through a specific method or an attack vector to target a certain asset (Pan *et al.*, 2017). Each attack flow would then result in a specific attack type, referred to as attack impact, which the authors relate back to the famous CIA-triad of IT security. In other words, this taxonomy by Pan *et al.* (2017) can be considered as an attempt to extend this triad to cyber-physical production systems. Each of those four elements was further broken down into even more categories (Pan *et al.*, 2017).

The second taxonomy is for quality inspection measures to counteract these attacks where quality inspection is defined as "measures aimed at checking, measuring, or testing of one or more product characteristics and to relate the results to the requirements to confirm compliance" (Pan *et al.*, 2017). For this taxonomy, Pan *et al.* (2017) discussed some measures or checks that can be applied in both the physical and cyber domains to detect the effects of cyber-physical attacks in manufacturing; however, the authors were not considering the possibility of those measures themselves being compromised by the attacks. Furthermore, although the use of Statistical Process Control (SPC) tools was mentioned along with quality inspection, there is certainly much more room to explore the area of improving QC tools in general, for more secure cyber-physical production systems.

The principal contribution I had for the effort discussed here was assisting with developing both of the cyber-physical security taxonomies. More specifically, I outlined the diverse types of manufacturing processes to which the first taxonomy could be applied and suggested some improvements during developing this cyber-physical attack taxonomy. For the second taxonomy, the main role I played was suggesting presenting it as a taxonomy for quality inspection measures which could only be used to check some of the product's characteristics, assuming they themselves have not been compromised.

2.3 Quality Control for Cyber-Security

From the discussion in the previous sub-section, it can be seen that most of the current research efforts in cyber-physical security in production systems do not consider the use of quality control tools or the role QC systems has during such attacks. Furthermore, no efforts have discussed the possibility of using QC tools to detect the effects of cyber-physical attacks or as a defense mechanism to reduce their occurrence in the first place. While there have still been a few efforts that incorporated the use of QC tools in cyber-security of CPS in general, they are limited to detecting changes in digital phenomena.

The majority of such research efforts were focused, however, on network intrusion detection; more specifically, on anomalous network behavior detection (Kammerdiner, 2014). Detecting the underlying sudden changes in the behavior patterns within networks usually involved the use of SPC tools. SPC related examples included proposing to use Shewhart control charts based on a chi-square distance metric (Ye *et al.*, 2001) or applying a variety of univariate control charts such as X-bar and S (Goonatilake *et al.*, 2011), Exponentially Weighted Moving Average (EWMA) (Ye *et al.*, 2002; Ye *et al.*, 2003), and cumulative sum (CUSUM) control charts (Park *et al.*, 2014). Other anomalous intrusion detection techniques used scan statistics with Markov models (Neil *et al.*, 2013).

Aside from network intrusion detection applications, additional QC-related research studies in network security were also performed. For instance, X-bar and R control charts were used for a distributed denial of service attack detection scheme within a network (Wu *et al.*, 2007). Furthermore, it was also suggested to use non-parametric CUSUM statistic to detect attacks on process control systems in a fashion similar to anomaly intrusion detection within networks

(Cárdenas *et al.*, 2011). However, as pointed out by Megahed and Jones-Farmer (2015), the implementation of a QC tool such as control charts in this field of cyber-physical security can be considered still in its infancy, which is also true for most other QC tools.

2.4 Existing Gaps and Proposed Research Areas

Despite the fact that the cyber-security related research efforts involving QC, which were discussed in the previous sub-section, actually included the use of some tools, these efforts did not focus on cyber-physical security and concentrated only on network-related issues in CPSs. So even with all those efforts, the use of QC approaches for improving cyber-physical security in production systems has not yet reached its full potential. This is because research efforts for the security of cyber-physical production systems have not included QC approaches; whereas cyber-security efforts that did include QC are not suitable to be implemented in manufacturing for ensuring cyber-physical security, as they focus on only detecting digital phenomena instead.

Consequently, there are opportunities to use QC approaches for the cyber-physical security of production systems. Through detecting physical phenomena more effectively, improved QC approaches could result in minimizing the instances of successful cyber-physical attacks or even discouraging attackers by reducing the amount of time needed to detect such attacks in manufacturing. To this end, the proposed research work will focus on the following specific areas, which are further discussed in the subsequent chapters and are aligned with the objectives mentioned in sub-section 1.3:

1) **Developing a Cyber-Physical Attack Taxonomy from a QC Perspective:** The focus of this research area is to analyze cyber-physical attacks in manufacturing within a QC taxonomy. More specifically, this research area is concerned with describing and analyzing diverse types of cyber-physical attack possibilities while taking into consideration the role of QC systems within these attacks.
2) **Identifying Existing Weaknesses in QC Tools and Their Effects:** The concern of this research area is to explore the adverse effects cyber-physical attacks can have on exploiting existing weaknesses in QC tools. Hence, the focus of this research area would be to illustrate how exactly these weaknesses could be exploited by attacks and the resulting consequences; along with presenting a few ideas for mitigation and best practices.
3) **Introducing Randomness to QC Tools:** With the objective of providing better QC tools, this research area is concerned with including randomness in tools implementation and

design. Such a practice is common in the field of IT and the usefulness of extending it to the field of QC will be the main emphasis of this third research area. Specifically, the focus will be on control charts as representatives of common QC tools and the key challenge would be introducing randomness without sacrificing performance or increasing associated costs.

3 A Cyber-Physical Attack Taxonomy for Production Systems: A Quality Control Perspective

As a first step toward the development of new QC tools, an attack taxonomy to better understand the relationships between QC systems, manufacturing systems, and cyber-physical attacks is proposed in this chapter. The proposed taxonomy is developed from a quality control perspective and accounts for the attacker's view point through considering four attack design consideration layers, each of which is required to successfully implement an attack. In addition, a detailed example of the proposed taxonomy layers being applied to a realistic production system is included in this chapter.

It should be noted that the work in this chapter was presented at both INFORMS 2015 (Elhabashy *et al.*, 2015b), and INFORMS 2016 (Elhabashy *et al.*, 2016) Annual Meetings. In addition, this chapter is published as a research article in the *Journal of Intelligent Manufacturing* (Elhabashy *et al.*, 2018a); and the material is reused with the journal's permission, as shown in Appendix A.

3.1 Background and Motivation

Over the past century, the manufacturing industry has relied heavily upon the use of Quality Control (QC) systems to detect quality losses, ensure production of high quality parts, and maintain stable processes. During this time, QC tools have been adapted and new theories/methods have been developed to handle practical problems faced by scientists and practitioners. For example, developments in the underlying statistical science behind QC have been adapted to account for increasing amounts of data and computational power available in many applications (Box and Woodall, 2012). As we progress further into the 21st century, a new need is beginning to surface; specifically, the need for QC tools and approaches designed to prevent and/or reduce vulnerabilities to cyber-enabled attacks against manufacturing.

3.1.1 Cyber-physical Attacks in Manufacturing Environments

Examples of cyber-attacks are many, spanning across a wide range of fields. Recent cyber-attacks targeted large retailers (Target, 2014; Kurtzman Carson Consultants, 2016), the entertainment industry (Lee, 2014), healthcare insurance companies (Anthem, 2016), power grids (Tuptuk and Hailes, 2016), nuclear plants (Steitz, 2016), technology corporations (Fahey and Wells, 2016; Castillo, 2017), and many more. For cyber-physical systems, such as advanced manufacturing, one of the most notable attacks to date is the infamous "Stuxnet" virus. Between late 2009 and early 2010, the Stuxnet virus was responsible for destroying as many as 1,000 high-speed uranium enrichment centrifuges in Iran (Albright *et al.*, 2010). Due to the integrated cyber-physical nature of these uranium enrichment systems, Stuxnet was able to periodically change centrifuge speeds, while simultaneously falsifying equipment readings to indicate normal operating conditions (Cherry, 2011). Despite high levels of both cyber and physical security in these enrichment facilities, the virus was still able to reach and successfully attack its target (Cherry, 2011); demonstrating that even highly secured cyber-physical systems can be infiltrated by cyber-attacks. This suggests that given enough time, computational resources, planning, intent, and luck, all systems can be breached, with production systems being no exception.

John Chambers, Cisco's executive chairman and former CEO, stated "I am convinced that there are only two types of companies: those who have been hacked, and those who don't yet know they have been hacked" (Chambers, 2015). This statement resonates across all industries, especially manufacturing. In fact, over the last few years, manufacturing has been the target of a variety of cyber-attacks seeking to acquire sensitive user information (Deloitte, 2014; Symantec, 2014; 2015; 2016; 2017). Manufacturers have also increasingly become the victims of ransomware (Teemu M., 2015; McGee, 2016; Liptak, 2017), where computer access is blocked (through encrypting its contents or locking it) until a ransom is paid. The critical manufacturing sector alone accounted for the majority (33%) of all security incidents reported to the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) between October 2014 and September 2015 (ICS-CERT, 2016). While most of these attacks concentrated on Intellectual Property (IP) theft (Pham, 2015; IBM X-Force Research, 2016; ICS-CERT, 2016), the cyber-physical nature of manufacturing opens the door to attacks that can maliciously affect the physical manufacturing environment.

Modern manufacturing enterprises consist of a highly heterogeneous assortment of integrated cyber and physical technologies, which often includes control systems (PCs, Supervisory Control and Data Acquisition (SCADA) systems, PLCs), machining equipment (lathes, mills, electric discharge machines), assembly equipment (pneumatics, welders, robotics), post-processing equipment (kilns, paint lines), and inspection equipment (digital cameras, coordinate measuring machines). Furthermore, the Information Technology (IT) structure of these systems is quite complex as the aforementioned technologies communicate across a multitude of architectures

(FireWire, RS-232, USB), protocols (FTP, HTTP), network technologies (Ethernet, LAN, Wi-Fi), and operating systems (Linux/Unix, MS-DOS, Windows 3.x, Windows CE, Windows 95/98, Windows XP, Windows 8, Windows 10). Given the multifaceted nature of these systems, reliance upon a mix of modern and legacy equipment/technologies, the lack of Original Equipment Manufacturer (OEM) support, and the infeasibility to implement software updates results in production systems being more susceptible to cyber-attacks.

The opportunities for attacks are further exacerbated by: 1) the Internet of Things (IoT), which has resulted in a rampant growth of networked devices (other than computers) across every sector (Evans, 2011), including manufacturing; 2) increased usage of internet-based Computer Aided Engineering (CAE) support tools, such as cloud computing and Software as a Service (SaaS); and 3) the current push in industry to easily extract and share manufacturing equipment/process data, which can be seen in the increasing adoption of emerging communication protocols such as MTConnect (MTConnect Institute, 2016). With the continual development of cyber-physical manufacturing environments, together with the recent increase in cyber-attack maliciousness and decrease in visibility (Bayuk *et al.*, 2011; Gendarmerie, 2011), there is a need for researchers to focus their attention on developing tools to secure production systems against cyber-physical attacks.

The successful implementation of cyber-physical attacks, which are cyber-attacks affecting the physical system components, against a manufacturing enterprise could be devastating. Potential attacks include, but are not limited to: 1) modifying product designs (Computer Aided Design (CAD) files, tolerances, original design criteria); 2) altering manufacturing processes (tool paths, tools used, machining parameters) (Wells *et al.*, 2014); and 3) manipulating process/product data (inspection results, machine maintenance indicators). Such attacks could cause the system to produce modified products; which if gone undetected, could adversely affect a product’s design intent, performance, quality, or perceived quality (Wells *et al.*, 2014). The results of such attacks could delay a product’s launch, increase warranty costs, or ruin customer trust. More importantly, these attacks could pose a safety risk to consumers exposed to faulty products (Wells *et al.*, 2014).

3.1.2 The Need for New QC Tools

Traditionally, manufacturing has relied on QC systems to detect changes and quality losses within a product and/or process. QC systems are based on making decisions concerning the different processes performance from collected process data over time, using appropriate tools. These QC tools are essential to both reducing variability within a process and monitoring its performance, such as control charts (Montgomery, 2009). QC tools could also include inspection equipment and measurement instruments.

Recently, Wells *et al.* (2014) and Vincent *et al.* (2015) pointed out that current QC approaches (and tools) are not designed to detect process changes caused by malicious cyber-physical attacks. Specifically, QC methods are based upon assumptions (e.g., sustained system shifts) and decisions on rational sub-grouping that may be invalid when considering the possibility of a cyber-physical attack (Wells *et al.*, 2014; Vincent *et al.*, 2015). Given the importance of maintaining the integrity of our manufacturing infrastructure, it becomes clear that new QC approaches need to be developed for current and future cyber-physical manufacturing environments that go beyond detecting non-malicious quality losses.

It was reported in 2015 that the median number of days between a security breach and its detection was over 200 days (Mandiant, 2015). In addition, 69% of these breaches were not detected by the victims, but by third parties, such as law enforcement agencies (Mandiant, 2015). In a manufacturing context, this could equate to 200 days of unknowingly producing modified or flawed products, until they begin to fail in-use. However, as discussed by Cárdenas *et al.* (2009); (2011), since cyber-physical systems affect the physical world, they offer an additional avenue for detecting attacks beyond traditional cyber-security techniques. Using new QC approaches to detect such physical phenomena would act as a "second line of defense" against malicious cyber-physical attacks, which would assist in potentially reducing cyber-physical breach detection times in manufacturing.

Quality control methods have been previously used in cyber-security. However, QC approaches have been used only to detect changes in digital phenomena, such as anomalous network behavior detection (Ye *et al.*, 2001; 2002; 2003; Goonatilake *et al.*, 2011; Neil *et al.*, 2013; Kammerdiner, 2014; Park *et al.*, 2014) and in various other network-related applications (Wu *et al.*, 2007; Cárdenas *et al.*, 2011; Megahed and Jones-Farmer, 2015). Research in cyber-physical security has traditionally grouped manufacturing in the generic area of security for critical infrastructures (Stamp, 2003). As its own unique research area, cyber-physical security for manufacturing researchers have not yet focused their attention toward QC and primarily presented case-studies demonstrating cyber-physical attacks for varying manufacturing processes (Wells *et al.*, 2014; Turner *et al.*, 2015; Belikovetsky *et al.*, 2016; Zeltmann *et al.*, 2016; Moore *et al.*, 2017b;

Sturm *et al.*, 2017), suggesting novel defense schemes (Vincent *et al.*, 2015; Chhetri *et al.*, 2016; Sturm *et al.*, 2016; Albakri *et al.*, 2017; Bayanifar and Kühnle, 2017; Belikovetsky *et al.*, 2017; Moore *et al.*, 2017a; Wegner *et al.*, 2017; Wu *et al.*, 2017), or proposing frameworks and classification approaches (Hutchins *et al.*, 2015; Yampolskiy *et al.*, 2016; Chhetri *et al.*, 2017; DeSmit *et al.*, 2017; Pan *et al.*, 2017; Wu and Moon, 2017; Yampolskiy *et al.*, 2017).

The first step in developing new QC approaches for cyber-physical security in manufacturing is to understand the relationships between QC systems, manufacturing systems, and cyber-attacks in the context of malicious process changes. In response to this need, we propose an attack taxonomy that encompasses these relationships from a QC perspective. This proposed taxonomy will allow for a systematic approach to: 1) describe cyber-physical manufacturing attack surfaces[8] involving QC systems, 2) analyze the effects of possible QC-related cyber-physical attack types, 3) explore the mechanisms through which QC tools could be exploited, 4) evaluate the weaknesses of specific QC tools, and 5) develop advanced QC approaches to prevent or minimize the effects of cyber-physical attacks against manufacturing. The proposed taxonomy is not, however, an exhaustive classification of all the attack possibilities in cyber-physical production systems; it is rather a taxonomy of how cyber-physical attacks operate with respect to quality control systems only.

It should be noted that Pan *et al.* (2017) proposed two taxonomies in the area of cyber-security for cyber-physical manufacturing systems. However, the taxonomies proposed by Pan *et al.* (2017) do not consider relationships between manufacturing and QC systems. The first taxonomy was for cyber-physical attacks against manufacturing processes, whereas the second was for quality inspection[9] measures to counteract the attacks; there was no overlap between these two taxonomies. Within the second (quality inspection) taxonomy, although the use of Statistical Process Control (SPC) tools was mentioned, Pan *et al.* (2017) only discussed measures or checks that can be applied in both the physical and cyber domains to detect the effects of cyber-physical attacks that are not all necessarily related to traditional QC approaches. Finally, Pan *et al.* (2017) did not account for the possibility of those quality tools themselves being compromised by the cyber-physical attacks, which is considered in the attack taxonomy proposed in this work.

[8] A system's attack *surface* is a term that represents all possible means through which an attacker can access the system and potentially cause harm (Manadhata and Wing, 2011).

[9] Quality inspection "*are measures aimed at checking, measuring, or testing of one or more product characteristics and to relate the results to the requirements to confirm compliance*" (Pan *et al.*, 2017).

This chapter is organized as follows; Section 3.2 contains the details of the proposed attack taxonomy. An example for implementing the proposed taxonomy in a realistic production system is provided in Section 3.3. Next, in Section 3.4 system-specific design considerations that need to be taken into account when implementing the attack taxonomy are discussed. Finally, the chapter is concluded with a general discussion in Section 3.5.

3.2 Cyber-Physical QC Attack Taxonomy for Manufacturing

In this section, the proposed taxonomy that governs the relationships between QC systems, manufacturing systems, and cyber-attacks in the context of malicious process changes is presented. Within cyber-physical production systems, several inter-connected sub-systems often exist to support manufacturing processes and assembly operations, such as production planning and control and quality control systems (Groover, 2010). For conciseness, the proposed taxonomy is developed around a simplified production system consisting of only manufacturing and QC sub-systems.

A common practice within the cyber-security community is to focus on understanding systems from the attacker's point of view (Intrusion Kill Chains (Assante and Lee, 2015), Penetration Tests (Underbrink *et al.*, 2012), White Hats, Hack-a-Thons, etc.), which often enables defenders to develop better defense strategies. Following suit, the proposed taxonomy is developed from an attacker's viewpoint and aims to understand the relationships between QC systems, manufacturing systems, and cyber-attacks. This proposed attack taxonomy could also be aligned with the manufacturing profile of the National Institute of Standards and Technology (NIST) Cybersecurity Framework's first core function "Identify", which is defined as "*develop the organizational understanding to manage cybersecurity risk to systems, assets, data, and capabilities*" (Stouffer *et al.*, 2017). The taxonomy would not fit, however, with other functions of this cybersecurity framework's core such as "Detect" and "Respond", which are concerned with developing and implementing activities to identify cybersecurity events and taking actions towards them, respectively (Stouffer *et al.*, 2017). In other words, the attack taxonomy is not intended to be used for attack detection or identification, but for realizing the existing vulnerabilities within cyber-physical production systems.

More specifically, the taxonomy includes four attack design consideration layers, referred to herein as just attack layers, required to successfully implement an attack; while taking into account a perspective focusing more toward QC systems. These four layers are: 1) attack objectives; 2) affected production sub-systems; 3) method of applying the attack; and 4) the actual attack location. These proposed taxonomy's four attack layers would mainly aid in understanding an attacker's decision-making process during early attack phases (such as planning and development (Assante and Lee, 2015)), assuming rational decision making. For defenders, understanding these layers provides crucial insights into the attackers' thought processes. With these insights, defenders have the opportunity to develop more secured production systems and appropriate future action plans for minimizing (if not eliminating) various attacker advancements.

It should be noted that these layers do not consider the processes of identifying cyber-security weaknesses or gaining access to a computer/network. In addition, system-specific attack design considerations are ignored within the four attack layers of the taxonomy, such as: 1) component accessibility; 2) attack pervasiveness; 3) attack impact; and 4) attack coordination. However, the importance of these system-specific design considerations is addressed in Section 3.4 to provide a more overall understanding of the system. Furthermore, attackers' motivation is also ignored within this taxonomy, and it is assumed that an attacker's objective is correctly aligned with their motivation. However, it should be noted that there are numerous examples of motivations for cyber-physical attacks against manufacturing, such as gaining a competitive advantage, revenge (e.g., disgruntled employees (Slay and Miller, 2008; Kravets, 2009; Poulsen, 2009; Brenner, 2013)), and terrorism. Detailed descriptions of these four layers are provided in the following sub-sections.

3.2.1 Attack Objectives

The first layer within the proposed taxonomy covers possible objectives for cyber-physical attacks against a production system. The proposed taxonomy divides possible objectives into three (not necessarily mutually exclusive) categories: 1) disrupting manufacturing operations; 2) reducing outgoing product quality; and 3) altering product design intent.

3.2.1.1 Disrupting Manufacturing Operations

Cyber-physical attacks may aim to cause disruption in the production system, resulting in unanticipated downtime, wasted production efforts, and/or ruined equipment. One way to fulfill this objective is by forcing the system to produce significantly nonconforming products. If these nonconforming products are detected, the production system should be brought to a halt. The system would remain offline until the cause for the nonconformities is identified and appropriate actions to recover from the attack are implemented.

This disruption objective could also be accomplished by destroying or forcing significant maintenance work on machines and/or equipment. An example of such an attack on cyber-physical equipment was accomplished in the Aurora project conducted at the Department of Energy's Idaho lab in 2007 (Meserve, 2007). In this project, a cyber-physical attack caused a diesel generator’s circuit breakers to rapidly cycle between being open and closed. This resulted in causing the generator to shake, smoke, and eventually break down (Meserve, 2007; Brenner, 2013; Tucker, 2014). Obviously, an attack of this nature would suspend a manufacturing operation until the affected resources are fixed or replaced.

Finally, to achieve the disruption objective, it may not be necessary for an attack to cause any physical changes to products or processes. Instead, an attack could take advantage of a system’s reliance upon automated product/process data collection systems to either make a fraction of products falsely appear nonconforming or incorrectly indicate that a specific piece of equipment needs unnecessary maintenance. Both these scenarios would result in unwanted or unscheduled system downtime while the respective system recovery processes are performed.

3.2.1.2 Reducing Outgoing Quality

In order to reduce a production system’s outgoing quality, an attack must either prevent the system from identifying and rectifying nonconforming products or force the system to produce nonconformities. In either case, a successful attack results in some portion of outgoing products from failing to satisfy specifications. One way to achieve this objective is to deceive the inspection process into classifying manufactured products as conforming, regardless of their actual quality. This form of attack would prevent manufacturers from adhering to product specifications; however, it requires the system to naturally produce nonconformities without any additional attack influence.

3.2.1.3 Altering Product Design Intent

Attacks with the third objective of altering a product’s design intent involves forcing a change in the production process, resulting in products that fail in-use, exhibit abnormal/reduced performance, or possess unintended functionalities. To accomplish such an attack objective, product characteristics would need to be specifically altered at some point in the production process to achieve a very deliberate product alteration. A synopsis of the three aforementioned cyber-physical attack objectives is detailed in Figure 3-1.

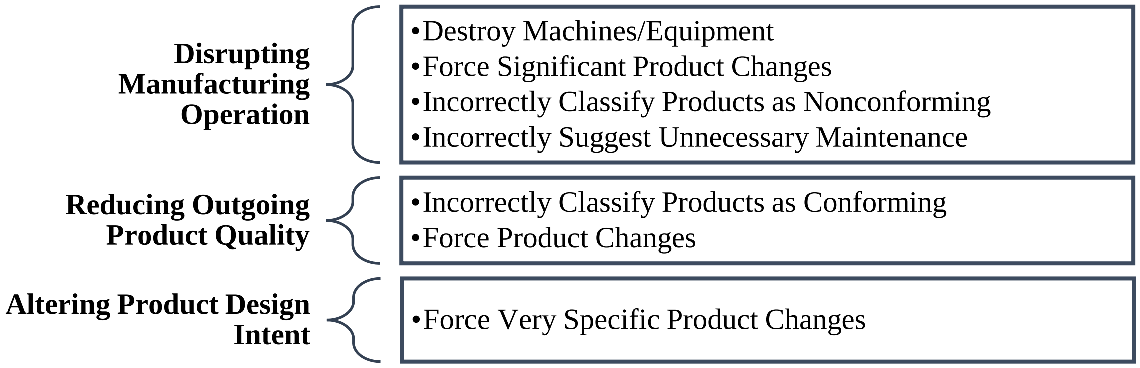
Disrupting Manufacturing Operation
•Destroy Machines/Equipment
•Force Significant Product Changes
•Incorrectly Classify Products as Nonconforming
•Incorrectly Suggest Unnecessary Maintenance
Reducing Outgoing Product Quality
•Incorrectly Classify Products as Conforming
•Force Product Changes
Altering Product Design Intent
•Force Very Specific Product Changes

Figure 3-1: Cyber-physical attacks objectives in manufacturing.

3.2.1.4 Objective Commonalities

Despite being unique objectives, reducing outgoing quality and altering product design intent share several characteristics. First, the outcome of an attack aimed at accomplishing either objectives must not be immediately detectable by inspection and/or quality control. If the effects of the attack are immediately detected, the compromised products would not be allowed to leave the facility and the objective of the attack would not be attained. However, this would inadvertently result in accomplishing the objective of disrupting manufacturing operations. If the effects of an attack are not immediately detected, some fraction of manufactured products will be compromised and may eventually enter into the hands of a customer. This suggests that attacks aimed at reducing outgoing quality or altering product design intent are not required to be persistent.

Second, these two objectives can both be accomplished by physically altering a product. For example, consider the two case studies recently conducted at Virginia Tech (Wells *et al.*, 2014; Sturm *et al.*, 2017), where cyber-physical attacks successfully caused a production system to produce manufactured products that are significantly different from their original designs. In the first case study, tool path files were altered for a machining process (Wells *et al.*, 2014) resulting in considerably incorrect part dimensions. It should be noted that this attack was designed and implemented in a manner that made its effects hidden from the inspection process; therefore, reducing outgoing quality. In the second case study, design files for an additive manufacturing process (Sturm *et al.*, 2017) were altered to produce an internal void within a product, which would be undetectable by the majority of modern manufacturing inspection systems. This internal void significantly increased stress concentrations in the product, which could have resulted in product failure under designed operating conditions, thereby altering its design intent.

3.2.2 Affected Production Sub-system

The second layer of the proposed taxonomy aids in determining which production sub-system(s) would be attacked to accomplish a specific objective. For this taxonomy, a production system is viewed as two concurrent sub-systems: the manufacturing system and the quality control system. The manufacturing system encompasses all activities required to realize the physical product, such as product design, manufacturing processes, and finishing operations. In contrast, the QC system covers all actions taken to ensure that manufactured products meet specifications and the production process remains stable. The QC system contains product inspection, SPC tools, process monitoring, etc. While these two systems operate in parallel, overlap exists by the fact that observations and decisions made through a QC system can affect the operations within a manufacturing system. For example, if a product inspected by a QC system is found to be nonconforming, it may be returned to the manufacturing system for rework. Conversely, product/process redesigns made within a manufacturing system may affect which data a QC system collects and how that data are analyzed. Together these two sub-systems result in transforming raw engineering materials into finished products within a production system, as shown in Figure 3-2.

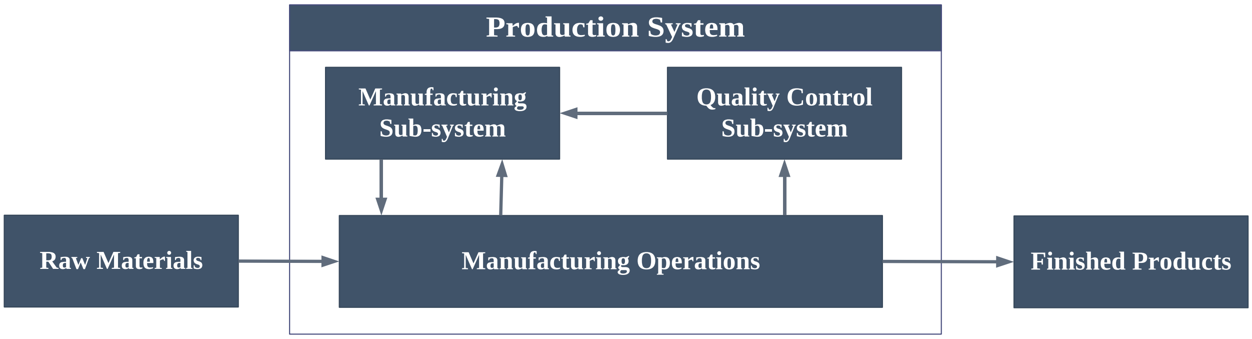

Figure 3-2: Overview of the production system components (adapted from Groover (2010)).

With respect to the objectives identified in Section 3.2.1, a successful attack requires attacking either the manufacturing, QC, or both systems; as can be seen in Table 3-1. For instance, an attack that aims to disrupt manufacturing operations through forcing significant product changes, would only need to affect the manufacturing system. Conversely, an attack that aims to incorrectly identify products as conforming to reduce product outgoing quality, only needs to affect the quality control system. While a cyber-physical attack that forces product changes and attempts to conceal the attack, needs to affect both the manufacturing and quality control systems.

Table 3-1: Examples of affected production systems.

Cyber-Physical Attack Objective	Affected System(s)
1) Disrupt Manufacturing Operations:	
a) Destroy machines/equipment	Manufacturing
b) Force significant product changes	Manufacturing
c) Incorrectly classify products as nonconforming	Quality Control
d) Incorrectly suggest unnecessary maintenance	Quality Control
2) Reduce Outgoing Quality:	
a) Incorrectly classify products as conforming	Quality Control
b) Force product changes	Manufacturing
c) Force product changes and attempting to conceal the attack	Both
3) Alter Product Design Intent:	
a) Force very specific product changes	Manufacturing
b) Force very specific product changes and attempting to conceal the attack	Both

Attacks against QC systems will be classified as *Direct Quality (DQ)* attacks. While these attacks are limited to activities within the QC system, they can be used to acquire QC implementation information, alter QC data, or affect how QC data are collected, used, stored, or analyzed. Similarly, attacks against manufacturing systems will be referred to as *Direct Manufacturing (DM)* attacks. Such attacks can physically alter the manufacturing process (which may require obtaining manufacturing data) to affect either equipment functionality and/or product quality (deviation from design intent).

Due to the physical manifestations that result from DM attacks, there is always the possibility that the effects of a DM attack could be detected by a QC system. Therefore, depending on the attack objective, it may be advantageous for the attacker to target both the QC and manufacturing systems. In attacks of this nature, referred to as *Joint* attacks, attacks against manufacturing and QC systems are designed and implemented together to reduce or eliminate attack detection. Joint attacks have potentially devastating consequences and are one of the main reasons behind the need to develop new or improve current QC tools and approaches to prevent and/or reduce vulnerabilities to cyber-enabled attacks. A summary of the affected systems within each attack type is shown in Figure 3-3.

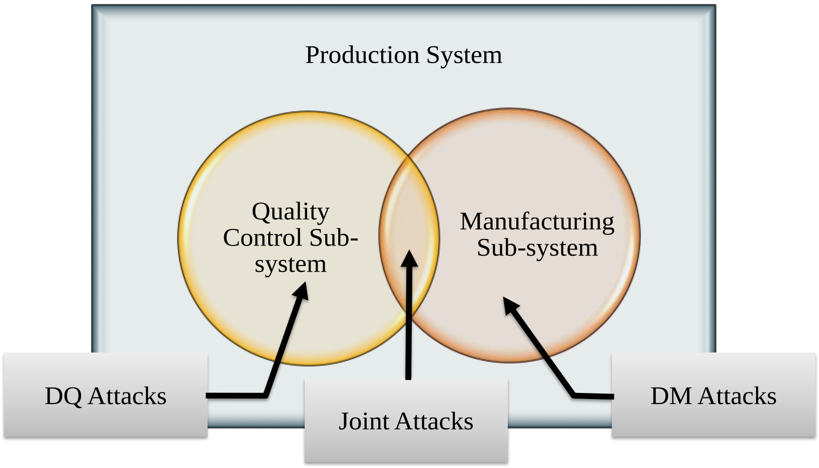

Figure 3-3: Summary of affected systems within each attack type on production systems.

Based upon the manner in which a QC system is attacked, Joint attacks are further classified into two types. The first Joint attack type, referred to as *Active Joint* or simply an *Active* attack, occurs when attacks against the manufacturing system and QC system result in a physically altered process/product and digitally altered QC data, respectively. For the second joint attack type, referred to as *Passive Joint* or simply a *Passive* attack, QC data is not altered. Instead, Passive attacks against the QC system acquires information regarding the implementation and use of QC tools and data. This acquired information is then used to design and execute a complementary manufacturing system attack to considerably decrease its probability of detection. In other words, Active attacks change both systems; whereas Passive attacks only change the manufacturing system but uses QC system information to make the change less obvious.

3.2.3 Attack Method and Attack Location

In this sub-section, the final two layers of the proposed taxonomy are discussed; namely, Attack Method and Attack Location. The Attack Method layer focuses on how to achieve the outcome a specific attack aims to cause, while the Attack Location layer determines the location(s) within the production system the attacks will take place. It should be noted that these two layers are presented together due to their close relationship with each other.

3.2.3.1 Attack Method

For any given production system, there is a myriad of possible attacks to achieve a desired outcome. The proposed taxonomy groups possible attacks into generalized categories, referred to as Attack Methods. Examples of possible Attacks Methods for a QC system may include; 1) the altering of part quality definitions, 2) the reporting of falsified QC data, and 3) the acquisition of QC implementation information. The latter method would be most often associated with Passive attacks, while the remaining two methods could occur in either DQ or Active attacks.

The desired outcome from attacking a manufacturing system is to produce a physical change in either a manufactured product or manufacturing equipment. Therefore, manufacturing system Attack Methods consist of attacks that alter either what product is to be manufactured or how a product is manufactured. The former type of Attack Methods could be thought of as methods that change a product’s design, such as: 1) modifying feature dimensions, 2) adding extra features, 3) removing features, and 4) changing material to be used. The latter type of Attack Method could be thought of as methods that change the process (that could ultimately result in a change of the product), such as: 1) altering machine/equipment parameters (e.g., feeds and speeds, temperatures, pressures, working envelopes) or 2) changing manufacturing set-ups (e.g., tooling, coolant). Regardless of the desired outcome, these Attack Methods can be associated with DM, Passive, or Active attacks.

3.2.3.2 Attack Location

Within either the QC or manufacturing systems there are numerous possible locations where attacks could be implemented. It must be noted that given the cyber-physical nature of this taxonomy, an attack location is not constrained to merely physical locations within a production system but can include “cyber locations” where data is entered, accessed, or stored. Furthermore, the location of the attack does not necessarily correspond to the location where the attack’s outcome occurs. The selection of attack location is the fourth and final layer of the proposed taxonomy.

While the decision of attack location is a unique layer, it is highly influenced by the choice of the attack method. Specifically, depending on the attack or production system scenario, the choice of the attack method constrains possible attack locations (and vice versa). For example, if the chosen attack method is to falsify QC data, the location of this attack will be limited to inspection equipment, QC stations, and QC databases. While this severely limits attack locations, multiple instances of prospective attack locations (e.g., multiple pieces of inspection equipment) may exist within a production system.

Attacks are not limited to only one location. In fact, the very definition of a Joint attack (Section 3.2.2) requires a minimum of two attack locations. Furthermore, to be successful, complex attacks or complex production systems may require attacks at multiple locations. As an example of the former, a complex attack designed to produce a very specific product alterations may need to occur through a series of many DM attacks at several different stations within the manufacturing system (e.g., attacks on a lathe followed by attacks on a mill). As an example of the latter, complex production systems with redundant quality inspection stations would require multiple DQ attacks.

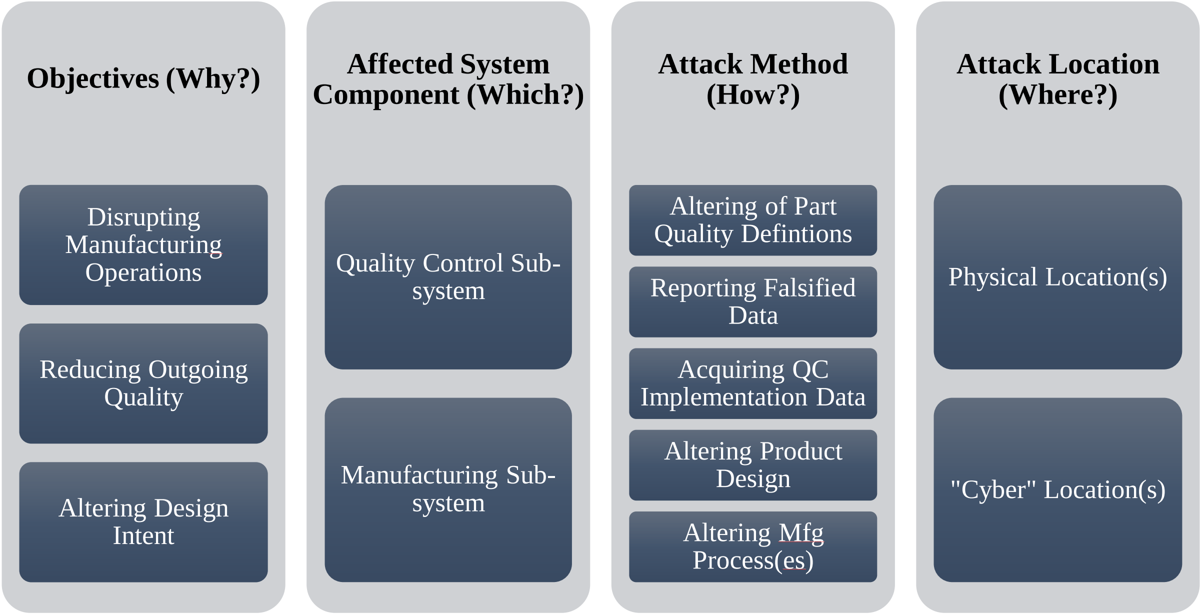

Figure 3-4: Overview of the proposed cyber-physical QC attack taxonomy, along with a summary of the details of each layer.

3.2.4 Taxonomy Overview

An overview of the proposed cyber-physical QC attack taxonomy along with the details of its four different layers is illustrated in Figure 3-4. The first layer (attack objectives) essentially addresses the question of *why* an attacker would target a specific production system, with the objectives being grouped into three (not necessarily mutually exclusive) categories. Layer two, of the taxonomy, addresses the question of *which* sub-system(s) within the production system needs to be targeted. The third layer is concerned with determining *how* the attack would be implemented. As shown in Figure 3-4, there are several possibilities for the method of the cyber-physical attack, depending on which of the two sub-systems is being affected. Finally, the last taxonomy layer answers the question of *where* the attack could be focused.

3.3 Taxonomy Implementation Example

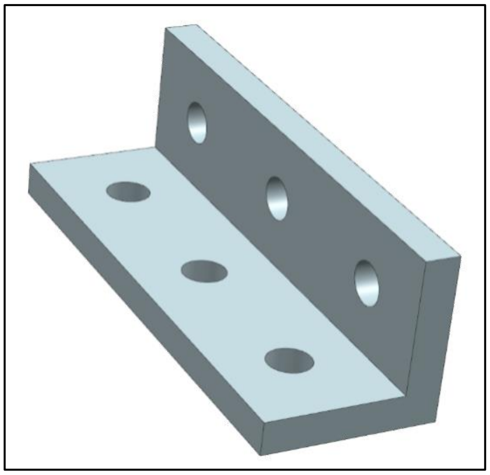

Figure 3-5: Implementation example part.

In this section, a unifying example is presented to demonstrate the cyber-physical QC attack taxonomy developed in Section 3.2. This example covers the production system involved in manufacturing the product shown in Figure 3-5. For this example, it is assumed that the product (features, material, etc.) and the process (machines, equipment, tooling, etc.) have already been appropriately designed to meet specifications. In addition, a flow diagram outlining the production of this part is given in Figure 3-6. Here dashed, dotted, and dashed-dotted lines respectively represent components within the manufacturing system, the QC system, or within both systems.

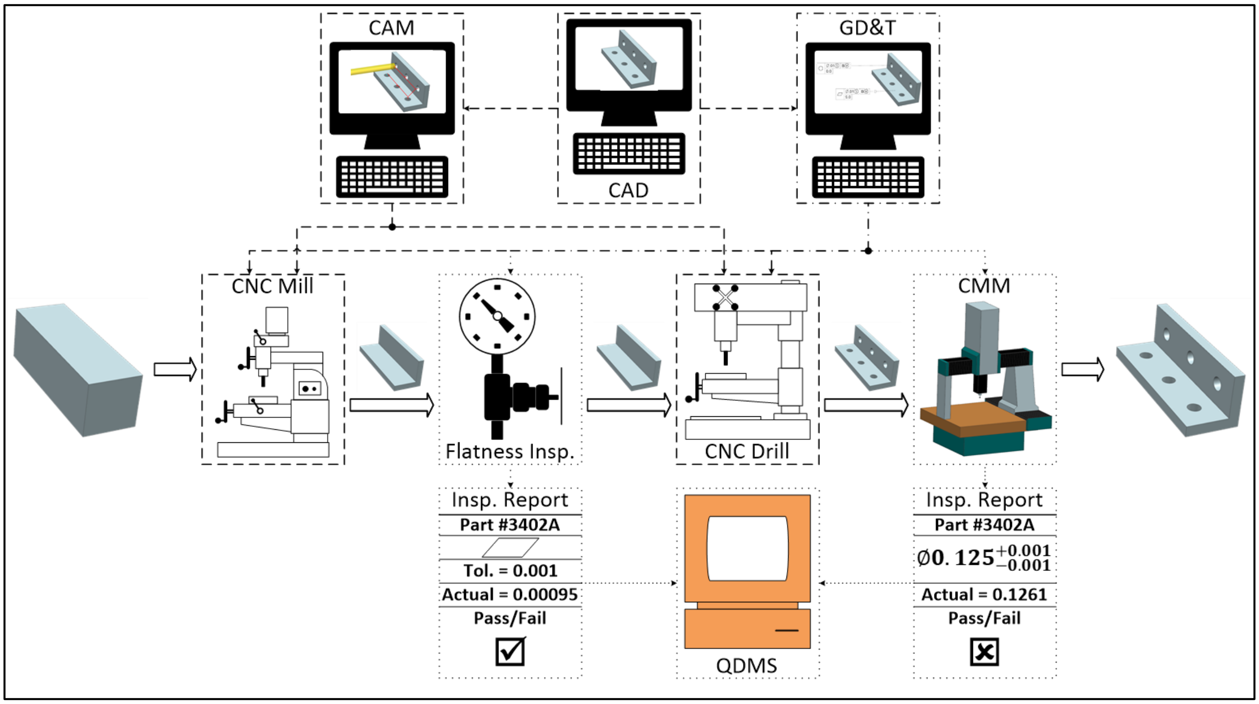

Figure 3-6: Production system example flow diagram

The first step to manufacture this product is to set-up the production system. This consists of generating a Computer Aided Design (CAD) rendering of the part, which will be used to develop both Computer Aided Machining (CAM) and Geometric Dimensioning and Tolerancing (GD&T) data. The CAM data (in the form of M and G-Codes) are sent to two Computer Numerical Control (CNC) machines, a milling and a drilling machine. Similarly, the GD&T data are sent to the CNC machines for process set-up. The GD&T data are also sent to a manual inspection station and a Coordinate Measurement Machine (CMM) for implementing QC.

Once system set-up has been completed, the part can be produced. The production process begins with bar stock entering the CNC Mill workstation. At this point, a surface milling operation is performed to produce the part's required cross-section. This milling operation is followed by a manual inspection (100%) of the part's overall flatness. This manual inspection is recorded in hard-copy form, which is later recorded into the Quality Data Management System (QDMS). After the inspection, six holes are drilled at the CNC Drill workstation. From here, hole sizes and locations are inspected 25% of the time (every fourth part) using the pre-programmed CMM. This inspection data is automatically uploaded to the QDMS and a quality report is automatically generated and presented to a CMM operator (in digital form).

This example will be used to illustrate the proposed QC attack taxonomy. Specifically, this will be accomplished by demonstrating several possible implementations of the taxonomy, starting from layer 1 - Attack Objectives.

3.3.1 Attack Objective: Disrupting Manufacturing Operations

While there exist several approaches to disrupting a manufacturing operation, this taxonomy implementation example will focus on incorrectly classifying products as nonconforming. Considering this proposed taxonomy's 2nd layer - Affected Production Sub-system, incorrectly classifying products as nonconforming could be accomplished through a DQ attack. Upon analyzing the production system given in Figure 3-6, it becomes apparent that a DQ attack can be implemented against: 1) GD&T Data Development, 2) Flatness Inspection Station, 3) CMM Station, or 4) the QDMS. When considering these four possible locations for a DQ attack with respect to this proposed taxonomy's 3rd layer - Attack Method, incorrectly classifying products as nonconforming could be accomplished by the following methods; 1) altering the definition of a nonconforming product, 2) altering the data collection process, 3) altering the data report process, or 4) altering the collected data. A discussion of this taxonomy's 4th layer - Attack Location, for each of the previously identified attack methods is presented next. It must be noted that the following discussion regarding attack location does not consider the difficulty in performing the attack.

3.3.1.1 Altering the Definition of a Nonconforming Product: Attack Locations

The following four cases will demonstrate the methods altering the definition of a nonconforming product at different locations within the targeted production system. For each of these four cases, the tolerance band for specific Key Quality Characteristics (KQCs) will either be moved or shifted:

1. The **Flatness Inspection Station** could be attacked by altering the digital file that is used to print the hard-copy inspection report. During the process of manually recording inspection data, operators would be alerted to false quality losses.
2. The **CMM Station** could be attacked by altering the digital files that generate the quality report for the CMM operator. This would alert the CMM operator to false quality losses.
3. The **QDMS** could be attacked by altering the digital files that contain KQC tolerance information, alerting the Quality Engineer who would be observing false quality losses.
4. The tolerance definition during **GD&T** itself could be attacked. Despite the fact that this attack occurs at the GD&T development phase, the outcomes of this attack would occur at the Flatness Inspection Station, the CMM Station, and the QDMS. It should be noted that this attack would have to occur during the initial production system set-up (i.e., before tolerance data were received by either inspection station). Once the system transitions to full production, operators at the Flatness Inspection Station/CMM Station or the Quality Engineer would react to false quality losses.

3.3.1.2 Altering the Data Collection Process: Attack Locations

The attack method of altering the data collection process could be applied at two different locations:

1. The **CMM Station** could be attacked and instructed to collect misleading data. For instance, the inspection parameters regarding the CMM ball probe diameter could be changed. This would result in all measurements being incorrect and would alert the CMM operator or downstream QC engineer to a nonexistent quality problem.
2. The **Flatness Inspection Station** could be attacked during the process of uploading the manually collected data to the QDMS. If incorrect data were uploaded to the system, the QC

engineer would be alerted to a problem. Despite the fact that the attack was initiated at the Flatness Inspection station the outcome of the attack occurred at the QDMS.

3.3.1.3 Altering the Data Report Process: Attack Locations

Altering the data reporting process as the attack method could also occur at two locations:

1. The **CMM Station** could be attacked by allowing the CMM to collect the correct data but report altered and nonconforming data. This would alert the CMM operator to a nonexistent problem.
2. The **QDMS** could be attacked by forcing the system to generate reports with incorrect data. This incorrect data could be pulled from a different part or feature, the data could undergo a transformation (e.g., linear scaling) when being accessed from the database, or the data could be randomly generated without ever querying the database.

3.3.1.4 Altering the Collected Data: Attack Locations

Altering the collected data could occur at only one location, which is the QDMS. The **QMDS** could be attacked by directly changing data collected and stored in it. This attack would cause the QC engineer to react to a nonexistent problem.

3.3.2 Attack Objective: Altering Product Design Intent

In this sub-section, the focus will be on implementing the proposed taxonomy when the objective is to alter the product's design intent through forcing very specific product changes and attempting to conceal the attack. According to the taxonomy's 2nd layer - Affected Production Sub-system, both the manufacturing and QC systems need to be affected through a Joint attack to achieve this objective. For this taxonomy implementation example, it will be assumed that a Passive Joint attack will be employed. The QC system itself will not be altered; however, to avoid detection, the manufacturing component of the attack will be designed with respect to QC implementation data acquired through the attack's QC component.

With respect to the taxonomy's 3rd layer - Attack Method, the attack objective of altering a product's design intent could be accomplished by 1) adding part features, 2) altering the KQC dimensions, or 3) removing part features. The details of the taxonomy's final layer - Attack Location, for these cases are discussed as follows.

3.3.2.1 Adding Part Features: Attack Locations

The following two cases will demonstrate the method of adding part features at different locations within the targeted production system. For each of these two cases, the QC component of the

Passive attack obtained information regarding the CMM Station. Specifically, the location of CMM hit points was acquired. With this information, the manufacturing component of the attack could be designed to add extra hole(s) without the risk of being detected by the CMM. In essence, since the CMM is programmed to check for only the 6 holes, there is no way for it to detect any additional holes. One could say that that the effect of the attack is visually quite obvious. However, if this product was produced in large volumes it could easily go unnoticed. In addition, as the complexity of a manufactured part increases or the size significantly decreases it becomes impossible to visually assess a part's features and one must solely rely on the use of measurement systems. The manufacturing component of this attack could be implemented at two locations:

1. An attack could alter the digital design files during the **CAM phase**, which would result in the generation of incorrect CNC code. It should be noted that this attack would have to occur during the initial production system set-up, which would result in this being a sustained attack.
2. The **CNC Drill Workstation** itself could be attacked by replacing or altering the CNC Code data files. In contrast to the previous case, it may be possible to deactivate and then reactivate this attack. This would allow for additional flexibility of this attack to remain concealed.

3.3.2.2 Altering the KQC Dimensions: Attack Locations

The CNC Mill Station is followed by a 100% manual inspection operation, while the CNC Drill Station is followed by an automated inspection operation that is performed for 25% (every fourth part) of the parts produced. With this QC knowledge, an attack could be designed to frequently alter the CNC code at the **CNC Drill Workstation** to change holes' diameters and locations. In this case, if the attacker ensured that attacks never coincided with an inspection, the attack would never be detected through inspection.

3.3.2.3 Removing Part Features: Attack Locations

Similarly, attackers can use the QC knowledge of the CMM inspection being done 25% of the time in their favor to remove part features without being detected. Again, this can only occur at the **CNC Drill Workstation**, because the CNC Mill Workstation is followed by a 100% manual inspection operation. For the CNC Drill Workstation, the attack in this case would be in the form of removing one of the six holes.

3.3.3 Attack Objective: Reducing Outgoing Quality

With the objective of reducing outgoing quality, this taxonomy implementation example will focus on forcing product changes and attempting to conceal the attack. Considering this proposed taxonomy's 2nd layer - Affected Production Sub-system, this objective could be accomplished through a Joint Active attack, affecting both production sub-systems. Due to the nature of Active attacks, the 3rd and 4th layers are very dependent on each other, since the attack requires a coordinated effort across the manufacturing and the QC systems. Specifically, the attack causes the manufacturing system to physically change the product, while also preventing the QC system from being able to detect these changes.

For the manufacturing system component, this proposed taxonomy's 3rd layer - Attack Method, could be implemented by 1) adding part features, 2) altering the KQC dimensions, or 3) removing features within the part. Regarding the taxonomy's 4th layer - Attack Location, the attack could be implemented against the 1) CAD Data Development, 2) CAM Data Development 3) CNC Mill Workstation, or 4) CNC Drill Workstation.

In order to conceal the effect of such an attack within the QC system, either 1) the definition of a conforming product, 2) the data collection process, 3) the data report process, or 4) the collected data needs to be altered as the proposed taxonomy's 3rd layer - Attack Method. With respect to the final layer - Attack Location, the attack could be implemented at either of these locations: 1) GD&T Data Development, 2) Flatness Inspection Station, 3) CMM Station, or 4) the QDMS.

There exist a large number of possible manufacturing system attack method/location and QC system attack method/location combinations possible to accomplish an Active attack. For brevity, in this taxonomy illustrative example it is assumed that the manufacturing system will be attacked to alter the parts dimensions (method) at CAM Data Development (location) by causing the CNC Drill Workstation to choose an incorrect (larger) drill bit from its tool turret. Therefore, the following discussion is of the taxonomy's 4th layer - Attack Location for each of the previously identified QC system attack methods is presented next, where the manufacturing system will be attacked to alter the parts dimensions (larger holes) at CAM Data Development.

3.3.3.1 Altering the Definition of a Nonconforming Product: Attack Locations

With this attack objective, there will be three cases demonstrating the method altering the definition of a nonconforming product:

1. The **CMM Station** could be attacked by altering the digital files that generate the quality report for the CMM operator. In this attack, the tolerances defined on the quality report could be widened to falsely inform the CMM operator that the measurements of out-of-tolerance hole diameters were within specifications.
2. The **QDMS** could be attacked by altering the digital files that contain hole diameter tolerance information. Similar to the previous attack, the Quality Engineer would be falsely led to believe that out-of-tolerance parts were within specification.
3. The tolerance definition during **GD&T** itself could be attacked. Despite the fact that this attack occurs at the GD&T development phase, the outcomes of this attack would occur at the CMM Station and the QDMS. It should be noted that this attack would have to occur during the initial production system set-up. Once the system transitions to full production, operators at the CMM Station or the Quality Engineer would never be alerted to quality losses.

3.3.3.2 Altering the Data Collection Process: Attack Locations

This attack method of altering the data collection process would only occur at one location. The **CMM Station** could be attacked and instructed to collect wrong data (i.e., wrong data that suggests good parts). For instance, the inspection parameters regarding the CMM ball probe diameter could be adjusted to compensate for the change in hole diameters. This would result in all holes' diameter measurements incorrectly considered to be within tolerance.

3.3.3.3 Altering the Data Report Process: Attack Locations

The attack method of altering the data reporting process could be applied at two different locations:

1. The **CMM Station** could be attacked by allowing the CMM to collect correct data but report completely different and conforming data. This could be done by running a subroutine through the CMM's programming that generated random hole diameters within the tolerances data.
2. The **QDMS** could be attacked by instructing the system to generate reports with incorrect data. For instance, incorrect conforming data could be randomly generated without ever querying the database.

3.3.3.4 Altering the Collected Data: Attack Locations

Altering the collected data could occur at only one location. The **QMDS** could be attacked by directly changing the collected data within the database as the attack method in this case.

Table 3-2 Summary of taxonomy details for the implementation example.

Layer	Choices	
Attack Objective: Disrupting Manufacturing Operations		
Specific Objective:	Incorrectly classifying products as nonconforming	
Affected Production Sub-system(s):	Quality Control (DQ Attack)	
Attack Method(s):	• Altering the definition of a nonconforming product • Altering the data report process	• Altering the data collection process • Altering the collected data
Attack Location(s):	• GD&T data development • CMM Station	• Flatness Inspection Station • The QDMS
Attack Objective: Altering Product Design Intent		
Specific Objective:	Forcing very specific products changes while attempting to conceal the attack	
Affected Production Sub-system(s):	Manufacturing *and* Quality Control (Passive Attack)	
Attack Method(s):	• Adding part features • Removing features within the part	• Altering the KQC dimensions
Attack Location(s):	• CAM Data Development	• CNC Drill Workstation
Attack Objective: Reducing Outgoing Quality		
Specific Objective:	Forcing products changes while attempting to conceal the attack	
Affected Production Sub-system(s):	Manufacturing *and* Quality Control (Active Attack)	
Attack Method(s):	• Removing features within the part • Altering the KQC dimensions • Adding part features	• Altering the definition of a nonconforming product • Altering data collection process • Altering data report process • Altering the collected data
Attack Location(s):	• CAD Data Development • CAM Data Development • CNC Mill Workstation • CNC Drill Workstation	• GD&T data development • Flatness Inspection Station • CMM Station • The QDMS

It should be noted that the aforementioned QC attacks may not necessarily guarantee a successful Active attack. For instance, attacking the data reporting process at the CMM station only prevents the CMM operator from detecting the effects of the attack. Since the actual holes' diameter measurements were altered, it would still be possible for an engineer to detect this attack through the QDMS. This suggests that Active attacks may require coordinated attacks across repetitive/redundant QC system components to be successful.

While this example illustrated only one case from a large set of possible combinations, it is important to note that with the QC system being actively attacked, a new set of attacks is now possible and some of the Attack Methods in the manufacturing system would have additional locations. For instance, it would be now possible to attack the CNC Mill Workstation when altering the KQC dimensions, since a corresponding DQ attack on the following Flatness Inspection Station could cover this one up. As another example, the CAD/CAM files generated in the Data Development phase could also be altered (or replaced) more freely now when removing features or altering the KQC dimensions, since the corresponding DQ attack on the following CMM Inspection Station would make sure that it goes undetected.

A summary of the mentioned taxonomy implementations for each of the three presented examples is given in Table 3-2. As can be seen, there is a wide range of potential attack methods that could be applied at almost every location within the production system.

3.4 System-Specific Attack Design Considerations

In this section, system-specific considerations required to successfully design an attack against a production system are discussed. As demonstrated in the taxonomy implementation example, there are numerous possible ways to accomplish a specific attack objective. The final attack design would depend upon a multitude of system-specific attack design considerations, which can be regarded as a 5th layer to the taxonomy. Those specific considerations are briefly described next, and include but are not limited to: 1) component accessibility; 2) attack pervasiveness; 3) attack impact; and 4) attack coordination.

3.4.1 Component Accessibility

Possibly the most important (and most constraining) system-specific design consideration is whether or not a specific production sub-system component is accessible or can be attacked. This consideration is a function of several factors, which include but are not limited to: 1) attacker's (individual or group) skill; 2) ability to access information regarding the targeted component; 3) time available/dedicated to accomplishing a specific attack; 4) system/sub-system/component security measures; 5) attacker's resources; and 6) number of alternative attacks that would accomplish the same objective. The large number of factors involved in system component accessibility could result in a myriad of varying sub-systems accessibilities for prospective attacks that would accomplish the same objective.

3.4.2 Attack Pervasiveness

Manufacturing and QC systems are typically composed of sequential stations (or operations), referred to herein as stages, that are required to transform raw materials into quality goods. Given this sequential nature, each stage is dependent upon it predecessor(s) and any information passed from one stage to the next is *trusted* to be correct. Additionally, the effects of a cyber-physical attack do not necessarily occur at the affected stage, but could appear downstream. Attacks occurring at earlier stages could result in more decisions downstream to be made incorrect, due to the progression of the malicious attack effects; e.g., the last stage in a production process can only impact itself (most likely). The earlier the "roots of trust" are contaminated, the more pervasive the attack would be. An attacker's desired attack pervasiveness could therefore limit possible attack locations.

3.4.3 Attack Impact

Different cyber-physical attacks made against a manufacturing system can have drastically different impacts. For instance, consider two attacks with the same objective (disrupting manufacturing operations) targeting the same production system (manufacturing) at the same location (milling machine) that are using two different attack methods. One attack method could cause unscheduled maintenance for a specific machine, resulting in hours of downtime. The other attack method could completely destroy a specific machine, resulting in days to weeks of possible downtime. An attacker's desired attack impact could therefore limits possible attack methods (or locations).

3.4.4 Attack Coordination

Attacks against manufacturing are not constrained to be individual attacks acting at one location within the production system. Joint Attacks, by definition, require the attack to occur at two separate locations, one in the manufacturing system and one in the QC system. Furthermore, to enhance their impact, DM or DQ attacks could be implemented at numerous locations throughout the production system. For attacks that occur at multiple locations in a production system, it may be necessary to ensure that these attack efforts are appropriately coordinated (with respect to time).

Consider an Active attack with a manufacturing system attack component that alters the dimensions of a part's features, which will eventually be inspected by a CMM. The QC system attack component will alter the parameters of the CMM to compensate (or offset) for the incorrect feature dimensions, so that out-of-tolerance parts will appear to be within specification. However, a time lag exists between when the first part affected by the attack's manufacturing component will be inspected by the CMM. If both attack components (manufacturing and QC) are implemented simultaneously, the attack on the CMM will compensate for parts that have not been subjected to the manufacturing system attack. This would result in actual within specification parts to appear out-of-tolerance, alerting the QC system to a quality loss. In order for such an attack to be successful, all individual attacks would have to be effectively coordinated, with respect to time. The inability or difficulty in coordinating such attacks would limit the attack design space.

3.5 Summary and Discussion

In this chapter, an attack taxonomy that governs the relationships between QC systems, manufacturing systems, and cyber-physical attacks in the context of malicious process changes has been proposed. The proposed taxonomy was developed from a QC perspective, while accounting for the attacker's point of view to help manufacturers better understand existing vulnerabilities and secure their production systems against cyber-physical attacks. In addition to describing the different design consideration layers incorporated into the attack taxonomy, an example of implementing the taxonomy in a practical situation was provided. Finally, specific system-related design considerations that need to be taken into account during implementing this attack taxonomy were also discussed.

One of the overarching goals of developing this taxonomy is to further raise awareness to the threat cyber-physical attacks pose to manufacturing. If any manufacturer who ignores the possibility of cyber-physical attacks against their production system falls victim to a cyber-attack, they will be ill-prepared to recover. An initial step toward enhancing a production system's resilience to cyber-physical attacks could be as simple as including cyber-physical attacks as possible failure modes or root causes when performing a Failure Mode and Effects Analysis (FMEA) or a Root Cause Analysis (RCA), respectively. Through the use of the proposed taxonomy, QC engineers together with production engineers, IT specialists, and equipment operators could determine the effects possible cyber-physical attacks would have on a production system. This knowledge added to an FMEA would help in the (re)design of products and/or processes to be less susceptible to attacks. When incorporated within an RCA, this information would aid in quick system recoveries when a cyber-physical attack occurs, increasing overall production system cyber-attack resiliency.

Future work in this particular research area will focus on applying this taxonomy in an actual manufacturing setting. Such an implementation will then enable describing the potential attack surfaces in a production system from a QC perspective and analyzing the different corresponding QC-related types of possible attacks. Additionally, the proposed attack taxonomy, along with the corresponding QC approaches that will be discussed in subsequent chapters, could be combined with other methods for detecting the physical effects of the attacks to form an all-inclusive security framework corresponding to cyber-physical threats on both manufacturing and QC systems.

4 Cyber-Physical Attack Vulnerabilities in Manufacturing Quality Control Tools

The next step toward the development of new QC tools is exploring the effects of exploiting existing weaknesses in QC systems, such as misusing specific QC tools, within cyber-physical production systems. Hence, opportunities where QC tools could be exploited by cyber-physical attacks within manufacturing are identified in this chapter; along with their possible subsequent adverse effects. Furthermore, best practices and potential mitigation guidelines for better cyber-physical security in manufacturing are also presented in this chapter. The idea of the work in this chapter was presented at the INFORMS 2015 Annual Meeting (Elhabashy *et al.*, 2015b) and further details were discussed at the IISE 2018 Annual Conference (Elhabashy *et al.*, 2018b). A corresponding journal article is currently being finalized for submission to the *Quality and Reliability Engineering International* (Elhabashy *et al.*, 2019a).

4.1 Introduction

4.1.1 Background

With the latest technological advancements, manufacturing systems are evolving into highly integrated Cyber-Physical Systems (CPSs) that rely on their cyber components as much as they do on their physical ones. Such systems are susceptible to cyber-physical attacks, where an attack targets a cyber component that ultimately results in affecting a physical component (Yampolskiy *et al.*, 2013). In manufacturing, such attacks can result in changed product designs, manipulated manufacturing equipment, and altered final products. While historically infrequent, several examples of cyber-physical attacks against manufacturing systems have occurred. One example is the infamous “Stuxnet” virus, which was responsible for destroying as many as 1,000 high-speed uranium enrichment centrifuges between late 2009 and early 2010 (Albright *et al.*, 2010). A more recent example was a spear-phishing attack against a German steel mill in 2014, which resulted in multiple cyber components failing and massive physical damage (Lee *et al.*, 2014).

One key reason that cyber-physical attacks are becoming more relevant in manufacturing can be attributed to the transition towards Industry 4.0. Manufacturers are placing more emphasis on the ability to automate equipment and easily exchange data, which increases the opportunities for cyber-physical attacks. For instance, the growth of the Industrial Internet of Things (IIoT) has resulted in more manufacturing equipment (e.g., machines, sensors, cameras, etc.) to be connected to the internet. Furthermore, for more efficient data extraction and sharing, additional communication protocols and cloud computing tools (such as MTConnect and Software as a Service (SaaS), respectively) have been incorporated into modern manufacturing systems. Such an increased connectivity has encouraged more reliance on creating almost everything within a virtual digital environment using Computer-Aided Engineering (CAE) support tools, such as Computer-Aided Design (CAD) and Computer-Aided Manufacturing (CAM) software.

Relying solely on traditional approaches for Information Technology (IT) security is insufficient for CPSs, such as manufacturing systems (Cárdenas *et al.*, 2009; 2011). The related literature has identified several reasons why IT security approaches, although necessary, are not sufficient. As reported by Cárdenas *et al.* (2009); (2011), these reasons include: 1) traditional IT security practices of patching and continuous updates are not well suited for CPSs; 2) CPSs often consist of legacy equipment with diverse assortment of operating systems; and 3) CPSs affect the physical world. One additional reason traditional IT security approaches may be insufficient toward manufacturing is that manufacturing systems tend to contain wide ranges of heterogeneous processes that result in complex IT structures (Elhabashy *et al.*, 2018a).

4.1.2 Motivation

In manufacturing, QC systems are typically used to ensure the production of high-quality parts and to maintain stable processes. With the evolving nature of manufacturing systems, the new security requirement for QC systems is to promptly detect cyber-physical attacks. Given that cyber-physical attacks can result in physical manifestations, these QC tools may be able to provide a second line of defense against attacks that have successfully bypassed traditional IT security. However, in their current state, QC tools may not be able to detect attacks in a timely manner (or perhaps even at all), which can be attributed to three reasons.

Firstly, QC tools themselves are often highly integrated components in manufacturing systems and are also susceptible to cyber-physical attacks. Secondly, current QC tools are simply not designed to detect process/product changes caused by a cyber-physical attack (Wells *et al.*, 2014; Vincent *et al.*, 2015). More specifically, QC tools are based upon assumptions that may no longer be valid with respect to a system being subjected to an attack (Wells *et al.*, 2014; Vincent *et al.*, 2015), which would render them less sensitive to malicious process changes. Thirdly, as mentioned by Wells *et al.* (2014), not only could assumptions be made invalid in the presence of an attack, but they could also be used against the QC system to make such attacks undetectable.

As part of their cyber-physical security taxonomy for production systems, Elhabashy *et al.* (2018a) showed how QC systems could be involved in cyber-physical attacks and separated these attacks into three categories: 1) Direct Quality (DQ): Attacks that only involve the QC systems, 2) Direct Manufacturing (DM): Attacks that only involve the manufacturing systems, and 3) Joint: Attacks that involve both the QC and manufacturing systems. The latter type, the Joint attacks, are quite dangerous for manufacturing systems. Unlike DQ or DM attacks, where attackers can alter a product and hope that the effects of their attacks are not detected, Joint attacks are implemented with higher chances of avoiding detection by QC tools while decreasing a product's quality, functional performance, reliability, or all of those together. Hence, this work will not consider DQ nor DM attacks, and will focus only on Joint attacks.

These joint attacks can be further broken down, as suggested by Elhabashy *et al.* (2018a), into two categories: Passive and Active. In a Joint Passive Attack, the QC system is not altered, but information regarding the QC systems is accessed, collected, and used in the design and/or implementation of the manufacturing system component of the Joint Attack. In Joint Active attacks, the QC system is altered in a way such that the effect of the manufacturing component of the attack is undetectable; e.g., modifying inspection data. Passive Joint attacks will be the main focus of this work because: 1) passive attacks may be easier to implement and therefore more prevalent and 2) there aren't many opportunities to design QC tools that are resilient to Active attacks, since they directly change QC information, unlike the Passive ones.

As an example of a Passive Joint attack, consider the case-study presented by DeSmit *et al.* (2017), which affected both the manufacturing and QC systems during the production of a jet engine bracket. The QC component of the attack involved: 1) knowing that the part was being inspected using a Coordinate Measuring Machine (CMM) and 2) obtaining the Geometric Dimensioning and Tolerancing (GD&T) information of the part and using it to determine how the part could be altered in a way to avoid detection during inspection. For the manufacturing system component of the attack, tool-path files for the milling operation to machine the jet engine bracket were modified, which resulted in three of part's features being altered. This Passive attack was successful because the induced part alterations were designed to not be detected by the CMM or even noticed visually. Despite being carried out in an experimental setting, this example showcased how a Passive Joint attack can result in a QC system not being able to detect maliciously induced quality losses and the need for developing QC tools that are more robust against cyber-physical attacks.

Despite the recent increased attention given to cyber-physical security in manufacturing, research efforts to secure product quality have mostly focused on developing entirely new approaches to detect cyber-physical attacks, rather than improving and/or adapting existing QC tools. These newly proposed approaches are based on using alternative side-channel measurements (features) for attack detection, such as the work by Vincent *et al.* (2015), Sturm *et al.* (2016), Chhetri *et al.* (2016), Albakri *et al.* (2017), Wu *et al.* (2017), Moore *et al.* (2017a), and Belikovetsky *et al.* (2018).

Furthermore, although references to the inability of modern QC tools to detect cyber-physical attacks against manufacturing systems exist within the literature (Wells *et al.*, 2014; Vincent *et al.*, 2015), there has been no work done to effectively categorize and understand the nature of QC tool vulnerabilities and how they could be exploited by Passive attacks. Therefore, more effort is needed to explore these weaknesses, which is purpose of this work. More specifically, the goal in this chapter is to identify such exploitable weaknesses in QC tools and categorize them effectively. This work also aims to demonstrate the harmful effects if such exploitations were to take place by cyber-physical attacks. This demonstration would help raising awareness amongst the manufacturing community of the significance of exploiting QC tools vulnerabilities.

Two different techniques through which QC tools vulnerabilities can be exploited are identified in this work. These two techniques are referred to herein as “exploitation classes” and are outlined in the following section. Each of these two exploitation classes is then described in details separately throughout Sections 4.3 and 0, where the effects of the exploitation by cyber-physical attacks are also shown. Section 4.5 afterwards contains a discussion about potential mitigation guidelines. Finally, the chapter is concluded in Section 4.6 with a summary.

4.2 Overview of the Exploitation Classes

The effective design and implementation of QC systems rely heavily on making decisions regarding: 1) how data is analyzed and interpreted and 2) how data is collected in the first place. An example of a data analysis decision is determining which type of control chart to use with a set of collected data. Alternatively, identifying appropriate part/process features to measure and how often to collect them is an example of a decision regarding how data is collected. It should be noted that how data is analyzed and interpreted is highly dependent upon how it was collected.

All these decisions are based on some type of assumptions. For example, during the process of analyzing and interpreting data, the choice of the most appropriate tool to use is done assuming the tool will be used properly. As another example, it is assumed that quality losses are only attributable to causes originating from within a manufacturing system (such as those due to changes in materials and processes) and are not due to malicious external causes (such as cyber-physical attacks) during data collection. Such assumptions, on which the decisions are based, could be incorrect, causing wrong decisions to be taken. When wrong decisions are made, vulnerabilities would exist in the used QC tools, which could then be exploited by Passive cyber-physical attacks.

This chapter identifies two types of “exploitation classes”, which are a result of wrong (or incomplete) decisions, described next:

1) **Exploiting bad decisions:** In this class, Passive attacks take advantage of the misuse of QC tools by users and is referred to as *exploiting bad decisions*. These exploitations occur when wrong decisions are taken during analyzing and interpreting the data the users had collected; due to assuming that QC tools are used properly, when they are not. In this case, the collected data could be adequate for users to detect cyber-physical attacks. However, this ability to detect such attacks is lost when the following analysis is not adequate or correct due to QC tool exploitation.
2) **Making decisions bad:** Attacks based on the second exploitation class exploit the vulnerabilities resulting from users not considering the possibility of cyber-physical attacks in the first place when taking decisions regarding data collection; this exploitation class is referred to as *making decisions bad*. The users’ decisions are not necessarily incorrect, they would have been appropriate, if it hadn’t been for cyber-physical attacks. The ability to detect attacks during data analysis is diminished due to the data collection process being no longer adequate as a result of QC tools exploitation. In this case, the collected data is not enough or the correct data needed to be able to detect a Passive cyber-physical attack.

The main characteristics of each of these two exploitation classes are summarized in Table 4-1. Each of the exploitation classes will be discussed in more details in the following two sections, respectively. It should be noted that the main focus of these two discussions will be primarily on Statistical Process Control (SPC) techniques (and especially control charts), as they are one of the predominant QC tools used in industry.

Table 4-1: Summary of the exploitation classes characteristics discussed in this chapter.

Exploitation Class	Is a Result of	Data Collection Process	Data Analysis Process
Exploiting Bad Decisions	Misusing QC tools	Adequate	Inadequate or Incorrect
Making Decisions Bad	Not considering the possibility of cyber-physical attacks	Inadequate or Insufficient	Adequate

4.3 Exploiting Bad Decisions

For this exploitation class, attackers take advantage of vulnerabilities resulting from the misuse of QC tools by users and industrial practitioners during the data analysis and interpretation process. Due to the misuse, the analysis done for the collected data is not adequate to detect a cyber-physical attack. Hence, the discussion in this section is focused on showing how QC tools can be misused, as a result of wrong decisions taken during the data analysis process.

There exist a number of reasons why QC tools are misused. For instance, the lack of a sufficient understanding regarding the importance of a QC tool's underlying statistical assumptions may lead to the improper use of a tool (i.e., tools are applied without ensuring statistical assumptions are valid). Additionally, QC tool users often have strong (and dangerous) beliefs regarding the infallibility of QC systems. Such beliefs in tool infallibility may prevent users from questioning the legitimacy of any of the results obtained. Misusing QC tools may lead QC practitioners to not realizing that they are using invalid or incorrect information about the production environment from the QC system. Such misuse can also prevent users from being able to perform proper assessment of their systems, leading to having false impressions about product quality.

In this work, QC tool misuse is separated to the three categories presented next. Then, a discussion about some of the effects of exploiting the weaknesses within QC tools in each of these categories on the tools' performance follows.

4.3.1 Misusing QC Tools Categories

It is important to realize that the misuse of QC tools is historically common in manufacturing. Based on a literature review, this sub-section highlights cases of misuse when utilizing quality control tools. From this review, we propose that QC tool misuse can be grouped into three categories: 1) Improperly Implementing QC Tools, 2) Unknowingly Violating QC-Related Statistical Assumptions, and 3) General Misconceptions.

4.3.1.1 Improperly Implementing QC Tools

The first category discussed is improperly implementing QC tools within a production system. As an example, consider the study performed by Bird and Dale (1994) at an automotive supplier specializing in manufacturing small high-precision metal components. In this study, the authors noted significant instances of SPC tool misuses. One thing they noticed was the fact that Gauge Repeatability and Reproducibility (R & R) studies were not conducted before applying other QC tools, or at all (Bird and Dale, 1994). This prevented the manufacturer from being able to differentiate between variation from the system and variation from the measurement process. Another thing they noticed was that the measurement tools themselves required frequent recalibration and there that were no standards used when measuring, with the procedure varying between operators (Bird and Dale, 1994).

In addition to these instances, Bird and Dale (1994) also noticed, at the same manufacturer, that the control charts used by the automotive supplier had control limits only during the day shift. Moreover, regardless of the reasoning, operators in the other shifts did not seem to find this lack of control limits in their shifts to be an issue. Not only was this omission of control limits odd, it made the operators believe that points plotted on the charts shouldn't be far from the center line, which led them to resetting the machine frequently, despite being in-control (Bird and Dale, 1994).

In another study, Hoyles *et al.* (2007) interviewed employees of four different companies, two of which were packaging companies, a pharmaceutical company, and an automotive manufacturer. The authors were able to gain interesting insights regarding applying SPC in practice within those industries. The authors noticed that if operators responded to control charts, they tended to react unnecessarily too quickly to data points that were far from the center line but still in-control. This overreaction caused a degradation in the overall production process since it introduced unnecessary variability and could have been just a single isolated case (Hoyles *et al.*, 2007). However, operators didn't respond often enough to trends existing in the data. The lack of response to trends caused expensive losses to one of the companies (Hoyles *et al.*, 2007).

Similar to the work of Hoyles *et al.* (2007), Wood and Preece (1992) presented case studies at three different companies. The authors observed in two of the companies that the constructed control limits were incorrect, which prevented the users from detecting and therefore investigating potential assignable cause variations. For one company, these incorrect control limits were caused by using incorrect equations. For the other company, these incorrect control limits were a result of trying to make the charts look more like textbook examples (Wood and Preece, 1992). In addition, the authors also noticed that some control charts were based on very little available data and that some of the users tended to use techniques from what they referred to as an “approved menu”, rather than appropriate tools, only due to pressure from within the company (Wood and Preece, 1992).

As a final example for this misuse category, consider the work of Ermer and Hurtis (1995), in which the authors presented a study on the application of SPC techniques to two soldering processes used to manufacture electronic card assemblies. The authors noticed that the majority of attribute control charts (c- and u-charts) used had upper control limits less than 1, meaning that zero defects are required for stability which is just not practical (Ermer and Hurtis, 1995). Such a situation of having a low upper control limit will probably lead to a significant number of false alarms, which would be costly in such a high production system.

4.3.1.2 Unknowingly Violating QC-Related Statistical Assumptions

QC tools can be misused by unknowingly violating their statistical assumptions during tool design and/or implementation. Examples of such statistical assumptions include, but are not limited to, 1) presuming that a certain tool is appropriate for the collected samples, 2) assuming underlying statistical distribution of observations, and 3) assuming observations are independent.

For a QC tool such as a control chart, the choice of the most appropriate one to use with a given set of observations needs to be based on specific assumptions regarding the process being monitored. Such assumptions should relate to the type of Key Quality Characteristic (KQC) of interest, the observations distribution, the type of process variations to be detected, and the size of shifts in mean and/or variance. For example, attribute control charts should be used to monitor the fraction of non-conforming products or number of defects, Exponentially Weighted Moving Average (EWMA) charts should be used to detect small mean shifts, and so on. Users may incorrectly select a control chart type for a given situation, such as using an EWMA chart to detect large mean shifts, due to inaccurate judgment which causes misleading results.

In addition to the assumptions mentioned in the previous paragraph, the choice of which control chart to use also depends on the distribution of the collected samples. So, another case of misuse in this category is using control charts designed for data following a certain distribution, when the actual observations follow another. For example, individual and moving range control charts should be only used with samples that are known to follow a normal distribution, as pointed out by Benneyan (2009). The author showed cases of incorrectly choosing to use the individual and moving range control charts with observations following binomial, Poisson, geometric, and exponential distributions, with all of them leading to inaccurate representations and decreased performance (Benneyan, 2009).

A similar, yet different, situation is assuming a certain distribution for the collected observations without knowing (or checking) the actual distribution. This is an important statistical assumption which needs to be always validated after first attaining stability in Phase I. Some control charts require the statistic being monitored to follow a normal distribution. However, users often do not perform any statistical tests or visualizations to confirm this behavior and just assume that it is true. This common mistake has been documented in applying control charts (Wood, 1995) and in performing process capability studies (Holmes and Mergen, 2012). Furthermore, the assumption of normality for control charts may not hold when the number of samples used is significantly small, as with automated processes (Cai *et al.*, 2001).

Another commonly violated statistical assumption is that the observations are independent. There are situations where the independence assumption will not hold, such as in an automated manufacturing environment where inspections are very frequent (Cai *et al.*, 2001). It is important to always consider checking for the independence assumption, as evident in the case-study discussed by Warner *et al.* (1993). In that case-study, it was assumed that the collected observations were independent and they were plotted directly on a control chart (Warner *et al.*, 1993). The existing autocorrelation, however, caused many out-control-signals, leading process engineers to mistakenly believe that there was something wrong with the process and immediate action was required (Warner *et al.*, 1993).

Other studies have also shown that when one or more of the underlying statistical assumptions concerning the collected samples is no longer valid, then the choice of the QC tool used will mostly likely be incorrect (Warner *et al.*, 1993; Wood, 1995) and the user may not aware of this. Lastly, the negative effects of the blind use of control charts with violated statistical assumptions, such as samples' normality, independence, and constant variance were shown by Sparks and Field (2000). Using Deming's well-known funnel experiment, the authors demonstrated the extent to which the results of different Shewhart control charts with those violated assumptions would be deceptive (Sparks and Field, 2000).

4.3.1.3 General Misconceptions

Finally, there are misconceptions (or incorrect beliefs) that some users have about QC tools which can contribute to misuse. One example is assuming that the product's specifications limits can be used instead of control chart limits and that they are the same, which is a misconception observed in multiple facilities (McCoy *et al.*, 2004; Hoyles *et al.*, 2007). The control and specifications limits are entirely different. The control limits are driven by natural process variability, whereas the specification limits are set externally by management, engineers, customers, or product developers/designers (Montgomery, 2009). Replacing control limits with specification limits could result in inaccurate conclusions about the processes being monitored. A similar misconception is believing that the purpose of SPC is only acceptance and rejection, rather than process control to understand and prevent future defects (Wood and Preece, 1992).

Other examples of misconceptions are when users assume that an out-of-control process implies not meeting customer requirements (Wood, 1995; Holmes and Mergen, 2012); or that if a process is stable, then it must be meeting customer specifications (Holmes and Mergen, 2012). What the customers want may not directly translate into the used control limits. The monitored process may be stable and in control, but still not meeting the customer requirements (Holmes and Mergen, 2012). Also, assuming that control charts are only needed to reach stability and then discontinuing using them once stability is attained (Bird and Dale, 1994) is another misconception by QC users.

Additional examples of misconceptions about the use of QC tools also include, but are not limited to: 1) seeing no need for involving the production personnel in running the QC tools (Bird and Dale, 1994); 2) believing that a "standard procedure" must be followed for constructing control charts without thoroughly considering the number of group sizes chosen or even the type control chart to be used (Caulcutt, 1995); 3) believing that SPC is only about having plotted control charts on walls to satisfy customers (Antony and Taner, 2003); and 4) believing there is no need for reaction plans to deal with out-of-control situations (Bird and Dale, 1994).

4.3.2 The Effects of Cyber-Physical Attacks on the Categories of QC Tools Misuse

Not only do these misuses in the three categories reduce QC tools effectiveness they can also be exploited by cyber-physical attacks. So, understanding the effects of such misuse exploitation in the context of cyber-physical attacks is particularly important with the current advancements in computer and internet technology in Industry 4.0 and the reliance on cyber-physical systems.

Cyber-physical Passive attacks that exploit misused QC tools can result in QC systems incorrectly presenting themselves as stable and in-control, leading to decreased tool performance and false sense of quality. This sub-section will contain examples of some of the effects of these Passive attacks and their impact on the misused tools. More specifically, it will be shown: 1) how cyber-physical attacks exploiting QC tools misuse can be designed to minimize detection in the manufacturing environment; and 2) how the performance of the QC tools will decrease in the presence of these Passive attacks, without the users even knowing. It should be just noted that the examples shown next are grouped into the three discussed misuse categories.

4.3.2.1 Misuse Category 1 Example – Using a Stand-alone X-bar Control Chart

The first example presented in this sub-section is a case of using just an X-bar control chart, without an accompanying Range (R) or Standard Deviation (S) charts, for manufacturing process monitoring. In this case, the X-bar control chart is used for Phase II monitoring of a certain product's KQC and contains five observations per sample. The decision was made that a stand-alone X-bar chart is enough and that there is no need to monitor the within samples variance. It could be argued that this is an extreme case that may not occur in an actual manufacturing setting. However, industrial practitioners who specialize in the development of QC software have confirmed to the authors that this particular situation is more common than one would expect.

Opting to use a stand-alone X-bar chart, without using an accompanying R or S chart, is a bad decision within the data analysis and interpretation process. Specifically, an X-bar chart alone does not take the within samples variability into account and the decisions taken in this process are no longer enough to prevent a carefully designed cyber-physical attack on that product's monitored KQC. The data collection process itself is, however, adequate in detecting a potential attack on that particular KQC.

To highlight the effects of a cyber-physical attack along with this bad decision, consider a set of randomly generated samples, each containing five observations, to represent the monitored product's KQC, which will be a victim of a cyber-physical attack. These observations follow a distribution of N(3,1) and are used to estimate the 3-sigma control limits[10] for Phase I, where any

[10] When estimated values are used for the process' mean and standard deviation, the traditional 3-sigma control limits for an X-bar control chart are evaluated from the equations: $\bar{\bar{x}} \pm A_2\bar{R}$; where $\bar{\bar{x}}$ is the mean estimator, $\bar{R}$ is the average range, and A_2 is a constant that equals 0.577 when the number of observations per sample (n) is 5 (Montgomery, 2009).

out-of-control points were excluded. After stability is obtained in Phase I and the practitioners start to use the chart for monitoring during Phase II, the attack is launched on a set of 50 samples.

More specifically, an attacker who has already accessed the cyber-physical manufacturing system, takes advantage of this misuse situation and implements the attack through: 1) deciding on the number of samples to attack; 2) choosing the specific samples to target; and 3) for each targeted sample, making physical changes to the KQC that only affect the sample's range. However, the attacker cannot know in advance the exact values each observation will have for the KQC. So, before monitoring starts in Phase II, the attacker generates a set of 50 random samples with observations following the same distribution and modifies the KQC values of the observations within the already targeted samples. The attacker then alters the manufacturing process to physically change the KQC for each observation to match its corresponding value in the newly generated dataset.

The goal of the attacker is to increase the value of two observations per a sample, decrease two other observations values per this sample, and leave the remaining observation value alone. With doing such changes, the range of the attacked samples will be affected more than its mean; an example of such a desired alteration in one sample can be seen in Table 4-2. Due to the manufacturing process' natural variation, the attacker will not be able ensure that the physically altered values and the generated ones are an exact match in practice. Some type of error will exist due to the difference between the two values, and the physical changes in the attacked samples will not be completely balanced out as shown in the last row of Table 4-2. However, since the error value is going to be very small, it could be neglected, and it would be assumed that the alterations in the observations in each sample do not cause a significant change in each attacked sample's mean.

Table 4-2: A numerical example to show the difference between physically altered KQC values for one sample during an attack and the values they would have taken in case of no attack implemented.

Observation	1	2	3	4	5	Mean	Range
KQC values assuming no attack is implemented	3.2	2.8	3	3.1	2.9	**3**	**0.4**
Altered KQC values after attack implementation	1.2	0.8	3	5.1	4.9	**3**	**4.3**

In this case, the changes caused by the cyber-physical Passive attack in the KQC across the different samples would not be detected by the stand-alone X-bar control chart, as can be seen in the following figure. More specifically, Figure 4-1 contains two identical-looking X-bar charts. The chart on the right (Figure 4-1b) shows the X-bar chart with the attack implemented and KQCs physically altered, whereas the one on the left shows the chart for the same dataset *without* the attack and the physical alterations. The charts, however, only look identical; as the mean of the attacked samples, designated by a magenta asterisk in Figure 4-1b, is almost the same as that of the corresponding samples in Figure 4-1a. Hence, the attacker was successful in hiding the effects of the attack due to the QC tool misuse, and the users would not be able to even realize that the monitored process mean is no longer stable and that the within samples variance has increased.

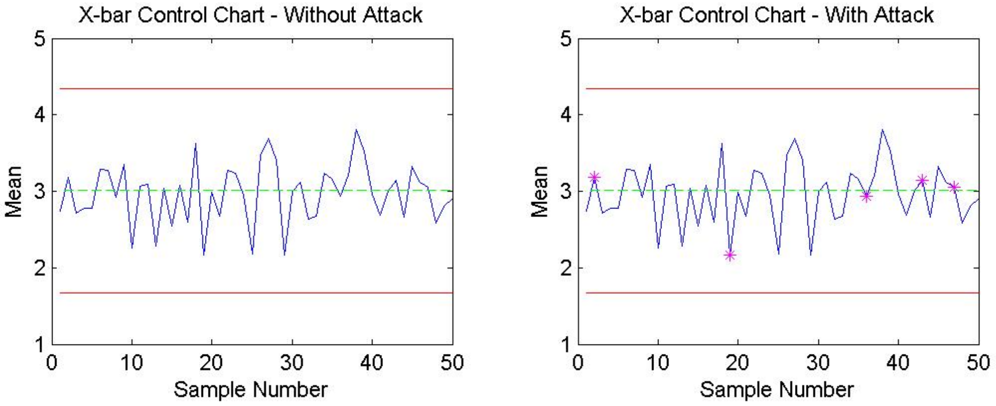

(a) (b)

Figure 4-1: The resulting X-bar chart without (a) and with (b) a cyber-physical attack, as an example of the effects of improperly using QC tools. Samples affected by the cyber-physical attack are designated with a magenta asterisk.

If an R chart[11] had been used as well, the out-of-control samples due to the cyber-physical attack would have been directly detected. As illustrated in Figure 4-2, a corresponding R chart detected all such samples and the users can then investigate into the causes of these process shifts. Hence, it is clear that relying on a standalone X-bar chart was a bad decision, as an X-bar chart was not able to detect that the within samples variance was not constant nor any of the situations where the KQC values were out-of-control due to a malicious attack; whereas an accompanying R chart would have been able to detect all of them (100%), in this case

[11] When estimated values are used for the process' mean and standard deviation, the centerline for an R chart equals to $\bar{R}$ and its upper and lower control limits are evaluated from the equations: $UCL = D_4\bar{R}$ and $LCL = D_3\bar{R}$, respectively; where $\bar{R}$ is the average range and both D_4 and D_3 are constants that are equal to 2.114 and 0, respectively, when the number of observations per sample (n) is 5 (Montgomery, 2009).

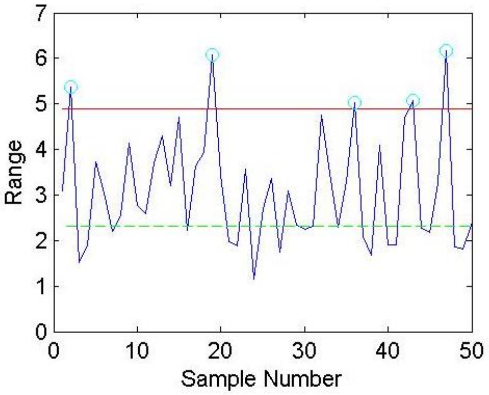

Figure 4-2: A corresponding R chart that would have been able to detect the effects of the cyber-physical attack in this case. Out-of-control points are designated with a cyan circle.

It should be just noted that this attack did not need to alter the QC system data, and only used knowledge from this system to physically alter the manufactured products. In addition, having an R chart is not only beneficial in the existence of process shifts caused by cyber-physical attacks, as demonstrated in this example, but also in detecting increasing variance within samples from other sources, which would not be picked up by the X-bar chart alone, that would cause the process to be out-of-control.

4.3.2.2 Misuse Category 2 Example – Assuming Independence within Samples

This second example illustrates a case where a statistical assumption is unknowingly violated; in particular, the assumption of independence. On one hand, if the collected samples are positively autocorrelated and plotted on a control chart, the limits are narrowed causing a large number of out-of-control signals (Montgomery, 2009). On the other hand, when negatively autocorrelated samples are plotted on a control chart, the limits become too wide and the chart may not be able to detect process mean shifts (Maragah and Woodall, 1992). So, it is crucial to always check for samples independence before using it; otherwise, users might be unnecessarily looking into the data itself to investigate non-existing assignable causes (due to positive autocorrelation) or using control charts with misleadingly wide control limits (as a result of negative autocorrelation).

For this example, the bad decision by the users is not checking for the independence of the collected samples before using it, and the attackers will take advantage of such a decision. Such a decision is taken during the data analysis process, making the process incapable of detecting a cyber-physical attack. To illustrate this case of QC tools misuse, consider a randomly generated dataset representing the collected samples of a product’s KQC (x), with a mean (μ) and standard deviation (σ) of 10 and 1 units, respectively. The generated samples are normally distributed, independent, and are first used to obtain the Phase I control limits of an I-chart. The QC users wait until all Phase I data are collected before starting the data analysis process and never check for the data independence at any point in time.

For the corresponding attack scenario, attackers have already accessed the cyber-physical production system and have full knowledge about the aforementioned data collection procedure. Knowing that QC users will not check for autocorrelation of the collected samples, attackers will launch a cyber-physical attack that has two objectives: 1) forcing the inspected product KQCs to be negatively autocorrelated during Phase I; and 2) introducing mean and standard deviation shifts of 2 units each to the inspected KQC during Phase II.

To achieve the first objective, the attackers start by generating negatively autocorrelated samples and then make physical changes to each manufactured product so that the KQC actual value would be very close from the corresponding generated one. These negatively autocorrelated samples, with the desired KQC value denoted as d, are generated according to the 1st order autoregressive model:

$$d_t = c + \phi d_{t-1} + \varepsilon_t \; \forall \; t > 1,$$

where c is a constant, ϕ is the autocorrelation parameter, t is the time index, and ε is the error term; such that $c = 2\mu$, $\phi = -0.9$, $d_1 = x_1$, and ε_t is a random variable following the standard normal distribution N(0,1). Similar to the previous example, the attackers cannot ensure a 100% match between the generated KQCs value and the actual modified ones, due to the manufacturing process natural variation. However, since the difference in these values is going to be very small, it would be neglected and both values will be assumed to be equal.

The originally generated independent samples, without the attack being implemented, are plotted on an I-chart as shown in Figure 4-3a, with 3-sigma control limits[12]. When the attack is applied in Phase I such that the samples become negatively autocorrelated, the corresponding I-chart can be seen in Figure 4-3b. Since the users never check for sample independence, they would not realize that the attack was designed in such a way to cause the control limits to be artificially increased, while physically altering the KQCs of almost all of the inspected products.

[12] When estimated values are used for the process' mean, the traditional 3-sigma control limits for an I-chart are evaluated from the equations: $\bar{x} \pm 3\overline{MR}/d_2$; where $\bar{x}$ is the mean estimator, $\overline{MR}$ is the moving range, and d_2 is a constant that equals 1.128 when the number of observations per sample (n) is 2 (Montgomery, 2009).

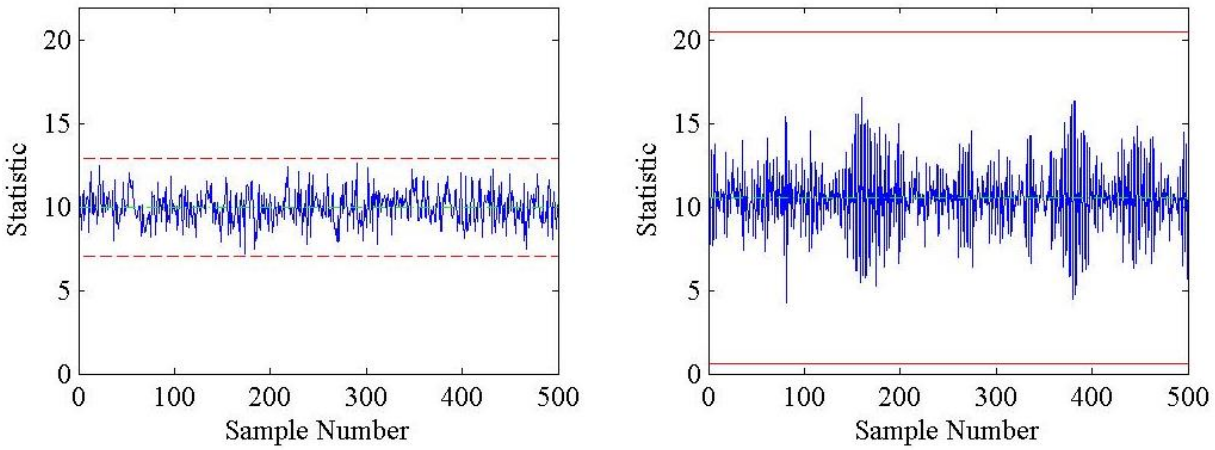

(a)

(b)

Figure 4-3: Individual (I) control charts for product's KQC during Phase I for a) independent samples without a cyber-physical attack and b) negatively correlated samples as a result of a cyber-physical attack; as an example of an attack taking advantage of the second misuse category. Red dashed lines represent the original control chart limits without the attack, red solid lines represent the control limits with the attack, whereas green dotted dashed lines represent the charts center lines.

With this attack affecting the resulting control limits established in Phase I, the quality of the manufactured products onwards would be questionable, without the manufacturers even knowing. So to accomplish the second objective of this attack, the attackers will just manipulate the manufacturing process to cause a shift of 2 units in both the mean and standard deviation of the samples during Phase II, without crossing the new control limits, as can be seen in Figure 4-4. With the new (widened) control limits, the effects of the attack were not detected; as all the samples are in-control according to these limits (solid red lines in Figure 4-4). However, if the original (correct) control limits (dashed red lines in Figure 4-4) were used, the changes in the manufactured products could have been detected faster. For this particular case, 95 out of 250 samples (38%) were actually out-of-control, designated by cyan circles in Figure 4-4, when using the new control limits.

It is important to note that although the charts limits were forced to change to avoid detection, this attack did not need to gain access nor make any changes in the QC system itself and only needed to know information about the QC tool used. Such a Passive attack only made direct physical changes within the manufacturing system, causing the QC tools to give misleading results about the process behavior. In addition, the same attack example could have been illustrated using an associated Moving Range (MR) chart just as well, with the same impacts; however, for the sake of brevity, such an attack is not demonstrated here.

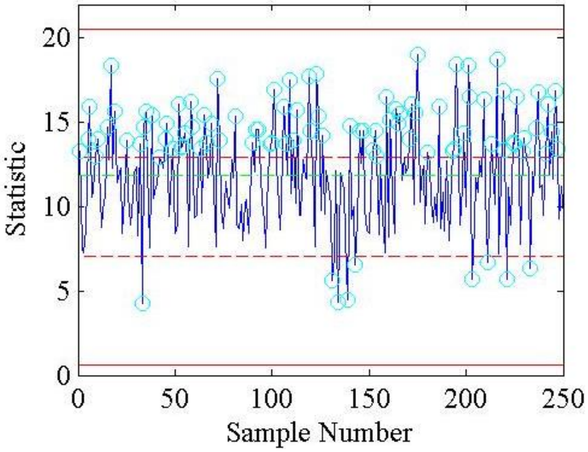

Figure 4-4: The I-chart during Phase II resulting from the cyber-physical attack for the autocorrelated samples, where the attacked samples outside the original limits are designated with cyan circles. Red dashed lines represent the original control chart limits when there was no attack, red solid lines represent the control limits with the attack, whereas green dotted dashed lines represent the new chart center lines.

4.3.2.3 Misuse Category 3 Example – Using Specification Limits instead of Control Limits

The final example presents a case of a general misconception of using the product's specification limits on control charts instead of the chart's actual control limits for Phase II monitoring. As pointed out in sub-section 4.3.1.3, this practice is actually quite common, and it is important to remember that specification limits are different from control limits and should not be used as a replacement. The decision of using the specification limits is a bad one, because using incorrect control chart limits makes the analysis done to the collected data not capable of detecting the effects of potential cyber-physical attacks. In addition, being within the specification limits does not imply that the process is stable.

For the sake of illustrating the potential consequences of this decision, consider a set of 50 randomly generated samples, containing five observations per sample with the observations being normally distributed according to N(5,1). This set of samples represents one of the product's inspected KQCs in Phase II monitoring, which has an Upper Specification Limit (USL) of 8 units, a Lower Specification Limit (LSL) of 2 units, and a process capability ratio C_p of 1. Unlike the example discussed in sub-section 4.3.2.1, the QC users do not estimate the traditional 3-sigma control limits within a Phase I for this set of collected samples nor do they ensure that stability is obtained first. The users just use the pre-specified specification limits for monitoring the KQC of interest.

The specification limits are often being used instead of the control limits due to the misconception that they are both equivalent. Also, with the observations' standard deviation being equal to one, users might incorrectly believe that the 3-sigma control limits are equal to the USL and LSL values, by adding/subtracting 3 standard deviations of 1 to/from the observations mean value of 5. However, since this is an X-bar chart, the control limits values need to take the sample size (n) value into account such that the limits equal either $\mu \pm 3\sigma/\sqrt{n}$ when the process parameters are known in advance or $\bar{\bar{x}} \pm A_2\bar{R}$ when the parameters are estimated, where A_2 is a constant that depends on the value of n. In either case, the control limits and specification limits are not equal. Figure 4-5 shows the X-bar chart plotted with both the specification and the traditional 3-sigma control limits (from estimated parameters) for the 50 generated samples; where the solid red lines represent the incorrect specification limits used, whereas the dashed red lines represent the correct control limits that should have been used. It is clear from the figure that both limits values are different.

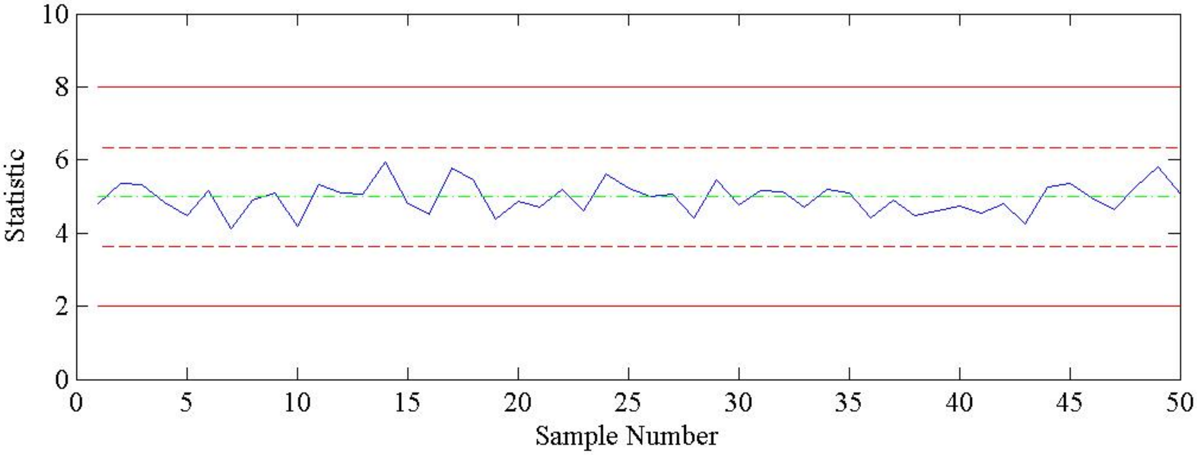

Figure 4-5: The resulting X-bar control chart with both limits for misuse example 3. Red solid lines represent the product's specification limits, red dashed lines represent the charts control limits, whereas green dotted dashed lines represent the chart's center line.

Consider an attacker, taking advantage of the usage of the specification limits and launches a disguised cyber-physical Passive attack. Since the actual control limits are not used, the attack will be designed such that the products KQCs mean values will be increased, while ensuring that such changes will not cause any of the X-bar statistic values to go beyond the specifications limits. After the first 50 products are manufactured, the attacker will modify the manufacturing process causing a mean shift of 1.5 units for the next 100 samples. The following figures demonstrate the detection capability using the different limits for these 100 new samples. The chart with the specification limit, on the left-hand side (Figure 4-6a), did not yield any out-of-control signals, despite the existing shift in the process mean. However, if the control (correct) limits were to be used, as in the plot in Figure 4-6b, the majority of the shifted samples would have been detected in this case, where the out-of-control samples are designated with a cyan circle.

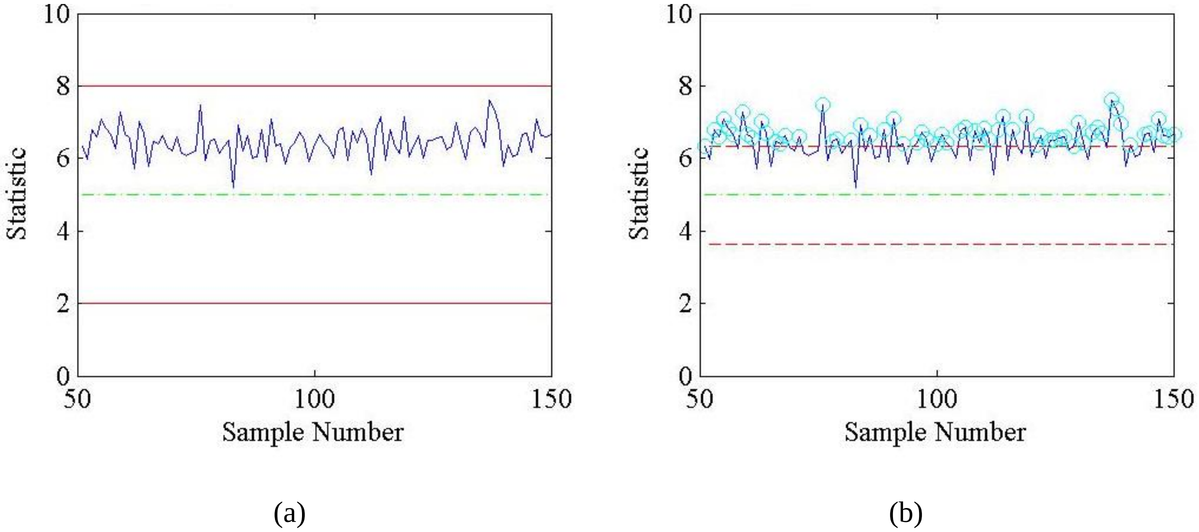

Figure 4-6: X-bar control charts with a cyber-physical attack using specification (a) and control (b) limits. Red solid lines represent the product's specification limits, red dashed lines represent the charts control limits, whereas green dotted dashed lines represent the charts center lines. Samples outside the control limits are designated with cyan circles.

In this case, the control chart with the specification limits will not be able to detect any (0%) of the shifted samples. However, using the control limits, 62% of the attacked samples would have been detected and the shift in the mean realized earlier for this particular case. This is rather significant because all the attacked samples were missed when the specification limits are incorrectly used, with the QC users not being even aware of the existence of a cyber-physical attack causing a mean shift. A summary of the bad decisions in each of the previously mentioned examples throughout this sub-section is given in Table 4-3.

Table 4-3: Summary of bad decisions within the different misuse category examples.

Misuse Category	The Bad Decision Taken	Why was the Decision Bad?
Category 1 Example	An X-bar chart alone is enough to detect KQC changes.	Not accounting for within process variance.
Category 2 Example	There is no need to check for collected samples independence before using it with an I-chart.	Not being able to detect process mean or standard deviation shifts.
Category 3 Example	The specification limits could be used instead of the control limits for an X-bar chart.	A process within the specification limits can still be out-of-control.

4.4 Making Decisions Bad

Unlike the exploitation class discussed in the previous section, this second class is not due to initially bad decisions, but is a result of taking decisions without considering the possibility of Passive cyber-physical attacks. In this exploitation class, the collected data is either not enough nor the correct ones for manufacturers to detect the effects of a cyber-physical attack. The subsequent collected data analysis, however, would be appropriate for detecting the effects of cyber-physical attacks, given that the collected data was appropriate to begin with.

Since this class of vulnerabilities exploitation relies on users not considering the possibility of cyber-physical attacks, the purpose of this section is to discuss the different levels at which the exploitation might occur in manufacturing. In addition, practical examples are presented in this section to highlight some of the effects of this second exploitation class on the performance of QC tools when the cyber-physical production system is under attack.

4.4.1 Making Decisions Bad Levels

Since the main distinction between the 2 classes is that the decisions involved in the second class are concerned with the data collection process, this exploitation class will be further classified into *levels* representing the type of exploitable knowledge used within this process. These exploitation levels are summarized in Table 4-4 and are further discussed next.

4.4.1.1 Level 1 – For the Feature

The first level considered in this class is the one where the knowledge of *how the collected data is used to assess the quality of a feature* within the QC system is exploited. How the quality of a feature is assessed is typically assumed sufficient, which may not be true in the presence of a cyber-physical attack. Level 1 exploitations utilize the fact that the used methods of assessing features' quality do not account for the possibility of a potential cyber-physical attack.

4.4.1.2 Level 2 – For the Part

The next level uses the knowledge of *which of the features' collected* data is used to assess the *quality of a part* within the QC system. Fundamentally, these types of Level 2 exploitations are due to the tendency in QC tools to focus only on KQCs, groups of features, or process parameters. When one assumes that a subset of system's outputs is sufficient, one becomes vulnerable to attacks on the reminder of that set or even to attacks introducing additional features within a part.

4.4.1.3 Level 3 – Along the Manufacturing Process

The final level uses knowledge of *how often data is collected* within the QC system. For instance, if the data collection frequency is not enough to detect a potential attack, then the quality of the whole manufacturing process would be affected. The basic idea here is to exploit the static design approach taken in QC tools development within a manufacturing process, such as a consistent sampling strategy. Table 4-4 provides a summary of the main issue with the decisions taken within the data collection process for each of the exploitation levels in this class.

Table 4-4: Class 2 exploitation levels summary.

Exploitation Class 2	Main Issue with Decisions Taken during Data Collection
Level 1	Collected data is not sufficient to assess the quality of a *feature*.
Level 2	Features identified through data collection are not enough to assess the quality of a *part*.
Level 3	Collected data is not enough to assess the quality of a manufacturing *process*.

4.4.2 The Effects of Cyber-Physical Attacks on Class 2 Exploitation Levels

After discussing the different levels of Class 2 exploitations, some examples of potential attacks within these levels are illustrated in this sub-section. Knowing the effects of such cyber-physical attacks is important for QC users and industrial practitioners; since they lead, for example, to decreased tool performance and a false sense of quality.

4.4.2.1 Level 1 Exploitation Example – Not Collecting Enough Data for a Feature

The first example presented here is of a Passive attack that was able to use knowledge about the collected dimensional data of a certain product to hides its effects on this feature. Consider the cylindrical product to be machined as shown in Figure 4-7a, via a turning process. The product's diameter needs to be equal to 50 mm and, to this end, the diameter of each part is inspected at three separate locations: midway across the product's length and at both ends. In this attack example, the QC engineer assumed that the data collected from these 3 inspection locations is sufficient to assess the part's diameter and did not consider the possibility of a cyber-physical attack altering it.

Hence, the attackers, who had already accessed the system, use this information and alter the tool-path used to manufacture the part on a Computer Numerically Controlled (CNC) machine at the locations not being inspected throughout its length, as shown in Figure 4-7b. More specifically, between every two diameter inspection locations, the part has a 2-degrees increasing and then decreasing taper, to give it a double barrel-like appearance; where the changes made to the part are symmetrical. This Passive cyber-physical attack renders the part not manufactured as intended while avoiding any detection of the attack at the same time.

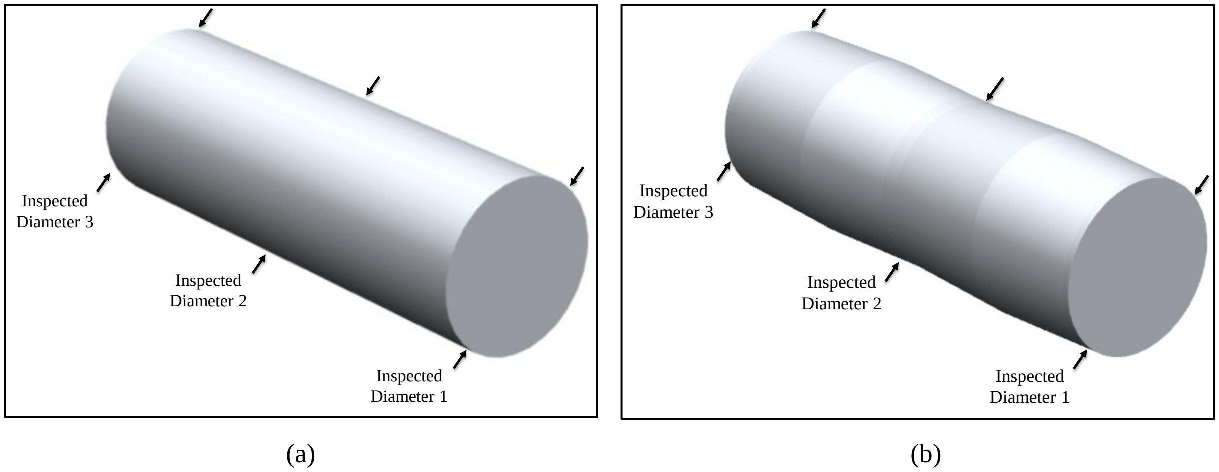

Figure 4-7: A manufactured product without an attack (a) and with a Passive cyber-physical attack (b), as an example of a Class 2 - Level 1 exploitation.

4.4.2.2 Level 2 Exploitation Example – Collecting Data for Only a Subset of Features

The next example presents a Passive attack taking advantage of the knowledge of which of the features collected data was used to assess the quality of the whole part. With such a knowledge, this attack simply adds a feature that should not exist in a manufacturing product, such as the manufactured bracket designed to have ten mounting holes shown in Figure 4-8a. In this case, a CMM is used to collect inspection data about the bracket's different dimensions. For the ten holes, the data collected pertains to their locations and sizes only. Using the knowledge about which of the features are assessed by the data collected from the CMM inspection, the attack simply places an 11th hole on the manufactured product (Figure 4-8b). Since an 11th hole should not exist on the product, it is never measured, resulting in the attack never being detected.

Although one could argue that such an attack is rather obvious, this may not necessarily be the case if the attacked products are mass-produced in an automated environment, the products are very complex, or the product's features are very small. The issue here with the manufactured bracket is deciding that the whole product quality can be assessed by just the 10 holes. Specifically, it was decided that inspecting all the holes was enough for part characterization, without considering that the product could be maliciously altered and assuming that this type of inspection was sufficient. Despite the fact that assuming the CMM can detect deviations in the locations and sizes of the different features is not a bad decision, it was not enough to prevent the attacker from using the available knowledge of the collected data in his/her favor to hide the corresponding attack.

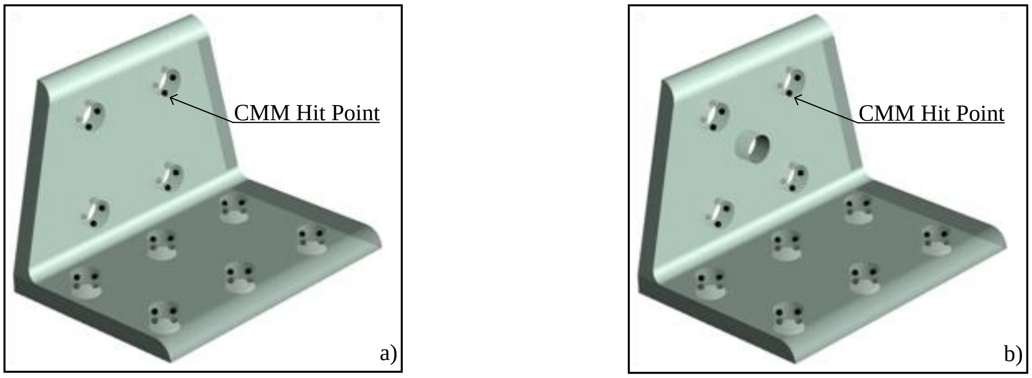

Figure 4-8: The manufactured bracket without an attack (a) and with a Passive cyber-physical attack adding an extra feature (b), as an example of a Class 2 - Level 2 exploitation.

4.4.2.3 Level 3 Exploitation Example – Data Collection Not Frequent Enough

The final example demonstrates a Passive attack that uses knowledge about the collected data within the QC system for the whole manufacturing process to potentially prevent attack detection. Consider the case of a manufacturing process being monitored by a control chart, where a sample is taken for inspection every certain number of products. The decision to collect inspection data in this manner is not a bad decision, since it is assumed that shifts in the process are sustained ones that can be eventually detected by this QC tool. A typical data inspection scheme can be seen in Figure 4-9, where the inspected products are the hollow ones.

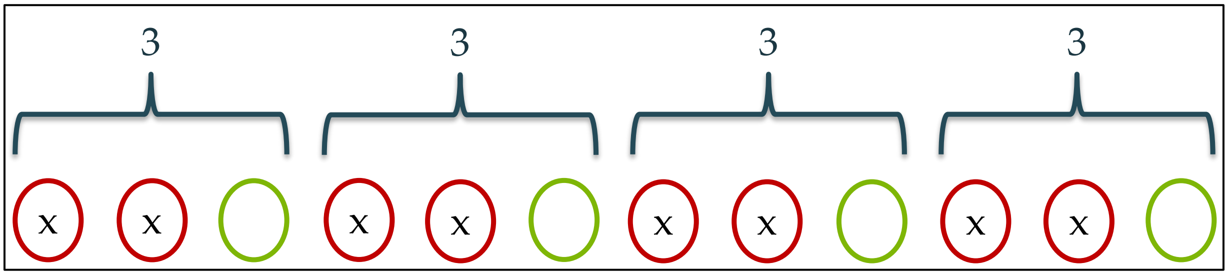

Figure 4-9: An example of a traditional data inspection decision of collecting data every fixed number (3) of products. The hollow products are the ones being inspected, whereas the solid products (also marked with an "x") are not inspected.

However, an attacker who has knowledge about such data inspection frequency can intentionally affect products between inspections only and do not tamper with the inspected ones, causing transient shifts instead. The products that are marked with an “x” in Figure 4-9 are the ones that are more vulnerable and could be targeted by the cyber-physical attacker in this case. With the attacker forcing such malicious transient process shifts, the collected data is no longer enough to assess the quality of the monitored manufacturing process. The collected data was only suitable when the considered shift was sustained, but not considering the possibility of cyber-physical attacks causing transient shifts has made this data collection decision a bad one.

4.5 Potential Mitigation Guidelines

After presenting, in the previous sections, the two different exploitation classes through which cyber-physical attacks could take advantage of the vulnerabilities in QC tools, potential mitigation guidelines are discussed here in this section. Mitigating the threats posed by both exploitation classes could be achieved through, but not limited to: 1) effectively applying QC best practices; 2) selecting alternate Key Quality Characteristics (KQCs) to monitor; and 3) developing new approaches that are meant to make QC systems more resilient to cyber-physical attacks.

4.5.1 Applying QC Best Practices

Existing systems can benefit from QC best practices to limit the vulnerability due to QC tools misuse; i.e., those exploiting bad decisions within the first exploitation class. Such best practices include, but are not limited to, ensuring that the most appropriate quality tools are selected in a given situation, the selected tools are correctly employed, and that the necessary assumptions for the application of the selected quality tools are continuously verified.

One example of QC best practice is applying multiple control charting strategies in conjunction to decrease the vulnerabilities of an attack. Use of an EWMA chart is effective for detecting small shifts occurring over time, whereas a Shewhart chart is able to detect significant changes in quality parameters. Together the combined Shewhart EWMA charts are able to monitor multiple shift magnitudes. Even if the in-control system rarely has large shifts, the dual charting strategy would be more capable of detecting a cyber-physical attack inducing a large alteration. These types of best practices can reduce the vulnerabilities posed by certain cyber-physical attacks.

4.5.2 Selecting Alternate KQCs

In order to limit the vulnerability within the second class that makes decisions bad, an alternate view of what Key Quality Characteristics (KQCs) are is now necessary. The selected KQCs should be thought of more as Key Security Characteristics (KSCs) that need to be monitored for the sake of detecting the effects of cyber-physical attacks within the manufacturing environment. Instead of monitoring traditional KQCs that are based on customer requirements, focusing on such KSCs should ensure that the product's quality has not been compromised by an attack.

An example of this idea can be illustrated by a simple part printed through additive manufacturing. Typical characteristics to monitor might include the physical dimensions and surface roughness. However, monitoring an additional characteristic such as the build time or weight can provide diagnostic information with minimal extra effort. Significant variation in the build time could point to any deviations from the nominal print pattern. Selecting the key characteristics to monitor should include the critical aspects to ensure a part's performance, as well as characteristics that may assist in the detection of a cyber-physical attack. This idea is similar to that of the work done to detect cyber-physical attacks in manufacturing using side channel analysis (Vincent *et al.*, 2015; Chhetri *et al.*, 2016; Sturm *et al.*, 2016; Albakri *et al.*, 2017; Moore *et al.*, 2017a; Wu *et al.*, 2017; Belikovetsky *et al.*, 2018).

4.5.3 Developing More Resilient QC Tools

Since one of the main reasons for cyber-physical attacks being able to exploit vulnerabilities in QC tools is the fact that the tools were designed without taking such attacks into account, the development of the next generation QC tools must include cyber-physical security as a key consideration to remove the vulnerabilities resulting from either exploitation classes. The IT domain contains several tactics which could be adapted to serve this purpose of developing more resilient QC tools to cyber-physical attacks.

One tactic that could be applied is introducing randomness to the design and implantation of QC tools, such as using a random sampling strategy. Another option is considering a more holistic approach to quality control by adding a unique "tag", such that quality is not defined by individual part features, but by the part as a whole. Such a tag, usually obtained from a hashing function (Sturm *et al.*, 2017), is then collected by a receiver that would match it with a reference value, ensuring that any changes to the processes by an attacker, no matter how minute, would indirectly change the process data and could be detected.

4.6 Summary

With the recent breakthroughs in computers and networking technology, along with the increasing threat of cyber-physical attacks in manufacturing environments, the work discussed in this chapter is a key step towards developing new or redesigning traditional QC tools in a way to increase the probability of being able to detect the effects of these types of attacks rather promptly. This work also demonstrated that there is an urgent need to be watchful to the effects of cyber-physical attacks exploiting existing vulnerabilities of QC tools within manufacturing and to start including cyber-physical security considerations in QC systems design.

Throughout this chapter, the relation between cyber-physical attacks and vulnerabilities existing within QC systems in manufacturing has been presented. First, the different techniques, referred to as exploitation classes, through which cyber-physical attacks can take advantage of the vulnerabilities existing within QC systems were discussed. Two distinct classes exist corresponding to the wrong decisions taken when implementing QC tools during either data analysis and interpretation or data collection. Then, for each of these classes, the different situations in which they could occur were presented, along with highlighting the additional negative effect cyber-physical attacks would have on each one of them. Finally, ideas for potentially reducing the effects of cyber-physical attacks and creating more secure QC systems were also provided.

5 Introducing Randomness into Control Chart Sampling Plans for Better Security against Cyber-Physical Attacks in Production Systems – Univariate Case

After discussing the role of Quality Control (QC) systems during cyber-physical attacks and identifying some of the issues existing within them in the previous two chapters, the focus of the next research area is on improving current QC tools to address the impacts of some of these issues for better cyber-physical security in manufacturing. More specifically, randomness is introduced into control chart sampling plans to make QC tools more resilient against these attacks. The work done in this research area includes both univariate and multivariate control charts and will be covered across this and the following chapter, respectively.

In this chapter, the details of introducing randomness into univariate control chart sampling plans to prevent attackers from knowing which products are being sampled for inspection is presented. Different random sampling plans are discussed and their performances under varying attack scenarios are evaluated, using a number of performance metrics. It should be just noted that control charts were used as representatives of common QC tools because they are one of the most frequently utilized tools in industry. In addition, the initial idea of the work in Chapter 5 was introduced at INFORMS 2015 Annual Meeting (Elhabashy *et al.*, 2015b) and more specific details about it were presented at INFORMS 2017 Annual Meeting (Wells *et al.*, 2017). A corresponding journal article will be submitted to the *Journal of Quality Technology* as soon as it is finalized (Elhabashy *et al.*, 2019b).

5.1 Background and Motivation

Recent technological advancements in cyber-physical systems have caused manufacturing systems to become more susceptible to cyber-physical attacks. Such attacks have been demonstrated in small-scale case studies in controlled settings, such as those by Meserve (2007); Wells *et al.* (2014); Turner *et al.* (2015); Zeltmann *et al.* (2016); Moore *et al.* (2017b); Sturm *et al.* (2017); and Belikovetsky *et al.* (2018), and have unfortunately occurred in actual production systems; such as nuclear facilities (Albright *et al.*, 2010; Cherry, 2011), steel mills (Lee *et al.*, 2014), and car manufacturers (Liptak, 2017). This work focuses specifically on cyber-physical attacks that attack a cyber element within a manufacturing system to physically alter a manufactured part in a specific way, referred to as Product-Oriented C2P (Cyber-to-Physical) attacks (Shafae *et al.*, 2018).

In manufacturing, Quality Control (QC) systems are employed to ensure the production of high-quality parts and to maintain stable processes. Unfortunately, QC systems are not immune to cyber-physical attacks, as they could be compromised during cyber-physical attacks (Wells *et al.*, 2014; Vincent *et al.*, 2015). As part of their cyber-physical attack taxonomy, Elhabashy *et al.* (2018a) identified that QC systems could be compromised either through: 1) Direct Quality Attacks: Altering QC systems without affecting a manufacturing system; 2) Joint Passive Attacks: Using knowledge regarding the implementation of a QC systems to intelligently design and implement a Product-Oriented C2P attack; or 3) Joint Active Attacks: Altering QC systems in conjunction with a Product-Oriented C2P attack.

With respect to Joint Passive Attacks, Elhabashy *et al.* (2019a) identified two categories to describe how knowledge regarding the implementation of a QC system could be used to intelligently design and implement a Product-Oriented C2P attack. The first category covers instances where adversaries have determined that QC tools are being incorrectly applied, and use this QC tool misuse to design their Product-Oriented C2P attack. The second category focuses on instances where QC tools are correctly used but the adverasary has determined that the data collected and used for QC are inadequate to either assess the quality of a feature, a part, or a process under the presence of an intelligently designed Product-Oriented C2P attack.

This work focuses on the latter of the three inadequacies in the second category; specifically, the situation where an adversary knows that 100% inspection of the manufactured parts is not performed and that sampling occurs at a fixed part interval. In this situation, the Product-Oriented C2P attack can be intelligently designed to affect parts that are not sampled for inspection. An example of such a situation is shown in Figure 5-1, where inspection occurs every 5 parts (colored in green). In this situation, a Product-Oriented C2P attack could be designed to affect some, if not all, of the parts colored in red.

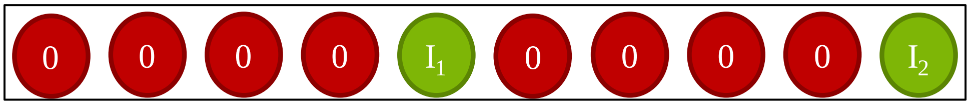

Figure 5-1: An example of a traditional inspection plan, where sampling occurs every 5 products. The products that are sampled for inspection are colored in green and designated with an "I", whereas the red products are the ones that are not inspected and could be affected by an intelligently designed Product-Oriented C2P attack.

Intelligently designed attacks of this nature have a zero probability of being detected by quality control tools and can be considered a unique type of transient shifts. Transient process shifts, as opposed to sustained shifts, only last for a finite duration (*l*) of time (Reynolds and Stoumbos, 2004a). The process parameters affected by such transient shifts would eventually return to their in-control value without any external intervention (Reynolds and Stoumbos, 2004a; 2005). Although previous related work in the field of QC has focused on shifts due to transient assignable causes, such as those by Reynolds and Stoumbos (2004b; 2004a; 2005), it has not addressed unique transient shifts caused by intelligently designed cyber-physical attacks.

The goal in this work is to develop and analyze the performance of more robust QC approaches that are capable of detecting these types of intelligently designed attacks. This goal will be achieved through developing new sampling plans that are based upon randomness, which eliminates the ability of an adversary to know which parts will be inspected. This randomness, in essence, prevents this type of an intelligently designed attack to be performed. It should be just noted that the benefits of the proposed random sampling plans will be demonstrated using Individual control charts (I-charts); yet, this work could be generalized to other types of univariate control charts.

This chapter is organized as follows; details of the proposed methodology, which includes two proposed random sampling plans and two intelligently designed attack models, are presented in the following section. In Section 5.3, the different performance metrics used for assessing the proposed methodology are discussed, along with how they are evaluated. The results of the methodology assessment are then shown in Section 5.4. Finally, the chapter is concluded with a summary in the last section.

5.2 Methodology

The proposed methodology aims to introduce randomness into sampling plans to increase cyber-physical security for situations where 100% inspection is not viable. In this manner, attackers would not be able to know which products are being inspected nor to intelligently design their cyber-physical attacks accordingly. Such an approach of deliberately introducing randomness has been frequently applied in the field of Information Technology (IT) to ensure process stability and security. For instance, it has been proposed to randomly simulate network node downtimes to improve network stability (Tseitlin, 2013). To extend this idea of intentionally introducing randomness to sampling plans, it is important to first understand the sampling plans that are currently employed in practice to establish a performance baseline, before discussing the newly-proposed ones.

This work considers processes that manufacture products at a constant rate, where MP_i is the i^{th} Manufactured Product. Inspection is to occur once every (N) products and the j^{th} Inspected Product is designated as IP_j. Another way of describing the inspection process is viewing the manufactured products as successive groups each of size N, such that $G_1 = \{MP_1, \dots, MP_N\}$; $G_2 = \{MP_{N+1}, \dots, MP_{2N}\}$; and $G_j = \left\{MP_{(j-1)N+1}, \dots, MP_{jN}\right\}$; where G_j is the j^{th} group. The products sampled for inspection in this case are $IP_j = G_j(N)\ ,\ \forall\, j \geq 1$. This developed notation helps in understanding the nature of the inspection process.

This sampling plan is just the base plan, where samples are chosen for inspection every N products, such as inspecting every 5 products illustrated in Figure 5-1. This is the plan where the sampling information can be used to intelligently design Passive Product-Oriented C2P attacks. Such a sampling plan will be referred to as **Plan 0 – No Randomness** and the indices of the sampled products are just the multiples of N, such that $IP_j = MP_{jN},\ \forall\, j \geq 1$.

There also exists two possible practical realizations of sampling Plan 0 that needs to be discussed. These other sampling plans, which might be also currently used in practice, are described as follows:

- **Plan 0a – Unintentional Randomness 1:** In this plan, the products are selected for inspection every certain period of time instead (e.g. every hour), which implies that inspection should also occur every N products when the production rate is constant. However, sometimes inspection would be done at a slightly earlier (or later) time, preventing the inspection from still occurring every N products in this plan. In this work, this unintentional randomness in sampling timing is assumed to be equivalent to a random variable D, such that $IP_j = MP_{jN+D_j},\ \forall\, j \geq 1$. This plan, however, is not considered in this work, since it is very similar to Plan 0.
- **Plan 0b – Unintentional Randomness 2:** Similarly, instead of sampling once every N products, sampling is done unintentionally a few products before or after. Again, this sampling randomness is assumed to be equal to the random variable D, such that

$$IP_j = MP_{jN+\,D'_j}, \forall\, j \geq 1 \; where\ D'_j = \sum_{s\,=\,1}^{j} D_s \tag{5-1}$$

Plan 0b differs from the previous one in two ways: 1) the index of the chosen product for inspection in Plan 0b greatly depends on that of the previously chosen sample and 2) zero, one, or two products could be inspected within one group of *N* products in Plan 0b.

5.2.1 Sampling Plans

In order to detect the effects of the intelligently designed Passive Product-Oriented C2P attacks causing transient shifts, two competing plans were developed in this work. Despite being developed for detecting transient shifts, the two sampling plans were chosen such that their ability to detect sustained shifts is not diminished. It is important for the proposed sampling plans to excel in detecting sustained shifts as well, or else applying such plans would be of limited usefulness. These two sampling plans are designed such that their in-control performance is comparable to the plans implemented in practice; hence, there would be no loss in performance when implementing them. The main difference between these two sampling plans is in how randomness is intentionally induced within each group of *N* products, such that a random product is sampled from every group.

The first of these random sampling plans will be referred to as **Plan 1 – Induced Randomness 1**, where just 1 product is chosen randomly within each group of *N* products, as shown in Figure 5-2. The value of the indices of the inspected products is set to be randomized from one group to the other; such that $IP_j = MP_{(j-1)N + RI_j}$, $\forall\, j \geq 1$, where *RI* is the Random Inspection variable representing the index of the randomly inspected product within a group of *N* products.

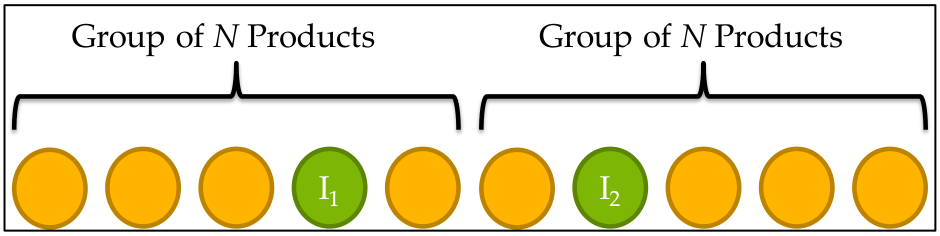

Figure 5-2: An example of sampling Plan 1, where the selected products for inspection, colored in green and designated with an "I", are randomly chosen within each group of N = 5 products.

In the second sampling plan of intentionally inducing randomness, each product within a group of *N* products could be selected for inspection according to a probability that equals to (1/N), as long as no shifts have been detected. This sampling plan will be referred to as **Plan 2 – Induced Randomness 2**, and it is important the note the distinction between it and the other random plan. The key difference between these two intentionally random plans, is that only one product can be chosen for inspection in Plan 1 within a group of N products. But with Plan 2, more than one product could be chosen for inspection within one group of *N* products; or no products at all.

5.2.2 Attack Scenarios

The performance of the sampling plans discussed in the previous sub-section needs then to be assessed under different intelligently-designed Passive Product-Oriented C2P attacks, referred to as simply attacks. For this purpose, two specific attack scenarios are identified. These attack scenarios are considered intelligently-designed because with the inspection known to occur every N products in the baseline Plan 0, the attackers make sure that the N^{th} product in each group is never attacked.

Only a single manufactured product within each group of N products is targeted in both scenarios, causing a transient shift in the manufacturing process mean of *l* equal to 1 product. The Attacked Product within each group is designated with an AP_q to represent the q^{th} attacked product, such that $AP_q \neq MP_{qN}$, $\forall\ q \geq 1$. The details of both attack scenarios are as follows:

1) **Single Fixed Attack Scenario:** In this first scenario, $AP_q = MP_{(q-1)N\ +FA}$, $\forall\ q \geq 1$, where FA is a Fixed Attack constant that doesn't equal to *N*. This constant represents the index of the attacked product within each group of *N* products.
2) **Single Random Attack Scenario:** Alternatively, the index of the attacked product within each group of N products is randomized in the second attack scenario; such that $AP_q = MP_{(q-1)N\ +RA_q}$, $\forall\ q \geq 1$, where RA is the Random Attack variable that is randomized every group of *N* products.

5.3 Assessment

In order to determine how well the random sampling plans fair in detecting the transient shifts during the different attack scenarios, appropriate performance metrics are needed. In this work, the monitored statistic by the I-charts would represent any of the product's Key Quality Characteristic (KQC) that the manufacturers are monitoring, with the in-control performance of the used control chart set to a designed value. When the process mean shifts from its original value due to an attack, the process then becomes out-of-control. This out-of-control performance, due to different sizes of process mean shift, is the one that would be evaluated for all the sampling plans and attack scenarios combinations.

The metrics considered in this work address questions regarding the process out-of-control performance in the existence of this unique type of transient shifts including, but not limited to: 1) how fast, in terms of number of inspections needed and manufactured products, can a shift be detected; 2) what is the probability of detecting a shift due to an actual attack; and 3) how many products are successfully attacked, before the attack effects can be detected. While many of the used metrics are common to most control charts, their values are estimated a little differently to account for the unique type of transient shifts caused by the intelligently designed Passive Product-Oriented C2P attacks. In addition, some new metrics will also be used to evaluate the performance of the control charts during these types of transient shifts.

5.3.1 The Used Performance Metrics

A range of performance metrics was used to be able to adequately assess the performance of the different random sampling plans when this unique type of transient is present. Additionally, the more performance metrics that can be used, the more accurate the comparison between the different sampling plans can be made. Each of the performance metrics provides a different piece of information regarding the sampling plans and is briefly described as follows:

- **Average Run Length (ARL):** The ARL is defined in this work as the average number of inspected products until a signal is obtained. A signal could occur due to either an actual shift in the process mean or due to the process natural variation (random error). The ARL could be also viewed as the average number of groups of *N* products to signal.
- **Average Time to Signal (ATS):** In this work, the ATS is treated as the number of manufactured products to signal since sampling has started, regardless of the cause of the signal.
- The **True Positive Probability (TPP)** is defined as the probability that the resulting signal is due to an attack and not natural variation. Conversely, the **False Positive Probability (FPP)** is defined as the probability that the resulting signal is due to natural variation only and not due to an attack.
- The number of successfully attacked products until a signal is a random variable of interest that is quantified by two new metrics:
  1) **The Number of Successful Attacks until Signal (NSAS):** The NSAS just represents the total number of products that were successfully attacked before an

out-of-control signal has occurred, regardless of this signal's source. The lower the value of the NSAS, the better.

2) **The Percentage of Successful Attacks until Signal (PSAS):** The PSAS just represents this number of successful attacks as a percentage of the total of number of attacked products.

- **Percentage Attacked products (PA):** This metric represents the percentage of products successfully attacked compared to the total number of products until signaling, regardless of the source of this signal.
- The **Detection Probability (DP)** is already a well-established performance metric when dealing with transient shifts. According to the related literature, the detection probability is typically defined as "the probability of the control chart to signal while the shift is present (or shortly after[13])" (Reynolds and Stoumbos, 2004b; 2004a; 2005; Reynolds and Lou, 2010). In this work, the detection probability is going to be defined, a little differently, as the percentage of attacked parts that have caused a signal.
- The final two new metrics are presented together, since they provide information regarding out-of-control statistics. The first is the **Percentage Out-of-Control (POC)** and it represents the percentage of products that would have been out-of-control if 100% inspection is done, whether due to a random error or a malicious attack. The second one, which is the **Percentage Attacked Out-of-Control (PAOC)**, is concerned with the products that would have been out-of-control due to an attack only.

It should be just noted that it may not be possible to obtain these last two metrics in an industrial setting, as it essentially requires inspecting all the products, as opposed to just a specified number of samples every certain period of time. More details about all of these metrics can be found in Appendix B.

5.3.2 Performance Metrics Evaluation

One way to obtain the performance metrics values is to derive theoretical formulas for each of them. Unfortunately, such a task is not simple for all the sampling plans, such as Plan 0b, nor necessarily possible for all the performance metrics, such as the POC and PAOC. Therefore, to avoid the complexity involved in such derivations, it was opted to use computer simulations for

[13] In one study, this was equal to 4 time periods after the shift existed (Reynolds and Lou, 2010).

all the different attack scenario and sampling plan combinations. It should be just noted that the theoretical derivations for some of the metrics across the different sampling plans for the 1st attack scenario are still included in Appendix B, whereas a comparison of a few metrics theoretical values to those obtained from computer simulations are presented in Appendix C. The latter appendix also demonstrates that the metrics values obtained from the simulations are repeatable and quite accurate, as the difference between them and the theoretical values is negligible.

The simulation procedure used, alternatively, is pretty straightforward and can be summarized in the following steps:

1) Setting the values for the different parameters used throughout, such as:
   - N = 10 products, for all attack scenario and sampling plan combinations;
   - $\mathrm{D_j} \sim \mathrm{U}\{-2,\ 2\}\ \forall\ \mathrm{j} \geq 1$, for sampling Plan 0b;
   - $\mathrm{RI_j} \sim \mathrm{U}\{1, \mathrm{N}\}\ \forall\ \mathrm{j} \geq 1$, for sampling Plan 1;
   - $\mathrm{FA} = 0.5\mathrm{N} = 5$, for the single fixed attack scenario; and
   - $\mathrm{RA_k} \sim \mathrm{U}\{1,\ \mathrm{N} - 1\}\ \forall\ \mathrm{k} \geq 1$, for the single random attack scenario.
2) Generating random data of size 1 following a normal distribution of mean and standard deviation values equal to 0 and 1, respectively, to represent the measured value of the current product's KQC;
3) If this product would be attacked, introducing a mean shift to the generated data when applying each of the different sampling plans in a manner that mimics both the considered attack scenarios;
4) If this product is going to be inspected, plotting its corresponding statistic value on the I-chart and checking if the statistic value is within the control limits;
5) Repeating steps 2-4 until the I-chart statistic value is out of the control limits and a signal is obtained, while keeping track of the values of the variables that are used to evaluate the metrics until a signal is obtained. For example, need to keep track of the number of inspected products to get the Run Length (RL) values;
6) Steps 2-5 are repeated for 1,000,000 simulation iterations, to obtain the different metric values across all of the iterations. For example, all the different RL values will be averaged to obtain the ARL value for this simulation.

5.4 Results and Discussion

In this section, the resulting values for the different metrics from the simulations are presented for all the attack scenario and sampling plan combinations. The results presented in this section help in assessing which of the proposed sampling plans is better against different potential cyber-physical attacks. Therefore, all the metric values are first shown for each attack scenario considered. Then, the most useful metrics from the proposed sampling plans are compared across the different attack scenarios.

5.4.1 For Each Attack Scenario

The metric values across the simulation iterations for the Single Fixed Attack Scenario are shown in Table 5-1, whereas those for the Single Random Attack Scenario are shown in Table 5-2. It should be just noted that for the metrics involving attacks, there will not be a corresponding value for Plan 0 and it is denoted by just a dash. For example, there are no corresponding values during Plan 0 for TPP, FPP, NSAS, PSAS, PA, DP, and PAOC metrics, since this plan cannot detect the shifts due to the intelligently designed attacks. Several observations can be made from these tables to assess the proposed random sampling plans and understand the attackers' behavior, as discussed next.

5.4.1.1 Scenario 1 – Single Fixed Attack

Considering first the ARL metric values, it was noticed that when the shift size increases, all the plans involving some type of randomness (0b, 1, and 2) have a decreasing value. Such a decrease in ARL values means that these sampling plans detect the presence of large shifts faster, unlike Plan 0. There is a significant difference between the ARL values for Plans 0 and 0b, with the latter being much smaller. Furthermore, the ARL metric values are even smaller for the 2 new sampling plans (Plans 1 and 2), especially with larger shift sizes. It can also be noticed that the ARL metric values across the different shift sizes are very close for Plans 1 and 2, which means that according to this metric they both have very similar performance, with both being more effective than any of the sampling plans currently used in practice. With the large shift sizes being more important to detect, both the proposed sampling plans are able to detect them promptly enough.

A pattern similar to the one for the ARL metric values can also be observed with both the ATS and NSAS metrics values. On one hand, the NSAS metric values are very close from those of the ARL value for all the plans, which is not surprising. On the other hand, the ATS metric values are closer to the ARL values multiplied by *N*, which is also expected (as discussed in Appendix B).

Table 5-1: Performance metrics values for the 1st attack scenario across the different sampling plans obtained from the simulation procedure with 1,000,000 iterations.

Shift Size	0	0.1	0.5	0.75	1	2.5	5	10
Average Run Length (ARL)								
Plan 0	370.2455	370.4883	370.5173	370.6745	371.0182	370.2872	370.3069	369.9276
Plan 0b	371.2555	368.6549	325.8660	274.2355	214.3650	36.3565	16.8522	16.6483
Plan 1	370.6418	368.4035	325.5718	273.3443	212.5604	30.0636	9.9496	9.7750
Plan 2	370.5610	368.8045	325.1569	272.6598	212.4421	29.6553	9.5048	9.2870
Average Time to Signal (ATS)								
Plan 0	3702.455	3704.883	3705.173	3706.745	3710.182	3702.872	3703.069	3699.276
Plan 0b	3712.579	3686.574	3258.710	2742.375	2143.663	363.569	168.516	166.483
Plan 1	3701.920	3679.536	3251.156	2728.814	2120.892	295.676	94.509	92.763
Plan 2	3705.864	3687.739	3251.472	2726.572	2124.461	296.580	95.105	92.879
True Positive Probabilities (TPP, %)								
Plan 0	-	0	0	0	0	0	0	0
Plan 0b	-	10.3461	20.6749	33.2351	47.6796	90.9939	95.7387	95.7737
Plan 1	-	10.3801	20.9006	33.6022	48.3557	92.757	97.5877	97.6315
Plan 2	-	10.4681	20.8915	33.6692	48.4129	92.8550	97.7286	97.7493
False Positive Probabilities (FPP, %)								
Plan 0	-	100	100	100	100	100	100	100
Plan 0b	-	89.6539	79.3251	66.7649	52.3204	9.0061	4.2613	4.2263
Plan 1	-	89.6199	79.0994	66.3978	51.6443	7.2430	2.4123	2.3685
Plan 2	-	89.5319	79.1085	66.3308	51.5871	7.1450	2.2714	2.25070
Number of Successful Attacks to Signal (NSAS)								
Plan 0	-	369.4883	369.5173	369.6745	370.0182	369.2872	369.3069	368.9276
Plan 0b	-	368.1067	325.3270	273.7002	213.836	35.8510	16.3482	16.1450
Plan 1	-	367.4051	324.5717	272.3442	211.5604	29.0636	8.9496	8.7750
Plan 2	-	368.2244	324.6014	272.1198	211.9164	29.1536	9.0086	8.7859
Percentage of Successful Attacks to Signal (PSAS, %)								
Plan 0	-	100	100	100	100	100	100	100
Plan 0b	-	99.8896	99.7569	99.5543	99.2310	94.8538	90.8515	90.7623
Plan 1	-	99.8335	99.6287	99.3052	98.7592	89.1471	74.9437	74.6698
Plan 2	-	99.8337	99.6287	99.3067	98.7871	89.1784	75.0175	74.6177

Shift Size	0	0.1	0.5	0.75	1	2.5	5	10
Percentage Attacked (PA, %)								
Plan 0	-	9.8398	9.8399	9.8406	9.8389	9.8398	9.8395	9.8399
Plan 0b	-	9.9110	9.9059	9.8979	9.8828	9.6724	9.4599	9.4545
Plan 1	-	9.8503	9.8364	9.8183	9.7880	9.1473	8.0252	8.0057
Plan 2	-	9.8503	9.8367	9.8213	9.7919	9.1523	8.03811	8.0003
Detection Probability (DP, %)								
Plan 0	-	0	0	0	0	0	0	0
Plan 0b	-	1.0669	1.1760	1.3410	1.6129	5.6556	9.5557	9.6453
Plan 1	-	1.6043	1.7766	2.0676	2.5660	11.7004	25.6756	25.9447
Plan 2	-	1.5886	1.7774	2.0591	2.5053	11.6543	25.5632	25.9667
Percentage Out-of-Control (POC, %)								
Plan 0	25.4553	25.4120	23.2879	20.5992	17.2754	3.7076	1.3442	1.3109
Plan 0b	25.4208	25.4721	25.3633	25.1605	24.8452	17.4404	9.7916	9.6497
Plan 1	25.5657	25.6564	25.5935	25.6175	25.5918	25.5934	25.6823	25.5522
Plan 2	25.5761	25.5425	25.6105	25.6091	25.5954	25.5560	25.5768	25.5904
Percentage Attacked Out-of-Control (PAOC, %)								
Plan 0	-	0	0	0	0	0	0	0
Plan 0b	-	7.2731	11.5274	15.0230	17.8219	16.4265	9.3732	9.2377
Plan 1	-	7.3333	11.7504	15.5006	18.6717	24.8624	25.4501	25.3302
Plan 2	-	7.3837	11.7362	15.5025	18.6246	24.8419	25.3731	25.3823

From both the TPP and FPP metrics, the proposed random sampling plans are doing a better job in signaling due to an actual attack; they are even more accurate than Plan 0b, which included some type of randomness. However, it can be observed from the DP metric that the best detection probability for any of the proposed random sampling plans does not exceed 26%. Although such a detection probability is rather low, it is better than having a 0% detection probability as with Plan 0 and not being able to detect any of the shifts. The same observation could also be noted from considering the PSAS metric values.

From the remaining three metrics, a few interesting facts can be deduced about the attack behavior for this scenario. The PA metric shows that the percentage of attacked parts is not affected that much by the sampling plans. The improvements from random sampling only improve this metric by a value less than 2%. Similarly, the POC and PAOC metrics values indicate that the sampling plans will also have almost no effect on the behavior of the attack, as it only helps with detecting it. For instance, the POC values are almost constant for both Plans 1 and 2; while the PAOC metric values are almost the same across Plan 0b, 1, and 2, except when there is a large shift size. Also, it has been observed that the POC values are much less for both the current plans when compared to the proposed ones. Such a low value is actually an indication of a larger number of out-of-control products, as discussed in Appendix B.

5.4.1.2 Scenario 2 – Single Random Attack

Overall, the metric values for the second attack scenario have the same patterns, and some even have very similar values, to those from the first scenario. However, one notable difference with this second attack scenario is the improved performance of the practical sampling Plan 0b. The values for the majority of the metrics, such as ARL, ATS, NSAS, TPP, and FPP, are now closer to those of the proposed plans (Plans 1 and 2). This improved performance of Plan 0b can be attributed to the fact that with the attack location being random in this scenario, this sampling plan has now a better chance to detect the shifts caused by the attacks faster. In addition, this sampling plan had a lower attack detection rate in the 1st scenario because the index of the attacked product was at the center of the group of *N* products. This behavior is also confirmed by the increased values of the DP, POC, and PAOC metrics and decreased values of the PA and PSAS metrics, especially at larger shift sizes. Yet, both of the two proposed random plans still outperform the current ones across all the metrics under this attack scenario too.

Table 5-2: Performance metrics values for the 2nd attack scenario across the different sampling plans obtained from the simulation procedure with 1,000,000 iterations.

Shift Size	0	0.1	0.5	0.75	1	2.5	5	10
Average Run Length (ARL)								
Case 0	370.4030	370.8892	370.3519	369.9275	369.6625	369.9335	370.3092	370.5586
Case 0b	371.0614	368.1579	325.5480	273.3640	212.3168	30.4488	10.3638	10.1508
Case 1	370.8554	368.6959	325.7909	273.6082	212.1650	30.0637	9.9741	9.7540
Case 2	370.7781	368.5986	325.6903	273.2249	212.1027	29.6631	9.4878	9.2766
Average Time to Signal (ATS)								
Case 0	3704.030	3708.893	3703.519	3699.275	3696.625	3699.335	3703.092	3705.587
Case 0b	3710.583	3681.540	3255.445	2733.640	2123.179	304.496	103.645	101.506
Case 1	3704.049	3682.460	3253.346	2731.447	2116.938	295.677	94.753	92.554
Case 2	3708.188	3685.912	3257.093	2731.961	2120.860	296.650	94.860	92.762

Shift Size	0	0.1	0.5	0.75	1	2.5	5	10
True Positive Probabilities (TPP, %)								
Case 0	-	0	0	0	0	0	0	0
Case 0b	-	10.3839	20.8395	33.4752	48.4052	92.5659	97.4942	97.5117
Case 1	-	10.4343	20.9213	33.6330	48.3962	92.6628	97.5723	97.6675
Case 2	-	10.4792	21.0234	33.6613	48.5062	92.8052	97.7292	97.7411
False Positive Probabilities (FPP, %)								
Case 0	-	100	100	100	100	100	100	100
Case 0b	-	89.6161	79.1605	66.5248	51.5948	7.4341	2.5058	2.4883
Case 1	-	89.5657	79.0787	66.3670	51.6038	7.3372	2.4277	2.3325
Case 2	-	89.5208	78.9766	66.3387	51.4938	7.1948	2.2708	2.2589
Number of Successful Attacks to Signal (NSAS)								
Case 0	-	396.8892	369.3519	368.9275	368.6625	368.9335	396.3092	369.5586
Case 0b	-	367.2068	324.6948	272.5320	211.5607	29.9133	9.8518	9.6380
Case 1	-	367.2995	324.4401	272.3142	210.9367	29.0322	8.9643	8.7447
Case 2	-	367.6443	324.8157	272.3655	211.3296	29.1297	8.9757	8.7656
Percentage of Successful Attacks to Signal (PSAS, %)								
Case 0	-	100	100	100	100	100	100	100
Case 0b	-	99.8512	99.6742	99.3990	98.9301	91.0857	80.4593	80.1699
Case 1	-	99.8346	99.6274	99.3068	98.7665	89.1688	74.9944	74.5810
Case 2	-	99.8298	99.6305	99.3086	98.7710	89.1716	74.9872	74.6208
Percentage Attacked (PA, %)								
Case 0	-	9.8402	9.8398	9.8388	9.8394	9.8397	9.8397	9.8396
Case 0b	-	9.8569	9.8499	9.8389	9.8177	9.3915	8.7207	8.6986
Case 1	-	9.8062	9.7958	9.7848	9.7513	9.1414	8.0591	8.0296
Case 2	-	9.8061	9.8001	9.7808	9.7550	9.1446	8.0601	8.0324
Detection Probability (DP, %)								
Case 0	-	0	0	0	0	0	0	0
Case 0b	-	0.1488	0.3258	0.6010	1.0699	8.9143	19.5407	19.8301
Case 1	-	0.1654	0.3726	0.6932	1.2335	10.8312	25.0056	25.4190
Case 2	-	0.1702	0.3695	0.6914	1.2290	10.8284	25.0128	25.3792
Percentage Out-of-Control (POC, %)								
Case 0	25.4709	25.3891	23.3211	20.6860	17.3066	3.7042	1.3445	1.3123
Case 0b	25.4517	25.5006	25.4719	25.4340	25.3778	23.4715	20.0743	19.9826

Shift Size	0	0.1	0.5	0.75	1	2.5	5	10
Case 1	25.5800	25.5607	25.5816	25.5372	25.6369	25.5955	25.6304	25.6427
Case 2	25.5627	25.5724	25.5538	25.5813	25.5925	25.5703	25.6187	25.5880
Percentage Attacked Out-of-Control (PAOC, %)								
Case 0	-	0	0	0	0	0	0	0
Case 0b	-	7.3152	11.6734	15.3784	18.5424	22.7986	19.9235	19.8301
Case 1	-	7.3509	11.7692	15.4573	18.6863	24.8427	25.3960	25.4190
Case 2	-	7.3883	11.7885	15.5182	18.6791	24.8587	25.4042	25.3792

5.4.2 Across the Attack Scenarios

In this sub-section, a few metric values are highlighted for both scenarios side-by-side to assess which of the proposed sampling plan is more effective across the different attack scenarios. The metrics considered are the ARL, PAOC, and DP, as shown in Table 5-3, Table 5-4, and Table 5-5, respectively. It should be noted that the values from Plan 0 and 0b are not displayed, since it has already been established in the previous sub-section that the proposed random plans outperform the ones used in practice across all the metrics.

For the ARL metric, the values for Plans 1 and 2 are quite similar across both scenarios. However, the ARL values resulting from Plan 2 remain the lowest across all shift sizes and can detect potential Passive Product-Oriented C2P attacks from both scenarios with a slightly faster rate than the other one.

Table 5-3: ARL metric values for both attack scenario across the proposed random sampling plans.

Shift Size	Fixed Attack Scenario		Random Attack Scenario	
	Plan 1	**Plan 2**	**Plan 1**	**Plan 2**
0	370.6418	370.5610	370.8554	370.7781
0.1	368.4035	368.8045	368.6959	368.5986
0.5	325.5718	325.1569	325.7909	325.6903
0.75	273.3443	272.6598	273.6082	273.2249
1	212.5604	212.4421	212.1650	212.1027
2.5	30.0636	29.6553	30.0637	29.6631
5	9.9496	9.5048	9.9741	9.4878
10	9.7750	9.2870	9.7540	9.2766

According to the PAOC metric, both Plan 1 and Plan 2 have values that are very close from each other under both attack scenarios. Due to this closeness in the metric values across both attack scenarios, it is not possible to conclude which of the induced randomness plans is superior to the other, based on this metric alone.

Table 5-4: PAOC (%) metric values for both attack scenario across the proposed random sampling plans.

Shift Size	Fixed Attack Scenario		Random Attack Scenario	
	Plan 1	Plan 2	Plan 1	Plan 2
0	-	-	-	-
0.1	7.3333	7.3837	7.3509	7.3883
0.5	11.7504	11.7362	11.7692	11.7885
0.75	15.5006	15.5025	15.4573	15.5182
1	18.6717	18.6246	18.6863	18.6791
2.5	24.8624	24.8419	24.8427	24.8587
5	25.4501	25.3731	25.3960	25.4042
10	25.3302	25.3823	25.4190	25.3792

Similarly, it can be seen from the DP metric values in Table 5-5 that it is not clear which of the intentional randomness plans (1 and 2) outperform the other. Moreover, it seems that the randomized nature of this second attack scenario make it more difficult for any of the random sampling plans to detect out-of-control signals at lower shift sizes.

Table 5-5: DP (%) metric values for both attack scenario across the proposed random sampling plans.

Shift Size	Fixed Attack Scenario		Random Attack Scenario	
	Plan 1	Plan 2	Plan 1	Plan 2
0	-	-	-	-
0.1	1.6043	1.5886	0.1654	0.1702
0.5	1.7766	1.7774	0.3726	0.3695
0.75	2.0676	2.0591	0.6932	0.6914
1	2.5660	2.5053	1.2335	1.2290
2.5	11.7004	11.6543	10.8312	10.8284
5	25.6756	25.5632	25.0056	25.0128
10	25.9447	25.9667	25.4190	25.3792

5.4.3 Conclusions

From all the metric values presented in this section, the following can be concluded: 1) incorporating any type of random sampling plan improves the chances of cyber-physical attacks being detected; 2) the performance of the unintentional randomness plan (0b) improves with a randomized attack when compared to a fixed attack at the center of groups of *N* products; 3) the performance of the induced randomness plans (1 and 2) is constantly better than that of Plans 0 and 0b; and 4) it is not possible to determine which of Plan 1 and 2 has the best performance across the different attack scenarios, based on the presented metrics alone. However, if inspections costs are to be considered alongside these metrics, then it would be recommended to use Plan 1 as a random sampling for inspection. This recommendation is based on the fact that Plan 1 doesn’t require any additional inspection costs than the base plan (Plan 0), since just one product is inspected within a sampling group. In Plan 2, however, more than one product could be inspected within one sampling group, despite the average being one inspected product.

5.5 Summary

In this chapter, the advantages of introducing randomness into the sampling plans of the univariate individual control charts for better security against cyber-physical attacks in production has been demonstrated. More specifically, two competing random sampling plans have been proposed, where randomness was intentionally induced such that their performance could be compared to the plans currently used in practice. The performance of these random plans has been assessed under two different scenarios of single attacks within a group of *N* products, using appropriate metrics. Some of these metrics are traditional metrics, such as ARL and ATS values, while others were specially developed to consider transient shifts resulting from intelligently designed Passive Product-Oriented C2P attacks. As a conclusion, using any of the proposed random sampling plans increases the chances of detecting these attacks.

6 Introducing Randomness into Control Chart Sampling Plans for Better Security against Cyber-Physical Attacks in Production Systems – Multivariate Case

In the previous chapter, the benefits of introducing randomness to univariate control chart sampling plans for dealing with the situation of not being able to perform 100% inspection has been discussed. Such an approach is extended to the case of multivariate control charts in this chapter to better defend against cyber-physical attacks in manufacturing, when 100% inspection is a viable option. More specifically, the focus in this chapter is on tackling the issue of selecting enough quality characteristics to inspect for a given product in a cost-effective way. This issue becomes more significant when the array of possible quality characteristics to monitor increases.

6.1 Background and Motivation

During the manufacturing of critical product components with complex designs, there is usually a number of quality characteristics whose production processes need to be closely monitored using multivariate control charts throughout. Manufacturers identify these quality characteristics (or variables) and then have the choice of either 1) selecting only a subset of Key Quality Characteristics (KQCs) to inspect from the start or 2) just inspecting all of them. Then, the inspected samples are all monitored together as a multivariate statistic.

On one hand, selecting only a subset of KQCs to inspect is more cost-effective, particularly when the total number of variables considered is significant for a given product. The choice of such a subset would be made under the assumption that monitoring this subset is going to be enough to ensure that products meet design specifications and quality requirements. However, if QC users and practitioners rely on only a subset of KQCs to ensure the resulting products' integrity and overall quality, they could become victims of cyber-physical attacks, even with 100% inspection of these products (Elhabashy *et al.*, 2019a). When attackers know in advance which KQCs are being monitored, they can intelligently design Passive joint cyber-physical attacks that target the other unmonitored KQCs and physically alter them. The effects of such Passive attacks, which only use QC tools information to make changes in the manufacturing sub-system (Elhabashy *et al.*, 2018a), would then be difficult to detect.

On the other hand, inspecting all the variables has the advantage of ensuring that the effects of such intelligently-designed Passive cyber-physical attacks are detected as soon as possible. However, the number of variables to monitor could be significantly large for a couple of reasons: 1) the product itself might be quite complex, requiring a larger number of variables to be inspected and 2) manufacturers might decide to inspect as many variables as possible to prevent cyber-physical attackers from being able to target any uninspected variables. Such a large number of variables means that this option would be more expensive and computationally intensive.

To rectify this issue, variable selection approaches could be used along with control charts to identify which of the monitored variables was targeted by the attack. In fact, some of these approaches select the variables before applying multivariate control charts (Peres and Fogliatto, 2018). In this case, only a subset of variables will be used for monitoring, in spite all of them being still measured; which makes these approaches not as computationally intensive and slightly less expensive. So, inspecting all the variables along with using a variable selection approach before monitoring starts, is the best option available when defending against potential intelligently-designed attacks, from a performance perspective.

Despite being able to detect the attack effects promptly, the combination of inspecting all the variables along with variable selection still suffer from a few shortcomings. Firstly, inspecting all available variables is obviously expensive, especially when the number of variables is considerable. Secondly, with the increase in the complexity of cyber-physical attacks and their potential attack surface becoming almost unlimited, it might not even be possible to enumerate all the variables needed for effective inspection. Finally, variable selection approaches in this field typically have theoretical assumptions that would not be valid in industrial settings.

Based on the aforementioned discussion, improved samplings plans are needed that: 1) are not as costly as plans that inspects all the variables; and 2) cannot be easily exploited by intelligently-designed cyber-physical attacks. Hence, the objective of this chapter is to propose sampling plans that are resilient to the increasing threat of intelligently-designed cyber-physical attacks in manufacturing systems dealing with multiple variables and, at the same time, can select the variables to inspect for a given product in a cost-effective manner. This objective is achieved by introducing randomness to the sampling plans of multivariate Hotelling T^2 control charts for individual observations. The key challenge of this work is to make sure that the proposed sampling plans have acceptable performance, while being cost-effective.

This chapter is organized as follows; the next section contains more details about the proposed approach and how the sampling plans were developed. In the section after that, the developed approach evaluation is discussed. In Section 6.4, the results of the approach evaluation are presented. Finally, this chapter is concluded with a summary.

6.2 Developed Approach

The proposed sampling plans are based on a newly-developed random variable selection approach. The idea behind this random variable selection approach is very simple: instead of inspecting all the available *p* variables for a given product, just selecting a random subset of *q* variables to inspect every time, where $q \leq p$. Randomly selecting variables to be sampled for inspection, prevents potential attackers from knowing which variables are monitored or being able to intelligently-design corresponding Passive cyber-physical attacks. In addition, this variable selection process occurs before control chart implementation and wouldn't require inspecting or monitoring all the *p* variables, which saves unnecessary sampling costs; especially when *p* gets considerably large.

This proposed random variable selection approach is the main contribution of this work and it has three distinct advantages: 1) it helps in securing manufacturing systems against intelligently-designed cyber-physical attacks; 2) it is not as expensive as an approach that inspects all *p* variables and 3) it was developed with many practical considerations in mind, unlike other similar variable selection approaches. In the forthcoming sub-sections, we discuss how exactly is the developed approach different from others, along with providing details about the random sampling plans considered and an attack scenario under which the proposed sampling plans are evaluated.

6.2.1 Differences from other Approaches

Although the topic of variable selection has been widely researched in multivariate Statistical Process Control (SPC), less than a handful of publications considered the case of Hotelling's T^2 control charts, such as the work by Wang and Tsung (2008) and Wang and Jiang (2009). These two research efforts had theoretical assumptions that would not necessarily hold true in practice. Most notably, both approaches assumed that the mean vector and covariance matrix of the monitored variables to be available beforehand.

In practical situations, however, the mean vector and covariance matrix are unlikely to be known and would need to be estimated. The problem with estimating them is that a large number of Phase I data of size (*m*) samples is typically required for an accurate estimation. Lowry and Montgomery (1995) recommended that *m* has to be at least equal to 250 samples to provide a good estimate in general. In a more recent study, Chen and Pan (2011) used simulations to provide *m* values recommendations for individual observations for T^2 charts, when *p* ranges from 2 to 10. As shown in Table 6-1, 1,200 sample values are needed when *p* is only equal to 10 variables.

Table 6-1: Recommended Phase I sample size (m) for T^2 control charts with individual observations, according to Chen and Pan (2011).

p	3	4	5	6	7	8	9	10
m	450	550	600	700	800	1,000	1,100	1,200

When the number of variables increases even more, as with many practical situations, the required *m* values will increase considerably and having such large number of data points for Phase I data is just not going to be feasible in an industrial setting. This implies that the case of knowing the values of the mean vector and covariance matrix (or even estimating them accurately) in practice is rather ideal and will probably not be achieved. Therefore, the proposed approach was developed assuming that the mean vector and covariance matrix are not known in advance and that the estimates used instead are not going to be accurate either, in order to mimic actual industrial settings.

6.2.2 Sampling Plans

This work considers processes that manufacture products at a constant rate, where P_i is the i^{th} product manufactured. Letting V_{ij} represent the variable *j* for product *i*, then each product is characterized by a set of *p* variables designated by $\boldsymbol{V}_i$ such that $\boldsymbol{V}_i = \{V_{ij}\}\ \forall\ \mathrm{i} \geq 1 \text{ and } \mathrm{j} = 1\!: \mathrm{p}$, or simply $\boldsymbol{V}_i = \{V_{i1}, \ldots, V_{ip}\}\ \forall\ \mathrm{i} \geq 1$. Every product manufactured is inspected (100% sampling) in the considered process, such that the index *i* also represents product inspections (or measurements).

Randomness is introduced to control chart sampling plans in the developed approach through selecting a random subset of *q* variables, where $q \leq p$. Letting IV_{ik} represent the Inspected Variable *k* for product *i*, then the set of these randomly selected variables is denoted by $\boldsymbol{IV}_i$ such that $\boldsymbol{IV}_i = \{IV_{ik}\}\ \forall\ \mathrm{i} \geq 1 \text{ and } \mathrm{k} = 1\!: \mathrm{q}$, or simply $\boldsymbol{IV}_i = \{IV_{i1}, \ldots, IV_{iq}\}\ \forall\ \mathrm{i} \geq 1$. Each element of $\boldsymbol{IV}_i$ is randomly chosen, without repetition, from $\boldsymbol{V}_i$, such that $\boldsymbol{IV}_i \subseteq \boldsymbol{V}_i$.

The first sampling plan developed in this work is the one where a fixed number of *q* variables is chosen every time sampling occurs, but the selected variables themselves are randomized from inspection to inspection. This plan will be referred to as **Plan 1 – Fixed-sized Random Subset**, where the size of the selected subset for each product *i* is fixed: $q_i = q\ \forall\ i \geq 1$.

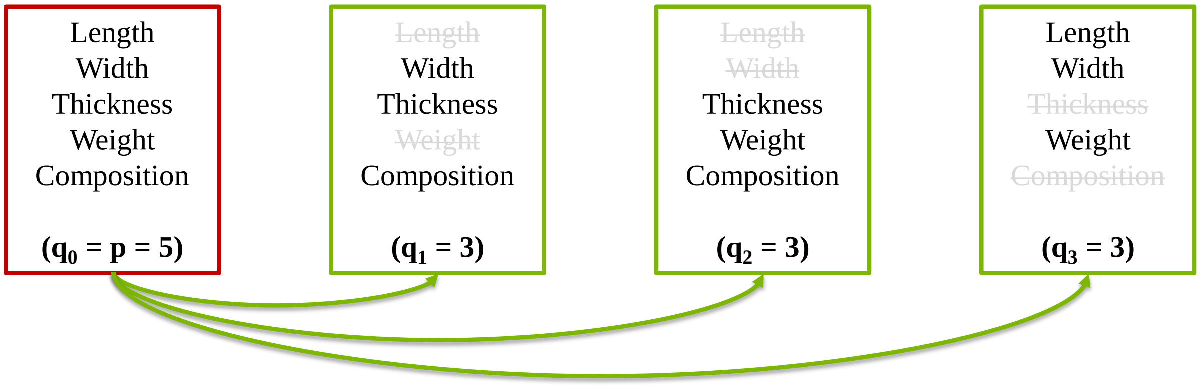

Figure 6-1: Examples of fixed-sized random subsets used in the sampling plan of multivariate T^2 control charts, instead of monitoring all the variables as shown on the far left.

For example, if an individuals T^2 control chart is used to monitor only p = 5 variables within every inspection, then the proposed sampling plan could be to monitor only a subset of q = 3 variables instead; where the 3 variables are randomly chosen every time, as shown in Figure 6-1. In this example, $\boldsymbol{IV}_1 = \{V_{12}, V_{13}, V_{15}\}$, $\boldsymbol{IV}_2 = \{V_{23}, V_{24}, V_{25}\}$, and $\boldsymbol{IV}_3 = \{V_{31}, V_{32}, V_{34}\}$.

Another alternative is having the number of inspected q variables to vary from inspection to inspection. In such a sampling plan, not only would the selected variables be randomly chosen within each subset of size q_i, but also the size of each subset would be randomized every time inspection occurs. This alternative plan will be referred to as **Plan 2 – Variable-Sized Random Subset**, where the size of the selected subset for each product i is variable: $q_i \sim U\{2, p\}\ \forall\ i \geq 1$.

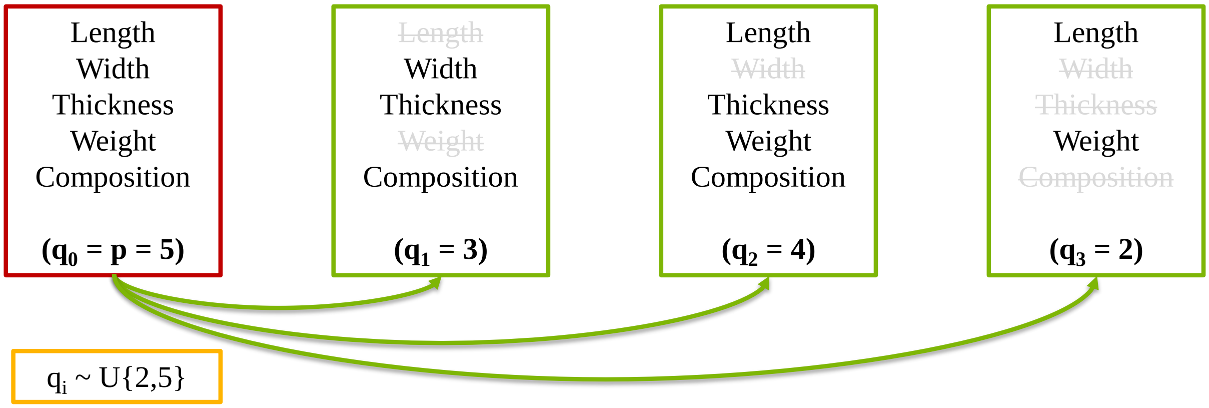

Figure 6-2: Examples of variable-sized random subsets used in the sampling plan of multivariate T^2 control charts, instead of monitoring all the variables as shown on the far left.

Considering again the T^2 control chart monitoring the same p = 5 variables. The value of the selected subset q_i variables could be either 2, 3, 4, or even 5 (i.e. all the variables), as seen in Figure 6-2. In this example, $\boldsymbol{IV}_1 = \{V_{12}, V_{13}, V_{15}\}$, $\boldsymbol{IV}_2 = \{V_{21}, V_{23}, V_{24}, V_{25}\}$, and $\boldsymbol{IV}_3 = \{V_{31}, V_{34}\}$.

6.2.3 Attack Details

Next is to identify the intelligently designed Passive cyber-physical attack scenario under which the proposed approach is evaluated. Each time a product is manufactured, the cyber-physical attack would cause a shift in the manufacturing process mean to affect a random number M_i of variables, where M is the upper bound of the number of attacked variables; such that $M_i \leq M \leq p$. For each product i, AV_{il} represents the randomly Attacked Variable l and the set of Attacked Variables is designated by an $\boldsymbol{AV}_i$ such that $\boldsymbol{AV}_i = \{AV_{il}\} \; \forall \; \mathrm{i} \geq 1 \text{ and } \mathrm{l} = 1{:}\, \mathrm{M_i}$, or simply $\boldsymbol{AV}_i = \left\{AV_{i1}, \dots, AV_{iM_i}\right\} \forall \; i \geq 1$, where $M_i \sim U\{2, M\} \; \forall \; i \geq 1$. Each element of $\boldsymbol{AV}_i$ is also randomly chosen, without repetition, from $\boldsymbol{V}_i$, such that $\boldsymbol{AV}_i \; \subseteq \; \boldsymbol{V}_i$.

6.3 Approach Evaluation

After discussing the specifics of the developed approach, the purpose of this section is to demonstrate how the developed approach would be evaluated when compared to an approach that inspects all the variables prior to selecting which one to monitor from them. With the goal being finding out which of the proposed random sampling plans is more cost-effective, the evaluation criterion used is first discussed. Then, the simulation procedure used for this evaluation is detailed. Finally, the idea behind a simplified version of one of the variable selection approaches in the literature is presented. This simplified approach would be the one against which the developed approach is compared.

6.3.1 Evaluation Criteria

Since the developed approach is based on monitoring random subsets of variables, its performance would not be the same as one that inspects all the variables and then selects the variables more likely to be out-of-control to monitor. The latter approach enables multivariate control charts to detect process shifts faster. In addition, the proposed approach, unlike other closely-related variable selection approaches, uses estimates of the mean and the covariance matrix that may be inaccurate, which further hinders its performance when compared to these other approaches. Therefore, it would not be appropriate to use common performance metrics, such as the Average Run Length (ARL) and Average Time to Signal (ATS), to evaluate the developed approach.

Instead, the focus would be on cost estimates as the main evaluation criterion of the developed approach. Obviously, since a lower number of random variables is inspected, the cost associated with the developed approach should be less than that of an approach that inspects all the variables;

especially when the total number of considered variables increases significantly. If such cost savings are not considerable, applying any of the proposed sampling plans of the developed approach would not be worthwhile. To estimate the associated costs, an economic cost model is needed.

Economic cost models for control charts have been discussed by many researchers, with the aim of optimally selecting a control chart's sample size, sampling frequency, and control limits (Montgomery, 2009). Despite the development of different economic cost models, these models suffer from some practical limitations, as discussed by Woodall *et al.* (1986). However, since our objective is only to compare the inspection costs of the different sampling plans and not to economically design control charts, using some aspects of these models should not be problematic.

A well-known economic cost model is the one proposed by Lorenzen and Vance (1986), where the authors presented a unified model that can be applied to many types of control charts. This cost model can also be extended to the case of Hotelling T^2 control chart (Chalaki *et al.*, 2016) and is discussed further in Appendix D. Since our main interest is only in cost comparisons of the different approaches, a simplified version of the model by Lorenzen and Vance (1986) is used to determine the inspection costs of the different sampling plans. Such a simplified model only considers the cost of sampling and the cost of manufacturing non-conforming products, with the latter referred to as "production cost", instead of all the other components in the original cost model. The idea behind such a model is similar to that of the one proposed by Chen and Pan (2011) to perform similar cost comparisons. The assumptions made in this simplified cost model along with its details are discussed next.

6.3.1.1 Assumptions

In addition to the assumptions made in the original model by Lorenzen and Vance (1986), the simplified cost model has an extra set of assumptions including:

- The sample size (n) is assumed to be equal to 1 (individual observations);
- The sampling frequency (h) is assumed to be equal to 1 hour, such that 100% sampling occurs when products are manufactured every hour;
- The quality characteristics monitored follow a p-variate normal distribution with an estimated mean vector and covariance matrix;
- Production is stopped to search for the cause of the potential assignable cause once a signal is obtained;

- The assignable causes are assumed to occur according to a Poisson process with an intensity of λ occurrences per hour (Montgomery, 2009). So, the in-control time can be considered to follow a negative exponential distribution with mean 1/λ (Lorenzen and Vance, 1986).
- All the following parameters/variables are not considered in the simplified cost model:
  - The expected time to sample and chart one product;
  - The time of repairing an assignable cause;
  - The time interval during which the assignable cause is detected and identified;
  - The hourly production cost due to non-conformities when the process is in-control;
  - The cost of signaling;
  - The fixed sampling cost.

6.3.1.2 Simplified Cost Model

The inspection cost in the simplified model will only account for the sampling and production costs. For the sampling cost, since the sample size *n* is assumed to be always equal to 1, then the average number of inspected variables (q_{average}) will be considered instead for cost estimation; as shown in equation (6-1). The reason for using an average *q* value, instead of the actual ones, is because the value of *q* is constantly changing in the random sampling Plan 2.

$$\mathrm{E(Sampling\ Cost)} = (a_2 q_{\text{average}})\ \mathrm{E(T)} \tag{6-1}$$

Where a_2 is the variable cost of sampling and testing per production cycle and *E(T)* is the expected cycle time, which can be obtained from equation (6-2):

$$\mathrm{E(T)} = \frac{1}{\lambda} + \mathrm{ARL}_1 - \tau \tag{6-2}$$

The ARL_1 is the out-of-control average run length, while τ is the expected time between assignable causes occurrence and the last in-control sample such that:

$$\tau = \frac{1 - (1 + \lambda)e^{-\lambda}}{\lambda(1 - e^{-\lambda})}.$$

The second cost considered, which is the production cost, will be obtained from equation (6-3), where C_1 is the hourly production cost due to non-conformities when the process is out-of-control.

$$\mathrm{E(Production\ Cost)} = C_1[\mathrm{ARL}_1 - \tau] \tag{6-3}$$

Finally, the total expected cost, E(C), would be the result of adding equations (6-1) and (6-3), resulting in a linear variable cost model:

$$\mathrm{E(C)} = \mathrm{a}_2\mathrm{q}_{\mathrm{average}}\left[\frac{1}{\lambda} + \mathrm{ARL}_1 - \tau\right] + C_1[\mathrm{ARL}_1 - \tau] \tag{6-4}$$

6.3.1.3 Needed Parameters

According to equation (6-4), there are three parameters that need to be specified beforehand for the simplified cost model: the variable cost of sampling and testing, a_2; the exponential distribution parameter, λ, for the assignable causes; and the hourly production cost due to non-conformities when the process is out-of-control, C_1. In this work, it is assumed that a_2 equals \$5 per variable to be sampled for a single product and that C_1 equals \$10 per hour.

To determine the value of λ, it is important to note that intelligently-designed cyber-physical attacks are considered as one of the possible assignable causes in this simplified cost model. In this case, the cyber-physical attacks are treated as just a single assignable cause that affects several product variables simultaneously. Given the nature of these intelligently-designed attacks, along with the fact that 100% sampling is carried out every hour, it will be assumed that process shifts are to occur on an hourly basis too. Hence, the value of λ used in the simplified model would be equal to 1.

Equation (6-4) has two variables that needs to be determined as well: the average number of subsets, $q_{average}$, and the values of the out-of-control average run length, ARL_1, considering different shift sizes in the process mean. In order to do so, different values of q need to be specified first, then the corresponding ARL values can be evaluated. One way to obtain ARL values is to use theoretical formulas for the error probabilities, such as those described in the last section of Appendix D. Using such formulas for in-control situations is straight-forward. However, since the mean vector and covariance matrix are assumed to not be known in advance and must be estimated first, similar formulas could not be used in the case of out-of-control situations. Hence, computer simulations are used in this work to determine the out-of-control ARL values for the different sampling plans under the attack scenario considered.

6.3.1.4 Simulation Procedure

Regardless of the sampling plan and attack scenario considered, the simulation procedure followed to evaluate the out-of-control ARL values is rather simple and can be summarized in the subsequent steps:

1. Specifying the values of the different parameters used throughout, such as:
   - Phase I data size: m = 500 samples;
   - Type I error probability: α = 0.005;
   - Total number of variables considered: p = 10, 20, 50, and 100 variables;
   - Number of randomly selected variables for inspection: $q = 0.1\mathrm{p} \times [1\ 3\ 6\ 10]$. For instance, if p equals 20, the corresponding q values considered are 2, 6, 12, and 20 variables;
   - The upper bound of the number of attacked variables: $M = [0.5\mathrm{p}\ \ \mathrm{p}]$;
   - Sizes of the shifts to the process mean start from 0.5 up to 4, with increments of 0.5.
2. Generating random data of size *m* following a *p*-variate normal distribution of mean vector and covariance matrix values equal to 0 and **1**, respectively. This data set would represent the measured value of all the variables during Phase I;
3. Estimating the mean vector and covariance matrix values of this Phase I data, to be used for Phase II monitoring;
4. Using the estimated mean and covariance matrix to generate a random data of size 1 following a *p*-variate normal distribution, representing the measured values of the variables for the current product;
5. For the current product's variables that would be affected by the attack, introducing mean shifts to the generated data of these variables in a manner that mimics the considered attack scenario;
6. Selecting the q variables to be monitored when applying the different sampling plans;
7. Plotting the corresponding T^2 statistic value for only the selected variables on the chart and checking if it is within the control limits;
8. Repeating steps 4-7 until a T^2 statistic value is outside of the control limits and a signal is obtained, while keeping track of the number of inspections to get the Run Length (RL) values;
9. Repeating steps 2-8 for 10,000 simulation iterations, to obtain the ARL values across all iterations. The obtained ARL values for both the random sampling plans are shown in Appendix E.

A few important implementation considerations need to be taken into account when determining the covariance matrix, Upper Control Limits (UCL), and the applied shift size used in the simulation procedure. Firstly, when estimating the covariance matrix, the formula used should be the one involving calculating the difference between successive pairs of observations, rather than the one containing a pooled estimator. Montgomery (2009) provides a brief discussion about the difference between the 2 commonly used estimators, along with a numeric example showing the different resulting T^2 charts from using each one of them for the same set of data.

Secondly, the UCL value for Phase II should be based on the F-distribution, as shown in equation (6-5), since other approximations for the Phase II UCL values have limitations to when they should be applied. However, it has been observed that with such a low *m* value of 500 samples, the in-control ARL values obtained from the simulation procedure when using this formula are not necessarily equal to the theoretical in-control ARL values; especially with large *p* values. So, with some experimentation, a more appropriate UCL value was obtained beforehand for each sampling plan to ensure that the in-control ARL value is accurate, while using the value provided by equation (6-5) as a starting point. Any UCL value providing an in-control value within ±2 units of the theoretical in-control ARL value was deemed acceptable.

$$\text{Phase II UCL} = \frac{p(m+1)(m-1)}{m^2 - mp} \mathrm{F}_{\alpha,\mathrm{p},\mathrm{m}-\mathrm{p}} \tag{6-5}$$

Finally, since more than one variable could be attacked at the same time and the number of attacked variables changes from product to product, it has to be ensured that the resulting shift is equivalent for all the attacked variables. In order to do that, the Mahalanobis distance was used to represent the magnitude of the mean shift (δ), such that the $\delta = shift\ size/\sqrt{M_i}$. For example, if the shift size is equal to 1 and only 1 variable is attacked, then $\delta = 1$; if 2 variables are attacked instead, then $\delta = 0.707$ for each variable. This calculated value of δ is the one used to represent the shifts by the cyber-physical attacks for each product.

6.3.2 Naïve Variable Selection Approach

The most closely-related variable selection approach in this research area is the one developed by Wang and Jiang (2009). So, it would be ideal if our proposed approach could be compared to their approach, which is quite effective in selecting variables that are more likely to be shifted. However, since this approach requires inspecting all the available *p* variables in advance and using their information, its computational speed increases drastically with the increase in the size of both *p* and *q* as it needs go through a lot of optimization routines. Hence, it was not possible to replicate their approach, using computer simulations, for the parameters specified in the previous sub-section in a reasonable amount of time and an alternative variable selection approach that is much less time-consuming was needed.

This alternative variable selection approach can be considered as a much simplified version of the approach by Wang and Jiang (2009) and would be referred to as the **Naïve Variable Selection Approach**. This naïve variable selection approach doesn't not involve any optimization routines and is much more computationally efficient with respect to run time. This approach, instead, selects the q variables with the largest measured values for each product i; and those are the only ones that are monitored and used in calculating the T^2 statistic value for this product. This variable selection approach can be represented mathematically as follows.

Let the measured values of the monitored variables be denoted by $\boldsymbol{Y}_i = \{Y_{ij}\} \; \forall \text{ i} \geq 1 \text{ and j} = 1\text{: p}$; where Y_{ij} represents the measured value of the variable j for product i. In other words, $\boldsymbol{Y}_i = \{Y_{i1}, \dots, Y_{ip}\} = \{F(V_{i1}), \dots, F(V_{ip})\} \; \forall \text{ i} \geq 1$, where $F(V_{ij})$ also represent the measured values of the variables V_{ij}. These Y_{ij}'s are then sorted descendingly to obtain the corresponding order statistics denoted by $\boldsymbol{O}_i = \{O_{is}\} \; \forall \text{ i} \geq 1 \text{ and s} = 1\text{: q}$, such that $O_{i1} > O_{i2} > \cdots > O_{iq}$ for product i. The first ordered statistic for product i, O_{i1}, is the variable corresponding to the largest Y_{ij}, i.e., to $\max(\boldsymbol{Y}_i)$; the 2nd ordered statistic, O_{i2}, corresponds to the next largest Y_{ij}; and so on until the q^{th} ordered statistic.

Aside from the process of selecting the variables, this naïve approach can be used to estimate the out-of-control ARL values in the same manner as the random variable selection approach. Hence, the simulation procedure would only need to account for selecting the variables differently when applied for the naïve approach. More specifically, the way the variables are selected in step 6 is the only thing that would need to change in the simulation procedure described in the previous sub-section.

Finally, it should be just noted that this naïve variable selection approach could be used to obtain an approximation of the upper bound of the variable selection approach of Wang and Jiang (2009). The ARL values for the naïve approach resulting from simulation (not shown here) are very close to the ones resulting from their approach for the same design parameters. Furthermore, the difference in the ARL values of both methods converges to zero with the increase in the values of p and q.

6.4 Results and Discussion

In this section, the random variable selection approach, referred to as just the “randomness” approach, is compared to the naïve variable selection approach, referred to as the “naïve” approach from a cost perspective. More specifically, the cost savings obtained from the simplified cost model for the randomness approach when compared to the naïve approach are the results presented in this section; the actual costs, however, associated with each approach are detailed in Appendix F. These results should help in assessing the usefulness of the randomness approach compared to another variable selection approach in reducing the accompanying inspection costs.

6.4.1 Sampling Plan 1 - Fixed-sized Subsets

The first set of cost comparisons are from using sampling Plan 1 when *p* = 10 variables for different combinations of attacks and variables subsets, as shown in Table 6-2. From this table, it can be observed that: 1) the randomness approach allows for more cost savings with small shift sizes; 2) the highest cost savings are obtained when *q* is equal to its lowest possible value and there are almost no cost savings when *q* is equal to its maximum possible value, which makes sense as both approaches would then involve inspecting all the variables; and 3) when the value of *M* increases, i.e., when more variables could be attacked, the cost savings decrease a little bit.

Table 6-2: Cost savings (%) resulting from both approaches, when p = 10 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 1.

Shift Size	**p = 10 and M = 5**				**p = 10 and M = 10**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	1.9900	0.6034	-0.0569	-0.0016	1.8038	0.4124	-0.1938	0.0152
1	1.2258	0.1587	-0.2848	0.0109	1.0887	-0.0277	-0.4356	0.0240
1.5	0.6521	-0.0896	-0.3906	0.0091	0.6675	-0.2218	-0.5048	-0.0233
2	0.3073	-0.1929	-0.4011	-0.0138	0.4161	-0.2842	-0.5242	0.0114
2.5	0.1088	-0.2322	-0.3533	0.0064	0.2314	-0.3103	-0.4943	0.0159
3	-0.0203	-0.2232	-0.3411	-0.0076	0.1419	-0.2913	-0.4304	0.0189
3.5	-0.0391	-0.1905	-0.2456	0.0005	0.0726	-0.2553	-0.3336	-0.0122
4	-0.0507	-0.1323	-0.1660	-0.0122	0.0640	-0.1603	-0.2213	-0.0169

When the total number of variables monitored increase to *p* = 100, the cost comparisons are shown in Table 6-3. It can be seen from this table that the same patterns just discussed are also present when the value of *p* increases too. More notably, the cost savings increase with the increase of the value of *p* when using the randomness approach compared to the naïve one. Finally, it is just important to point out that the actual costs associated with both approaches has increased significantly (as shown in Appendix F), which means that the cost savings now are even more substantial.

Table 6-3: Cost savings (%) resulting from both approaches, when p = 100 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 1.

Shift Size	**p = 100 and M = 50**				**p = 100 and M = 100**			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0.5	5.3291	1.0791	0.0279	-0.0088	4.8915	0.7910	-0.1344	-0.0056
1	3.7552	0.4133	-0.2915	-0.0218	2.9668	0.0743	-0.4680	0.0028
1.5	2.6196	0.0398	-0.4868	-0.0201	1.8297	-0.2900	-0.6597	-0.0270
2	1.7497	-0.2430	-0.5982	-0.0133	1.1367	-0.4796	-0.7571	-0.0155
2.5	1.1036	-0.3783	-0.6718	0.0143	0.6323	-0.6007	-0.7968	0.0141
3	0.7016	-0.4637	-0.7025	-0.0183	0.3408	-0.6503	-0.8154	-0.0157
3.5	0.4206	-0.5136	-0.7172	-0.0088	0.1240	-0.6714	-0.8163	-0.0029
4	0.1973	-0.5301	-0.7026	-0.0150	0.0169	-0.6714	-0.8016	-0.0232

6.4.2 Sampling Plan 2 - Variable-sized Subsets

A couple of examples of the resulting cost savings when applying the second sampling can be seen in Table 6-4 and Table 6-5 for an increasing value of *p*. Similar patterns to that of the previous sampling plan are observed with this one as well. One difference, though, is that there is a small amount of cost savings when *q* equals *p* in this sampling plan. Furthermore, the cost savings seems to be more in Plan 2 than that of Plan 1. However, comparing both the plans directly for the same value of *q* wouldn't be a fair comparison, since the value of q_i is randomized for each product *i* being inspected in Plan 2. Instead, the comparisons should be based on the $q_{average}$ values, which were used in the simplified cost model evaluation.

Table 6-4: Cost savings (%) resulting from both approaches, when p = 10 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 2.

Shift Size	**p = 10 and M = 5**				**p = 10 and M = 10**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	2.0680	0.9908	0.4451	0.3433	1.8225	0.8142	0.2373	0.2543
1	1.2447	0.4593	0.0504	0.1618	1.1790	0.2789	-0.1369	0.0491
1.5	0.6461	0.1987	-0.1391	0.0605	0.7093	0.0301	-0.2919	-0.1158
2	0.3386	0.0395	-0.1623	0.0125	0.4166	-0.0711	-0.3446	-0.1497
2.5	0.1165	-0.0649	-0.2114	-0.0043	0.2522	-0.1101	-0.3150	-0.1498
3	0.0131	-0.0823	-0.1793	-0.0147	0.1295	-0.1359	-0.2925	-0.1338

Shift Size	p = 10 and M = 5				p = 10 and M = 10			
	q = 1	q = 3	q = 6	q = 10	q = 1	q = 3	q = 6	q = 10
3.5	-0.0430	-0.0676	-0.1270	0.0126	0.0943	-0.0940	-0.1978	-0.0651
4	-0.0707	-0.0138	-0.0563	0.0558	0.0460	-0.0183	-0.0841	0.0171

Table 6-5: Cost savings (%) resulting from both approaches, when p = 100 variables across different shift sizes and attack values for a range of *q* values, using sampling Plan 2.

Shift Size	p = 100 and M = 50				p = 100 and M = 100			
	q = 10	q = 30	q = 60	q = 100	q = 10	q = 30	q = 60	q = 100
0.5	9.7202	3.1347	1.1561	0.9415	8.7726	2.5819	0.8050	0.9542
1	7.3425	1.9621	0.4839	0.9511	6.1860	1.3904	0.1120	0.9392
1.5	5.2994	1.2015	0.0766	0.9345	4.4089	0.6269	-0.2742	0.8647
2	4.1455	0.7106	-0.1857	1.0512	3.1758	0.1882	-0.4732	0.9102
2.5	3.1467	0.3679	-0.3267	1.1119	2.2328	-0.0376	-0.5692	0.9150
3	2.3845	0.2173	-0.3999	1.2382	1.8074	-0.1248	-0.6017	0.8924
3.5	1.8915	0.1032	-0.4394	1.3868	1.3613	-0.2419	-0.6153	0.9679
4	1.5085	0.0033	-0.4373	1.4392	1.1587	-0.2331	-0.5926	1.0034

6.4.3 Discussion

From these results, it can be seen that applying either of the random sampling plans provides a cost-effective alternative to an approach that involves variables selection after all the variables have already been inspected. In general, the cost savings from applying the randomness approach are at its greatest value when a small subset of *q* variables is being used, along with a small shift size. Furthermore, when the total number of considered variables *p* increases considerably, the cost savings from the randomness approach, when compared to the naïve one, are even more (regardless of which of the sampling plans is being applied). Finally, comparing both of the proposed sampling plans from a cost perspective, or any other one, should be only based on the average value of *q* for a given product.

6.5 Summary

In this chapter, the usefulness of introducing randomness into the sampling plans of multivariate control charts to design more resilient QC tools to cyber-physical attacks have been discussed. More specifically, two sampling plans were developed, based on a random variable selection approach that does not require inspecting all considered variables, which were still able to detect intelligently designed cyber-physical attacks. Using a simplified cost model, the proposed random variable selection approach was evaluated under one specific attack scenario when compared to a “naïve” variable selection approach in which one inspects all the variables but monitors only a subset *q* of them. The cost comparisons have shown that the random variable selection approach allows for more cost savings when 1) the size of the mean shift is small, 2) the subset size *q* is much lower than the original *p* variables, and 3) the number of the total *p* variables increases considerably. These results confirm that the developed approach is cost-effective and well-suited for industrial applications where *p* is quite large.

7 Contributions and Future Work

This final chapter summarizes the contributions made in this dissertation and discusses future work ideas in the research topic of QC tools for cyber-physical security of production systems.

7.1 Contributions

The proposed work here in the field of quality control allows quality control systems to be a more effective second line of defense for manufacturers against the increasing threat of cyber-physical attacks. This work has a set of overall expected contributions that are discussed in the first sub-section; whereas Chapters 3 through 6 have unique expected contributions that are discussed in the subsequent sub-sections, respectively.

7.1.1 Overall Contributions

The first overall contribution of this work is **developing QC techniques that are more resilient to cyber-physical attacks** within manufacturing. While the scope of this work is only on detecting the effects of cyber-physical attacks, unlike that of IT security that aims to detect the attacks themselves, these techniques are more effective in preventing, or at least minimizing, the occurrences of such attacks against production systems. These novel techniques, such as the proposed taxonomy and the QC approaches, were designed with cyber-physical attacks in mind, making it more difficult for attackers to exploit existing weaknesses and easier for defenders to better appreciate the capabilities and limitations of their systems.

Another overall contribution is **making production systems more secure against cyber-physical attacks** through significantly reducing the frequency of the associated attacks or even eliminating them completely. Having more secure production systems will enable manufacturers to avoid fabricating faulty products unintentionally or dangerous products unknowingly, unnecessary production delays, warranty or recall issues, and dissatisfied customers' concerns. This would also lead to minimizing safety risks for all humans involved and, ultimately, to better and more trustworthy products.

The third overall contribution is **raising the general awareness among manufacturers** in different industries about cyber-physical attacks in manufacturing. This would allow manufacturers not only to realize the possibility of being targets of cyber-physical attacks, but also the serious effects these attacks could have on them. In addition, manufacturers would be able to have a better understanding to the means through which traditional QC approaches could be exploited to generate attacks that may not be easily detected. Manufactures would then recognize the need for better securing production systems and developing appropriate action plans in the event of cyber-physical attacks. This is important, since cyber-physical attacks may not be considered as common threats nor even be a cause of concern in production systems by many manufacturers (NDIA, 2014; Wells *et al.*, 2014).

7.1.2 Cyber-physical Attack Taxonomy from a QC Perspective

The main contribution of the presented work in this research area is **creating a systematic approach to describe the cyber-physical attack surface components** in manufacturing through the developed taxonomy. More specifically, the taxonomy makes it possible to analyze the different layers involved in an attack and to define the resulting types of cyber-physical attacks in a systematic and repetitive manner. As shown in Chapter 3, with the taxonomy, it is also now possible to assess the impact of each different layer within the system on the characteristics of cyber-physical attacks.

Another advantage of the proposed taxonomy is understanding the relationships between cyber-physical attacks, QC systems, and manufacturing in the context of malicious process changes. Hence, an additional contribution in this research area is **identifying the role QC systems play in cyber-physical attacks within production systems**. For instance, compromising QC systems can either mask the effects of an attack to reduce outgoing product quality and/or alter the product's design intent or disrupt manufacturing operations through incorrectly indicating that something is wrong.

In addition, having such a taxonomy also helps in **understanding the potential thought process and different considerations involved during a cyber-physical attack**, which is the last contribution in this research area. Obtaining a potential outline for how a cyber-physical attack is most likely to be carried out in manufacturing could be quite useful. Similar to methods used in software systems development, such as penetration tests, looking at the system from an attacker's potential point of view allows for better defense mechanisms to be implemented for counterattacking and decreasing attacks occurrences and/or impact.

7.1.3 Cyber-Physical Attack Vulnerabilities in QC Tools

The first contribution of this research area would be **discussing the different vulnerabilities existing in QC tools**. Previous research work, such as those by Wells *et al.* (2014) and Vincent *et al.* (2015), have pointed out that current QC tools are based on specific assumptions that may no longer be valid during cyber-physical attacks, but didn't discuss these assumptions. Our work, however, have actually presented how these assumptions could be invalidated during cyber-physical attacks and the resulting vulnerabilities.

The second contribution in this research area is **showing the effects of exploiting QC tools vulnerabilities within manufacturing by cyber-physical attacks**. In other words, the contribution lies in demonstrating how exploiting these vulnerabilities by cyber-physical attacks can lead to a decreased performance of QC tools. This is vital because it should act as an incentive for manufacturers to be more aware of such negative consequences and be more careful when using these tools.

Demonstrating that the different types of vulnerabilities can be categorized systematically in manufacturing systems is the third contribution within this research area. This categorization would help researchers in gaining a better understanding about the different types of vulnerabilities present within the tools used. Such an understanding should also help manufacturers to be able to assess the impact of potential exploitations, make better decisions regarding the use of those tools, and better secure their cyber-physical production systems.

7.1.4 Introducing Randomness to Control Charts

For this research area, which was covered by both Chapters 5 and 6, the principal contribution is **developing novel approaches for more effective QC tools within production systems**. These approaches, which were based on introducing randomness in the sampling plan of different types of control charts, were developed to deal with different issues in both univariate and multivariate situations. A key advantage of these novel approaches is allowing for faster detection of cyber-physical attacks effects in the univariate case and providing a cost-effective variable selection approach during 100% sampling for the multivariate case.

Another important contribution is **being able to model actual cyber-physical attack scenarios in manufacturing systems**. Throughout the previous two chapters, different attack scenarios were implemented to assess the developed approaches. Modelling these different attacks helped with designing corresponding solutions to decrease the impact these attacks might have on manufacturing systems. As previously mentioned, considering issues from the attacker's perspective, helps in devising better solutions for defenders.

Furthermore, **creating unique evaluation criteria for assessing the usefulness of the proposed approaches** against cyber-physical attacks in manufacturing is an additional contribution within this research area. For the univariate case, it was necessary to create new performance metrics to deal with this unique type of transient shift caused by cyber-physical attacks, since the traditional metrics may not take the possibility of cyber-physical attacks into consideration adequately enough, and to learn even more information about the tool considered. For the multivariate case, a simplified cost model approach was used to determine whether the newly-proposed random variable selection approach was cost-effective or not.

As a final contribution in this research area, this work **presented a different way for dealing with the issue of cyber-physical security in manufacturing**. Firstly, an idea from the IT security domain has been extended for use within the cyber-physical security domain. Secondly, new evaluation criteria have been developed to assist in evaluating the usefulness of the proposed approaches. These evaluation criteria played an essential role in ensuring that the ARL performance of the proposed approach is comparable to existing methods and, at the same time, the approaches are cost-effective.

7.2 Future Work

The work in this dissertation has paved the way for developing new and/or redesigning existing QC tools for better security in manufacturing systems, which should be the main direction of future work in this research field. Introducing randomness to QC tool sampling plans have proven to be cost-effective, but both of the corresponding approaches presented in the previous two chapters could still use further improvement and is one future research avenue. Furthermore, additional QC tools that are resilient to cyber-physical attacks could be also developed. Some of the potential future work ideas are highlighted in the following sub-sections.

7.2.1 Introducing Randomness to Univariate Control Charts

With the promising results of the effort discussed in Chapter 5, there are more future work ideas to consider. One idea is to assess the effects of changing the location of the attack within the groups of *N* products during the fixed attack scenario; it is suspected that such changes would affect the performance of Plan 0b only. Another idea is to consider more attack scenarios and random sampling combinations; for example, introducing multiple attacks or multiple random plans per a group of *N* products. Similarly, attacks of varying shift sizes could also be considered. Finally, since this work covered only the individuals control chart, future efforts would also include extending this work to other types of univariate control charts.

7.2.2 Introducing Randomness to Multivariate Control Charts

With regards to the random variable selection approach for multivariate control charts sampling plans, future work ideas include, but are not limited to: 1) developing more sampling plans, 2) considering more attack scenarios; and 3) improving the simplified cost model. An interesting sampling plan to consider, for example, is a plan that constantly selects one or two variables every time inspection occurs, while randomly selecting the remaining q variables. Such a sampling plan could be applied in conjunction with either the fixed-sized or varying-sized random subset sampling plans.

The work discussed in Chapter 6 only had one attack scenario that targeted multiple variables randomly for each product. Alternate attack scenarios could also be used to further assess the proposed sampling plans, such as a scenario that attacks a fixed number of variables or attacks a mix between fixed and random variables per product. Finally, the simplified cost model treated the attacks as assignable causes following a Poisson process, such an assumption could be generalized to other distributions and the cost model could also include assignable causes that are not due to cyber-physical attacks.

7.2.3 Improving QC Tools to be More Resilient

As suggested in Chapter 4, more resilient QC tools against cyber-physical attacks are needed and the IT domain contains several approaches that could be used for the development of such tools. One approach that was already discussed in Chapter 4 is using a unique “tag” for ensuring that any changes made by the attacker to the process can be quickly detected. Another approach would be employing some type of encryption to the common language used along the different QC tools; doing so would drastically reduce the ability of an attacker to interpret the states of the QC system. Finally, it could be also possible to implement a mixed off-line/on-line QC system. The goal of such a mixed system would be minimizing the overall system vulnerability while maximizing the amount of data available for system analytics to reduce the likelihood of successful cyber-physical attacks in manufacturing.

References

Abdelhamid, S. E., Elhabashy, A. E., Reid, K. J. and Kuhlman, C. J. (2015) Studying the Impact of Social Interactions on Students' Academic Retention using Agent-based Simulation, in *the 7th Annual First Year Engineering Experience (FYEE) Conference*.

Al Faruque, M. A., Chhetri, S. R., Canedo, A. and Wan, J. (2016a) Acoustic Side-channel Attacks on Additive Manufacturing Systems, in *the Proceedings of the 7th International Conference on Cyber-Physical Systems (ICCPS '16)*, IEEE Press, pp. 1-10.

Al Faruque, M. A., Chhetri, S. R., Faezi, S. and Canedo, A. (2016b) Forensics of Thermal Side-channel in Additive Manufacturing Systems, University of California, Irvine.

Albakri, M., Sturm, L., Williams, C. B. and Tarazaga, P. (2015) Non-destructive Evaluation of Additively Manufactured Parts Via Impedance-based Monitoring, in *Proceedings of the 26th Annual International Solid Freeform Fabrication Symposium*, pp. 1475-1490.

Albakri, M., Sturm, L., Williams, C. B. and Tarazaga, P. (2017) Impedance-based Non-destructive Evaluation of Additively Manufactured Parts. *Rapid Prototyping Journal*, **23**(3), 589-601.

Albright, D., Brannan, P. and Christina, W. (2010) Did Stuxnet Take Out 1,000 Centrifuges at the Natanz Enrichment Plant?, in Institute for Science and International Security (ISIS).

Anthem (2016) How to Access & Sign Up For Identity Theft Repair & Credit Monitoring Services, Available from https://www.anthemfacts.com/cyber-attack, [Last accessed on 9 April 2017].

Antony, J. and Taner, T. (2003) A Conceptual Framework for the Effective Implementation of Statistical Process Control. *Business Process Management Journal*, **9**(4), 473-489.

Assante, M. J. and Lee, R. M. (2015) The Industrial Control System Cyber Kill Chain, The SANS Technology Institute.

Bayanifar, H. and Kühnle, H. (2017) Enhancing Dependability and Security of Cyber-Physical Production Systems, in *Technological Innovation for Smart Systems: 8th IFIP WG 5.5/SOCOLNET Advanced Doctoral Conference on Computing, Electrical and Industrial Systems, DoCEIS 2017, Costa de Caparica, Portugal, May 3-5, 2017, Proceedings*, Camarinha-Matos, L. M.Parreira-Rocha, M. and Ramezani, J. (Eds) Springer International Publishing, Cham, pp. 135-143.

Bayuk, J. L., Cavit, D., Guerrino, E., Mahony, J., McDowell, B., Nelson, W., Snevel, R. and Staarfanger, P. (2011) Malware Risks and Mitigation Report, BITS Financial Services Roundtable, Washington, DC.

Belikovetsky, S., Yampolskiy, M., Toh, J. and Elovici, Y. (2016) Dr0wned-Cyber-Physical Attack with Additive Manufacturing. *arXiv preprint* **arXiv:1609.00133**.

Belikovetsky, S., Solewicz, Y., Yampolskiy, M., Toh, J. and Elovici, Y. (2017) Detecting Cyber-Physical Attacks in Additive Manufacturing using Digital Audio Signing. *arXiv preprint* **arXiv:1705.06454**.

Belikovetsky, S., Solewicz, Y., Yampolskiy, M., Toh, J. and Elovici, Y. (2018) Digital Audio Signature for 3D Printing Integrity. *IEEE Transactions on Information Forensics and Security*, 1-1.

Benneyan, J. (2009) Misuse of XmR Quality Control Charts for Common Single Parameter Probability Distributions, in *Proceedings of the 2009 International Conference on Computers & Industrial Engineering*, pp. 1291-1296.

Bird, D. and Dale, B. G. (1994) The Misuse and Abuse of SPC: A Case Study Examination. *International Journal of Vehicle Design*, **15**(1-2), 99-107.

Box, G. E. P. and Woodall, W. H. (2012) Innovation, Quality Engineering, and Statistics. *Quality Engineering*, **24**(1), 20-29.

Brenner, J. F. (2013) Eyes Wide Shut: The Growing Threat of Cyber Attacks on Industrial Control Systems. *Bulletin of the Atomic Scientists*, **69**(5), 15-20.

Cai, D. Q., Xie, M. and Goh, T. N. (2001) SPC in an Automated Manufacturing Environment. *International Journal of Computer Integrated Manufacturing*, **14**(2), 206–211.

Cárdenas, A., Amin, S., Sinopoli, B., Giani, A., Perrig, A. and Sastry, S. (2009) Challenges for Securing Cyber Physical Systems, in *Proceedings of the Workshop on Future Directions in Cyber-physical Systems Security*, DHS, pp. 5.

Cárdenas, A., Amin, S., Lin, Z.-S., Huang, Y.-L., Huang, C.-Y. and Sastry, S. (2011) 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)*, pp. 355-366.

Castillo, M. (2017) Yahoo's Hack Warning Comes from Third breach, The Company Says, Available from http://www.cnbc.com/2017/02/15/yahoo-sends-new-warning-to-customers-about-data-breach.html, [Last accessed on 9 April 2017].

Caulcutt, R. (1995) The Rights and Wrongs of Control Charts. *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, **44**(3), 279-288.

Chakraborty, R. S., Narasimhan, S. and Bhunia, S. (2009) Hardware Trojan: Threats and emerging solutions, in *Proceedings of the 2009 IEEE International High Level Design Validation and Test Workshop*, pp. 166-171.

Chalaki, K., Saghaei, A. and Moghadam, M. B. (2016) A comparison study of effectiveness and robustness of robust economic designs of T^2 chart using genetic algorithm. *Communications in Statistics - Theory and Methods*, **45**(11), 3383-3396.

Chambers, J. (2015) What Does the Internet of Everything Mean for Security?, Available from https://www.weforum.org/agenda/2015/01/companies-fighting-cyber-crime/?utm_content=bufferb0881&utm_medium=social&utm_source=facebook.com&utm_campaign=buffer, [Last accessed on 30 October 2015].

Chen, S.-C. and Pan, J.-N. (2011) Determining Optimal Number of Samples for Constructing Multivariate Control Charts. *Communications in Statistics - Simulation and Computation*, **40**(2), 216-228.

Cherry, S. (2011) Sons of Stuxnet, Available from http://spectrum.ieee.org/podcast/telecom/security/sons-of-stuxnet, [Last accessed on 15 December 2014].

Chhetri, S. R., Canedo, A. and Al Faruque, M. A. (2016) KCAD: Kinetic Cyber-Attack Detection Method for Cyber-Physical Additive Manufacturing Systems, in *Proceedings of the International Conference On Computer Aided Design (ICCAD '16)*, IEEE/ACM.

Chhetri, S. R. and Al Faruque, M. A. (2017) Side-Channels of Cyber-Physical Systems: Case Study in Additive Manufacturing. *IEEE Design & Test*, **PP**(99), 1-1.

Chhetri, S. R., Wan, J. and Al Faruque, M. A. (2017) Cross-Domain Security of Cyber-Physical Systems, in *Proceedings of the 22nd Asia and South Pacific Design Automation Conference (ASP-DAC 2017)*, IEEE, pp. 200-205.

Dastoorian, R., Elhabashy, A. E., Tian, W., Wells, L. J. and Camelio, J. A. (2018) Automated Surface Inspection Using 3D Point Cloud Data in Manufacturing: A Case Study, in

Proceedings of the ASME 2018 13th International Manufacturing Science and Engineering Conference (MSEC2018), ASME, pp. V003T02A036.

Deloitte (2014) Global Cyber Executive Briefing - Manufacturing, Available from http://www2.deloitte.com/global/en/pages/risk/articles/Manufacturing.html#, [Last accessed on 16 August 2015].

DeSmit, Z., Elhabashy, A. E., Wells, L. J. and Camelio, J. A. (2016) Cyber-physical Vulnerability Assessment in Manufacturing Systems. *Procedia Manufacturing*, **5**, 1060-1074.

DeSmit, Z., Elhabashy, A. E., Wells, L. J. and Camelio, J. A. (2017) An Approach to Cyber-physical Vulnerability Assessment for Intelligent Manufacturing Systems. *Journal of Manufacturing Systems*, **43, Part 2**, 339-351.

Elhabashy, A. E., Abdelhamid, S. E., Reid, K. J. and Camelio, J. A. (2015a) Factors Affecting Better Use of Laboratory Courses in Engineering, in *the 7th Annual First Year Engineering Experience (FYEE) Conference*.

Elhabashy, A. E., Wells, L. J., Camelio, J. A. and Woodall, W. H. (2015b) Designing Quality Control Tools for Enhanced Cyber-Security in Manufacturing, INFORMS 2015 Annual Meeting, Philadelphia, PA, USA.

Elhabashy, A. E., Wells, L. J., Camelio, J. A. and Woodall, W. H. (2016) A Cyber-physical Vulnerabilities Framework for Manufacturing Systems: A Quality Control Perspective, INFORMS 2016 Annual Meeting, Nashville, TN, USA.

Elhabashy, A. E., Wells, L. J., Camelio, J. A. and Woodall, W. H. (2018a) A Cyber-Physical Attack Taxnomy for Production Systems: A Quality Control Perspective. *Journal of Intelligent Manufacturing*.

Elhabashy, A. E., Wells, L. J., Woodall, W. H. and Camelio, J. A. (2018b) Misuse of Quality Control Tools in Manufacturing, IISE 2018 Annual Conference and Expo, Orlando, FL, USA.

Elhabashy, A. E., Wells, L. J., Woodall, W. H. and Camelio, J. A. (2019a) Cyber-Physical Attack Vulnerabilities in Manufacturing Quality Control Tools. *Quality and Reliability Engineering International*, (To be Submitted).

Elhabashy, A. E., Wells, L. J., Woodall, W. H. and Camelio, J. A. (2019b) Introducing Randomness into Univariate Control Chart Sampling Plans for Better Security against

Cyber-Physical Attacks in Production Systems. *Journal of Quality Technology (JQT)*, (To be Submitted).

Ermer, D. S. and Hurtis, G. M. (1995) Advanced SPC for Higher-Quality Electronic Card Manufacturing. *Quality Engineering*, **8**(2), 283-299.

Evans, D. (2011) The Internet of Things: How the Next Evolution of the Internet is Changing Everything, Cisco Internet Business Solutions Group (IBSG).

Fahey, M. and Wells, N. (2016) Yahoo Data Breach is among the Biggest in History, Available from http://www.cnbc.com/2016/09/22/yahoo-data-breach-is-among-the-biggest-in-history.html, [Last accessed on 9 April 2017].

Finkle, J. (2013) Malicious Virus Shuttered Power Plant: DHS, Available from http://www.reuters.com/article/us-cybersecurity-powerplants-idUSBRE90F1F720130116, [Last accessed on 6 April 2017].

Gendarmerie, N. (2011) Prospective Analysis on Trends in Cybercrime From 2011 to 2020.

Giraldo, J., Sarkar, E., Cárdenas, A., Maniatakos, M. and Kantarcioglu, M. (2017) Security and Privacy in Cyber-physical Systems: A Survey of Surveys. *IEEE Design & Test*, **34**(4), 7-17.

Goonatilake, R., Bachnak, R. and Herath, S. (2011) Statistical Quality Control Approaches to Network Intrusion Detection. *International Journal of Network Security & Its Applications (IJNSA)*, **3**(6), 115-124.

Groover, M. P. (2010) *Fundamentals of Modern Manufacturing: Materials, Processes, and Systems*, 4th, J. Wiley & Sons, Hoboken, NJ.

Hojjati, A., Adhikari, A., Struckmann, K., Chou, E., Nguyen, T. N. T., Madan, K., Winslett, M. S., Gunter, C. A. and King, W. P. (2016) Leave Your Phone at the Door: Side Channels that Reveal Factory Floor Secrets, in *the Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*, ACM, pp. 883-894.

Holmes, D. and Mergen, A. E. (2012) Control Limits versus Acceptance Limits - Which Limits are Appropriate for your Task?, in *Proceedings of the Northeast Decision Sciences Institute 41st Annual Meeting*.

Hoyles, C., Bakker, A., Kent, P. and Noss, R. (2007) Attributing Meanings to Representations of Data: The Case of Statistical Process Control. *Mathematical Thinking and Learning*, **9**(4), 331-360.

Huang, S., Zhou, C.-J., Yang, S.-H. and Qin, Y.-Q. (2015) Cyber-physical system security for networked industrial processes. *International Journal of Automation and Computing*, **12**(6), 567-578.

Huang, Y.-L., Cárdenas, A. A., Amin, S., Lin, Z.-S., Tsai, H.-Y. and Sastry, S. (2009) Understanding the physical and economic consequences of attacks on control systems. *International Journal of Critical Infrastructure Protection*, **2**(3), 73-83.

Hutchins, M. J., Bhinge, R., Micali, M. K., Robinson, S. L., Sutherland, J. W. and Dornfeld, D. (2015) Framework for Identifying Cybersecurity Risks in Manufacturing. *Procedia Manufacturing*, **1**, 47-63.

IBM X-Force Research (2016) 2016 Cyber Security Intelligence Index Infographic for Manufacturing, IBM Corporation.

ICS-CERT (2016) ICS-CERT Monitor Newsletters: November-December 2015, Available from https://ics-cert.us-cert.gov/sites/default/files/Monitors/ICS-CERT_Monitor_Nov-Dec2015_S508C.pdf, [Last accessed on 26 May 2016].

Kammerdiner, A. R. (2014) Statistical Techniques for Assessing Cyberspace Security, in *Dynamics of Information Systems - Computational and Mathematical Challenges*, Vogiatzis, C.L., W. J. and Pardalos, P. M. (Eds) Springer International Publishing, pp. 161-177.

Kaspersky (2015) What is Spear Phishing? - Definition, Available from http://usa.kaspersky.com/internet-security-center/definitions/spear-phishing#.WOuQM_nyvX5, [Last accessed on 22 November 2015].

Keefe, M. J., Loda, J. B., Elhabashy, A. E. and Woodall, W. H. (2017) Improved Implementation of The Risk-Adjusted Bernoulli CUSUM Chart to Monitor Surgical Outcome Quality. *International Journal for Quality in Health Care*, **29**(3), 343-348.

Kessem, L. (2017) WannaCry Ransomware Spreads Across the Globe, Makes Organizations Wanna Cry About Microsoft Vulnerability, Available from https://securityintelligence.com/wannacry-ransomware-spreads-across-the-globe-makes-organizations-wanna-cry-about-microsoft-vulnerability/, [Last accessed on 25 May 2017].

Kravets, D. (2009) Feds: Hacker Disabled Offshore Oil Platforms' Leak-Detection System, Available from www.wired.com/2009/03/feds-hacker-dis, [Last accessed on 15 December 2014].

Kurtzman Carson Consultants (2016) Home Depot Breach Settlement, Available from http://www.homedepotbreachsettlement.com/, [Last accessed on 15 June 2016].

Larson, S. (2017) Massive Cyberattack Targeting 99 Countries Causes Sweeping Havoc, Available from http://money.cnn.com/2017/05/12/technology/ransomware-attack-nsa-microsoft/, [Last accessed on 25 May 2017].

Lee, R. M., Assante, M. J. and Conway, T. (2014) German Steel Mill Cyber Attack, in Industrial Control Systems, SANS Institute.

Lee, T. B. (2014) The Sony Hack: How it Happened, Who is Responsible, and What we've Learned, Available from www.vox.com/2014/12/14/7387945/sony-hack-explained, [Last accessed on 15 December 2014].

Liptak, A. (2017) Renault Shut Down Several French Factories after Cyberattack, Available from https://www.theverge.com/2017/5/14/15637472/renault-nissan-shut-down-french-uk-factories-wannacry-cyberattack, [Last accessed on 15 May 2017].

Lorenzen, T. J. and Vance, L. C. (1986) The Economic Design of Control Charts: A Unified Approach. *Technometrics*, **28**(1), 3-10.

Lowry, C. A. and Montgomery, D. C. (1995) A review of multivariate control charts. *IIE Transactions*, **27**(6), 800-810.

Manadhata, P. K. and Wing, J. M. (2011) An Attack Surface Metric. *IEEE Transactions on Software Engineering*, **37**(3), 371-386.

Mandiant (2015) M-Trends 2015: A View from The Front Line, Mandiant.

Maragah, H. D. and Woodall, W. H. (1992) The effect of autocorrelation on the retrospective X-chart. *Journal of Statistical Computation and Simulation*, **40**(1-2), 29-42.

McCoy, T., Yearout, R. and Patch, S. (2004) Misrepresenting Quality Data through Incorrect Statistical Applications – A Statistical Quality Control (SQC) Case Study. *International Journal of Industrial Engineering: Theory, Applications and Practice*, **11**(1), 66-73.

McGee, B. (2016) Move over Healthcare, Ransomware Has Manufacturing in Its Sights, Available from https://blog.fortinet.com/2016/06/06/move-over-healthcare-ransomware-has-manufacturing-in-its-sights, [Last accessed on 3 August 2016].

Megahed, F. M. and Jones-Farmer, L. A. (2015) Statistical Perspectives on “Big Data”, in *Frontiers in Statistical Quality Control 11*, Knoth, S. and Schmid, W. (Eds) Springer, pp. 29-47.

Meserve, J. (2007) Mouse Click Could Plunge City into Darkness, Experts Say, Available from http://www.cnn.com/2007/US/09/27/power.at.risk/index.html, [Last accessed on 22 February 2016].

Montgomery, D. C. and Klatt, P. J. (1972) Economic Design of T^2 Control Charts to Maintain Current Control of a Process. *Manage. Sci.*, **19**(1), 76-89.

Montgomery, D. C. (2009) *Introduction to Statistical Quality Control*, 6th, Wiley, Hoboken, N.J.

Moore, S. B., Gatlin, J., Belikovetsky, S., Yampolskiy, M., King, W. E. and Elovici, Y. (2017a) Power Consumption-based Detection of Sabotage Attacks in Additive Manufacturing. *arXiv preprint arXiv:1709.01822*.

Moore, S. B., Glisson, W. B. and Yampolskiy, M. (2017b) Implications of Malicious 3D Printer Firmware, in *Proceedings of the Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, pp. 6089-6098.

MTConnect Institute (2016), Available from www.mtconnect.org, [Last accessed on 11 April 2016].

NDIA (2014) Cybersecurity for Advanced Manufacturing.

Neil, J., Hash, C., Brugh, A., Fisk, M. and Storlie, C. B. (2013) Scan Statistics for the Online Detection of Locally Anomalous Subgraphs. *Technometrics*, **55**(4), 403-414.

Pan, Y., White, J., Schmidt, D. C., Elhabashy, A., Sturm, L., Camelio, J. and Williams, C. (2017) Taxonomies for Reasoning About Cyber-physical Attacks in IoT-based Manufacturing Systems. *International Journal of Interactive Multimedia and Artificial Intelligence, Special Issue on Advances and Applications in the Internet of Things and Cloud Computing*, **4**(3), 45-54.

Park, Y., Baek, S. H., Kim, S.-H. and Tsui, K.-L. (2014) Statistical Process Control-Based Intrusion Detection and Monitoring. *Quality and Reliability Engineering International*, **30**(2), 257-273.

Peres, F. A. P. and Fogliatto, F. S. (2018) Variable selection methods in multivariate statistical process control: A systematic literature review. *Computers & Industrial Engineering*, **115**, 603-619.

Pham, T. (2015) Authentication-Based Attacks Target Energy & Critical Manufacturing Industries, Available from https://duo.com/blog/authentication-based-attacks-target-energy-and-critical-manufacturing, [Last accessed on 17 August 2015].

Poulsen, K. (2009) Ex-Employee Fingered in Texas Power Company Hack, Available from www.wired.com/2009/05/efh/, [Last accessed on 15 December 2014].

Privault, N. (2018) Discrete-Time Markov Chains, in *Understanding Markov Chains: Examples and Applications*(Eds) Springer Singapore, Singapore, pp. 89-113.

Reynolds Jr., M. R. and Stoumbos, Z. G. (2004a) Should Observations be Grouped for Effective Process Monitoring? *Journal of Quality Technology*, **36**(4), 343-366.

Reynolds Jr., M. R. and Stoumbos, Z. G. (2004b) Control Charts and the Efficient Allocation of Sampling Resources. *Technometrics*, **46**(2), 200-214.

Reynolds Jr., M. R. and Stoumbos, Z. G. (2005) Should Exponentially Weighted Moving Average and Cumulative Sum Charts Be Used With Shewhart Limits? *Technometrics*, **47**(4), 409-424.

Reynolds Jr., M. R. and Lou, J. (2010) An Evaluation of a GLR Control Chart for Monitoring the Process Mean. *Journal of Quality Technology*, **42**(3), 287-310.

Rost, J. and Glass, R. L. (2011) *The Dark Side of Software Engineering: Evil on Computing Projects*, Wiley-IEEE Computer Society Press.

Saleh, N. A., Mahmoud, M. A., Keefe, M. J. and Woodall, W. H. (2015) The Difficulty in Designing Shewhart $\bar{X}$ and X Control Charts with Estimated Parameters. *Journal of Quality Technology*, **47**(2), 127-138.

Shafae, M. S., Wells, L. J. and Purdy, G. T. (2018) Defending against Product-Oriented Cyber-Physical Attacks on Intelligent Machining Systems. *IEEE Transactions on Automation Science and Engineering*, (Submitted).

Slay, J. and Miller, M. (2008) Lessons Learned from the Maroochy Water Breach, in *Critical Infrastructure Protection*, Goetz, E. and Shenoi, S. (Eds) Springer US, Boston, MA, pp. 73-82.

Song, C., Lin, F., Ba, Z., Ren, K., Zhou, C. and Xu, W. (2016) My Smartphone Knows What You Print: Exploring Smartphone-based Side-channel Attacks Against 3D Printers, in *the Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*, ACM, pp. 895-907.

Sparks, R. S. and Field, J. B. F. (2000) Using Deming's Funnel Experiment to Demonstrate Effects of Violating Assumptions Underlying Shewhart's Control Charts. *The American Statistician*, **54**(4), 291-302.

Stamp, J., Dillinger, J., Young, W., and DePoy, J. (2003) Common Vulnerabilities in Critical Infrastructure Control Systems, Sandia National Laboratories.

Steitz, C., and Auchard, E. (2016) German Nuclear Plant Infected with Computer Viruses, Operator Says, Available from http://www.reuters.com/article/us-nuclearpower-cyber-germany-idUSKCN0XN2OS, [Last accessed on 15 June 2016].

Stouffer, K., Zimmerman, T., Tang, C., Lubell, J., Cichonski, J. and McCarthy, J. (2017) Cybersecurity Framework Manufacturing Profile, National Institute of Standards and Technology (NIST).

Sturm, L. D., Williams, C. B., Camelio, J. A., White, J. and Parker, R. (2014) Cyber-physical Vulnerabilities in Additive Manufacturing Systems, in *Proceedings of the 25th Annual Solid Freeform Fabrication Symposium*.

Sturm, L. D., Albakri, M., Williams, C. B. and Tarazaga, P. (2016) In-situ Detection of Build Defects in Additive Manufacturing via Impedance-Based Monitoring, in *the Proceedings of the 27th Annual International Solid Freeform Fabrication Symposium – An Additive Manufacturing Conference*, pp. 1458-1478.

Sturm, L. D., Williams, C. B., Camelio, J. A., White, J. and Parker, R. (2017) Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the .STL file with human subjects. *Journal of Manufacturing Systems*, **44, Part 1**, 154-164.

Symantec (2014) Internet Security Threat Report 2014, Volume 19, Symantec Corporation.

Symantec (2015) Internet Security Threat Report 2015, Volume 20, Symantec Corporation.

Symantec (2016) Internet Security Threat Report 2016, Volume 21, Symantec Corporation.

Symantec (2017) Internet Security Threat Report 2017, Volume 22, Symantec Corporation.

Target (2014) Data Breach FAQ, Available from https://corporate.target.com/about/shopping-experience/payment-card-issue-faq, [Last accessed on 28 October 2015].

Teemu M. (2015) 3 Key Learnings: Ransomware Hits A Concrete Manufacturer, Available from https://business.f-secure.com/3-key-learnings-ransomware-hits-a-concrete-manufacturer/, [Last accessed on 3 August 2016].

Tseitlin, A. (2013) The Antifragile Organization. *Communications of the ACM*, **56**(8), 40-44.

Tucker, P. (2014) Forget the Sony Hack, This Could Be the Biggest Cyber Attack of 2015, Available from www.defenseone.com/technology/2014/12/forget-sony-hack-could-be-he-biggest-cyber-attack-2015/101727/, [Last accessed on 22 February 2016].

Tuptuk, N. and Hailes, S. (2016) The Cyberattack on Ukraine's Power Grid is a Warning of What's to Come, Available from https://theconversation.com/the-cyberattack-on-ukraines-power-grid-is-a-warning-of-whats-to-come-52832, [Last accessed on 15 June 2016].

Turner, H., White, J., Camelio, J. A., Williams, C., Amos, B. and Parker, R. (2015) Bad Parts: Are Our Manufacturing Systems at Risk of Silent Cyberattacks? *IEEE Security & Privacy*, **13**(3), 40-47.

Underbrink, A., Potter, A., Jaenisch, H. and Reifer, D. J. (2012) Application Stress Testing Achieving Cyber Security by Testing Cyber Attacks, in *Proceedings of the IEEE Conference on Technologies for Homeland Security (HST)*, pp. 556-561.

Vincent, H., Wells, L., Tarazaga, P. and Camelio, J. (2015) Trojan Detection and Side-channel Analyses for Cyber-security in Cyber-physical Manufacturing Systems. *Procedia Manufacturing*, **1**, 77-85.

Wang, K. and Tsung, F. (2008) An adaptive dimension reduction scheme for monitoring feedback-controlled processes. *Quality and Reliability Engineering International*, **25**(3), 283-298.

Wang, K. and Jiang, W. (2009) High-Dimensional Process Monitoring and Fault Isolation via Variable Selection. *Journal of Quality Technology*, **41**(3), 247-258.

Warner, D. J., Lindsay, C. E. and Sansom, C. L. (1993) The Use and Misuse of Statistical Process Control in GaAs MMIC Manufacture, in *Proceedings of the Gallium Arsenide Integrated Circuit (GaAs IC) Symposium, 15th Annual*, IEEE, pp. 131 - 133.

Wegner, A., Graham, J. and Ribble, E. (2017) A New Approach to Cyberphysical Security in Industry 4.0, in *Cybersecurity for Industry 4.0: Analysis for Design and Manufacturing*, Thames, L. and Schaefer, D. (Eds) Springer International Publishing, Cham, pp. 59-72.

Wells, L. J., Camelio, J. A., Williams, C. B. and White, J. (2014) Cyber-physical Security Challenges in Manufacturing Systems. *Manufacturing Letters*, **2**(2), 74-77.

Wells, L. J., Elhabashy, A. E., Woodall, W. H. and Camelio, J. A. (2017) Introducing Randomness into Control Chart Sampling Strategies to Better Protect against Cyber-Physical Attacks on Production Systems, INFORMS 2017 Annual Meeting, Houston, TX, USA.

Whitman, M. E. and Mattord, H. J. (2012) Introduction to Information Security, in *Principles of Information Security*(Eds) Cengage Learning, pp. 1-37.

Wood, M. and Preece, D. (1992) Using Quality Measurements: Practice, Problems and Possibilities. *International Journal of Quality & Reliability Management*, **9**(7), 42-53.

Wood, M. (1995) Three Suggestions for Improving Control Charting Procedures. *International Journal of Quality & Reliability Management*, **12**(5), 61-74.

Woodall, W. H. (1986) Weaknesses of the Economic Design of Control Charts. *Technometrics*, **28**(4), 408-410.

Wu, M. and Moon, Y. B. (2017) Taxonomy of Cross-Domain Attacks on CyberManufacturing System. *Procedia Computer Science*, **114**, 367-374.

Wu, M., Song, Z. and Moon, Y. B. (2017) Detecting Cyber-physical Attacks in CyberManufacturing Systems with Machine Learning Methods. *Journal of Intelligent Manufacturing*, 1-13.

Wu, Q., Zhang, H. and Pu, J. (2007) Mitigating Distributed Denial-of-Service Attacks using Network Connection Control Charts, in *Proceedings of the Proceedings of the 2nd International Conference on Scalable Information Systems (InfoScale '07)*, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), pp. 1-4.

Yampolskiy, M., Horvath, P., Koutsoukos, X. D., Xue, Y. and Sztipanovits, J. (2013) Taxonomy for Description of Cross-domain Attacks on CPS, in *the Proceedings of the 2nd ACM International Conference on High Confidence Networked Systems (HiCoNS '13)*, ACM, pp. 135-142.

Yampolskiy, M., Skjellum, A., Kretzschmar, M., Overfelt, R. A., Sloan, K. R. and Yasinsac, A. (2016) Using 3D Printers as Weapons. *International Journal of Critical Infrastructure Protection*, **14**, 58-71.

Yampolskiy, M., King, W., Pope, G., Belikovetsky, S. and Elovici, Y. (2017) Evaluation of Additive and Subractive Manufacturing from the Security Perpsective, in *Critical Infrastructure Protection XI: 11th IFIP WG 11.10 International Conference, ICCIP 2017, Arlington, VA, USA, March 13-15, 2017, Revised Selected Papers*, Rice, M. and Shenoi, S. (Eds) Springer International Publishing, Cham, pp. 23-44.

Ye, N., Emran, S. M., Li, X. and Chen, Q. (2001) Statistical Process Control for Computer Intrusion Detection, in *Proceedings of the DARPA Information Survivability Conference & Exposition II, 2001 (DISCEX'01).* IEEE, pp. 3-14.

Ye, N., Borror, C. and Zhang, Y. (2002) EWMA Techniques for Computer Intrusion Detection through Anomalous Changes in Event Intensity. *Quality and Reliability Engineering International*, **18**(6), 443-451.

Ye, N., Vilbert, S. and Qiang, C. (2003) Computer Intrusion Detection through EWMA for Autocorrelated and Uncorrelated Data. *IEEE Transactions on Reliability*, **52**(1), 75-82.

Yu, Z., Ouyang, J., Li, S. and Peng, X. (2017) Formal Modeling and Control of Cyber-physical Manufacturing Systems. *Advances in Mechanical Engineering*, **9**(10), 1-12.

Zeltmann, S. E., Gupta, N., Tsoutsos, N. G., Maniatakos, M., Rajendran, J. and Karri, R. (2016) Manufacturing and Security Challenges in 3D Printing. *The Journal of The Minerals, Metals & Materials Society (JOM)*, **68**(7), 1872-1881.

Appendix A *Journal of Intelligent Manufacturing* Copyright Permission

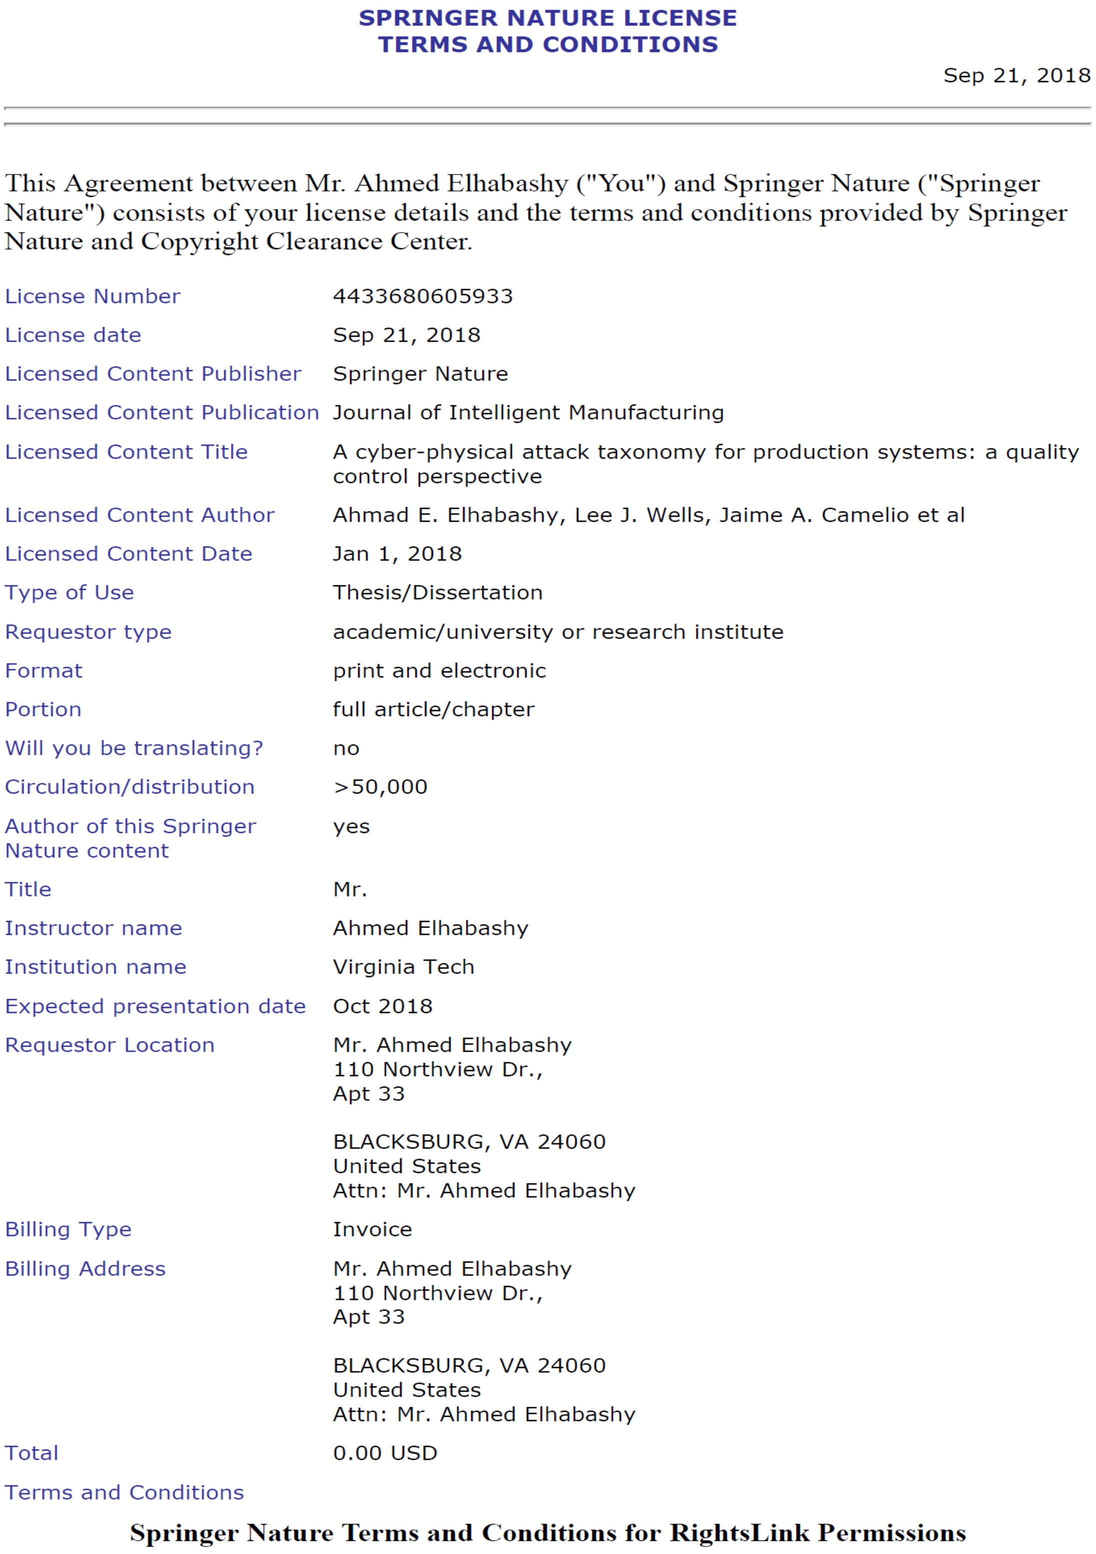

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

Sep 21, 2018

This Agreement between Mr. Ahmed Elhabashy ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	4433680605933
License date	Sep 21, 2018
Licensed Content Publisher	Springer Nature
Licensed Content Publication	Journal of Intelligent Manufacturing
Licensed Content Title	A cyber-physical attack taxonomy for production systems: a quality control perspective
Licensed Content Author	Ahmad E. Elhabashy, Lee J. Wells, Jaime A. Camelio et al
Licensed Content Date	Jan 1, 2018
Type of Use	Thesis/Dissertation
Requestor type	academic/university or research institute
Format	print and electronic
Portion	full article/chapter
Will you be translating?	no
Circulation/distribution	>50,000
Author of this Springer Nature content	yes
Title	Mr.
Instructor name	Ahmed Elhabashy
Institution name	Virginia Tech
Expected presentation date	Oct 2018
Requestor Location	Mr. Ahmed Elhabashy 110 Northview Dr., Apt 33 BLACKSBURG, VA 24060 United States Attn: Mr. Ahmed Elhabashy
Billing Type	Invoice
Billing Address	Mr. Ahmed Elhabashy 110 Northview Dr., Apt 33 BLACKSBURG, VA 24060 United States Attn: Mr. Ahmed Elhabashy
Total	0.00 USD
Terms and Conditions	

Springer Nature Terms and Conditions for RightsLink Permissions

Springer Nature Customer Service Centre GmbH (the Licensor) hereby grants you a non-exclusive, world-wide licence to reproduce the material and for the purpose and requirements specified in the attached copy of your order form, and for no other use, subject to the conditions below:

1. The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of this material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).

   If the credit line on any part of the material you have requested indicates that it was reprinted or adapted with permission from another source, then you should also seek permission from that source to reuse the material.

2. Where **print only** permission has been granted for a fee, separate permission must be obtained for any additional electronic re-use.

3. Permission granted **free of charge** for material in print is also usually granted for any electronic version of that work, provided that the material is incidental to your work as a whole and that the electronic version is essentially equivalent to, or substitutes for, the print version.

4. A licence for 'post on a website' is valid for 12 months from the licence date. This licence does not cover use of full text articles on websites.

5. Where **'reuse in a dissertation/thesis'** has been selected the following terms apply: Print rights of the final author's accepted manuscript (for clarity, NOT the published version) for up to 100 copies, electronic rights for use only on a personal website or institutional repository as defined by the Sherpa guideline (www.sherpa.ac.uk/romeo/).

6. Permission granted for books and journals is granted for the lifetime of the first edition and does not apply to second and subsequent editions (except where the first edition permission was granted free of charge or for signatories to the STM Permissions Guidelines http://www.stm-assoc.org/copyright-legal-affairs/permissions/permissions-guidelines/), and does not apply for editions in other languages unless additional translation rights have been granted separately in the licence.

7. Rights for additional components such as custom editions and derivatives require additional permission and may be subject to an additional fee. Please apply to Journalpermissions@springernature.com/bookpermissions@springernature.com for these rights.

8. The Licensor's permission must be acknowledged next to the licensed material in print. In electronic form, this acknowledgement must be visible at the same time as the figures/tables/illustrations or abstract, and must be hyperlinked to the journal/book's homepage. Our required acknowledgement format is in the Appendix below.

9. Use of the material for incidental promotional use, minor editing privileges (this does not include cropping, adapting, omitting material or any other changes that affect the meaning, intention or moral rights of the author) and copies for the disabled are permitted under this licence.

10. Minor adaptations of single figures (changes of format, colour and style) do not require the Licensor's approval. However, the adaptation should be credited as shown in Appendix below.

<u>Appendix — Acknowledgements:</u>

For Journal Content:
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g.

Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

For **Advance Online Publication papers:**
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)

For Adaptations/Translations:
Adapted/Translated by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

Note: For any republication from the British Journal of Cancer, the following credit line style applies:

Reprinted/adapted/translated by permission from [**the Licensor**]: on behalf of Cancer Research UK: : [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT]** (year of publication)

For **Advance Online Publication** papers:
Reprinted by permission from The [**the Licensor**]: on behalf of Cancer Research UK: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj. [JOURNAL ACRONYM])

For Book content:
Reprinted/adapted by permission from [**the Licensor**]: [**Book Publisher** (e.g. Palgrave Macmillan, Springer etc) [**Book Title**] by [**Book author**(s)] [**COPYRIGHT**] (year of publication)

Other Conditions:

Version 1.1

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

Appendix B Theoretical Evaluation of Some of the Performance Metrics – Individual Control Charts

The purpose of this appendix is to provide more information regarding the different performance metrics discussed in Chapter 5 and how they can be estimated in general. This appendix also contains the technical details and derivations behind evaluating the theoretical values for some of these performance metrics during the Single Fixed Attack scenario, where only one product will be attacked (at *FA* = 5) within each group of *N* products.

Average Run Length (ARL)

The Average Run Length (ARL) is perhaps the most commonly used performance metric for control charts, which would represent in this work the average number of inspected products until a signal occurs. Typically, the ARL values can be obtained directly using the probabilities of getting an out-of-control signal (p) within a group of *N* products. In this appendix, however, it will be shown how to use p to obtain the ARL values only for plans 0, 1 and 2. For plan 0b, using p to evaluate the ARL values is not a straightforward task and Discrete Time Markov Chains (DTMCs) will be used instead. Each of these approaches to obtain the ARL values will be discussed separately next.

Using the Probabilities of Getting an Out-of-Control Signal within a Group of *N* products

Plans 0 and 1

For the individuals control charts (and even x-bar charts) the Run Length (RL) behaves as a geometric random variable when the in-control parameters values are known, with the ARL value equal to its mean (Montgomery, 2009):

$$\text{ARL} = {}^{1}\!/_{p} \tag{B-1}$$

where p is the probability of getting an out-of-control signal within a group of *N* products. When there is no shift, the in-control probability (p_0) equals 0.0027 and the ARL is equal to approximately 370.4. This value also corresponds to both the in-control and out-of-control ARL values for Plan 0, since the shift will never be detected in any of the different attack scenarios.

Getting an out-of-control signal within a group of *N* products could be due to either an actual process shift or a random error (false alarm). Traditionally, the probability of getting such a signal can be written as: p = P(Monitored Statistic > UCL) + P(Monitored Statistic < LCL), where UCL and LCL are a chart's Upper and Lower Control Limits, respectively. We have

$$p = 1 - \Phi(L - \delta) + \Phi(-L - \delta), \tag{B-2}$$

where Φ is the standard normal distribution Cumulative Density Function (CDF), L is typically set to 3 (to achieve a 3-sigma control limit when sigma equals 1), and δ is the shift size. In other words, p is a function of both L and the shift size (δ). If no shift exists, then this probability is the in-control probability p_0, which equals 0.0027 and is denoted by α (Montgomery, 2009). It should be just noted that the probability p is conditional on the considered product being inspected within a group of N products.

However, when transient shifts exist, such as the ones due to Passive Product-Oriented C2P attacks on specific products, equation (B-2) is no longer an accurate representation for the probability of getting an out-of-control signal within a group of N products. Now, more factors need to be considered to obtain the out-of-control signal, including: a) the probability that the signal occurs when an attacked product is inspected (due to an actual transient process shift) and b) the probability that the signal occurs when a non-attacked product is inspected (due to a random error). In other words, the probability of getting an out-of-control signal, within a group of N products, during inspection is

$$\mathrm{p} = \mathrm{P(Getting\ a\ signal\ due\ to\ the\ Attack)} + \mathrm{P(Getting\ a\ signal\ due\ to\ Random\ Error)}$$

$$= \mathrm{P(Attacked\ product\ inspected\ \&\ signals)} + \mathrm{P(Nonattacked\ products\ inspected\ \&\ signal)}$$

$$= \mathrm{P}(\mathrm{S_{Att}} \cap \mathrm{I}) + \mathrm{P}(\mathrm{S_{RE}} \cap \mathrm{I}), \text{where I denots Inspection.}$$

$$\therefore \mathrm{p} = \mathrm{P}(\mathrm{S_{Att}} \mid \mathrm{I}) \times \mathrm{P_I} + \mathrm{P}(\mathrm{S_{RE}} \mid \mathrm{I}) \times \mathrm{P_I},$$

where $\mathrm{P_I}$ is the probability of a part being inspected within the group of N products and $\mathrm{P}(\mathrm{S_{Att}} \mid \mathrm{I})$ is the probability of getting a signal from an attacked product, if it gets inspected.

To get these two components of p, we have:

$$\mathrm{P}(\mathrm{S_{Att}} \mid \mathrm{I}) = [1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})], \text{and}$$

$$\mathrm{P}(\mathrm{S_{RE}} \mid \mathrm{I}) = [1 - \Phi(L - \delta_0) + \Phi(-L - \delta_0)] = \alpha$$

where L is set to 3 and δ_{FA} is the shift size caused by the attack. For Plan 1, there is only one attacked product and *N-1* non-attacked ones; with $\mathrm{P_I}$ equal to *1/N* (regardless if this inspected product was attacked or not). So, the probabilities p would be

$$\mathrm{p} = \mathrm{P_I} \times [1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})] + \mathrm{P_I} \times \sum_{k=1}^{N-1} \alpha_k$$

$$\text{i.e.}, \mathrm{p} = \mathrm{P_I} \times [1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})] + \mathrm{P_I} \times (\mathrm{N} - 1)\,\alpha$$

$$\therefore\ \mathrm{p} = {}^{1}\!/_{\mathrm{N}} \times [1 - \Phi(\mathrm{L} - \delta_{\mathrm{FA}}) + \Phi(-\mathrm{L} - \delta_{\mathrm{FA}})] + (\mathrm{N} - 1) \times {}^{\alpha}\!/_{\mathrm{N}} \qquad \text{(B-3)}$$

The second column in Table B-1 shows the resulting *p* values from equation (B-3) corresponding to different values of transient shift sizes caused by malicious Passive Product-Oriented C2P attacks for this sampling plan and attack scenario combination. If the possibility of the attack is not considered, the *p* values for all shift sizes would always be equal to about 0.0027, since the attack portion in equation (B-3) would not be accounted for and all the *N* products would be treated as non-attacked ones. This attack portion in equation (B-3) can actually be considered as the true alarm probability; i.e., the probability of getting a signal due to an attack. Such true alarm probabilities would depend on both the inspection probability and the shift size.

Table B-1: Probabilities of getting an out-of-control signal and true alarms, along with the ARL values, corresponding to different transient shift sizes caused by a cyber-physical attack for the fixed attack scenario – Plan 1.

Shift Size (δ)	*p*	True Alarm Probability	ARL
0	0.002699796	0.00026998	370.398
0.1	0.002713158	0.000283342	368.574
0.5	0.003074046	0.000644229	325.304
0.75	0.003661105	0.001231289	273.142
1	0.004707997	0.00227818	212.405
2.5	0.033283572	0.030853756	30.045
5	0.100154803	0.097724987	9.985
10	0.102429816	0.1	9.763

The third column of Table B-1 shows the corresponding values of these true alarm probabilities for the same attack scenario and sampling plan combination. While the random error portion in equation (B-3) can be considered as the false alarm probability. This false alarm probability is the probability of getting a signal *not* due to the attack and only depends on P_I; i.e., it would be equal to 0.00243 in this plan. Finally, the 4th column of the table shows the corresponding ARL values obtained in this plan, using equation (B-3).

Plan 2

This plan is a little different than Plan 1, although the value for the probability of a part being inspected within the group of *N* products P_I remains unchanged. The probabilities of getting a signal follow a Binomial distribution in this plan, unlike Plan 1. Another key difference between Plan 2 and Plan 1 is, in Plan 2 no products could be chosen for inspection at all or more than one product could be inspected first before getting an out-of-control signal, *as long as* these other chosen products don't cause any signals themselves.

As an example, consider a signal obtained when inspecting the attacked product; this means that no matter how many (if not all) of the previous four products were inspected, they all didn't cause a signal *and* the out-of-control signal was only caused by the attacked product inspection. So, we have

$$\mathrm{P(S_{Att} \cap\ I) = P(Getting\ a\ signal\ due\ to\ the\ Attack)}$$

$$= \mathrm{P}(\text{Attacked product inspected \& signals } \textit{and} \text{ None of prevoiusly inspected products signal})$$

$$= P(S_{Att}\mid I) \times P_I = P_I \times [1 - \Phi(L-\delta_{FA}) + \Phi(-L-\delta_{FA})] \times \prod_{k=1}^{FA-1}(1-P_I\alpha_k).$$

$$\therefore P(S_{Att} \cap\ I) = P_I \times [1 - \Phi(L-\delta_{FA}) + \Phi(-L-\delta_{FA})] \times (1-P_I\alpha)^{FA-1}$$

Similarly, for the non-attacked products, all the previously inspected products need to be considered.

$$\mathrm{P(S_{RE} \cap\ I) = P(Getting\ a\ signal\ due\ to\ Random\ Error)}$$

$$= \mathrm{P}(\text{A nonattacked product inspected \& signals } \textit{and} \text{ None of prevoiusly inspected products signal})$$

For easier evaluation, the non-attacked products will be grouped into 2 sets:

1) The first non-attacked products *before* the attacked one, whose indices within the group of *N* products are from 1 to *FA – 1*. The probabilities for these products are somewhat similar to the attacked one, the difference is just in not having a shift:

$$\mathrm{P\left(S_{RE_{Before}} \cap\ I\right) = \ P_I \times P\left(S_{RE_{Before}} \mid I\right)}$$

$$= P_I \times \sum_{k=1}^{FA-1}\left[\prod_{l=1}^{k-1}(1-P_I\alpha_l) \times [1 - \Phi(L-\delta_0) + \Phi(-L-\delta_0)]\right]$$

$$= P_I\alpha \times \sum_{k=1}^{FA-1}\left[\prod_{l=1}^{k-1}(1-P_I\alpha_l)\right] = P_I\alpha \times \sum_{k=1}^{FA-1}(1-P_I\alpha)^{k-1}$$

2) The non-attacked products *after* the attacked one, whose indices within the group of *N* products are from *FA + 1* to *N*. The probabilities for these products are a combination of the previous ones and that for the attacked product.

$$\mathrm{P}\left(\mathrm{S}_{\mathrm{RE}_{\mathrm{After}}} \cap \mathrm{I}\right) = \mathrm{P}_{\mathrm{I}} \times \mathrm{P}\left(\mathrm{S}_{\mathrm{RE}_{\mathrm{After}}} \mid \mathrm{I}\right)$$

$$\begin{aligned}
&= P_I \times \sum_{k=FA+1}^{N} \left[[1 - \Phi(L - \delta_0) + \Phi(-L - \delta_0)] \prod_{l=1}^{FA-1} (1 - P_I \alpha_l) \right. \\
&\qquad \left. \times \left[1 - P_I\left(1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})\right)\right] \times \left[1 + \prod_{l=1}^{N-k} (1 - P_I \alpha_l)\right] \right] \\
&= P_I \alpha \times (1 - P_I \alpha)^{FA-1} \times \left[1 - P_I\left(1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})\right)\right] \\
&\qquad \times \sum_{k\,=\,0}^{N-(FA+1)} [\,(1 - P_I \alpha)^k\,]
\end{aligned}$$

$$\begin{aligned}
\therefore\ p = {} & P_I \alpha \times \sum_{k=1}^{FA-1} (1 - P_I \alpha)^{k-1} + P_I \times [1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})] \times (1 - P_I \alpha)^{FA-1} \\
& + P_I \alpha \times (1 - P_I \alpha)^{FA-1} \times \left[1 - P_I\left(1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})\right)\right] \\
& \times \sum_{k\,=\,0}^{N-(FA+1)} [\,(1 - P_I \alpha)^k\,]
\end{aligned}$$

For plan 2, the value of the probability of getting an out-of-control signal is shown in equation (B-4), where P_I has the same value as in Plan 1:

$$\begin{gathered}
p = P_I \times \left[\alpha \textstyle\sum_{k=1}^{FA-1} (1 - P_I \alpha)^{k-1} + [1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})] \times \right. \\
(1 - P_I \alpha)^{FA-1} + \alpha (1 - P_I \alpha)^{FA-1} \times \left[1 - P_I\left(1 - \Phi(L - \delta_{FA}) + \Phi(-L - \delta_{FA})\right)\right] \times \\
\left. \textstyle\sum_{k\,=\,0}^{N-(FA+1)} [\,(1 - P_I \alpha)^k\,]\,\right]
\end{gathered} \tag{B-4}$$

Obtaining the ARL values for this plan is not as straightforward though, since using equation (B-1) will not provide accurate results. Equation (B-1) is only correct for sampling Plans 0 and 1, since the expected number of inspected products in each group of *N* products is always equal to one, unlike in the other 2 plans. In Plan 2, the inverse of the out-of-control probability doesn't translate directly to the ARL. This inverse represents the expected numbers of groups until a signal occurs, given inspection, and will be referred to as E[NGS], where

$$\text{Expected Number of Groups to Siganl: E[NGS]} = {}^{1}\!/_{p} \tag{B-5}$$

The ARL can also be thought of as the expected number of complete groups of *N* products to a signal (all the previous groups before the one being currently considered) in addition to the expected number of inspections to signal within the current group, given that the current group has a signal:

$$\mathrm{ARL} = \mathrm{E[NGS]} - 1 + \mathrm{E[NIS|\, S_G]} \tag{B-6}$$

where $\mathrm{E[NIS|\, S_G]}$ is that expected number of inspections to signal within a group, given a group signal. It should be noted that equation (B-1) is just a special case of equation (B-6) where the expected number of inspections to signal within a group is always going to be equal to one, such as in Plans 0 and 1. For Plan 2, $\mathrm{E[NIS|\, S_G]}$ would be equal to the expected number of parts to signal in this group, $\mathrm{E[NPS|\, S_G]}$:

$$\mathrm{E[NPS|\, S_G]} = \sum_{i=1}^{N} \frac{(i \times P[\mathrm{S_i}|\, \mathrm{S_G}])}{N} \tag{B-7}$$

where $\mathrm{P[S_i|\, S_G]}$ is the probability of getting an individual signal, given a group signal and it equals:

$$\mathrm{P[S_i|\, S_G]} = \frac{\mathrm{P[S_i \cap S_G]}}{\mathrm{P[S_G|I]}} = \frac{\mathrm{P[S_i|I]}}{\mathrm{P[S_G|I]}} = \frac{\mathrm{P[S_i|I]}}{\sum_{i=1}^{N} \mathrm{P[S_i|I]}} \tag{B-8}$$

where $\mathrm{P[S_i|I]}$ is the probability of getting an individual signal from a product, given inspection and $\mathrm{P[S_G|I]}$ is the probability of getting a signal within a group of *N* products, given inspection of this group; which just equals *p*.

Using Discrete Time Markov Chains (DTMCs)

Plan 0b

In Plan 0b, the indices of the chosen product for inspection within each group of *N* products is obtained from equation (5-1), where they are dependent on previously selected ones. So, it is important to first understand how exactly product inspection occurs and why it would be difficult to derive an appropriate formula to calculate *p* in this case. As can be seen from equation (5-1), the selection of a product for inspection depends on the indices of the previously chosen ones, whether it is in the previous group of *N* products or the current one. Table B-2 shows this dependency and the range of possible parts to be inspected next for each specific product. For example, if part 1 is currently inspected within one group of *N* products, the next part to be inspected would range from part 9 ($1 + \ N - 0.2N$) in this same group to the third part in the next group of *N* products ($1 + \ N + 0.2N$), when each group has *N* = 10 parts.

Table B-2: Possibilities for the next part to be inspected after a specific part has been currently chosen for inspection within sampling Plan 0b. An “I” indicates that the part could be inspected within groups of size 10.

	Part 1	**Part 2**	**Part 3**	**Part 4**	**Part 5**	**Part 6**	**Part 7**	**Part 8**	**Part 9**	**Part 10**
Part 1	I	I	I						I	I
Part 2	I	I	I	I						I
Part 3	I	I	I	I	I					
Part 4		I	I	I	I	I				
Part 5			I	I	I	I	I			
Part 6				I	I	I	I	I		
Part 7					I	I	I	I	I	
Part 8						I	I	I	I	I
Part 9	I						I	I	I	I
Part 10	I	I						I	I	I

Accordingly, the probability of getting an out-of-control signal for a specific product is a function of not getting a signal in all the previously inspected products. For instance, to get a signal when inspecting the attacked product, all the previously inspected products should not signal *and* the out-of-control signal would be caused by the inspected attacked product only.

However, Table B-2 only shows the system during a steady-state condition that it would reach after a while from starting inspection. Initially, products 8, 9, 10, 11, or 12 only can be chosen; which means that during the very 1st group of *N* products only the last three products could be selected. In the second group of *N* products, the 1st two products could be chosen for inspection and/or products 6-10, depending on which product was chosen in the 1st group of *N* products and did not cause a signal. This initial behavior continues until the system reaches that steady state, given no signal occurs during inspection.

So, to evaluate the probability of getting an out-of-control signal, this initial behavior has to be considered first, until the steady-state condition is reached and then its behavior would be considered too. Unfortunately, such a task is not straightforward and deriving a formula for evaluating the probability of getting an out-of-control signal within a group of *N* products would not be carried out here. Instead, the ARL value would be obtained using Discrete Time Markov Chains (DTMCs).

DTMCs are stochastic processes that take values in discrete-time state space and have the Markov property, such that the probability distribution of a given state depends on that of the previous state only (Privault, 2018). In order to use DTMCs to evaluate the ARL values, all the possible states occurring need to be defined first. For all the attack scenarios and sampling plan combinations, the DTMCs used would have a total of (N + 2) states. These states are described in Table B-3 and are summarized as follows:

- State (0): The initial state at which sampling first starts. This is a transient state that is never returned to again.
- State (N + 1): The state representing obtaining an out-of-control signal during inspection, regardless of the source of the signal; this is an absorbing state.
- Finally, for each product there is one potentially recurring transient state. For a given product, this state represents that this product has been inspected, but did not cause a signal.

Table B-3: Description of the states within the used DTMCs.

State #	Description
0	Initial state at the beginning of sampling
1	Product *1* did not signal when inspected
2	Product *2* did not signal when inspected
.	.
.	.
.	.
N - 1	Product *N-1* did not signal when inspected
N	Product *N* did not signal when inspected
N + 1	Getting an out-of-control signal

So, the transitional probability matrix (**P**) would end up being of size (N + 2) by (N + 2) and could vary from one sampling plan and attack scenario combination to the other. However, there are some common terms in each of these probability matrices, such as:

- pns: The probability of NOT getting a signal due to a false alarm (random error), where:

$$\text{pns} = \Phi(L-\delta_0) - \Phi(-L-\delta_0) = 1 - \alpha$$

- ps: The probability of getting a signal due to a false alarm (random error), where:

$$\text{ps} = 1 - \text{pns} = 1 - \Phi(L-\delta_0) + \Phi(-L-\delta_0) = \alpha$$

- pns1: The probability of NOT getting a signal due to an intentional process shift (actual attack), where:

$$\text{pns1} = \Phi(L-\delta_{FA}) - \Phi(-L-\delta_{FA})$$

- ps1: The probability of getting a signal due to an intentional process shift (actual attack), where:

$$\text{ps1} = 1 - \text{pns1} = 1 - \Phi(L-\delta_{FA}) + \Phi(-L-\delta_{FA})$$

For Plan 0b, the complete transition probability matrix is shown next:

$$\mathbf{P} = \begin{bmatrix}
0 & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & ps \\
0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & ps \\
 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & 0 & 0 & & \frac{2pns}{N} & ps \\
 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns1}{N} & 0 & 0 & 0 & 0 & 0 & \frac{8ps}{N} + \frac{2ps1}{N} \\
 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns1}{N} & \frac{2pns}{N} & 0 & 0 & 0 & 0 & \frac{8ps}{N} + \frac{2ps1}{N} \\
\vdots & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns1}{N} & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & 0 & \frac{8ps}{N} + \frac{2ps1}{N} \\
 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns1}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & \frac{8ps}{N} + \frac{2ps1}{N} \\
 & 0 & 0 & 0 & 0 & \frac{2pns1}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & 0 & \frac{8ps}{N} + \frac{2ps1}{N} \\
 & 0 & 0 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & ps \\
 & \frac{2pns}{N} & 0 & 0 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & ps \\
0 & \frac{2pns}{N} & \frac{2pns}{N} & 0 & 0 & 0 & 0 & 0 & \frac{2pns}{N} & \frac{2pns}{N} & \frac{2pns}{N} & ps \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}$$

Once a probability transition matrix is determined, the ARL values could be easily obtained. Since there is only one absorbing state (the last one) in this DTMC, then the following is true:

$$\text{The transition rate matrix: } \mathbf{Q} = \mathbf{P}_{(N+1)\times(N+1)} \text{ and } \mathbf{I} = p_{(N+2)\times(N+2)} = 1.$$

$$\text{Also, the mean recurrance time: } \mathbf{M} = \text{Inverse}(\mathbf{I} - \mathbf{Q}),$$

where each element of the matrix **M** is the expected number of times the chain is in state j, given that it starts in state i. Hence, the desired ARL can be obtained from the 1st row of this matrix **M**, since it represents the excepted number of times the DTMC will be in each of the *N + 1* states, given it starts in state 0, so:

$$\text{ARL} = \text{sum}\,[\mathbf{M}(1, 1{:}\,N + 1)] \qquad \text{(B-9)}$$

It should be just noted that this approach could also be used for the other plans too. The steps needed to obtain the ARL values, using equation (B-9), are the same, but each attack scenario and sampling plan combination has its unique transition probability matrix. Hence, the following sub-sections will only show the transition probability matrix for the other two plans within this attack scenario, with the probability of not getting a signal due to the attack (pns1) would be always considered at state number *N/2* only.

Plan 1

$$\mathbf{P} = \begin{bmatrix} 0 & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns1}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{N}-1}{\mathrm{N}}(\mathrm{ps}) + \frac{\mathrm{ps1}}{\mathrm{N}} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns1}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{N}-1}{\mathrm{N}}(\mathrm{ps}) + \frac{\mathrm{ps1}}{\mathrm{N}} \\ 0 & 0 & & & & \cdots & & & & & 0 & 1 \end{bmatrix}$$

Plan 2

$$\mathbf{P} = \begin{bmatrix} 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & & & \cdots & & & & & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & & & \cdots & & & & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & & & \cdots & & & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & & \cdots & & & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns1}}{\mathrm{N}} & 0 & \cdots & & & 0 & \frac{\mathrm{ps1}}{\mathrm{N}} \\ \vdots & \vdots & & & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & \cdots & & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & \vdots & \vdots & & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & \cdots & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & & & \vdots & \vdots & & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ & & & & & & \vdots & 0 & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & \frac{\mathrm{ps}}{\mathrm{N}} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & \frac{\mathrm{ps}}{\mathrm{N}} \\ 0 & \frac{\mathrm{N}-1}{\mathrm{N}} + \frac{\mathrm{pns}}{\mathrm{N}} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & & \frac{\mathrm{ps}}{\mathrm{N}} \\ 0 & 0 & & & & \cdots & & & & & & 1 \end{bmatrix}$$

It should be noted that for this particular plan, the values obtained from the DTMCs need to be divided by N to obtain the corresponding ARL values. In other words, equation (B-9) will be modified as follows for just this plan:

$$\mathrm{ARL} = \mathrm{sum}\frac{[\mathbf{M}(1, 1\!:\mathrm{N}+1)]}{\mathrm{N}}$$

Average Time to Signal (ATS)

Since the signal could occur here due to either an attack or a random error, the ATS value would depend on 1) the average number of times an inspection has occurred, which equals to the average number of product groups considered, and 2) the number of products in each of these groups, which is fixed to *N*; as shown in equation (B-10):

$$\mathrm{ATS} = \mathrm{E[NGS]} \times \mathrm{N} \tag{B-10}$$

However, since randomness is now introduced, the last group of products would probably signal before *N* products are manufactured. So, for this group only, the number of products to signal, shown in equation (B-7), needs to be considered instead of *N*. Hence, equation (B-10) for estimating ATS can be re-written as follows:

$$\mathrm{ATS} = \ (\mathrm{E[NGS]} - 1) \times N + \ \mathrm{E[NPS|\,S_G]} \tag{B-11}$$

True Positive and False Positive Probabilities

The True Positive Probabilities (TPP) could be calculated as follows:

$$\mathrm{TPP} = \ \frac{\text{Probability of getting a true alarm}}{\text{Probability of getting an out of control signal in general}} \tag{B-12}$$

The denominator of equation (B-12) is just the probability of getting an out-of-control signal either from an actual attack (true alarm) or a random error (false alarm), which can be obtained from formulas, such as those in the previous sub-section. The numerator of equation (B-12) is just a portion of the probabilities used to get an out-of-control signal that accounts for true alarms only. In other words, the numerator is equivalent to the attack portion of the formulas in the denominator, as discussed previously for equation (B-3). The sum of the TPP and FPP values should be equal to one. Hence, the random error portion is the one used instead in the numerator of equation (B-12) to obtain the FPP values, where

$$\mathrm{FPP} = \ \frac{\text{Probability of getting a false alarm}}{\text{Probability of getting an out of control signal in general}} \tag{B-13}$$

A set of example values for the ARL, ATS, TPP, and FPP metrics for Plan 1, with different shift sizes caused by the attack is shown in Table B-4. The fact that TPP and FPP values add up to one is also confirmed in this table.

Table B-4: ARL and ATS values, along with TPP and FPP values (as percentages), corresponding to different transient shift sizes caused by an attack, for the fixed attack scenario – Plan 1.

Shift Size	ARL	ATS	TPP (% Signal due to attack)	FPP (% Signal due to error)
0	370.398	3,699.483	0	100
0.1	368.574	3,681.239	10.443	89.557
0.5	325.304	3,248.481	20.957	79.043
0.75	273.142	2,726.784	33.632	66.368
1	212.405	2,119.332	48.390	51.610
2.5	30.045	295.489	92.700	7.300
5	9.985	94.859	97.574	2.426
10	9.763	92.641	97.628	2.372

Number of Successfully Attacked Products until Signal Metrics

Intuitively, the NSAS can be thought of as the average number of times an inspection has occurred (E[NGS]) multiplied by the average number of attacked products between inspections; as follows

$$NSAS = \mathrm{E[NGS]} \times \text{ Avg number of attacked products between inspections} \qquad \text{(B-14)}$$

However, equation (B-14) would only hold if the inspection causing a signal occurs after a part is attacked, such as in Plan 0. If a signal is obtained due to inspection taking place before an attack occurs or at the attacked product, then the NSAS can be estimated instead as:

$$\mathrm{NSAS} = (\mathrm{E[NGS]} - 1) \times \text{ Avg number of attacked products between inspections} \qquad \text{(B-15)}$$

When the attack scenario involves only a single attack per a group of *N* products, the average number of attacked products between inspections is equal to 1 and the NSAS value would be very close (if not equal) to the ARL value.

Conversely, the PSAS depends on the source of the signal. If the signal was due to a random error and not because of an attack, then this means that all the attacks were successful and the PSAS value is 100%. However, if the attack did cause a signal, then it is the ratio between the NSAS and the total number of attacked parts. For instance, if the 4th attacked sample caused a signal, then the PSAS value is 3/4 = 75%. This metric's dependency on the signal's source can be seen in the following equation:

$$\text{PSAS} = \begin{cases} 1, \text{if the signal is due to a random error} \\ \dfrac{\text{NSAS}}{\text{NSAS} + 1}, \text{if the signal is due to an attack.} \end{cases} \quad \text{(B-16)}$$

Percentage of Attacked Products

A closely related metric is the Percentage of Attacked products (PA), which represents the percentage of products successfully attacked compared to the total number of products until signaling, regardless of the source of this signal. This metric would just be the ratio between the NSAS value to the total number of products until a signal is obtained, as follows:

$$\text{PA} = \frac{\text{NSAS}}{\text{ATS}} \quad \text{(B-17)}$$

Detection Probability

The next metric, the attack detection probability, is defined as the percentage of attacked parts that have caused a signal. So, it measures the number of detected parts by the control chart (which is always going to be equal to 1 only) as a percentage of the total number of parts attacked. For instance, if a signal occurs and 10 samples were attacked, then the DP is 1 out of 10 = 0.1. This is only true, however, if the resulting signal is due to an actual attack and not to a random error. If the resulting signal was not due to an attack, then the detection probability would actually be equal to zero. The formula that can be used for estimating it is stated in the following equation:

$$\text{DP} = \begin{cases} 0, \text{if the signal is due to a random error} \\ \dfrac{1}{\text{NSAS}+1}, \text{if the signal is due to an attack} \end{cases} \quad \text{(B-18)}$$

Out-of-Control Metrics

For the two out-of-control metrics considered, we have

$$\text{POC} = \frac{1}{\text{Samples that are out of control due to any reason}} \quad \text{(B-19)}$$

$$\text{PAOC} = \frac{1}{\text{Samples that are out of control due to an attack only}} \quad \text{(B-20)}$$

as the formulas used to estimate POC and PAOC, respectively. A low value of the POC metric reflects a higher number of out-of-control samples in general, whereas a low value of the PAOC metric indicates a higher number of samples out-of-control due to an attack. As an example, consider a situation where a signal is obtained when there is a total of 4 products out-of-control, with 2 out of those 4 are out-of-control due to an attack. In this situation, the POC value is just 1 out of 4 = 25%; whereas the PAOC value is 1 out 2 = 50%.

Appendix C Comparison of the Metrics Theoretical and Simulated Values – Individual Control Charts

In this appendix, comparisons of a few metrics theoretical values vs those obtained from the simulation procedure are presented for the first attack scenario discussed in Chapter 5. However, to get a more accurate estimate of the simulated values, the whole simulation procedure, steps 2-6 in sub-section 5.3.2 is repeated for 20 replications (using 100,000 iterations only per replication). It should be noted that now each of the metrics has a mean and a standard deviation value across the different replications, as discussed by Saleh *et al.* (2015).

More specifically, the metrics compared are the ARL, ATS, TPP, and FPP, as demonstrated in the forthcoming five tables. Each of these tables have the theoretical values and the mean of the simulated values of the metrics displayed side-by-side, along with the standard deviation of the simulated values between parentheses. In addition, the percentage difference from the theoretical values is also shown for each of the metrics.

Table C-1: Comparison of the ARL metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the current sampling plans.

Shift Size	Theoretical Values		Values obtained from Simulation		Difference (%)	
	Plan 0	Plan 0b	Plan 0	Plan 0b	Plan 0	Plan 0b
0	370.3983	370.3983	370.3722 (1.1893)	370.6906 (1.1404)	0.0071	- 0.0789
0.1	370.3983	368.5915	370.5615 (1.4650)	369.1260 (0.8726)	- 0.0441	- 0.1450
0.5	370.3983	325.7822	370.3297 (0.9538)	325.5011 (0.9208)	0.0185	0.0863
0.75	370.3983	274.3004	370.2190 (1.1420)	274.0253 (0.7205)	0.0484	0.1003
1	370.3983	214.5283	370.2459 (1.0448)	214.3117 (0.6252)	0.0411	0.1009
2.5	370.3983	36.1704	370.7891 (1.1316)	36.3744 (0.0903)	- 0.1055	- 0.5641
5	370.3983	16.6509	370.1729 (0.7782)	16.8286 (0.0334)	0.0609	- 1.0671
10	370.3983	16.4352	370.1438 (1.3251)	16.6312 (0.0432)	0.0687	- 1.1926

Table C-2: Comparison of the ARL metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans.

Shift Size	Theoretical Values		Values obtained from Simulation		Difference (%)	
	Plan 1	Plan 2	Plan 1	Plan 2	Plan 1	Plan 2
0	370.3983	370.3983	370.1951 (1.2687)	370.0802 (1.1727)	0.0549	0.0859
0.1	368.5742	368.5739	368.7550 (1.3610)	368.8599 (0.9372)	- 0.0491	- 0.0776
0.5	325.3042	325.2914	325.1103 (1.2587)	325.5533 (0.9384)	0.0596	- 0.0805
0.75	273.1415	273.0974	273.0665 (0.7014)	273.4113 (1.0622)	0.0275	- 0.1149
1	212.4046	212.3014	212.4261 (0.5272)	212.3199 (0.6754)	- 0.0101	- 0.0087

Shift Size	Theoretical Values		Values obtained from Simulation		Difference (%)	
	Plan 1	Plan 2	Plan 1	Plan 2	Plan 1	Plan 2
2.5	30.0449	29.6187	30.0122 (0.870)	29.6350 (0.0799)	0.1087	- 0.0553
5	9.9845	9.5095	9.9752 (0.0304)	9.5166 (0.0263)	0.0935	- 0.0750
10	9.7628	9.2871	9.7563 (0.0319)	9.2929 (0.0306)	0.0665	- 0.0619

For the ARL metric, the comparison is shown for both the sampling plans used currently in practice (Plans 0 and 0b) as well as the newly proposed ones (Plans 1 and 2). For the other metrics, however, the comparison is only shown for the proposed sampling plans. Overall, it can be seen from the comparisons that the values of all the metrics in this attack scenario are almost identical and the margin of the error due to simulation is insignificantly small. These comparisons confirm that the simulation procedure was effective and provides an adequate representation of the sampling plans performance discussed in this work.

Table C-3: Comparison of the ATS metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans.

Shift Size	Theoretical Values		Values obtained from Simulation		Difference (%)	
	Plan 1	Plan 2	Plan 1	Plan 2	Plan 1	Plan 2
0	3,699.4835	3,703.9835	3700.2416 (12.6848)	3705.1466 (11.7775)	- 0.0205	- 0.0314
0.1	3,681.2392	3,685.7391	3679.9469 (13.6115)	3683.3035 (9.4507)	0.0351	0.0661
0.5	3,248.4810	3,252.9144	3251.6226 (12.5874)	3251.3105 (9.6698)	- 0.0967	0.0493
0.75	2,726.7842	2,730.9741	2728.8471 (7.0124)	2734.0434 (10.4577)	- 0.0757	- 0.1124
1	2,119.3324	2,123.0137	2122.8558 (5.2707)	2122.3439 (6.7623)	- 0.1663	0.0315
2.5	295.4891	296.1865	295.1977 (0.8698)	296.1471 (0.7334)	0.0986	0.0133
5	94.8589	95.0946	94.9208 (0.3042)	95.1266 (0.2648)	- 0.0653	- 0.0336
10	92.6410	92.8714	92.5452 (0.3195)	92.8715 (0.3572)	0.1034	- 0.0001

Table C-4: Comparison of the TPP metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans.

Shift Size	Theoretical Values (%)		Values obtained from Simulation (%)		Difference (%)	
	Plan 1	Plan 2	Plan 1	Plan 2	Plan 1	Plan 2
0	0.0000	0.0000	0.0000 (0.0000)	0.0000 (0.0000)	0.0000	0.0000
0.1	10.4432	10.4447	10.4233 (0.0908)	10.4539 (0.0944)	0.1914	- 0.0880
0.5	20.9571	20.9630	20.9495 (0.1288)	21.0042 (0.1256)	0.0363	- 0.1966
0.75	33.6316	33.6469	33.6488 (0.1287)	33.6819 (0.1798)	- 0.0509	- 0.1039
1	48.3896	48.4212	48.3805 (0.1398)	48.3336 (0.1330)	0.0188	0.1810
2.5	92.6997	92.8157	92.6728 (0.0897)	92.8314 (0.0623)	0.0290	- 0.0168
5	97.5739	97.7026	97.5659 (0.0552)	97.7108 (0.0536)	0.0083	- 0.0084
10	97.6278	97.7566	97.6179 (0.0396)	97.7633 (0.0568)	0.0102	- 0.0068

Table C-5: Comparison of the FPP metric theoretical values vs. mean of simulated values, with the standard deviation values from simulation between parentheses, for the proposed random sampling plans.

Shift Size	Theoretical Values (%)		Values obtained from Simulation (%)		Difference (%)	
	Plan 1	Plan 2	Plan 1	Plan 2	Plan 1	Plan 2
0	100	100	100 (0.0000)	100 (0.0000)	0.0000	0.0000
0.1	89.5568	89.5553	89.5768 (0.0908)	89.5461 (0.0944)	- 0.0223	0.0103
0.5	79.0429	79.0370	79.0506 (0.1288)	78.9958 (0.1256)	- 0.0096	0.0522
0.75	66.3684	66.3531	66.3513 (0.1287)	66.3182 (0.1798)	0.0258	0.0527
1	51.6104	51.5788	51.6195 (0.1398)	51.6665 (0.1330)	- 0.0176	- 0.1699
2.5	7.3003	7.1843	7.3272 (0.0897)	7.1687 (0.0623)	- 0.3678	0.2174
5	2.4261	2.2974	2.4342 (0.0552)	2.2892 (0.0536)	- 0.3334	0.3588
10	2.3722	2.2434	2.3821 (0.0396)	2.2368 (0.0568)	- 0.4183	0.2971

Appendix D Details of the Economic Cost Model Proposed by Lorenzen and Vance (1986)

The original economic cost model proposed by Lorenzen and Vance (1986) is detailed in this appendix. The authors of the model break down the costs into three distinct categories. But before discussing these cost categories, the assumptions they built their model upon needs to be mentioned first, along with the different components used in estimating the expected quality cycle time.

Model Assumptions

- Assignable causes are assumed to occur according to a Poisson process with an intensity of λ occurrences per hour (Montgomery, 2009). So, the in-control time can be considered to follow a negative exponential distribution with mean $1/\lambda$ (Lorenzen and Vance, 1986);
- Only a single assignable cause is assumed, which results in a change in the process mean;
- "Drifting" processes that affect process variability are not considered here and the covariance matrix used is assumed to be constant over time;
- The process is not "self-correcting" and can only be returned to an in-control condition by external intervention;
- The quality cycle is assumed to follow a "renewal reward" process, where it starts in an in-control state and continues until the process is repaired after an out-of-control signal is reached;
- The *p* quality characteristics monitored follow a multivariate normal distribution with mean vector (μ) and covariance matrix (Σ).

Quality Cycle Time

The quality cycle is divided into four intervals of time: 1) the interval during which the process is in-control; 2) the interval during which the process is out-of-control, but still undetected; 3) the interval during which the assignable cause is detected and identified, referred to as T_1; and 4) the interval during which the assignable cause is repaired, referred to as T_2.

$$\mathrm{E(In - control\ Time)} = \frac{1}{\lambda} + \left(1 - \gamma_1 \frac{\mathrm{sT_0}}{\mathrm{ARL_0}}\right) \tag{D-1}$$

where ${}^{1}/_{\lambda}$ is the mean of the Poisson process, *s* is the expected number of samples while the process is in-control, *T₀* is the expected time spent to search for false alarms, *ARL₀* is the average run length for an in-control process, and γ_1 is a binary variable indicating if production continues when searching for the assignable cause; such that:

$$\gamma_1 = \begin{cases} 1, \text{if production continues with searches,} \\ \qquad\qquad\qquad 0, \text{otherwise;} \end{cases}$$

$$\mathrm{s} = \frac{\mathrm{e}^{-\lambda \mathrm{h}}}{1 - \mathrm{e}^{-\lambda \mathrm{h}}}, \text{where h is the sampling frequency; and}$$

$$\mathrm{ARL}_0 = \frac{1}{\alpha}, \text{where } \alpha \text{ is the type I error.}$$

$$\mathrm{E(Out-of-control\ Time)} = \mathrm{h(ARL_1)} - \tau + \mathrm{nE} + \mathrm{T_1} + \mathrm{T_2} \qquad \text{(D-2)}$$

where τ is the expected time between assignable causes occurrence and the last in-control sample, *ARL₁* is the out-of-control average run length, *n* is the sample size, *E* is the expected time to sample and chart 1 item, *T₁* is the time interval during which the assignable cause is detected and identified, and *T₂* is the time interval during which the assignable cause is repaired; such that:

$$\tau = \frac{1 - (1 + \lambda \mathrm{h})\mathrm{e}^{-\lambda \mathrm{h}}}{\lambda(1 - \mathrm{e}^{-\lambda \mathrm{h}})} \text{ and } \mathrm{ARL}_1 = \frac{1}{1 - \beta}, \text{where } \beta \text{ is the type II error.}$$

$$\therefore \text{The expected cycle time is E(T)}$$

$$= \frac{1}{\lambda} + (1 - \gamma_1 s T_0 \alpha) + \frac{\mathrm{h}}{1 - \beta} - \tau + \mathrm{nE} + \mathrm{T_1} + \mathrm{T_2} \qquad \text{(D-3)}$$

Quality Costs

The three main costs considered in this model are: 1) the sampling cost; 2) the cost of producing non-conformities while the process is operating, referred to as “production cost”; and 3) the cost of signaling of an out-of-control condition, which includes investigating the cause behind the signal and repairing it.

$$\mathrm{E(Sampling\ Cost)} = (\mathrm{a_1} + \mathrm{a_2 n})\frac{\mathrm{E(T)}}{\mathrm{h}} \qquad \text{(D-4)}$$

where *a₁* is the fixed cost of sampling and testing, *a₂* is the variable cost of sampling and testing per production cycle, and *n* is the sample size.

$$\mathrm{E(Production\ Cost)} = \frac{\mathrm{C_0}}{\lambda} + \mathrm{C_1}[\mathrm{h(ARL_1)} - \tau + \mathrm{nE} + \gamma_1 \mathrm{T_1} + \gamma_2 \mathrm{T_2}] \qquad \text{(D-5)}$$

where *C₀* is the hourly production cost due to non-conformities when the process is in-control, *C₁* is the hourly production cost due to non-conformities when the process is out-of-control and

$$\gamma_2 = \begin{cases} 1, \text{if producting continues during process repair,} \\ \quad 0, \text{otherwise.} \end{cases}$$

$$\text{E(Signaling Cost)} = a_3 + \frac{s a'_3}{ARL_0} \qquad \text{(D-6)}$$

where a_3 is the cost of locating and repairing an assignable cause and a'_3 is the cost of investigating each false alarm. It should be noted that

$$\frac{s}{ARL_0} = \frac{\alpha e^{-\lambda h}}{1 - e^{-\lambda h}} \text{ is just the expected number of false alarms.}$$

$$\therefore \text{ The expected net income per cycle is } E(C)$$

$$= (a_1 + a_2 n)\frac{E(T)}{h} + \frac{C_0}{\lambda} + C_1\left[\frac{h}{1-\beta} - \tau + nE + \gamma_1 T_1 + \gamma_2 T_2\right] + a_3 + \frac{a'_3 \alpha e^{-\lambda h}}{1 - e^{-\lambda h}} \qquad \text{(D-7)}$$

Typically, the expected net income per hour is obtained by dividing E(C)/E(T) to find the optimal design parameters, using some type of optimization technique(s). However, our work only focuses on cost comparisons and not parameters optimization.

Related Error Probabilities

According to Montgomery and Klatt (1972), the Type I and II errors for a Hotelling's T-square control chart can be estimated as shown next.

$$\alpha = \text{Type I Error Probability}$$

$$\alpha = \text{Prob of Getting an out} - \text{of} - \text{control Signal}$$

$$\alpha = \text{Prob(Statistic} > \text{UCL)} + \text{Prob(Statistic} < \text{LCL)}$$

For T^2 charts, only an upper control limit exists such that $\alpha = \text{Prob}(T_0^2 > \text{UCL})$

$$\therefore \alpha = P\left(T_0^2 > T^2_{\alpha,p,m-p}\right) = \int_{T^2_{\alpha,p,m-p}}^{\infty} f(T_0^2)\,dT_0^2 \qquad \text{(D-8)}$$

where $f(T_0^2)$ is the T^2 distribution with *p* and *m-p* degrees of freedom, and the UCL is based on the F-distribution. Typically, α is set to 0.005, yielding an in-control ARL of 200.

$$\beta = \text{Type II Error Probability}$$

$$\beta = \text{Probability of NOT detecting the shift in the } 1^{st} \text{ subsequent sample.}$$

$$\mathrm{\beta = Prob(LCL \leq Statisitc \leq UCL|A\ Shift\ exisits)}$$

$$\mathrm{\beta = Prob(T_1^2 \leq UCL) = 1 - Prob(T_1^2 > UCL)}$$

$$\therefore 1 - \ \mathrm{\beta = 1 - [1 - \ Prob(T_1^2 > UCL)] = Prob(T_1^2 > UCL)}$$

$$\therefore 1 - \ \beta = \ \mathrm{P}\left(T_1^2 > T_{\alpha,p,m-p}^2\right) = \ \int_{T_{\alpha,p,m-p}^2}^{\infty} f(T_1^2)dT_1^2 \tag{D-9}$$

where $f(T_1^2)$ is the T^2 distribution with a non-centrality parameter for the corresponding non-central F-distribution; i.e., it is the T^2-distribution when a shift exists.

Appendix E ARL Values Obtained from Simulation for Both of the Proposed Random Sampling Plans – Individual T^2 Control Charts

In this appendix the ARL values obtained from simulation for both of the proposed random sampling plans discussed in Chapter 6 are presented side-by-side for all the considered combinations in the following tables. It should be just noted that comparing the plans for the same value of q wouldn’t be a fair comparison, since the value of q_i is randomized for each product i being inspected in Plan 2. Instead, the comparisons should be based on the $q_{average}$, which are shown as the last row of each of the following tables.

In addition, with a significance level (α) of 0.005, the in-control ARL is set to be equal to 200. That is the same value that the out-of-control ARL would take, if the cyber-physical attacks were intelligently-designed to avoid the inspected variables every time one occurs. Since all the of the out-of-control values presented in this appendix are less than 200, this implies that the proposed random sampling plans work better against these types of cyber-physical attacks. Not only would the prospective attackers no longer know which variables are being monitored at any given point of time, they would not also know the best variables to target during a Passive cyber-physical attack; and, hence, the attackers would not be able to implement their attack without risking a higher chance of detection

Table E-1: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0	200.8425	198.7685	201.84	198.4467	198.4204	201.7546	199.5644	201.2787
0.5	178.8567	172.3648	167.6589	161.2548	179.5759	177.1843	171.0444	169.5557
1	133.6214	113.4078	104.3707	94.0451	132.9467	124.1484	114.004	106.1567
1.5	85.0439	66.3658	55.7643	46.0909	85.2499	73.5709	66.9088	57.4919
2	51.7749	35.7545	27.5721	21.8193	51.4480	41.6056	33.4297	28.9004
2.5	32.4001	19.7677	13.7189	10.5042	32.7384	24.5821	19.0239	14.8669
3	21.4547	11.6314	7.9477	5.5742	20.9083	15.112	11.0047	8.4467
3.5	14.5806	7.5609	4.6688	3.2466	14.8046	9.9627	6.954	5.0914
4	10.9054	5.3416	3.175	2.1316	11.0831	7.0439	4.8729	3.4145
q_{avg}	**1**	**3**	**6**	**10**	**1**	**2**	**3.5**	**5.5**

Table E-2: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0	199.1095	199.8519	200.6945	201.2149	201.9465	201.1101	201.9876	201.5172
0.5	178.7853	171.6352	169.045	161.4649	179.4035	178.7934	172.2888	169.1163
1	137.6768	118.902	105.7238	91.8619	132.3913	125.9291	118.2279	106.156
1.5	91.0207	70.0026	55.6344	47.0072	91.994	79.2774	68.3218	59.7606
2	57.9106	38.9386	28.4968	21.6963	57.6536	46.7008	37.4673	30.1572
2.5	37.2194	22.3509	14.9743	10.3587	37.2905	27.5238	20.5584	15.9562
3	24.4105	13.251	8.3629	5.4907	24.8278	17.5503	12.4247	9.0502
3.5	17.3275	8.6035	5.0388	3.2306	17.1289	11.3756	7.7435	5.4204
4	12.7152	5.9134	3.2676	2.157	12.8252	8.1167	5.2601	3.5987
q_{avg}	**1**	**3**	**6**	**10**	**1**	**2**	**3.5**	**5.5**

Table E-3: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
0	198.1945	201.074	199.2113	198.5343	198.5912	198.616	201.2675	201.9556
0.5	187.6743	185.6358	178.6435	178.109	187.0999	185.2841	182.2234	178.0561
1	153.987	139.4145	129.1684	120.1645	155.3859	147.2473	140.6281	136.5477
1.5	113.6157	93.6096	80.3238	68.4443	118.5191	104.8779	94.4414	84.8002
2	76.5541	56.8654	45.0254	36.3876	81.6715	67.3562	57.9491	48.9582
2.5	50.7238	33.6972	24.6105	18.8809	56.1796	42.2426	34.0739	27.5469
3	33.5393	20.2144	13.7746	10.1823	37.5525	26.5899	19.8501	15.306
3.5	22.7896	12.6104	7.9132	5.5942	26.0459	17.6934	12.3669	9.2739
4	16.5594	8.4542	5.1722	3.4035	19.3586	11.9564	8.0733	5.7875
q_{avg}	**2**	**6**	**12**	**20**	**1.5**	**3.5**	**6.5**	**10.5**

Table E-4: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
0	199.6248	200.3249	198.1576	201.3853	200.775	198.7419	199.1624	201.1077
0.5	188.6504	181.2779	177.6605	172.6214	183.121	185.024	182.6458	176.3979
1	155.4806	139.1985	133.5092	120.5125	156.9893	150.8964	141.0687	133.0632
1.5	117.2697	97.1235	82.6026	69.7205	119.355	110.1176	96.4663	86.5614
2	82.0641	60.6049	47.7136	36.8784	87.2095	73.7581	60.8138	50.7307
2.5	56.6537	36.7743	26.0527	19.291	62.5279	46.3694	36.4098	28.6168

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
3	38.6698	22.6731	15.0125	10.158	44.4348	30.6525	21.7917	16.3277
3.5	27.0509	14.292	8.7654	5.6235	31.2723	20.322	13.807	9.614
4	19.6368	9.3619	5.4432	3.4573	22.6335	13.9431	8.8186	6.0685
q_{avg}	**2**	**6**	**12**	**20**	**1.5**	**3.5**	**6.5**	**10.5**

Table E-5: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 5	**q = 15**	**q = 30**	**q = 50**	**q = 5**	**q = 15**	**q = 30**	**q = 50**
0	199.837	201.6302	200.324	201.546	199.325	199.8187	199.54	198.7464
0.5	190.119	191.3216	186.781	186.0721	194.518	190.6931	190.324	188.7168
1	171.835	165.0978	159.589	148.6277	176.796	170.3988	168.29	159.3015
1.5	146.488	130.5249	117.464	106.5735	154.317	140.0592	131.224	121.9083
2	113.364	95.2961	79.3714	68.9085	123.896	109.8971	96.4896	84.5802
2.5	85.365	65.352	50.8791	42.271	96.6029	77.7454	65.8308	57.7389
3	61.9819	43.0458	31.9702	25.1333	74.3884	55.6199	43.5013	35.743
3.5	44.3122	28.2325	19.6899	14.3463	54.8338	38.6193	28.7982	21.955
4	32.2665	18.6682	12.4054	8.9558	41.2895	27.0052	19.2744	14.2445
q_{avg}	**5**	**15**	**30**	**50**	**3**	**8**	**15.5**	**25.5**

Table E-6: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 5	**q = 15**	**q = 30**	**q = 50**	**q = 5**	**q = 15**	**q = 30**	**q = 50**
0	201.371	201.553	200.573	199.6868	198.389	200.687	200.606	200.4968
0.5	188.203	187.939	190.54	184.9767	193.513	190.17	191.323	186.1552
1	172.93	164.873	157.149	152.618	180.286	170.382	166.072	159.1841
1.5	148.704	134.65	116.512	104.7928	154.344	143.509	130.794	121.4437
2	117.92	96.0018	79.4951	68.0468	127.689	111.519	98.3672	84.7388
2.5	92.8059	67.5689	52.962	42.9581	103	83.7805	67.1287	57.4054
3	68.1954	47.3396	33.1986	24.986	80.5362	59.9193	46.5702	36.4896
3.5	50.198	31.1485	20.4721	14.7564	61.5848	42.2529	30.5184	22.6091
4	37.386	20.6618	12.6597	8.8602	47.3509	30.1437	20.2453	14.2197
q_{avg}	**5**	**15**	**30**	**50**	**3**	**8**	**15.5**	**25.5**

Table E-7: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0	201.245	201.711	198.696	200.328	201.438	199.401	197.986	198.925
0.5	198.222	192.634	194.913	191.593	198.225	194.654	192.093	189.495
1	184.737	179.523	173.194	168.423	186.284	184.527	173.488	170.749
1.5	162.866	151.845	142.824	134.418	173.49	161.041	150.433	147.052
2	141.902	124.877	107.32	101.189	151.756	135.988	125.48	113.901
2.5	117.202	94.8351	79.1904	69.3827	127.72	111.096	95.6967	86.2345
3	93.0835	70.4002	56.5179	48.1165	106.34	84.8039	71.8099	61.1008
3.5	71.2353	50.3099	37.9942	31.5503	85.7389	64.2548	52.1299	42.5417
4	55.3865	36.0942	25.7877	20.2552	68.2175	49.2752	36.2739	29.269
q_{avg}	**10**	**30**	**60**	**100**	**5.5**	**15.5**	**30.5**	**50.5**

Table E-8: ARL values obtained from simulation for both of the proposed random sampling plans, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values.

Shift Size	**Proposed Random Sampling Plan 1**				**Proposed Random Sampling Plan 2**			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0	198.756	200.475	200.454	199.482	201.493	199.347	200.576	199.774
0.5	191.569	194.402	192.495	189.222	197.99	198.239	195.908	189.302
1	183.838	179.134	167.157	164.174	185.755	182.708	175.872	170.305
1.5	165.123	154.358	143.38	136.338	172.188	160.262	154.093	145.332
2	141.687	122.735	110.357	101.61	151.875	139.807	126.777	114.785
2.5	121.124	96.0988	78.7938	70.0077	135.556	113.183	98.5294	85.7688
3	96.4571	70.5753	56.0535	48.2474	110.772	86.7749	71.7359	63.8949
3.5	76.6207	52.2795	39.1493	31.6228	93.6842	69.2261	52.3758	44.0967
4	59.7923	38.2859	26.6804	20.477	73.4536	51.9485	37.9703	29.5318
q_{avg}	**10**	**30**	**60**	**100**	**5.5**	**15.5**	**30.5**	**50.5**

Appendix F Cost Values Obtained from the Simplified Cost Model – Individual T^2 Control Charts

In this appendix the cost values obtained from the simplified cost model for both the variable selection approaches (the randomness and naïve approaches discussed in Chapter 6) are presented side-by-side for all the considered combinations in the following tables. It should be noted that the last row of each table represents the actual number of variables (q_{actual}) that are inspected per production cycle. In the naïve approach, this value is always going to be equal to *p*, whereas in the randomness approach this value could be either just the *q* value (Plan 1) or the average *q* value (Plan 2).

Table F-1: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	2,681.58	4,313.67	6,719.64	9,700.21	8,017.80	6,916.38	6,337.38	9,685.14
1	2,003.05	2,839.74	4,188.11	5,667.62	4,458.31	3,290.30	2,995.48	5,729.24
1.5	1,274.39	1,663.69	2,243.85	2,790.37	2,105.45	1,514.65	1,367.48	2,815.76
2	775.35	898.41	1,116.16	1,334.08	1,013.59	725.10	668.49	1,315.63
2.5	484.73	498.74	562.04	655.17	537.46	382.96	363.46	659.35
3	320.55	295.33	331.19	359.37	314.04	229.40	218.21	356.63
3.5	217.44	193.57	200.03	219.71	208.93	156.70	150.89	219.82
4	162.31	138.09	140.28	152.81	154.07	119.81	116.99	150.95
q_{actual}	**1**	**3**	**6**	**10**	**10**	**10**	**10**	**10**

Table F-2: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	2,680.51	4,295.43	6,775.08	9,712.81	7,515.69	6,066.80	5,462.17	9,860.20
1	2,063.88	2,977.10	4,242.23	5,536.63	4,310.87	2,894.50	2,394.43	5,669.35
1.5	1,364.04	1,754.61	2,238.66	2,845.35	2,274.55	1,365.49	1,108.58	2,779.15
2	867.39	978.01	1,153.15	1,326.70	1,228.32	700.04	548.63	1,341.80
2.5	557.02	563.32	612.25	646.44	685.90	388.51	309.61	656.74
3	364.89	335.82	347.80	354.36	416.66	238.01	198.10	361.04
3.5	258.64	219.64	214.83	218.75	277.42	163.55	143.15	216.09
4	189.46	152.38	143.98	154.34	201.58	127.96	112.12	151.73
q_{actual}	**1**	**3**	**6**	**10**	**10**	**10**	**10**	**10**

Table F-3: Cost values obtained (in $) for both approaches, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
0.5	3,755.13	7,438.71	12,535.78	19,646.01	16,805.31	13,381.13	12,621.74	19,178.07
1	3,081.38	5,589.86	9,072.53	13,272.11	9,480.29	7,205.84	6,522.78	13,582.03
1.5	2,273.95	3,757.66	5,653.40	7,582.89	4,103.26	3,703.83	3,311.59	7,730.83
2	1,532.72	2,287.90	3,182.52	4,056.65	1,713.27	1,859.85	1,681.02	4,117.96
2.5	1,016.12	1,361.17	1,753.47	2,130.92	807.10	1,012.63	917.56	2,162.06
3	672.43	821.86	994.96	1,174.07	463.79	594.22	544.66	1,156.35
3.5	457.43	517.70	584.66	669.38	303.64	386.35	361.94	674.34
4	332.83	351.45	392.79	428.40	227.78	279.62	264.85	430.59
q_{actual}	**2**	**6**	**12**	**20**	**20**	**20**	**20**	**20**

Table F-4: Cost values obtained (in $) for both approaches, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
0.5	3,774.65	7,264.40	12,466.97	19,042.37	17,044.71	11,481.13	10,687.48	19,136.28
1	3,111.25	5,581.22	9,376.38	13,310.39	9,644.16	5,824.03	4,972.82	13,177.74
1.5	2,347.03	3,898.22	5,812.92	7,723.27	4,465.19	3,006.04	2,374.43	7,733.58
2	1,642.92	2,437.48	3,370.69	4,110.64	1,985.50	1,558.54	1,228.34	4,143.50
2.5	1,134.71	1,484.25	1,854.43	2,176.03	917.00	888.39	709.36	2,122.19
3	775.04	920.20	1,081.61	1,171.40	512.41	547.10	447.26	1,157.41
3.5	542.66	584.96	644.32	672.60	327.94	373.20	313.30	678.04
4	394.38	387.76	411.76	434.32	239.97	275.14	243.48	431.22
q_{actual}	**2**	**6**	**12**	**20**	**20**	**20**	**20**	**20**

Table F-5: Cost values obtained (in $) for both approaches, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 5	**q = 15**	**q = 30**	**q = 50**	**q = 5**	**q = 15**	**q = 30**	**q = 50**
0.5	6,664.54	16,301.80	29,968.09	48,520.06	37,979.87	32,786.52	31,631.79	48,102.40
1	6,024.60	14,072.78	25,617.37	38,784.52	25,085.79	19,366.96	17,802.62	39,084.32
1.5	5,137.45	11,134.08	18,877.37	27,850.42	16,051.73	11,101.22	9,904.73	27,330.92
2	3,978.09	8,139.64	12,782.54	18,057.52	9,430.52	6,400.50	5,617.10	17,888.52
2.5	2,998.14	5,594.39	8,223.77	11,131.77	5,446.77	3,703.13	3,228.09	11,147.14
3	2,179.74	3,698.36	5,198.35	6,675.97	3,244.36	2,219.57	1,978.34	6,589.39
3.5	1,561.30	2,439.23	3,233.50	3,871.35	2,001.93	1,432.55	1,273.95	3,835.32
4	1,139.70	1,626.27	2,067.98	2,469.82	1,338.35	973.47	908.16	2,395.41
q_{actual}	**5**	**15**	**30**	**50**	**50**	**50**	**50**	**50**

Table F-6: Cost values obtained (in $) for both approaches, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	Randomness Approach				Naïve Approach			
	q = 5	**q = 15**	**q = 30**	**q = 50**	**q = 5**	**q = 15**	**q = 30**	**q = 50**
0.5	6,597.47	16,014.29	30,569.55	48,235.26	34,064.42	27,995.97	26,421.75	47,729.24
1	6,062.92	14,053.64	25,226.94	39,821.99	21,488.04	15,087.88	13,576.76	38,828.87
1.5	5,214.99	11,484.74	18,725.04	27,387.44	13,123.11	8,022.43	6,690.12	27,594.32
2	4,137.56	8,199.62	12,802.33	17,833.48	8,026.65	4,513.79	3,657.74	18,108.74
2.5	3,258.58	5,782.82	8,557.04	11,310.42	5,096.03	2,687.49	2,146.17	10,928.90
3	2,397.21	4,063.33	5,394.89	6,637.67	3,144.34	1,721.46	1,377.95	6,543.37
3.5	1,767.30	2,687.09	3,358.65	3,977.98	2,057.83	1,152.19	955.97	3,905.78
4	1,318.88	1,795.72	2,108.67	2,444.97	1,399.51	844.43	723.01	2,414.83
q_{actual}	**5**	**15**	**30**	**50**	**50**	**50**	**50**	**50**

Table F-7: Cost values obtained (in $) for both approaches, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	Randomness Approach				Naïve Approach			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0.5	11,918.24	30,904.57	60,593.32	97,999.29	75,432.30	64,252.79	62,281.80	97,140.81
1	11,109.14	28,806.83	53,860.65	86,182.33	52,826.14	40,711.45	38,162.01	84,307.78
1.5	9,796.88	24,378.33	44,445.73	68,840.19	35,460.34	25,348.51	22,807.64	67,457.28
2	8,539.01	20,063.40	33,439.67	51,893.15	23,479.93	15,187.84	13,435.73	51,202.35
2.5	7,057.04	15,256.73	24,719.44	35,671.99	14,844.96	9,485.17	8,112.76	36,183.36
3	5,609.93	11,347.15	17,690.96	24,826.22	9,545.65	6,085.05	5,263.24	24,370.95
3.5	4,299.04	8,132.70	11,948.61	16,377.46	6,107.29	3,955.49	3,379.40	16,233.59
4	3,348.11	5,858.19	8,164.60	10,616.96	4,008.84	2,753.02	2,428.25	10,457.38
q_{actual}	**10**	**30**	**60**	**100**	**100**	**100**	**100**	**100**

Table F-8: Cost values obtained (in $) for both approaches, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values, using sampling Plan 1.

Shift Size	Randomness Approach				Naïve Approach			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0.5	11,519.03	31,187.40	59,843.77	96,790.18	67,864.15	55,855.54	51,800.17	96,244.12
1	11,055.20	28,744.49	51,989.11	84,015.34	43,853.76	30,880.69	27,657.28	84,246.58
1.5	9,932.32	24,780.32	44,618.09	69,818.93	28,105.22	17,593.81	15,183.91	67,934.74
2	8,526.13	19,720.72	34,380.99	52,107.86	18,217.85	10,261.95	8,351.23	51,300.07
2.5	7,292.36	15,458.92	24,596.49	35,990.74	11,903.33	6,173.48	4,998.60	36,496.55
3	5,812.34	11,375.16	17,547.00	24,892.98	7,793.45	3,977.63	3,238.38	24,502.32
3.5	4,622.16	8,447.84	12,306.70	16,414.44	5,195.15	2,775.81	2,260.76	16,366.29
4	3,612.46	6,208.86	8,441.34	10,730.08	3,673.67	2,040.34	1,674.47	10,480.79
q_{actual}	**10**	**30**	**60**	**100**	**100**	**100**	**100**	**100**

Table F-9: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 5 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	2,692.37	3,545.33	4,709.73	6,370.16	8,260.12	7,057.98	6,805.81	8,557.21
1	1,992.93	2,484.61	3,141.11	3,992.70	4,473.58	3,625.82	3,299.37	4,638.88
1.5	1,277.48	1,473.06	1,846.00	2,167.77	2,102.82	1,765.82	1,589.18	2,298.89
2	770.45	833.75	925.32	1,095.59	1,031.33	866.69	775.15	1,109.28
2.5	489.81	493.28	529.16	569.33	546.88	461.28	417.31	566.89
3	312.35	303.88	308.63	328.58	316.45	278.87	253.30	323.73
3.5	220.80	200.89	197.24	202.75	211.31	187.32	172.18	205.31
4	164.98	142.52	140.01	139.87	153.31	140.55	132.13	147.68
q_{actual}	**1**	**2**	**3.5**	**5.5**	**10**	**10**	**10**	**10**

Table F-10: Cost values (in $) obtained for both approaches, when p = 10 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 1	**q = 3**	**q = 6**	**q = 10**	**q = 1**	**q = 3**	**q = 6**	**q = 10**
0.5	2,689.78	3,577.51	4,743.95	6,353.69	7,592.01	6,490.37	5,869.74	7,969.20
1	1,984.60	2,520.22	3,257.27	3,992.67	4,324.39	3,223.10	2,811.41	4,188.65
1.5	1,378.64	1,587.19	1,884.85	2,252.85	2,356.55	1,634.95	1,334.63	1,991.90
2	863.53	935.66	1,036.36	1,142.72	1,223.25	869.11	679.26	971.63
2.5	558.09	552.12	571.36	610.18	698.86	491.33	391.40	518.80
3	371.15	352.65	347.68	351.21	419.21	304.73	245.99	304.20
3.5	255.66	229.15	218.95	215.09	279.77	207.60	175.63	201.08
4	191.11	163.97	150.66	146.78	199.90	160.98	137.99	149.29
q_{actual}	**1**	**2**	**3.5**	**5.5**	**10**	**10**	**10**	**10**

Table F-11: Cost values (in $) obtained for both approaches, when p = 20 variables and M = 10 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 2	**q = 6**	**q = 12**	**q = 20**	**q = 2**	**q = 6**	**q = 12**	**q = 20**
0.5	3,274.43	5,101.32	7,759.23	11,154.88	15,696.96	14,348.58	12,829.29	17,508.56
1	2,719.44	4,055.31	5,991.43	8,560.60	9,893.73	8,390.93	7,037.46	11,345.06
1.5	2,074.27	2,890.15	4,028.49	5,326.39	5,527.68	4,357.86	3,773.54	6,672.60
2	1,429.44	1,858.30	2,477.57	3,086.26	3,034.73	2,364.85	2,003.49	3,536.58
2.5	983.33	1,167.68	1,462.87	1,748.05	1,681.61	1,308.13	1,112.86	1,902.91
3	657.35	737.23	858.36	983.00	961.03	785.74	660.83	1,034.27
3.5	455.99	492.57	540.33	605.99	620.78	502.40	442.75	636.52
4	338.96	334.81	357.85	388.09	433.10	363.17	320.67	425.34
q_{actual}	1.5	3.5	6.5	10.5	20	20	20	20

Table F-12: Cost values (in $) obtained for both approaches, when p = 20 variables and M = 20 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	Randomness Approach				Naïve Approach			
	q = 2	q = 6	q = 12	q = 20	q = 2	q = 6	q = 12	q = 20
0.5	3,204.80	5,094.16	7,777.18	11,051.24	14,382.68	12,647.73	10,992.34	16,635.52
1	2,747.50	4,155.66	6,010.15	8,342.82	9,152.36	6,933.15	5,552.90	10,203.30
1.5	2,088.90	3,034.24	4,114.55	5,436.46	5,414.58	3,767.74	2,868.84	5,645.99
2	1,526.35	2,034.35	2,599.32	3,197.04	3,181.39	2,099.80	1,593.47	2,978.68
2.5	1,094.42	1,281.16	1,562.15	1,814.92	1,873.74	1,236.79	917.20	1,575.87
3	777.79	848.95	940.88	1,046.85	1,154.58	786.65	586.36	916.96
3.5	547.45	564.86	601.53	627.25	761.04	523.79	410.48	565.22
4	396.27	389.44	389.52	405.65	542.91	387.44	315.22	391.08
q_{actual}	**1.5**	**3.5**	**6.5**	**10.5**	**20**	**20**	**20**	**20**

Table F-13: Cost values (in $) obtained for both approaches, when p = 50 variables and M = 25 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	Randomness Approach				Naïve Approach			
	q = 5	q = 15	q = 30	q = 50	q = 5	q = 15	q = 30	q = 50
0.5	4,867.50	9,553.75	16,694.27	26,018.58	39,065.91	34,595.76	27,731.55	46,785.60
1	4,424.45	8,539.04	14,766.25	21,973.98	27,251.07	21,921.33	16,866.18	39,281.53
1.5	3,862.48	7,022.06	11,523.00	16,832.41	18,422.28	13,527.52	9,913.93	28,887.23
2	3,101.95	5,513.95	8,483.76	11,699.80	11,708.66	8,353.91	6,057.61	20,238.64
2.5	2,419.62	3,906.37	5,801.12	8,009.12	7,204.58	5,065.56	3,648.56	12,872.68
3	1,864.26	2,800.09	3,847.29	4,984.68	4,413.48	3,170.76	2,321.47	8,095.34
3.5	1,375.39	1,950.06	2,560.77	3,088.83	2,902.02	2,086.11	1,588.42	4,955.21
4	1,036.79	1,369.36	1,727.43	2,028.64	1,930.95	1,430.73	1,128.07	3,081.52
q_{actual}	**3**	**8**	**15.5**	**25.5**	**50**	**50**	**50**	**50**

Table F-14: Cost values (in $) obtained for both approaches, when p = 50 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	Randomness Approach				Naïve Approach			
	q = 5	q = 15	q = 30	q = 50	q = 5	q = 15	q = 30	q = 50
0.5	4,842.37	9,527.57	16,781.68	25,666.36	35,457.17	30,730.03	23,436.43	45,636.58
1	4,511.71	8,538.18	14,572.23	21,957.84	23,611.72	17,757.25	12,950.66	36,388.46
1.5	3,863.15	7,194.54	11,485.42	16,768.53	15,611.47	10,563.86	7,198.86	25,458.63
2	3,196.78	5,595.06	8,648.05	11,721.61	10,265.01	6,288.91	4,217.26	16,612.65
2.5	2,579.54	4,208.12	5,914.68	7,963.26	6,666.43	3,992.82	2,626.24	10,258.64
3	2,017.95	3,015.06	4,115.82	5,087.34	4,454.14	2,603.88	1,729.86	6,169.78
3.5	1,544.17	2,131.74	2,711.28	3,178.77	2,976.59	1,810.85	1,235.71	3,780.46
4	1,188.32	1,526.28	1,812.39	2,025.23	2,120.28	1,320.91	959.61	2,398.09
q_{actual}	**3**	**8**	**15.5**	**25.5**	**50**	**50**	**50**	**50**

Table F-15: Cost values (in $) obtained for both approaches, when p = 100 variables and M = 50 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0.5	7,445.27	17,073.15	31,299.63	49,885.13	79,815.04	70,592.55	67,486.50	96,853.32
1	6,997.47	16,186.99	28,276.35	44,964.33	58,376.52	47,947.79	41,959.72	87,731.36
1.5	6,517.68	14,132.03	24,529.98	38,743.97	41,057.69	31,111.62	26,410.03	74,950.96
2	5,702.67	11,939.86	20,475.01	30,041.83	29,343.04	20,424.72	16,673.72	61,621.14
2.5	4,801.32	9,761.85	15,635.28	22,779.33	19,909.52	13,353.26	10,526.74	48,107.42
3	3,999.57	7,461.26	11,753.68	16,181.73	13,536.35	9,082.27	7,052.93	36,217.58
3.5	3,227.03	5,663.22	8,555.68	11,309.97	9,330.89	6,247.94	4,796.28	26,994.13
4	2,569.98	4,352.50	5,979.08	7,825.88	6,446.79	4,367.06	3,364.20	19,088.83
q_{actual}	**5.5**	**15.5**	**30.5**	**50.5**	**100**	**100**	**100**	**100**

Table F-16: Cost values (in $) obtained for both approaches, when p = 100 variables and M = 100 variables across different shift sizes for a range of *q* values, using sampling Plan 2.

Shift Size	**Randomness Approach**				**Naïve Approach**			
	q = 10	**q = 30**	**q = 60**	**q = 100**	**q = 10**	**q = 30**	**q = 60**	**q = 100**
0.5	7,436.46	17,386.82	31,919.65	49,834.47	72,673.30	62,278.02	57,615.76	97,386.37
1	6,977.62	16,027.90	28,663.80	44,847.70	50,141.04	38,313.22	31,873.10	86,967.12
1.5	6,468.86	14,063.83	25,124.62	38,292.37	34,989.15	22,880.22	18,236.56	71,402.84
2	5,707.12	12,274.01	20,685.77	30,273.86	23,831.77	14,583.43	10,898.07	57,830.47
2.5	5,095.17	9,944.39	16,095.60	22,657.08	16,471.66	9,570.95	6,933.43	43,387.88
3	4,165.76	7,633.73	11,741.65	16,915.18	11,694.85	6,681.14	4,677.04	32,009.73
3.5	3,524.98	6,098.21	8,595.64	11,718.15	8,323.44	4,622.88	3,307.13	23,060.25
4	2,766.33	4,586.42	6,254.74	7,894.87	5,971.57	3,517.35	2,548.10	15,816.72
q_{actual}	**5.5**	**15.5**	**30.5**	**50.5**	**100**	**100**	**100**	**100**