OAK RIDGE
NATIONAL LABORATORY
MANAGED BY UT-BATTELLE
FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2010/30

Cybersecurity through Real-Time Distributed Control Systems

February 2010

Prepared by

R. A. Kisner
W. W. Manges
L. P. MacIntyre
J. J. Nutaro
J. K. Munro
P. D. Ewing
M. Howlader
P. T. Kuruganti
R. M. Wallace
M. M. Olama

UT-BATTELLE
ORNL-27 (4-00)

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL/TM-2010/30

Measurement Science & Systems Engineering Division

CYBERSECURITY THROUGH REAL-TIME DISTRIBUTED CONTROL SYSTEMS

R. A. Kisner

W. W. Manges
L. P. MacIntyre[*]
J. J. Nutaro[*]

J. K. Munro
P. D. Ewing
M. Howlader

P. T. Kuruganti[*]
R. M. Wallace[†]
M. M. Olama[*]

[*] ORNL Computational Sciences & Engineering Division
[†] ORNL Measurement Science & Systems Engineering Division

Date Published: February 2010

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

ABSTRACT

Critical infrastructure sites and facilities are becoming increasingly dependent on interconnected physical and cyber-based real-time distributed control systems (RTDCSs). A mounting cybersecurity threat results from the nature of these ubiquitous and sometimes unrestrained communications interconnections. Much work is under way in numerous organizations to characterize the cyber threat, determine means to minimize risk, and develop mitigation strategies to address potential consequences. While it seems natural that a simple application of cyber-protection methods derived from corporate business information technology (IT) domain would lead to an acceptable solution, the reality is that the characteristics of RTDCSs make many of those methods inadequate and unsatisfactory or even harmful. A solution lies in developing a defense-in-depth approach that ranges from protection at communications interconnect levels ultimately to the control system's functional characteristics that are designed to maintain control in the face of malicious intrusion. This paper summarizes the nature of RTDCSs from a cybersecurity perspective and discusses issues, vulnerabilities, candidate mitigation approaches, and metrics.

1. INTRODUCTION

Cyber-critical infrastructure is the juncture of control systems and cyber systems. Control systems can be as simple as a self-contained feedback loop, or a very complex, networked system of interdependent, complex, hierarchical control systems with multiple components physically distributed over a wide area (miles, counties, states, or larger). The key word in the prior description is "networked." In its truest sense, the term means "an interconnected or interrelated group of nodes." The consequences of control failure and damage potential are proportional to the systems under direct control. Control systems must perform their critical functions without interruption. Real-Time Distributed Control Systems (RTDCSs) integrate computing and communication capabilities with monitoring and control of entities in the physical world. These systems are usually composed of a set of networked agents, including sensors, actuators, control processing units, and communication devices as described in Fig. 1. While some forms of RTDCSs are already in use, the widespread growth of wireless embedded sensors and actuators is creating several new applications in areas such

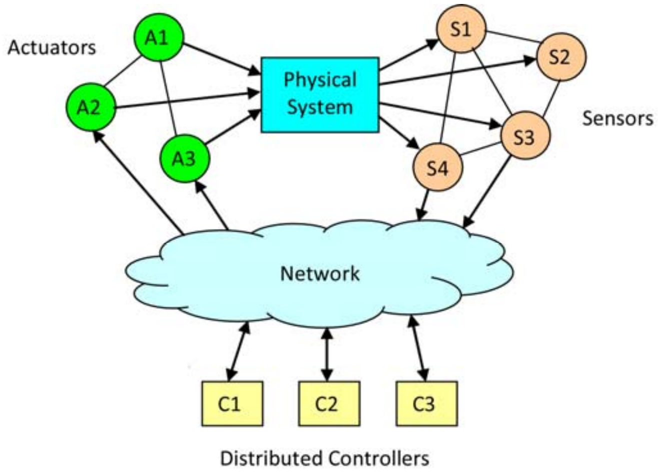

Fig. 1. The general architecture of RTDCSs.[1]

as medical devices, autonomous vehicles, and smart structures, as well as increasing the role of existing ones such as Supervisory Control and Data Acquisition (SCADA) systems.

Currently, RTDCSs are ill prepared for the highly interconnected communications environment that is becoming standard practice. Originally, the systems use was on standalone networks in physically protected locations without threat of subversion. With the use of data collection and control activation systems being set in remote, unattended locations connected to a public or shared network, this exposure allows intrusion if not properly protected from both a perimeter aspect, and more importantly, a resilient component aspect.

Many of RTDCSs are safety critical: their failure can cause irreparable harm to the physical system being controlled and to the people who depend on it. SCADA systems, in particular, perform vital functions in national critical infrastructures, such as electric power distribution, oil and natural gas, water and wastewater distribution systems, and transportation systems. The disruption of these control systems could have a significant impact on public health and safety and lead to large economic losses. While most of the effort for protecting RTDCSs (and SCADA in particular) has been done in reliability (i.e., protection against random failures), there is an urgent growing concern for protection against malicious cyber attacks.[2–34]

Methods derived from a corporate business information technology (IT) domain would lead to an acceptable solution if the physical loss were limited to just data. The reality is that the characteristics of RTDCS make many of those methods inadequate and unsatisfactory or even harmful. A solution lies in developing a defense-in-depth approach ranging from protection of communication interconnect levels to the control system functional characteristics designed to ensure proper control under malicious intrusion or for an analog of fail-safe that includes intrusion tolerant capabilities that ensure critical functionality and survivability. This paper provides a synopsis of the problem domain, a framework for defense in depth, mitigation methods, and metrics that codify RTDCS resilience to intrusion. We conclude that while the current various fields used to solve the problem (using elements from information security, sensor network security, and control theory) can give necessary mechanisms for the security of control systems, these mechanisms alone are not sufficient for the security of RTDCSs.

2. SYNOPSIS OF PROBLEM

Historically control systems are in manned, protected environments and under constant monitoring. Such perimeter isolation, or “fence-and-gate,” views of protection are impractical as control systems are frequently located at unmanned, unmonitored installations. Security of these sites is by a literal fence and lock. Such security is easily subverted by a well informed intruder who can gain physical access undetected and consequently leave such remote systems subject to control by hostile intruders. Extending perimeter security may be impractical, if not impossible. Furthermore, it is entirely possible that a trusted insider can become an adversary,[5] which raises the risk of danger to the greater control system as well as the equipment under its control, or both.

RTDCS have an additional complication of being responsible for operating critical infrastructures and facilities of great economic or strategic value. Examples include electric power distribution, telecommunications, public transportation, water supply and sewage, chemical plants, oil and gas pipelines, and military vessels.

Cyber control is considered fast, accurate, and able to optimize resources (e.g., energy efficiency) and delivery of services while minimizing overall cost. These advantages drive networked implementation. A recent example is the synchrophasor,[6] which captures time-accurate current and voltage (phase) at critical points on the electric grid. Unprecedented knowledge of power flow and stability is obtained from this information. Installation of RTDCS elements in the power system base improves the information from the “Smart Grid” and if designed properly (e.g., attack tolerant) improves the cybersecurity of the conglomerate of networked devices that make-up the Smart Grid.

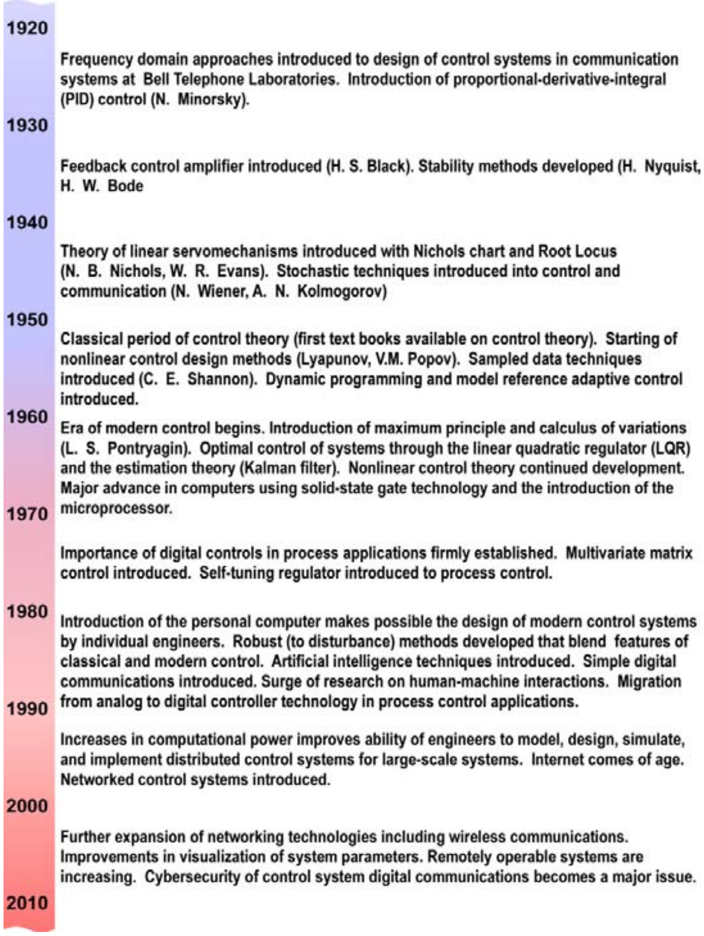

1920

Frequency domain approaches introduced to design of control systems in communication systems at Bell Telephone Laboratories. Introduction of proportional-derivative-integral (PID) control (N. Minorsky).

1930

Feedback control amplifier introduced (H. S. Black). Stability methods developed (H. Nyquist, H. W. Bode

1940

Theory of linear servomechanisms introduced with Nichols chart and Root Locus (N. B. Nichols, W. R. Evans). Stochastic techniques introduced into control and communication (N. Wiener, A. N. Kolmogorov)

1950

Classical period of control theory (first text books available on control theory). Starting of nonlinear control design methods (Lyapunov, V.M. Popov). Sampled data techniques introduced (C. E. Shannon). Dynamic programming and model reference adaptive control introduced.

1960

Era of modern control begins. Introduction of maximum principle and calculus of variations (L. S. Pontryagin). Optimal control of systems through the linear quadratic regulator (LQR) and the estimation theory (Kalman filter). Nonlinear control theory continued development. Major advance in computers using solid-state gate technology and the introduction of the microprocessor.

1970

Importance of digital controls in process applications firmly established. Multivariate matrix control introduced. Self-tuning regulator introduced to process control.

1980

Introduction of the personal computer makes possible the design of modern control systems by individual engineers. Robust (to disturbance) methods developed that blend features of classical and modern control. Artificial intelligence techniques introduced. Simple digital communications introduced. Surge of research on human-machine interactions. Migration from analog to digital controller technology in process control applications.

1990

Increases in computational power improves ability of engineers to model, design, simulate, and implement distributed control systems for large-scale systems. Internet comes of age. Networked control systems introduced.

2000

Further expansion of networking technologies including wireless communications. Improvements in visualization of system parameters. Remotely operable systems are increasing. Cybersecurity of control system digital communications becomes a major issue.

2010

Fig. 1. Timeline of major developments in RTDCS.

Development of the mathematics, hardware components, and communications foundations of today's RTDCSs has evolved over the last century. Figure Fig. 1 shows the history in brief of major control system technological developments. With the addition of cybersecurity as a relatively recent occurrence in the timeline, it adds to the complexity of controls.

The parameter estimation and control algorithms used in RTDCSs are designed to satisfy certain operational goals, such as closed-loop stability, safety, responsiveness, or the optimization of a performance function. Intuitively, our security goal is to protect these operational goals from a malicious party attacking our cyber infrastructure. Security, however, also needs to deal with non-operational goals. For example, if the measurements collected by the sensor network contain sensitive private information, we must ensure that only authorized individuals can obtain this data.

The need to secure high-value national infrastructures against remote, external cyber threats and internal agents is intensifying. The risk of infiltration/compromise opportunities increases as these infrastructures become more dependent on interconnected cyber communications. Current considerations also include protecting against the consequences of unintentional attacks resulting from non-hostile or naive trusted entities (people or devices) in the system. Disruptive events are inevitable as the system becomes more complex and dispersed. The goal of a control system is to ride through the attack without serious financial or productivity cost or loss of human life.

2.1 SUMMARY OF VARIOUS TYPES OF ATTACKS

A general abstraction of RTDCS can be seen in Fig. 2. Let *y* represent the sensor measurements and *u* the control commands sent to the actuators. A controller can usually be divided in two components: an estimation algorithm to track the state of the physical system given *y* and the control algorithm which selects a control command *u* given the current estimate.

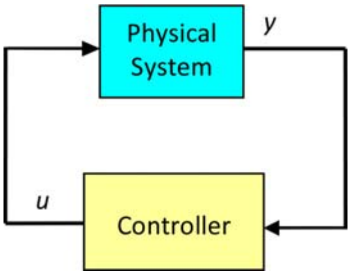

Fig. 2. Abstraction of an RTDCS.

Attacks to an RTDCS can be summarized as follows (see Fig. 3):[7] A1 and A3 represent deception attacks, where the adversary sends false information $\tilde{y} \neq y$ or $\tilde{u} \neq u$ from (one or more) sensors or controllers. The false information can include (1) an incorrect measurement, (2) the incorrect time when the measurement was observed, or (3) the incorrect sender identification (ID). The adversary can launch these attacks by obtaining the secret key or by compromising some sensors (A1) or controllers (A3). A2 and A4 represent denial of service (DoS) attacks, where the adversary prevents the controller from receiving sensor measurements. To launch a DoS, the adversary can jam the communication channels, compromise devices and prevent them from sending data, attack the routing protocols, etc.

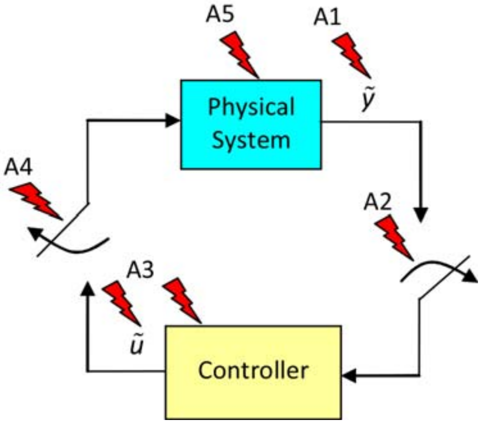

Fig. 3. Various types of cyber attacks.[7]

A5 represents a direct attack against the actuators or an external physical attack on the plant. From an algorithmic perspective, we cannot provide solutions to these attacks (other than detecting them). Therefore, significant efforts must be placed on deterring and preventing the compromise of actuators and other direct attacks against the physical system by, for example, securing the physical system, monitoring cameras, etc. Although these attacks are more devastating, we believe that a risk-averse adversary will launch cyber attacks A1–A4 because it is more difficult to identify and prosecute the perpetrators, it is not physically dangerous for the attacker, and the attacker may not be constrained by geography or distance to the network.

2.2 NATURE OF REAL-TIME DISTRIBUTED CONTROL

A real-time control system is an automatic device that maintains a system within a set of parameters. The real-time control system may range from a simple mechanical device to a complex computer system. Control of any system by a real-time control system is possible. Most control systems currently in use have little or no presence in cyberspace. This is changing due to the increased emphasis on a "Smart-Grid"[8] implementation where any device powered by electricity is connected to a public or shared network.

Control systems have unique properties. The term "real-time" carries connotations that can confuse practitioners as they attempt to discuss their requirements for applications and vendors of hardware and software. The term "real-time" use follows the IEEE definition of the actual time during which a physical process transpires or pertains to the performance of a computation during the actual time of related physical processing in order that results of computation guide the physical process.[9] A real-time control system becomes a distributed one by assembling many controllers to achieve coordinated control of a large-scale system. A RTDCS is a set of computational devices (e.g., sensors, controllers, and actuators) that run several tasks, sequentially or simultaneously, and communicate data across a network, nominally a digital communication network. Figure Fig. 4 shows the main functions of a single controller building block in a distributed system. The primary functions are measuring sensor inputs, computing output values, and sending those values to actuators. The figure details additional functions of distributing data to a corporate business-level system, communicating with other devices in the distributed control system (e.g., other controllers), and interacting with human operators through a control workstation. Threats to mission success come from many sources in which the RTDCS operates such as natural events, random failure, human error, physical intrusion, and cybersecurity breaches.

Typical SCADA configurations have the RTDCS element controllers, as shown in Fig. 5, at the lowest level in a given hierarchy with governance and coordination having both horizontal and vertical components. A subsystem at a given level controls or coordinates those on the level below it and is controlled or coordinated by the unit on the level immediately above it. Information may be passed laterally between subsystems within a level, as shown in Fig. 5.

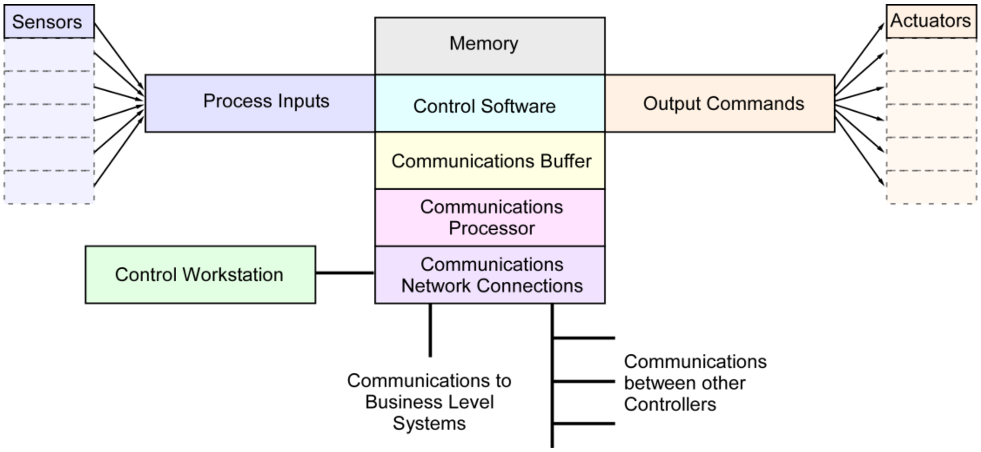

Fig. 4. Distributed control system block elements.

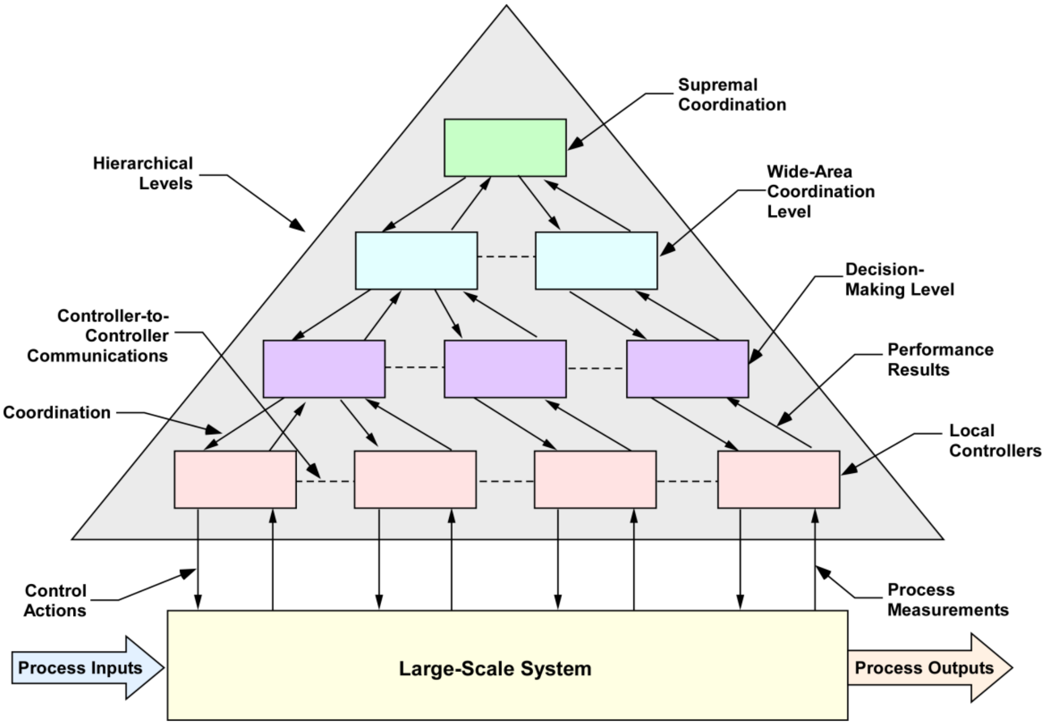

Fig. 5. Hierarchical multilevel control configuration for distributed (large-scale) system.[10]

Two distinct control archetypes (dynamic and discrete-event control) presently exist for RTDCS. Integrating the functions of both constitutes hybrid control.

1. Dynamic (continuous variable): Control of a dynamic system in which the variables such as flow and temperature are continuous. The functional objective is to maintain process variables within a specified uncertainty at a set-point value or on a trajectory. This form of control is typical of feedback control.
2. Discrete event (state): For system control involving discrete states, the functional objective is to place the system in a specific state by a sequence of actions that brings about a planned transition such as initiating specific processes while terminating others.

In Fig. 6, all the electronics and software are collocated with sensors and actuators at the point of control. There is no centrally located supervisory control in this configuration with physically dispersed or collocated components for the controlled process. Each of the local control modules, while responsible for control over their proximal process, possesses a complete image of the entire system and the governance criteria for that system. This architecture is beginning to replace the traditional Purdue Model in Fig. 7 that has been the mainstay of process control architectures for 50 years.[11] It becomes clear that cybersecurity must be addressed through layers of defense to determine the means to minimize risk and develop mitigation strategies. Not shown in the figure is the real-time communication network that interconnects all sense and control modules of the system.

Revisiting the deep-rooted assumptions made in the Purdue Model, the business drivers begin to appear on the factory floor. Data generated on the floor are no longer restricted to the floor. Indeed, more configuration implementations in which data are from lowest level in the hierarchy are turning up in the boardroom and in near real time.

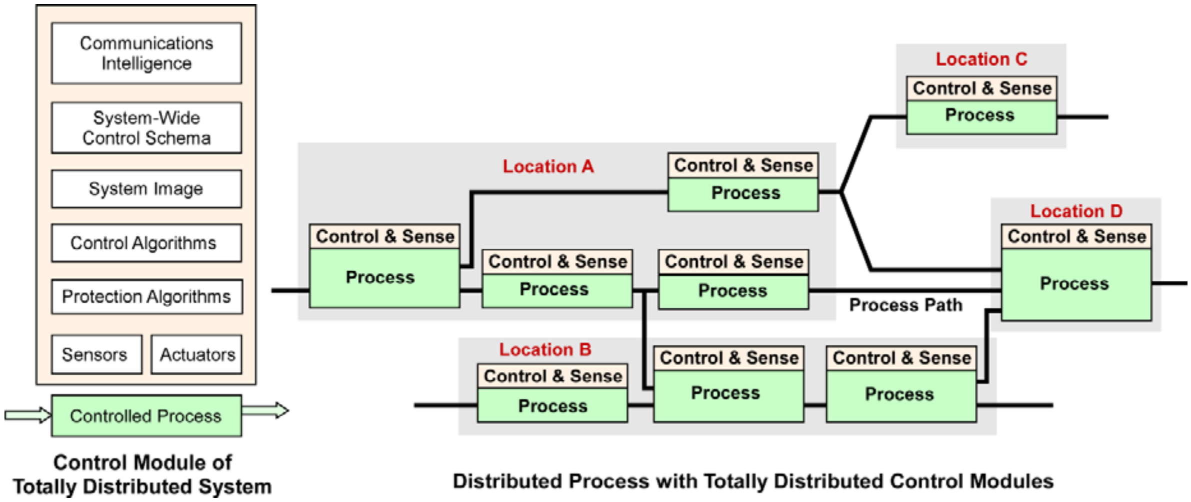

Fig. 6. Fully distributed control system having no (physical) central control supervisor.

Level 5	Enterprise
Level 4	Site Business Planning and Logistics
Level 3	Plant-Wide Operations and Control
Level 2	Area Operations
Level 1	Basic Control Safety Critical
Level 0	Process

Fig. 7. Levels in the Purdue Model levels.

Note that in the Purdue Model, lower level data must pass through intermediate layers for presentation to higher layers. Current business practices (e.g., RF tags) are making this approach obsolete.

2.3 INDUSTRIAL CONTROLS AND INFORMATION TECHNOLOGY

Briefly, data—regardless of origination—is important. The intertwining of financial data, operational data, and instruction for controls results in the need to protect these data equally. The difference between data for controls and data for business lies in the fact that RTDCS control equipment has the potential for death and destruction. At this time, many of these systems are poorly protected from both physical and cyberspace attacks, as they are low-value targets where the effect that might be gained by attacking any one of them is low. Isolation and separation are seemingly adequate protection in combination with the low value assigned to these targets. Connection of these ostensibly low-value targets makes them accessible, with the resulting aggregated target becoming a high-value asset and consequently requires a higher level of protection due to their aggregation.

2.4 ISSUES IN DISTRIBUTED REAL-TIME CONTROL

A defining characteristic of an RTDCS is the need for very frequent transmission of input and output signals. These transmissions may be contained within the components of a feedback loop, or they may extend between modules of separate but interacting loops.

Traditional cyber vulnerability issues for real-time control systems are as shown in Table 1.[12] Addressing these vulnerabilities during the planning, design, installation, and operation phases helps to mitigate their effects. Other reports further catalog cyber vulnerabilities and recommended mitigation approaches as they apply to real-time control systems.[13,14]

Table 1. Control systems vulnerabilities

1. Inadequate policies, procedures, and culture governing control system security
2. Inadequately designed networks with insufficient defense in depth
3. Remote access without appropriate access control
4. Separate auditable administration mechanisms
5. Inadequately secured wireless communication
6. Use of a non-dedicated communications channel for command and control
7. Lack of easy tools to detect/report anomalous activity
8. Installation of inappropriate applications on critical host computers
9. Inadequately scrutinized control system software
10. Unauthenticated command and control data

2.4.1 Reliability, Resiliency, and Security

The concepts of reliability, resiliency, and security constitute trustworthiness, which sustains the major functions of a control system. To achieve system reliability, the components that constitute the control system must exhibit a low failure rate. The resilient aspect adds tolerance to degraded and failed conditions that permits continued performance of critical functions. In the event of significant system failure, a resilient system may even reconfigure process streams and control parameters to meet new functional objectives including establishing new operational priorities such as shutting down low-priority processes in order to direct remaining resources to higher-priority ones.

A control system must meet its control objective despite external and internal disturbances such as noise, component failure, process variation, and communication degradation. A significant list of measures of utility should be considered when developing a control system, as shown in Table 2.[15] The measures include factors related to human operators, stability, communications, and resource use. By satisfaction of these utility measures, the objectives of reliability, resiliency, and security are accomplished, as are the objectives of operability, the ability to be secured, availability, and maintainability.

2.4.2 Known Vulnerabilities

Several pathways may be available for a cyber intrusion, as shown in Table 3. (See Ref. 16 for additional communications vulnerabilities.) The obvious course of action is to block that pathway to unauthorized intrusion. However, following a defense-in-depth philosophy, one must assume that communication blockage was ineffective and subsequent defense mechanisms are required to recognize that an intrusion is in progress and ultimately to know what signals or instructions are reasonable and allowable at the individual controller level.

Table 2. Measures of utility for control systems (adapted from Ref. 15)

1. Compatibility with human operators
   a. Meaningfulness of information
   b. Understandability
2. Real-time quantitative performance and stability
   a. Dynamic performance and stability
   b. Frequency domain characteristics
   c. Static performance (accuracy and precision of results)
3. Reliability of results or conclusions
   a. Opportunity for branching to incorrect path
   b. Repeatability of decision
4. Tolerance to degraded conditions and robustness
   a. Modeling errors
   b. Noise corruption
   c. Process parameter variation
   d. Sensor and actuator failure
5. Interactions with nearby components and subsystems
   a. Actuators
   b. Subsystems
6. Ability to tune in the field
   a. Ability to verify controller tuning
   b. Complexity of tuning process
   c. Disruption of the controlled process
7. Security
   a. Communications
   b. Physical
   c. Code control
8. Resource requirements
   a. Real-time computational requirements
   b. Sensor count, uncertainty, and bandwidth requirements
   c. Communication network requirements
9. Development considerations
   a. Design resources and staff effort
   b. Verification, validation, and testing
10. Long-term considerations
   a. Flexibility to alter and upgradability
   b. Maintainability
   c. Compatibility

Table 3. Realization of possible pathways for cybersecurity intrusions

Failure source (vulnerability)	Example
Change software code to achieve a new control objective	Requires access ability to change code. For firmware implementation, change cannot be made without physically replacing integrated circuit or circuit module. For code stored in rewritable memory, access must be restricted by authentication. Firmware implementation is preferred for critical, high-value assets.
Introduce incorrect (spoofing) input signals	With access to the digital communications network, it becomes possible to trick a controller into incorrect action by generating erroneous process information. Signal types that may be affected are set points, command functions, go/no-go (interlocks), general data transfers, and system status information. Various defensive mechanisms can be applied to protect the controller at several layers including signal validation.
Generate incorrect output values or commands	In a manner similar to incorrect inputs, erroneous output values can be sent to network connected actuators and other controllers. Various defensive mechanisms can be applied to protect against erroneous outputs including command validation.
Insert messages to indicate incorrect operational status of parts of system	Messages can be posted for other controllers or corporate networked computers to read that incorrectly indicate operational status. Such tactics can be used to spoof a maintenance action or force an unnecessary shutdown.

Table 3. (continued)

Failure source (vulnerability)	Example
Collect operational information (data, set points)	Simply by tapping into digital data streams, it is possible to determine operating parameters and states that can be used detrimentally by an adversary as part of a larger, more complex cyber attacks.
Interrupt or corrupt communications between control system components	Rather than directly introducing erroneous process signals for controllers to act on, one can interfere with communications and disrupt stability process. The following error types may be created depending on the network type and configuration: • Corruption • Unintended Repetition • Incorrect Sequence • Loss • Unacceptable Delay • Insertion • Masquerade • Addressing • Broadcast Storm (Denial of Service) • Babbling Idiot (Commission Fault) • Inconsistency (Byzantine Generals' Problem) • Excessive Jitter • Collision

Other failure vulnerabilities exist besides those that are considered as part of cybersecurity, as shown in Table 4. These vulnerabilities are always a core consideration in control systems design. Although these failures do not directly result from cyber attack, it is possible to mask a cyber attack by mimicking a natural failure.

Table 4. Failure sources derived from natural vulnerabilities

Failure source (vulnerability)	Example
Random component failure	Electronic and mechanical components can fail at any time, especially when under stress. Failure rate data is available from many sources.[17] Redundancy plays a crucial role in defense of random failure.
Common cause failure	Component failures result from a single shared cause and coupling mechanism. Component redundancy is negated by common cause failure.[18] Hardware and software diversity is used to mitigate common cause failure.
Latent fault (hardware or software)	The latent fault is representative of a mistake made in the design-fabrication process. Faults could be in hardware or software. A specific trigger makes the fault generate an error and a corresponding failure.[19]
Incorrect input (signal) values	Sensors or the communication chain from sensor to controller can drift or fail to generate incorrect values. Signal and sensor validation as well as redundancy can be used to mitigate this vulnerability.
Incorrect application (software)	Control software may be installed at the factory as firmware or as software. Field changes are sometimes necessary. The firmware implementation prevents changes via network connection. Engineering control methods are used to identify correct application software.
Incorrect operating parameters (set point, alarm limits, etc.)	Engineering and operations staff determine operating points, often in consultation with established standards. These data are subject to change and therefore usually not contained in firmware. Engineering controls are used to track changes.

2.4.3 Control System Failure Modes

A control system can fail according to one of three models.[20] The designer must determine the appropriate model to use according to the process under control.

Fail Arbitrary—control system without any dedicated error detection or fault-tolerance techniques other than memory protection, detection of attempts to execute illegal instructions, and detection of some arithmetic exceptions like overflow and divide by zero.

Fail Silent—control system generates either correct outputs or no commands at all, indicating that it has failed. Methods to achieve fail silent include redundancy (e.g., voted outputs or comparison with diverse controller) and output command validation through independently observing system. The procedure is to place the physical application into a safe state if a disagreement occurs.

Fail Bounded—control system can generate an incorrect output, but its output is constrained to remain within a specified range. A system is thus said to be Fail-Bounded if it (a) generates apparently correct output, (b) stops generating outputs after detecting some error, or (c) generates wrong outputs, but the errors have a boundary defined by the output assertions whose execution is guaranteed.

With these models as a basis, individual, independently operating control loops can fail in one of several modes:

1. off-line condition, in which the controller is dead and the actuators are left unpowered
2. degraded condition, which though functional, provides something less than specification performance
3. erratic output, for which the control output signals are intermittent or bizarre
4. alien, where the controller's mission has been altered and new operating parameters are inserted either through inadvertent or malicious actions

The failure modes of individual loops are likewise applicable to distributed control systems; however, addition modes emerge:

1. failure propagation through coupled process—aberrant control of one system can affect performance and stability of other connected systems through their shared energy or material flows
2. failure propagation through shared signals—controllers whose signals are interconnected or cascaded can pass along erroneous data
3. common cause failure—a common source can reprogram an entire fleet of controllers

Control system failures can exhibit several temporal failure modes.[21] A permanent failure is one in which a control system has failed catastrophically and will remain in a failed condition until major external action is performed (e.g., through repair or replacement). Temporary failures exhibit a failure due to a fault corrected through either prompt human intervention or automatic system action. Transient failures have a failed condition that rapidly returns to normal, perhaps without any external corrective action. Recurring failures are intermittent and may be either random or periodic.

Several models have developed to characterize the attack stages. One such model generally characterizes according to the type of attacker, the method of entry, and the objective of the attack, as illustrated in Fig. 8. The Attack-Vulnerability-Damage (AVD) Model[22] is more detailed and slightly different as shown in Table 5. Understanding the possible stages of an attack and access path allows system designers to better engineer layers of defense.

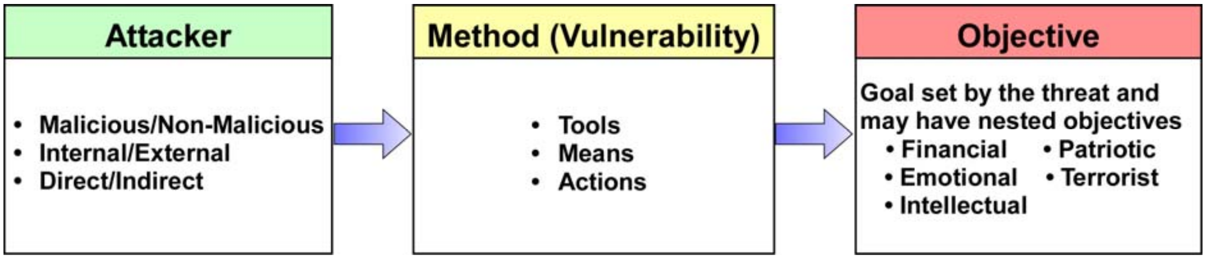

Fig. 8. Basic Cybersecurity Threat Model.

Table 5. Cybersecurity Attack-Vulnerability-Damage Model (adapted from Ref. 22)

Attack			Vulnerability weakness	Damage		
Origin	**Action**	**Target**		**State effect**	**Performance effect**	**Severity**
Local	Probe	Network	Configuration	None	None	None
Remote	Scan	Process	Specification	Availability	Timeliness	Low
	Flood	System	Implementation	Integrity	Precision	Medium
	Authenticate	Data		Confidentiality	Accuracy	High
	Bypass	User				
	Spoof					
	Eavesdrop					
	Misdirect					
	Read/Copy					
	Terminate					
	Execute					
	Modify					
	Delete					

2.4.4 The Effects of Time Jitter

An important aspect of in a failure mode is time jitter, or just jitter, in the communication of real-time process data. Factors contributing to jitter include the network capacity limits and network load. Packet-switched networks potentially have high jitter. Consistency of transmission timing is as important as the transmission timing itself in maintaining system stability.[23,24]

2.5 COMMUNICATIONS VULNERABILITIES

Communicating systems are potentially vulnerable to intrusion. This vulnerability can be due to improper design or to an inadequate implementation of the design.

2.5.1 Wired-versus-Wireless Vulnerabilities

Wired and wireless networks face similar vulnerabilities to attacks but do have differences in the attack vectors. Wired networks can have distant attackers but can have adequate perimeter defenses. Wireless networks—given only wireless access—must be within the range of the attacker’s wireless device and are accessible without a perimeter defense.

2.5.2 Passive-versus-Active Attacks

Violations of the desired security properties typically arise through known attack mechanisms. A taxonomy developed by the National Institute of Standards and Technology (NIST) is segregated into *passive attacks*, which require nothing more than an ability to eavesdrop on wireless communications, and *active attacks*, which require active interference. Passive attacks are difficult to detect, as they involve no alteration or introduction of data. Both passive attacks enumerated in the taxonomy are attacks on confidentiality.

Passive attacks:

Eavesdropping
By passive interception of information transactions, an attacker acquires data. If encryption is used, cracking the encryption and decrypting the traffic counts as a passive eavesdropping attack.

Traffic Analysis
Deduction of certain properties about information transactions based on the participants, duration, timing, bandwidth, and other properties that are difficult to disguise in a packet-encrypted wireless environment allow an attacker to examine a network by observing its transmissions.

Active attacks allow an attacker to be more intrusive. Active attacks include the following:

Masquerade
An attacker fraudulently impersonates an authorized entity to gain access to information resources. A “man-in-the-middle” attack involves a double masquerade—the attacker convinces the sender that she is the authorized recipient, and convinces the recipient that she is the intended sender. Man-in-the-middle attacks on Wi-Fi networks using a counterfeit AP are common. Successful masquerades can compromise all aspects of security.

Replay
An attacker is able to rebroadcast a previous message and elicit a reaction. This reaction either allows the attacker to force the information system into a vulnerable state (e.g., a system reset) or to collect information to enable further attacks (such as WEP encrypted packets). Replays are most directly a compromise of integrity but also compromise authentication, access control, and non-repudiation. Selected replay attacks can also impinge on availability and confidentiality.

Message modification
Modification of transmitted packets by delaying, inserting, reordering, or deleting en-route changes a message. In a wireless network, “man-in-the-middle” attacks are the most direct route to message modification. Message modification is a violation of integrity but can potentially affect all aspects of security.

Denial-of-service
Denial-of-service occurs when an attacker compromises the availability of an information system. In a wireless environment, the most direct routes to DoS are to disable one of the communications partners or to jam the wireless channel itself.

2.5.3 Jamming

Traditionally, the term jamming refers to the disruption of communications systems by the use of intentional electromagnetic interference. Jamming targets to corrupt the desired signals from expected users or to block communications between users by keeping the communications medium busy. Jamming can originate from a single attacker or multiple attackers in coordination and can target a specific user or the entire shared medium. The result is a DoS. DoS attacks can vary from simple to sophisticated. An attacker can send a signal with considerably higher signal strength than the usual signal levels in the system, and then flood the channel so that no user can communicate through it. The more sophisticated way is for the attacker to gain access to the system and violate the network protocol for sending packets, thereby causing many more packet collisions. In the context of electric power grids, jamming can result in a security breach in the form of DoS for communications systems by blocking the on and off activation of remote generating sites or the opening and closing of transmission line switches in response to load demands. In particular, wireless communications systems are more vulnerable to jamming because of their potential for access from covert locations.

3. MITIGATING STRATEGIES

The nuclear industry, and especially Oak Ridge National Laboratory, has built reliable, secure, and fail-safe control systems.[25–,26,27,28,29,30,31,32,33] The methods and techniques developed for these systems are directly applicable to cybersecurity, specifically the following:

1. Authentication
2. Redundancy and diversity
3. Design and Analysis Principles
4. Specification and design of continuously available secondary systems
5. Distributed, federated systems that do not depend on a central system as used in the Purdue Model
6. System recovery of critical functions for fail-safe or “safe mode” end state
7. Robust networked control systems
8. Defense in depth where de-perimeterized protection is distributed throughout the control system

It is an anticipation of unexpected actions for autonomous, continuously available RTDCS that must be the foundation for mitigating strategies of RTDCSs. The processes of intrusion, detection, patch, and reboot results in unacceptable downtime for all concerned. The nuclear industry had no choice but to start fresh to address the safety issues as the consequences of not doing so were deadly. This now must be done with cybersecurity for RTDCS. Regression testing is not the answer as it only anticipates the known and is not economically or practically capable of exhaustively testing for the unknown. A more robust analysis must be done at the very beginning of the control system design process as has been described.

Control system failure compensation and mitigation mechanisms fall into categories of redundant components and functions, independent observers, diversity, customized mitigation systems, and human interaction:

- Redundant fault-tolerant controller configuration with discrete voting logic
- Redundant controllers with weighted output values
- Redundant communication channels
- External observer with dynamic model (signal and command validation)
- External observer with static limit conditions (signal and command validation)
- Diversity in controller configuration and implementation
- Engineered mitigation systems independent and separate from controlled system
- Human operator intervention

Engineering deep defense against cyber attack should involve more than mere protection of the communications network. Table 6 refers to the existing resources from various government agencies on vulnerabilities, best practice guides, and security mechanisms for industrial control systems.

3.1 AUTHENTICATION

An important tool for securing distributed systems is authentication. Authentication schemes prevent humans and devices from impersonating another entity in the system. Access control prevents unauthorized access to the system: it prevents outsiders (unauthenticated principals) from gaining access to the network, while imposing and enforcing proper restrictions on what insiders (authenticated principals) can do. Accountability can be maintained by keeping audit logs of the actions by authenticated entities. Secure communication between two honest entities is achieved

Table 6. List of existing cybersecurity resources

Name	Description
DHS Catalog	Catalog of Control Systems Security: Recommendations for Standards Developers[34]
DHS CS2SAT	Control System Cyber Security Self-Assessment Tool (CS2SAT)[35]
NIST SP 800-82	DRAFT Guide to Industrial Control Systems (ICS) Security[36]
NIST SP 800-30	Risk Management Guide for Information Technology Systems[37]
NIST SP 800-53 Rev. 3	Recommended Security Controls for Federal Information Systems and Organizations[38]
NIST SP 800-94	Guide to Intrusion Detection and Prevention Systems (IDPS)[39]
AMI SEC SSR	AMI System Security Requirements[40]
ISA99	Industrial Automation and Control System Security[41]
ISA100	Wireless standard for industrial automation

with the help of Message authentication codes or digital signatures (they can detect when messages have been tampered by a third party). Message freshness can also be guaranteed by the use of timestamps (which require secure time-synchronization protocols) or by challenge and response mechanisms. Additionally, verification tools and software security can test the correctness of the system design and implementation, thereby limiting the number of vulnerabilities, and the separation of privilege principle is a design guideline to limit the amount to privileges that a corrupted entity can have.

3.2 REDUNDANCY AND DIVERSITY

Redundancy is a means to prevent a single-point of failure. Diversity is a way to prevent that a single attack vector can compromise all the replicas (the added redundancy). Communication channel redundancy for safety actuation signals is a foundational concept of nuclear power plant safety systems. 10 CFR 50 Appendix A[42] Criterion 21 requires that redundancy and independence be designed into the protection system so that (1) no single failure results in loss of the protection function and (2) removal from service of any component or channel does not result in loss of the required minimum redundancy. Oak Ridge National Laboratory recently completed an in-depth study of topics related to diversity and defense-in-depth for the U.S. Nuclear Regulatory Commission.[43] Implementation of channel redundancy is accomplished by voting logic—frequently two out of three. The primary protection that redundant voting systems offer is defense against single component failure.

A secondary benefit of redundant voting, which offers defense against cyber attack, is the requirement that messages must match by the defined majority (e.g., two out of three). An errant message in only one channel (whether arising from malicious or accidental causes) is therefore disregarded using this voting method. An attacker must gain access to at least two communication channels.

A robust implementation of communication redundancy would be multiple, independent network communication channels. A less robust but practical implementation would be a black channel approach—the communications channel carries RTDCS messages but is not itself redundant or safety grade.[44] The black channel makes use of a safety communication layer (SCL) that is present at both black channel end-points. The SCL performs protective transmission functions and checks the communication to ensure that the integrity of the link meets its requirement. Having detected a problem, the SCL corrects it or, failing that, puts the system into a safe state (e.g., by rendering the subsystem to a safe state). To achieve redundancy using black channels, SCLs corresponding to each redundant channel would be required at both sending and receiving ends. In the event that any SCL detected a communications error, there would be an attempt to correct the error (e.g., retransmit a lost message). The final message is derived by voting the redundant (SCL) channels. Failure to arrive at a majority vote signals that an unsafe condition exists, and a safe-state command is issued as determined by the design.

3.3 DESIGN AND ANALYSIS PRINCIPLES

In security research, when we say that a system is secure, we usually mean that the system is secure as long as our adversary model and trust assumptions are satisfied in practice. In general, the adversary model is a way of restricting the scope of the problem. A careful balance must be kept when defining the adversary model. On one hand, restrictive adversary models, such as assuming that an attacker will follow a Bernoulli distribution when performing DoS attacks, will limit the applicability of our analysis (Why would an adversary select such a distribution? What is the incentive?) On the other hand, sometimes these restrictive assumptions are useful to start modeling the adversary, in the hopes of giving us better insights into the nature of the problem, and of how to start obtaining better models in time. As long as the adversary assumptions are explained clearly, we believe that defining a problem with a restrictive adversary is a reasonable first step.

An essential part of security analysis is also in identifying the entities or systems that we trust. Trust is generally defined as accepted dependence;[45] that is, trusted systems are systems we rely on. For example, if in Fig. 3 we do not trust the actuators, there is very little we can do to secure the system. A human, device, or system is trustworthy if we have evidence to believe it can be trusted.

3.4 SPECIFICATION AND DESIGN OF CONTINUOUSLY-AVAILABLE SECONDARY SYSTEMS

An RTDCS is considered trustworthy and goes beyond cybersecurity by including the qualities of reliability and resiliency. The cybersecurity component of trustworthiness has more in common with traditional IT security. Reliability and resiliency issues for RTDCSs are different for office automation and business systems. Security is often measured by how well the control system withstands attempts to interfere with proper operation, such as recording the number of times attempts were made by someone to gain unauthorized access to some part of the control system, and of these attempts, how many were successful.

Reliability is measured by how long the equipment operates correctly before failure and by how long the software works correctly before encountering a defect that disrupts the desired system behavior. Correct transmission of data and commands are part of reliability, so losses in transmission must be detected and monitored. Reliability is also that aspect of trustworthiness most affected by activities of humans: operators, maintenance workers, system and subsystem engineers, and IT staff. Therefore, logging human activity related to the communications and control system operation should be a part of the larger data acquisition activity. This information becomes critical to identifying and diagnosing problems (i.e., the forensics part of metrics).

Configuration control is an important part of attaining and maintaining high reliability. It is also essential to system security because it provides reference points that can be used to detect the appearance of problems such as a virus or inadvertent introduction of a defect to the software or data by comparing the current version that has the problem to a previous version used before the problem appeared.

Applying software patches and upgrades to the control system is part of reliability and security of operation. These software changes often imply measurements that can be made to ensure that the changes improved the system trustworthiness. The same considerations apply to changes made to the control system and plant hardware.

Resiliency has several possible measures including how quickly the system recovers from a disruption, the degree of robustness given an incorrect command or a data entry error (e.g., a set point entry error). For many recovery processes, human operators need to be in the loop. However, an expected trend is automated recovery systems for which human intervention may not be required at least in the short term.

Identification of trustworthiness requirements are at the conceptual design stage of a project, especially when wireless network communications are part of the control system. System architecture should support trustworthiness requirements. Where appropriate, support for trustworthiness requirements should show up clearly in the functional description documentation to write appropriate test plans and performing testing demonstrating that the system meets trustworthiness requirements.

Experience with alarm-handling requirements for digital control systems should provide a useful starting and reference point for doing the kind of analysis of trustworthiness requirements that have not been getting the attention they need. It is by no means sufficient to carry over IT security practices from the areas of office automation and business systems to the area of RTDCSs.

3.5 DISTRIBUTED FEDERATED SYSTEMS

System components may be typically located in open media and may be limited in transmission power and memory. This motivates the need for designing distributed algorithms that can perform a global task with local information exchange and limited computation at nodes. Research in distributed estimation, which falls in the more general area of consensus problems, addresses these problems.[46]

The security of RTDCSs also depends on sensor network security.[47] Most of the efforts for the security of sensor networks have focused on designing a secure communication infrastructure in the presence of malicious insiders. The main results include efficient algorithms for (1) bootstrapping security associations and key management to build a trusted infrastructure,[48] (2) secure communication,[49] and (3) secure routing protocols.[50]

One example of secure communication is anti-jamming techniques, which have existed since the 1950s and implemented mostly with signal processing techniques at the physical layer.[51,52] Other anti-jamming defenses include interference cancellation techniques at the physical layer, orthogonal-multiple-access mitigation techniques at the link level, and secure routing techniques at the network layer.[53,54] Techniques using smart antennas have also been introduced to avoid interference with directional communications transmissions, and these make the communications link more robust against jamming.[55] Recent developments include milestone interference-cancellation techniques based on theoretical bounds as compared to feasibility of implementation.[55,56]

3.6 SYSTEM RECOVERY OF CRITICAL FUNCTIONS

In general, sensors and actuators are vulnerable to random failures. To enable desired operation under failure modes, appropriate redundancies need to be introduced at the design stage. Such techniques also aim at reconfigurable control and graceful performance degradation in the event of failure thus limiting the negative effects that failure can cause. Research in fault tolerant control addresses these issues.[57]

3.7 ROBUST NETWORKED CONTROL SYSTEMS

The architecture of RTDCSs in Fig. 1 indicates a spatially distributed system in which the system, sensors, actuators, and controllers coordinate their operation over a communication network to achieve some performance goal. A typical problem in control theory is to design a control policy to ensure that under the feedback-loop, an open-loop unstable system remains stable. The nature of such systems imposes several constraints on the design of control algorithms. For example, constraints imposed by communication networks such as limited capacity, random delay, packet loss, and intermittent network connectivity can cause DoS. Under DoS the actuator may fail to receive certain packets from the controller that are critical to stabilize an open-loop unstable system. As a result the system may enter a state from which it might be impossible to stabilize it. If the information content of measurement and/or control packets is compromised, it may lead to implementation of incorrect control policies. These factors strongly indicate the need to incorporate network characteristics in the design of control algorithms. Such problems are studied in robust networked control systems.[58]

3.8 DEFENSE IN DEPTH

A system designer must place fencing between the outermost accessible portions of the systems and the inner working parts. With such partitioning, layers of defense against cyber attack are built. Typically, the approach taken concentrates on the communication pathways as shown in Fig. 9. The illustration shows defensive layers for the corporate network on the left and the RTDCS network on the right. The RTDCS network is accessed from deep within the corporate network, which affords some degree of protection.

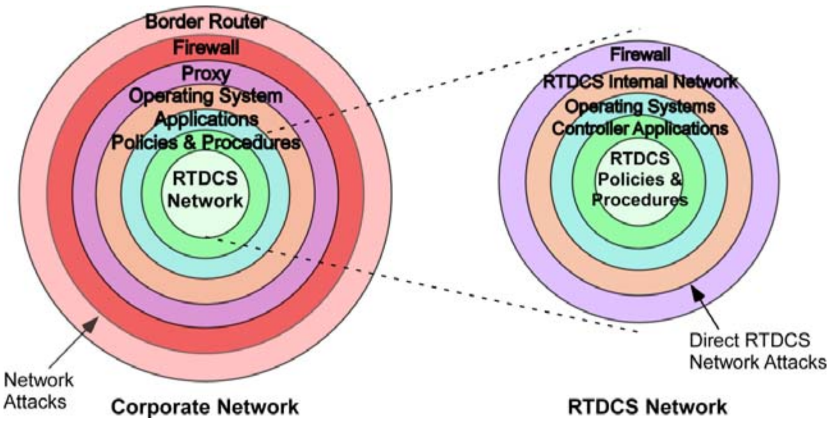

Fig. 9. Traditional concept of layers of defense applied to corporate and RTDCS networks.

The principle of defense in depth can be extended beyond communication networks to include all aspects of the controlled system, as shown in the illustrative example of Fig. 10. The first layer of defense from the system-wide perspective is the underlying physics of the processes followed by the system design and subsystem and component engineering. Protection of communication networks, although important, is not the only defensive mitigation means. Protective functions that are normally placed at the local control level and at regional control levels can include detection of and response to cyber attack objectives.

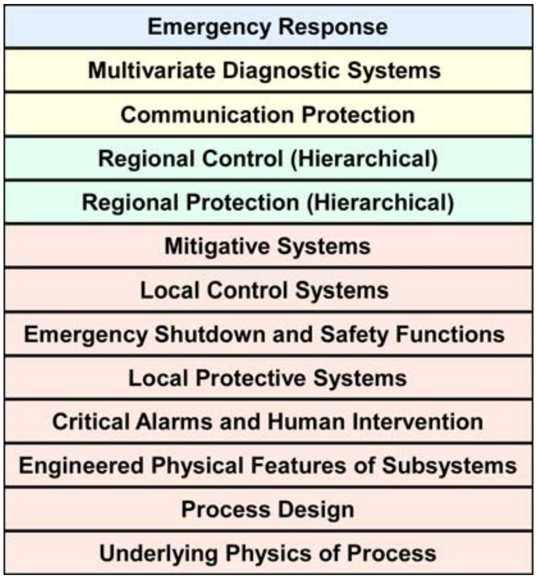
Emergency Response
Multivariate Diagnostic Systems
Communication Protection
Regional Control (Hierarchical)
Regional Protection (Hierarchical)
Mitigative Systems
Local Control Systems
Emergency Shutdown and Safety Functions
Local Protective Systems
Critical Alarms and Human Intervention
Engineered Physical Features of Subsystems
Process Design
Underlying Physics of Process

Fig. 10. Defense in depth for a system under real-time distributed control.

4. SECURITY METRICS

Historically, approaches to security metrics come from several different points of view:

1. as a means to assess how well financial consequences of a security problem are minimized or avoided, that is, the business impact;
2. to indicate how successfully the control system avoids problems that jeopardize desired operation or behavior of the system, that is, to quantify effectiveness of operations;
3. as a measure of how well quality assurance goals are met, that is, how effectively security flaws are detected; and
4. as a means to document how well the control system satisfies/complies with security requirements.

The point of view taken in this discussion is to identify all measurements that will support the goals of trustworthy behavior of the control system: reliability, security, and resiliency—the ability of the control system to recover from a failure or disruption to desired operation. Financial impacts, program planning activities, productivity, and quality assurance goals, though important, are not the primary basis for the kinds of measurements that need to be made.

System behaviors and frequencies of occurrence must be measured and recorded for both real-time monitoring and event reconstruction following a serious or fatal disruption. Times between carefully chosen events must be measured, and the number of times things that happen must be recorded. Carefully chosen measurements will not only provide an indication of how well the system is operating, they will also help to identify where improvements to trustworthiness need to be made.

4.1 METHODS OF MEASURING CYBERSECURITY

There are, in practice, three types of security metrics: designed-based metrics, policy-based metrics, and performance-based metrics.

4.1.1 Design-Based Metrics

Measuring the design and implementation of the system rather than its performance or operation is an up-front method. Two notable documents define design-based metrics.

1. TCSE was created by the US Department of Defense in 1985[59] to evaluate operating systems until the Common Criteria, an international standard, was created in 1999.[60] Both of these documents have gone through a series revisions and updates.[61] The Common Criteria security evaluation follows an international standard, ISO/IEC 15408.
2. DO-178B[62] Software Considerations in Airborne Systems and Equipment Certification was produced by the Radio Technical Commission for Aeronautics Federal Advisory Committee.[63] This document defines five levels (A–E) of critical software; with level "A" being the most critical level and therefore requiring the most effort to achieve compliance.

4.1.2 Performance-Based Metrics

These metrics indicate how often the security system was successful in repelling an attack and conversely how often the security system did not succeed in repelling an attack. These metrics are difficult to employ, as it is not possible to determine with certainty each time that an attack is attempted, and it is even less possible to know with certainty when an attack is successful, because if the attack was detected, it should have been possible to defeat the attempt. The absolute performance-based metric would be to measure the number of vulnerabilities present in a system, but there is always the possibility that unknown vulnerabilities exist, so this metric is of limited real value.

Because performance-based metrics are so difficult to use in practice, common practice is to fall back on policy-based metrics. Policy-based metrics measure such things as the timeliness of keeping defensive systems updated with knowledge of known attack methods and the speed with which alarms from intrusion detection systems receive response. These metrics do not actually measure the security of the system but rather the degree of effort put forth by the IT staff to maintain an adequate defensive posture. Often use is made of policy-based metrics for no other reason than to verify that the IT staff follows the generally accepted best practices for securing a system.

4.1.3 Ideal-Based Metric

New metrics need to be established based on the inabilities of the design-based metric to establish a measure, and performance-based metrics having little success due to their cumbersomeness in practice. The ideal-based metrics are agreements on the attributes of an ideal cybersecurity system and then assessing how closely the considered system approaches the ideal.[64]

Using the known approaches in Sect. 3, MITIGATING STRATEGIES, and ideal-based metrics, one can make a positive statement-of-measure for cybersecurity protection. This is in contrast to the use of *argumentum ad ignorantiam** where just because you have no evidence of a protection breach implies a fully protected RTDCS. This combination of mitigating strategies and a known scalar measurement system for RTDCS is the correct method of establishing the known level of protection.

*Literally, "appeal to ignorance." The fallacy that a proposition is true simply on the basis that it has not been proven false or that it is false simply because it has not been proved true.

5. CONCLUSIONS

RTDCSs have a significant influence on any system. As cybersecurity becomes an issue for RTDCSs, using ideal-based metrics will allow an ability to detect and prevent protection problems associated with the three main aspects of trustworthiness (reliability, resiliency, and security). Closing the loop with security metrics, approached in the manner described in this paper, gives a concrete, reachable goal to count, record, and monitor behavior or operation of the system as desired. The metrics are functioning correctly when the level of trustworthy operation characterized by a set of measurements shows control system protection and wellness with reproducible results.

We also conclude that while the several mitigation strategies discussed in Sect. 3 can give necessary mechanisms for the security of control systems, these mechanisms alone are not sufficient for the security of RTDCSs. In particular, computer security and sensor network security have focused on prevention mechanisms but do not address how a control system can continue to function when under attack. Control systems, on the other hand, have strong results on robust and fault-tolerant algorithms against well-defined uncertainties or faults, but there is very little work accounting for faults caused by a malicious adversary. Therefore, we conclude the paper by outlining some challenges (adapted from Ref. 65) in secure control.

Challenge 1: In the design and analysis of secure control algorithms, we need to introduce a trust analysis of the RTDCS architecture and realistic, rational adversary models that can launch deception or DoS attacks against RTDCSs.

Challenge 2: We must design new proactive algorithms and architectures that are robust against a given adversary model and that provide provable performance bounds (to understand the limits of the resiliency of the algorithms).

Challenge 3: We must design reactive algorithms and architectures for real-time detection and response for a given adversary model.

Challenge 4: In the design of these new algorithms, we need to study how attacks affect the performance of the estimation and control algorithms—and ultimately, how they affect the real world—by incorporating the dynamical models of the systems being monitored and controlled.

6. REFERENCES

1. A. Cardenas, S. Amin, and S. Sastry, “Secure Control: Towards Survivable Cyber-Physical Systems,” pp. 495–500 in *Proceedings of the 28th International Conference on Distributed Computing Systems Workshops*, June 2008.
2. J. Eisenhauer, P. Donnelly, M. Ellis, and M. O’Brien, “Roadmap to Secure Control Systems in the Energy Sector,” Energetics Incorporated, sponsored by the U.S. Department of Energy and the U.S. Department of Homeland Security, January 2006.
3. U. S. G. A. Office, Critical infrastructure protection: Multiple efforts to secure control systems are under way, but challenges remain, Technical Report GAO-07-1036, Report to Congressional Requesters, 2007.
4. R. J. Turk, Cyber incidents involving control systems, Technical Report INL/EXT-05-00671, Idaho National Laboratory, October 2005.
5. L. Lamport, R. Shostak, and M. Pease, “The Byzantine Generals Problem,” *ACM Trans. Programming Languages and Systems* **4**(3), 382–401 (July 1982).
6. J. Sykes, K. Koellner, W. Premerlani, B. Kasztenny, and M. Adamiak, “Synchrophasors: A primer and practical applications,” Power Systems Conference: Advanced Metering, Protection, Control, Communication, and Distributed Resources, 2007. PSC 2007, pp.213–240, March 13–16, 2007.
7. S. Amin, A. Cardenas, and S. Sastry, “Safe and Secure Networked Control Systems Under Denial-of-Service Attacks,” Hybrid Systems: Computation and Control, Lecture Notes in Computer Science. *Springer Berlin/Heidelberg*, **30**, pp. 31–45, April 2009.
8. “The Smart Grid: An Introduction,” prepared for the U.S. Department of Energy by Litos Strategic Communication under contract No. DE-AC26-04NT41817, Subtask 560.01.04, http://www.oe.energy.gov/DocumentsandMedia/DOE_SG_Book_Single_Pages.pdf (checked 9/21/2009).
9. IEEE 100, *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition, IEEE, 2000.
10. M. Jamshidi, *Large-Scale Systems*, Series Volume 9, North-Holland Series in System Science and Engineering, Elsevier Science Publishing, Inc., pp. 103–104, 1983.
11. H. J. Reekie and R. J. McAdam, *A Software Architecture Primer*, Angophora Press, Sydney, Australia, 2006.
12. “Top 10 Vulnerabilities of Control Systems and their Associated Mitigations—2006,” North American Electric Reliability Council, Control Systems Security Working Group, U.S. Department of Energy, National SCADA Test Bed Program, March 16, 2006.
13. K. Stouffer et al., “Guide to Industrial Control Systems (ICS) Security,” National Institute of Standards and Technology, U.S. Dept. of Commerce, Special Publication 800-82, Draft, September 2008.
14. *Common Cybersecurity Vulnerabilities Observed in Control System Assessments by the INL NSTB Program*, INL/EXT-08-13979, Idaho National Laboratory, November 2008.
15. R. Kisner, “A Framework for Selecting Appropriate Control Technologies for Nuclear Power Plant Systems,” pp. 405–418 in *Control—Theory and Advanced Technology*, Vol. 8, No. 3, 1992.
16. R. Kisner et al., Design Practices for Communications and Workstations in Highly Integrated Control Rooms, NUREG/CR-6991, September 2009.
17. *Reliability Prediction of Electronic Equipment*, MIL-HDBK-217F, Notice 2, Feb. 28, 1995.
18. L. Xie et al., “Data Mapping and the Prediction of Common Cause Failure Probability,” *IEEE Trans. on Reliability* **54**(2), June 2005.
19. J. Gray and D. Siewiorek, “High Availability Computer Systems,” *IEEE Computer*, 1991, Draft.

20. J. C. Cunha et al., “A Study of Failure Models in Feedback Control Systems,” The International Conference on Dependable Systems and Networks (DSN), Göteborg, Sweden, 1–4 July 2001.

21. B. E. Ossfeldt, “Maintaining Permanent and Temporary Faults in a Communications, System,” *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 7, October 1986.

22. T. Fleury et al., “Towards a Taxonomy of Attacks Against Energy Control Systems,” Proceedings of the IFIP International Conference on Critical Infrastructure Protection, March 2008.

23. P. Marti et al., “Jitter Compensation for Real-Time Control Systems,” Real-Time Systems Symposium, 2001 (RTSS 2001) Proceedings, 22nd IEEE, Dec. 3–6, 2001.

24. P. Marti et al., “An Integrated Approach to Real-time Distributed Control Systems Over Fieldbuses,” pp. 177–182 in *8th IEEE International Conference on Emerging Technologies and Factory Automation*, 2001 Proceedings, Vol. 1, 2001.

25. R. E. Battle et al., “Reactor Protection System Design Using Application Specific Integrated Circuits,” 2nd Annual ISA/EPRI Joint Controls and Instrumentation Conf., Kansas City, MO, June 1–3,1992.

26. J. K. Munro, Jr., R. A. Kisner, and S. C. Bhatt, “Verification and Validation of Control System Software,” Proc. American Power Conf., Chicago, April 29–May 1, 1991, CONF-9104106-4, 1991.

27. G. V. S. Raju, J. Zhou, and R. A. Kisner, “Fuzzy Logic Controller to a Steam Generator Feedwater Flow,” vol. 2, pp. 1491–92 in Proc. American Control Conf., San Diego, May 23–25,1990, ACC No. 90CH2896-9, 1990.

28. E. Eryurek and B. R. Upadhyaya, “Fault-Tolerant Control and Diagnostics for Large-Scale Systems,” *IEEE Control Systems* **15**(5), 34–43 (October 1995).

29. R. C. Berkan, B. R. Upadhyaya, L. Tsoukalas, R. A. Kisner, and R. L. Bywater, “Advanced Automation Concepts for Large Scale Systems,” *IEEE Control Systems Magazine* **11**(6), 4–12 (October 1991).

30. G. V. S. Raju, J. Zhou, and R. A. Kisner, “Hierarchical Fuzzy Control,” *Int. J. Control* **54**(5), 1201–1216 (November 1991).

31. R. A. Kisner and G. V. S. Raju, *Automating Large-Scale Power Plant Systems: A Perspective and Philosophy*, ORNL/TM-9500, December 1984.

32. H. Basher and J. S. Neal, *Autonomous Control of Nuclear Power Plants*, ORNL/TM-2003/252, October 2003.

33. R. W. Winks, T. L. Wilson, and M. Amick, “B&W PWR Advanced Control System Algorithm Development,” in *Proceedings: Advanced Digital Computers, Controls, and Automation Technologies for Power Plants*, edited by S. C. Bhatt, EPRI TR-100804, Electric Power Research Institute, Palo Alto, California, 1992.

34. http://www.us-cert.gov/control_systems/pdf/Catalog_of_Control_Systems_Security_Recommendations.pdf.

35. http://csrp.inl.gov/Self-Assessment_Tool.html.

36. http://csrc.nist.gov/publications/drafts/800-82/draft_sp800-82-fpd.pdf.

37. http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf.

38. http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/800-53-rev3_final-markup_final-publicdraft-to-final-updt.pdf.

39. http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf.

40. http://osgug.ucaiug.org/utilisec/amisec/Shared Documents/1. System Security Requirements/AMI System Security Requirements - v1_01 - Final.doc.

41. http://www.isa.org/MSTemplate.cfm?MicrositeID=988&CommitteeID=6821.

42. *Code of Federal Regulations*, Title 10 CFR Part 50, Appendix A.

43. R. T. Wood et al., *Diversity Strategies for Nuclear Power Plant Instrumentation and Control Systems*, NUREG/CR-7007, January 2010.

44. EC 62280-1, “IEC 62280-1 Railway applications communication, signaling and processing systems Part 1: Safety-related communication in closed transmission systems,” 1st Ed., 2002.

45. A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, “Basic concepts and taxonomy of dependable and secure computing,” *IEEE Transactions on Dependable and Secure Computing* **1**(1), 1–32 (2004).

46. R. Olfat-Saber, “Distributed Kalman filter with embedded consensus filter,” Proceedings of CDC and ECC, Seville, Spain, 2005.

47. A. Perrig, J. A. Stankovic, and D. Wagner, “Security in wireless sensor networks,” *ACM Journal of Communications* **47**(6), 53–57 (2004).

48. L. Eschenauer and V. Gligor, “A key-management scheme for distributed sensor networks,” pp. 41–47 in Proceedings of the 9th ACM Conference on Computer and Communications Security, 2002.

49. M. Luk, G. Mezzour, A. Perrig, and V. Gligor, “Minisec: A secure sensor network communication architecture,” in Sixth International Conference on Information Processing in Sensor Networks (IPSN 2007), April 2007.

50. C. Karlof and D. Wagner, “Secure routing in sensor networks: Attacks and countermeasures,” pp. 293–315 in Ad Hoc Networks, vol. 1, issues 2–3 (Special Issue on Sensor Network Applications and Protocols), Elsevier, Sept. 2003.

51. J. L. Massey, *Towards an Information Theory of Spread-Spectrum Systems*, Kluwer Academic Publishers, 1995.

52. R. C. Dixon, *Spread Spectrum Systems with Commercial Applications*, John Wiley & Sons, 3rd ed., 1994.

53. S. Verdu, *Multiuser Detection*, Cambridge Press, 1998.

54. A. Mishra, Security and Quality of Service in Ad Hoc Wireless Networks, Cambridge University Press, 2008.

55. P. Patel and J. Holtzman, “Analysis of a simple successive interference cancellation scheme in a DS/CDMA system,” *IEEE Journal on Selected Areas in Communications* **12**, 796–807 (June 1994).

56. A. J. Viterbi, “Very low rate convolutional codes for maximum theoretical performance of spread-spectrum multiple-access channels,” *IEEE Journal on Selected Areas in Communications* **8**, 641–649 (May 1990).

57. M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and fault-tolerant control*, Springer-Verlag, September 2003.

58. L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. S. Sastry, “Foundations of control and estimation over lossy networks,” pp. 163–187 in Proceedings of the IEEE, Special Issue on the Emerging Technology of Networked Control Systems, vol. 95, no. 1, January 2007.

59. “Trusted Computer System Evaluation Criteria,” DoD 5200.28-STD, Department of Defense, December 1985, http://csrc.nist.gov/publications/history/dod85.pdf.

60. *Common Criteria for Information Technology Security Evaluation*—Part 1: Introduction and general model; Part 2: Security function requirements; Part 3: Security assurance requirements, ISO/IEC 15408, Version 2.0, 1999.

61. See http://www.niap-ccevs.org/cc-scheme/cc_docs/ for the international standard ISO/IEC 15408 Common Criteria; for links to “Rainbow Series” for NSA and DOD, see links available at http://www.fas.org/irp/nsa/rainbow.htm. See also http://www.niap-ccevs.org/cc-scheme/nstissam_compusec_1-99.pdf.

62. DO-178B, “Software Considerations in Airborne Systems and Equipment Certification”; see also http://en.wikipedia.org/wiki/DO-178B and references.

63. For current information about RTCA activities, see http://www.rtca.org/.

64. W. Boyer and M. McQueen, “Ideal Based Cyber Security Technical Metrics for Control Systems,” CRITIS’07 2nd International Workshop on Critical Information Infrastructures Security, October 3–5, 2007.

65. A. Cardenas, S. Amin, and S. Sastry, “Research Challenges for the Security of Control Systems,” Proceedings of the 3rd USENIX Workshop on Hot topics in security, USENIX, Article 6, 25 July 2008.

APPENDIX A

APPENDIX A. CONSENSUS PRACTICES

A.1   WIRELESS CONSENSUS PRACTICES

Many wireless technologies are currently in use in industrial environments, including utility control rooms. Wireless applications vary in range (e.g., short-range wireless sensor networks, local-area data acquisition systems, long-range distributed control systems). Consequently, wireless networks are typically defined by their nominal transmission distances, with wireless personal area networks (WPANs) operating over a coverage area of a few tens of meters, wireless local area networks (WLANs) operating over a coverage area of hundreds of meters, wireless metropolitan area networks (WMANs) covering several kilometers, and wireless wide area networks (WWANs) covering hundreds of kilometers. Much of the success of wireless networks can be directly attributed to the successful development and adoption of the Institute of Electrical and Electronics Engineers (IEEE) 802 standards. Figure A.1 lustrates the relationship of the IEEE 802 wireless standards and their associated technologies.

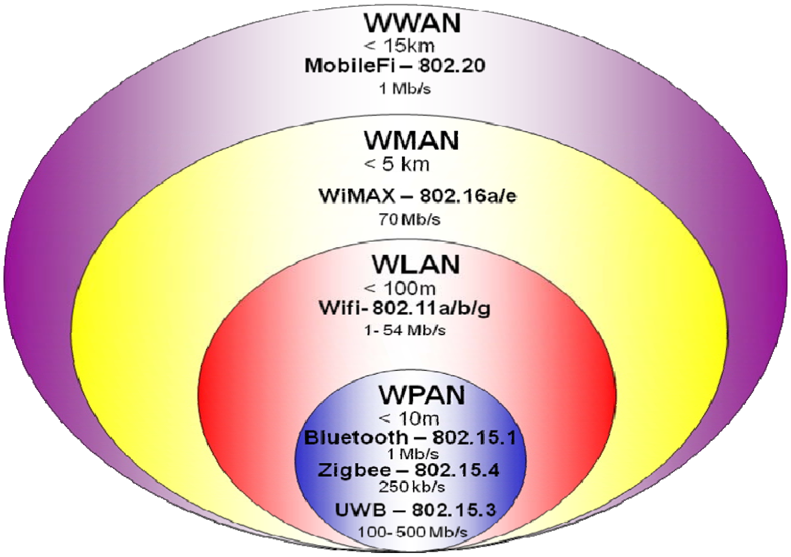

Fig. A.1. Wireless standards and associated technologies.

WPANs are covered by the IEEE 802.15 series of standards and include the Bluetooth (IEEE 802.15.1), ZigBee (802.15.4), and UltraWideband (UWB) (802.15.3) technologies. Bluetooth is a technology that was developed for short-range cable replacement. Table A.1 describes the various 802.15 networks. A consortium of companies with similar needs, known as the Bluetooth Special Interest Group,* collaborated and decided to generate a new and universal mode for which data transfer could be accomplished without wires and without sacrificing the speed of the transfer. The cornerstone for Bluetooth-compliant devices to date has been their ability to communicate with a personal computer. Bluetooth products include keyboards, mice, printers, and devices that can be used in conjunction with computers, such as personal digital assistants and cell phones. Bluetooth has a data rate of 1 Mb/s and operates in the 2.4 GHz Industrial, Scientific, and Medical (ISM) frequency band.

*www.bluetooth.org

Table A.1. 802.15 network descriptions

Name	Description
802.15.1	WPAN based on Bluetooth (1.1 and 1.2)
802.15.2	Co-existence of WPAN with other systems in the 2.4 GHz band
802.15.3	High rate WPAN (11-54 Mbps)
802.15.4	Low-rate WPAN (12-250 Kbps)
ISA100	Wireless standard for industrial automation
Wireless HART	Wireless protocol for process measurement and control

A.2 APPROACHES

When considering cybersecurity for RTDCSs, concerns for confidentiality, integrity, and availability are not of primary importance; violations of these properties are significant because, and only if, they can adversely affect the operation of the manufacturing facility, electrical generator, or whatever machinery is under the system's control. The distinguishing characteristic of security for control systems is the exclusive focus on protecting a dynamic process. Other concerns for the theft, loss, or manipulation of data that do not impinge directly on the dynamic process are not within the scope of cybersecurity for control systems.

A security analysis that examines a control process should consider three facets of operation:

1. *Safety*—Can a malfunctioning or compromised network, computer, or security asset cause loss of life, limb, or the destruction of critical machinery?
2. *Operational availability*—Can a malfunctioning or compromised network or computer prevent the control system from operating as intended?
3. *Operational risk*—Can a security measure, by introducing an unknown or unquantifiable dynamic or by requiring resources beyond the capabilities of the computers and networks responsible for operations, prevent the control system from operating as intended?

This perspective creates tension between security controls imposed for the sake of the dynamic process and security controls imposed for the protection of information. In designing the former, simplicity is the first concern: complicated security controls may impose greater risks to the controlled process by introducing new, unquantifiable dynamics. The capability of a compromised asset to disrupt operations must be weighed against the additional complexity, and therefore risk, of ensuring that the asset is not compromised.

Simple, automatic controls such as dead-man switches, fuses, pressure-relief valves, and other devices will be considered first; these are often the most reliable protections against catastrophic failure. Supervisory controls that locally (i.e., without network access or within a wholly isolated network) monitor operating conditions to ensure safety are next. More traditional security measures are considered last because they are the most difficult to quantify. How does encryption interact with a task schedule? How does key distribution affect delivery guarantees in the control network? Will an operating system patch disrupt operations? Although difficult, answers are needed to determine whether the control system will work correctly alongside more advanced security mechanisms; that is, they prevent malicious misuse but, potentially, at the cost of increased system unreliability.

Consequently, where process control data must be protected for business reasons, a tradeoff must be made; is additional protection of the data in the process control system justified by the risks that the protection mechanisms pose to its proper operation? The traditional concerns for confidentiality, availability, and integrity must therefore be complemented by concerns for the physical quantities that are monitored and manipulated by the computer assets, that is, for safety and for operational availability. For a security analysis focused on maintaining proper operation of the control system, the security risk associated with an asset should be based on two criteria:

1. The potential for the asset to affect one or more of the physical quantities under consideration
2. The potential impacts of disturbances of these quantities by a compromised or malfunctioning asset

With an understanding of these effects, it is possible to derive security controls by working backwards from the physical asset at risk in three steps. First, what can be reasonably done to mitigate dangerous action by, for instance, installing supervisory controllers, automatic protection systems, and other non-networked mechanisms? These failsafe devices are essential because they are the most reliable method for preventing catastrophic failures due computer errors and security breaches.

Second, what computer-based protections (digital signatures, encryption, intrusion detection, etc.) can be reasonably installed on the controllers or in the control network to prevent illegal operations? At this point, safety should have been ensured by protections in Step 1; networked control which is exposed to computer attacks should not, if at all possible, be responsible for safe operations. Rather, ensuring availability is the goal of Steps 2 and 3; the security risks averted in these steps should be carefully weighed against the potential risks to control system performance that are incurred by its addition.

Third, what other security controls can be reasonably enacted to prevent errors, accidents, and intentional maliciousness from disrupting operations? Physical protection, personnel screenings, etc., might be considered in this third stage as preventions against insider threats and physical, external hazards.

ZigBee is a collection of major corporations committed to standardizing cost-effective, low-power, wirelessly networked monitoring and control products based on an open global standard. IEEE 802.15.4 is supported by the ZigBee Alliance* and targets applications that do not need high data speeds or share large amounts of data. In return, ZigBee devices do not consume large amounts of power. ZigBee devices operate in the 2.4 GHz ISM frequency band at a data rate of 250 kb/s. ISA100 and Wireless Hart are two emerging standards for wireless sensors based on the IEEE 802.15.4 radio. ISA100, developed by ISA,† a leading global nonprofit organization of industrial automation professionals, allows the deployment of a single integrated wireless infrastructure platform that can simultaneously communicate over existing application protocols (e.g., HART Foundation Fieldbus, Modbus, and Profibus). Wireless HART combines the well-established HART communication protocol with IEEE 802.15.4 radios and is supported by the HART Communication Foundation,‡ an independent not-for-profit organization providing worldwide support for the HART technology.

IEEE 802.15.3 uses UWB technology for low-cost, low-power, high-speed wireless multimedia applications for portable consumer electronic devices. These applications include wireless connections to surround-sound speakers, portable video displays, flat panel displays, digital video cameras, and digital still cameras. UWB devices also operate in the 2.4 GHz ISM frequency band but at data rates from 100 to 500 Mb/s. The benefits of WPAN include ubiquitous sensing and enhanced process visibility. Denial of service remains the biggest risk or concern for these low-power devices. With careful implementation, the devices can respond to a denial of service attack by self-locating interference sources and rerouting messages through mesh networking.

WLANs are covered by the IEEE 802.11 series of standards. They are typically called the Wireless Fidelity (Wi-Fi) standards and are supported by the Wi-Fi Alliance.§ Table A.2 describes the various 802.11 networks. Three of the Wi-Fi standards are enjoying widespread use today: 802.11a, 802.11b, and 802.11g. The most prominent of the three IEEE 802.11 protocols is IEEE

*See www.zigbee.org

†See www.isa.org

‡www.hartcomm.org

§www.wi-fi.org

Table A.2. Amendments to IEEE 802.11

Number	Description
802.11a	Phy. layer for the 5 GHz ISM band 6–54 Mbps
802.11b	Phy. layer for the 2.4 GHz ISM band, 5.5 and 11 Mbps
802.11c	Supplement to support MAC bridge operation
802.11d	Specification for operation in different regulatory domains
802.11e	Enhancements for Quality of Service (QoS)
802.11f	Inter access point protocol
802.11g	Phy. layer for operation in 2.4 GHz band (OFDM)
802.11h	Spectrum and power management operations to 802.11a
802.11i	Security enhancements
802.11j	Enhancement to 802.11a for operation in 4.9–5.0 GHz in Japan
802.11k	Radio resource management
802.11m	Technical corrections and classifications
802.11n	High-throughput enhancement (OFDM, MIMO)

Abbreviations:
MAC: Media Access Control
OFDM: Orthogonal Frequency Division Multiplexing
MIMO: Multiple Input Multiple Output

802.11b, which has been successfully deployed in business offices, university buildings, and homes around the world for many years. IEEE 802.11b can transmit data at rates up to 11 Mb/s and operates in the ISM frequency band at 2.4 GHz. IEEE 802.11a offers a fivefold increase in data rate over IEEE 802.11b by transmitting up to 54 Mb/s. To increase its output bit rate, IEEE 802.11a takes advantage of the 300 MHz of bandwidth available in the 5 GHz Unlicensed National Information Infrastructure (UNII) band. IEEE 802.11g is the most recent standard, and products have been appearing in the marketplace for the last few years. It is capable of maintaining IEEE 802.11a-type data rates up to 54 Mb/s and is essentially a version of 802.11a (with slight differences) placed in the 2.4 GHz ISM band.

In the industrial environment, Wi-Fi networks are regularly used for sensor data acquisition, Internet connectivity, and enterprise-wide connectivity. All laptops used in the field or within control centers are likely equipped with any or all of the WLAN types. While providing mobile/ instantaneous Internet access for authorized users within the facility, using WLAN technology poses the biggest risk for unauthorized access to the enterprise or control center networks. A benefit of WLAN technology is rapid Internet connectivity for non-stationary authorized users (e.g., field engineers assembled in control centers during a crisis). The risk includes the potential for unauthorized access to a control center’s enterprise network and possible access into the ESP. Careful implementation of defense-in-depth is required to separate authorized stationary users, authorized non-stationary users, and unauthorized users to reduce the risk of a wireless attack on a control center network.

The IEEE 802.16 standards enable the development of WMANs by incorporating broadband wireless access technology. This technology is typically referred to as Worldwide Interoperability for Microwave Access (WiMAX). The proliferation of WLAN hotspots based on the IEEE 802.11 standards is driving the demand for broadband connectivity back to the Internet, with the term “broadband” simply meaning that the wireless system is capable of delivering a transmission rate greater than 1.5 Mb/s. Originally, the WMAN was intended to be a fixed wireless access system capable of providing the desired last-mile broadband access. WMAN has since developed into broadband access for hard-to-reach areas for wired infrastructure or where high installation costs make broadband access prohibitive. The IEEE 802.16 standards now include both fixed and mobile wireless broadband technology and are supported by the WiMAX Forum.* IEEE 802.16a addresses fixed non-line-of-sight point-to-multipoint transmissions in the 2 to 11 GHz band, and IEEE 802.16e addresses portable applications in the 2 to 6 GHz band. Looking toward the future, an emerging IEEE 802.20 working group has been tasked with developing standards for mobile broadband wireless systems designed to be used in WWANs that cover hundreds of kilometers.

*www.wimaxforum.org