Cyber-safety Analysis of an Industrial Control System for Chillers using STPA-Sec

Shaharyar Khan, Stuart Madnick, and Allen Moulton

Working Paper CISL# 2018-06

July 2018

Cybersecurity Interdisciplinary Systems Laboratory (CISL)
Sloan School of Management, Room E62-422
Massachusetts Institute of Technology
Cambridge, MA 02142

Working Paper – July 2018

Cyber-safety Analysis of an Industrial Control System for Chillers using STPA-Sec

Shaharyar Khan, Stuart Madnick, Allen Moulton

***Abstract*-- As Industrial Control Systems (ICS) become increasingly software-intensive and more complex, the traditional approaches to cybersecurity that undertake a narrow, static technical view of the system are proving to be increasingly inept in the face of new threat vectors and vulnerabilities. To date, most attacks on Energy Systems have targeted either the IT infrastructure (e.g., the Aramco Shamoo attack) or Circuit breakers of Operational Technology (e.g., the Ukraine attack.). In such cases, recovery is usually rather fast – either by rebooting computers or resetting breakers. But, if the Operation Technology equipment, especially the important, large, customized equipment, is physically damaged, recovery can take weeks or even months. In this paper, we demonstrate the use of *Systems-Theoretic Process Analysis (STPA)* to identify cyber vulnerabilities that have the potential to cause physical damage in industrial control systems using the MIT Central Utilities Plant as a use-case. It is shown that the method provides a well-guided and structured analysis process to unveil new cyber vulnerabilities that span not only technical aspects but also the broader socio-organizational system. The method ties system-level losses to violation of constraints at both the component-level as well as the process level and provides recommendations to make the system more resilient by defining additional constraints to control vulnerabilities in the system.**

I. INTRODUCTION

Cyber-Physical systems are electronic control systems that control physical processes and machines such as, motors and valves, in an industrial plant using Information and Communication Technologies (ICT). They can be thought of as the central nervous system of a plant that enable monitoring and control of all operations of a plant. The advances in computing power and network transmission speeds, coupled with a decrease in hardware cost, has enabled new applications of ICT in industrial settings to improve efficiency of the underlying physical processes. The resulting displacement of traditional analog and mechanical devices with complex, software-intensive Industrial Control Systems (ICS), has inadvertently intertwined the architecture of physical processes with cyberspace; thus, exposing them to new threat vectors and vulnerabilities.

ICSs monitor and control industrial processes across a wide spectrum of industries; from critical infrastructures such as electric grids, nuclear power plants, gas and water distribution pipelines and oil refineries to standalone cogeneration power plants and Building Management Systems (BMS) in hospitals, universities, malls and commercial buildings. Despite the diversity of scale and application across industries, their system architecture is fairly identical. Typically, these control systems rely on sensors, limit switches and measuring devices to acquire data from controlled processes, which is then fed back to Programmable Logic Controllers (PLC) in conjunction with some kind of a Supervisory Control and Data Acquisition (SCADA) system, to control the physical processes through actuators, motors and valves.

While security (of data) has been a primary concern for traditional Information Technology (IT) systems since their inception, it is a rather recent phenomenon for ICSs; the traditional top priority for ICS being the *reliability* and *availability* of physical devices. This lack of urgency or attention to security risks exposes

ICS to potential cyberattacks that can cause actual physical damage or disruption of critical infrastructure or services. The 2009 Stuxnet cyberattack that partially destroyed a third of the centrifuges at a uranium enrichment facility in Natanz, Iran, demonstrated the unprecedented capabilities of such attacks on ICS, ushering a new era in cyber warfare.

Current approaches to examining cybersecurity of cyber-physical systems are often based on analysis of ICT protocols or network configurations; they undertake a narrow technical view that is biased by information security concerns [7]. In reality, security (and by extension cybersecurity), like safety, is an *emergent* property of a system where the interactions of simple components produce complex behaviors which cannot be predicted by linearly analyzing individual components. Instead, a top-down, *systems thinking* approach is required that examines not only the components on their own but also holistically considers the functional interactions between components, people and management as a whole.

System Theoretic Accident Model & Processes (STAMP) is an accident causality model originally developed to address *safety* of complex systems. The actual method based on the STAMP accident model is called System Theoretic Process Analysis (STPA). Young and Levenson [2] adapted the STPA method to cybersecurity; the new method is called STPA-Sec. In this paper, we perform a limited STPA-Sec analysis on the MIT Central Utilities Plant (CUP) to demonstrate the use of the method for an archetypal industrial control system.

In this paper, the centrifugal chiller subsystem is selected over other subsystems (such as the electricity distribution subsystem) due to easy accessibility to information about the chiller design and processes. The chiller also provides an illustrative example of a modern-day, software intensive, cyber-physical system that is small enough in scope, to enable a thorough STPA-Sec analysis in a limited timeframe. In this paper, a robust analysis of a single control loop (the chiller capacity control loop) is used to illustrate how a cyberattack at a component level can propagate into system-level losses (such as catastrophic failure of compressor gear assembly or compressor motor burnout etc.) and how such losses can be mitigated through implementation of control measures both at the component level as well as at the organizational level.

II. BACKGROUND

The MIT Central Utilities Plant (CUP) contains a 21 MW Siemens ABB (GT10) gas turbine generator that provides electricity to the MIT campus. Waste heat from the turbine is directed to a Heat Recovery Steam Generator (HRSG) to produce steam along with other gas/oil-fired water-tube boilers. The steam is used for campus heating as well as for driving steam-driven chillers. The plant operates sixteen (16) steam and electric-driven chillers to provide chilled water and air conditioning to the campus buildings.

The gas turbine meets about 60 percent of the campus electricity demand; the plant is connected to the local electric grid and its generation capacity is throttled to most economically supply power based on fluctuating electricity and natural gas prices [1, 3]. The plant has been designed to provide near 100 percent reliability through maintaining standby units at all times, as the steam, chilled water and electrical power generated is used to maintain critical research facilities, laboratories, classrooms and dormitories [3]. Figure 1 shows the energy flow diagram of the MIT CUP.

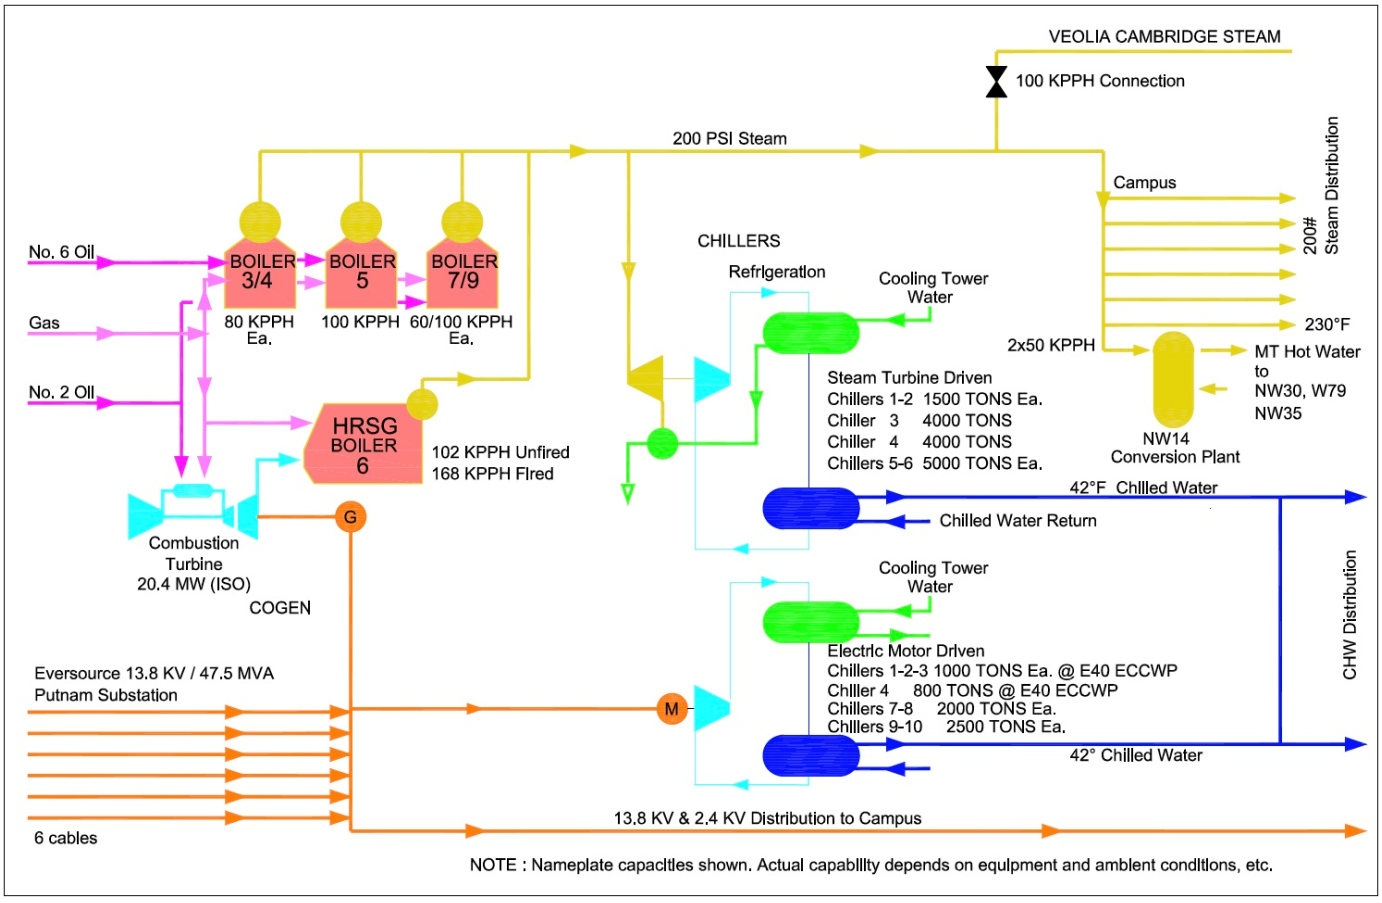

Figure 1 - Energy Flow Diagram for MIT Central Utilities Plant

The plant equipment, including boilers, chillers, gas turbine and other ancillary equipment, is monitored and controlled via a Distributed Control System (Emerson Ovation DCS) by operators who man the station 24/7. A DCS is integrated as a control architecture containing a supervisory level of control that consists of geographically distributed control elements over the control area. Process control is achieved by deploying feedback or feedforward control loops whereby key process conditions are automatically maintained around a desired set point using Programmable Logic Controllers (PLCs) [8].

It differs from centralized control system wherein a single controller at a central location handles the control function; in DCS each process element or machine or group of machines is controlled by a dedicated controller. DCS consists of a large number of local controllers in various sections of plant control area that are connected via a high-speed communication network. Furthermore, there is a subtle distinction between SCADA and DCS in that DCS's typically control equipment in the same geographical location.

DCS uses MODBUS TCP/IP protocols to communicate with the various PLCs, I/O modules and gateways; MODBUS is a serial communications protocol for PLCs which is the de facto standard communication protocol for connecting industrial control devices. A high-level system architecture implementation for a generic DCS is shown in Figure 2; it shows several controllers (PLCs, machine controllers, process controllers etc.) connected to an integrated supervisory control system annotated as the 'Main HMI (Human Machine Interface)'.

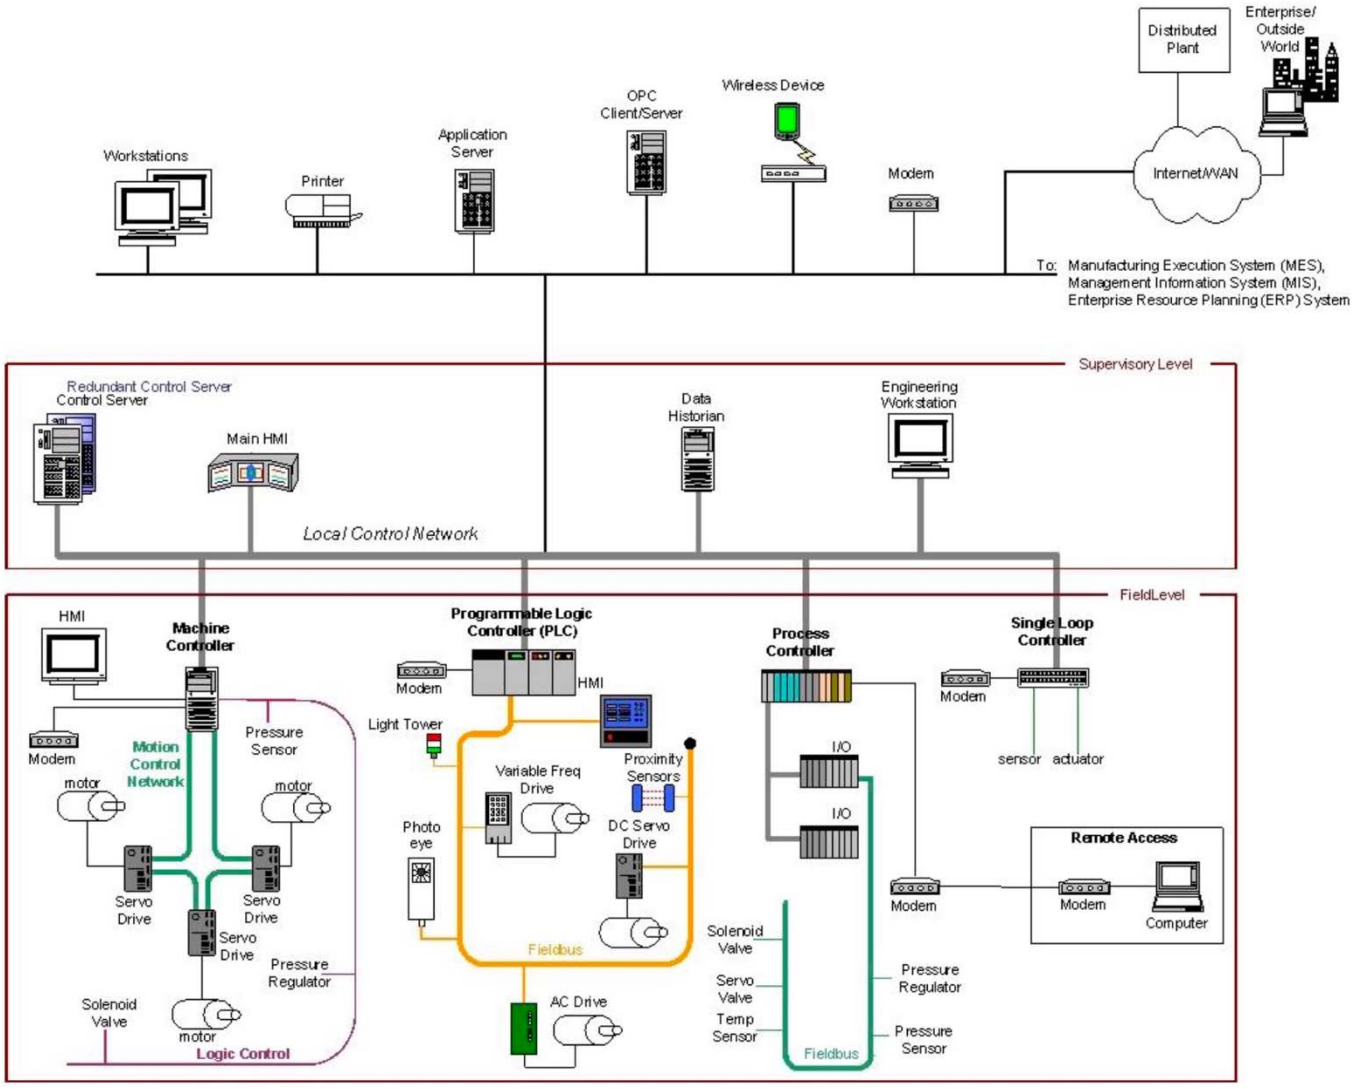

Figure 2 - DCS System Architecture Example [8]

III. MOTIVATION

The MIT Central Utilities facility exhibits a fairly representative control architecture for industrial control systems which makes it an excellent use-case for cybersecurity analysis of ICS using the STPA-Sec methodology. It is important to appreciate the significance of this analysis; a cursory internet search revealed that several universities in the United States operate cogeneration facilities (including Princeton, Stanford, Berkeley etc.). In fact, precluding electricity generation, all large buildings, hospitals, malls, hotels that have bulk heating and cooling requirements, have some form of building utility facilities that operate boilers and chillers to maintain environmental control.

As we found during our research, the specifics of the control architecture of a particular facility are not very difficult to excavate. In fact, some universities and colleges willingly share details about operations of their SCADA and DCS systems online. Figure 3 shows the operator screenshots for the MIT CUP DCS as available on the internet [9]. As can be observed from the figure, an attacker has easy access to critical information about the points of vulnerability within the control architecture at his disposal (including valve and pump numbers and configuration etc.). In one instance, it was found that a college put training material online about how to operate the boiler and chiller plant from the SCADA system, including the web address to remotely log in to the plant control system [10]. In this extreme case, except for login credentials, the college has unsuspectingly provided full operational insights about the inner workings of their building management systems to a potential attacker.

Investigating further, it was discovered that this problem was much more serious than originally envisaged; in some cases, SCADA and building management systems were accessible over the internet without password protection. Search engines designed exclusively for Internet of Things (IoT) devices and systems exist that readily expose unprotected devices, building management and SCADA systems. Two such IoT search engines are *Shodan* and *Censys.* These can be queried with common search terms such as *SCADA, Building Management System* or Equipment manufacturer name to expose unprotected devices. As shown in Figure 4, querying the *Censys* IoT search engine with *'Trane'* which is one of the chiller manufacturers used at the CUP, resulted in close to 3500 hits. Following one such link, we unsuspectingly acquired access to a building's front-desk camera and lighting control without any user credentials (we were unsuccessful in accessing Building Management or SCADA systems in our cursory search). It is worth mentioning that in several cases, with negligible effort, we found default login credentials for industrial chillers in online forums as well as publicly available operating manuals. The point is that detailed information about plant specifics in some cases is readily available which makes mounting a cyberattack, a fairly trivial task to implement.

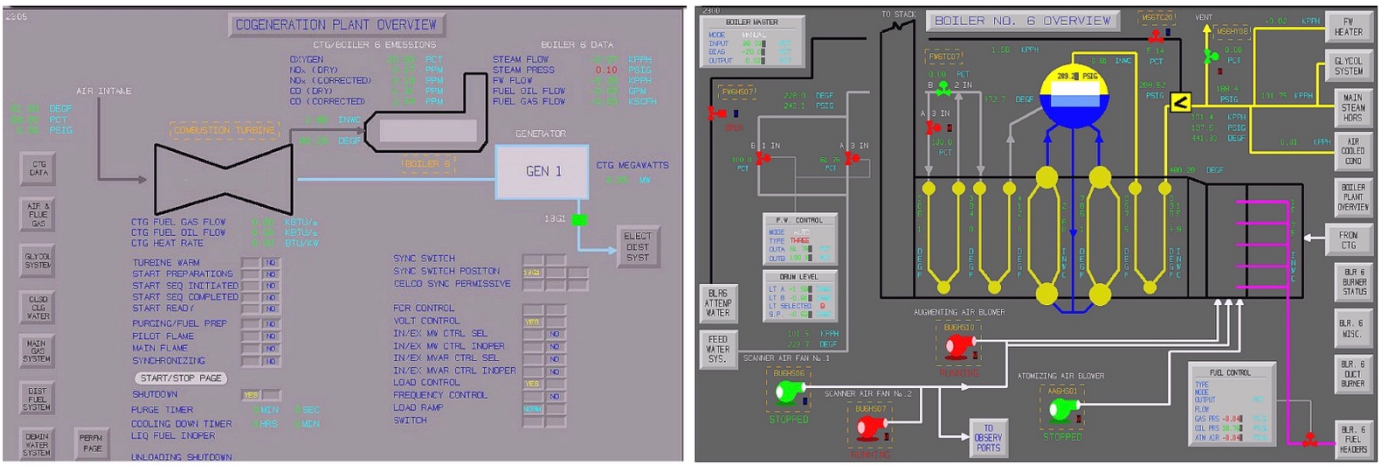

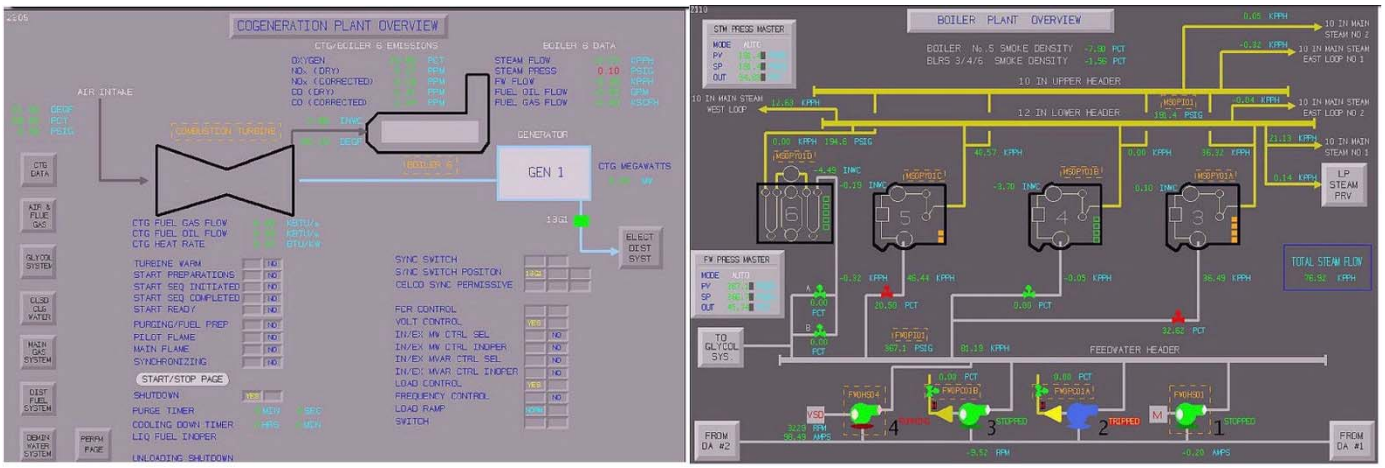

Figure 3 - Screenshots of MIT Central Utilities Plant DCS System [9]

Figure 4 - CENSYS results for querying 'Trane'

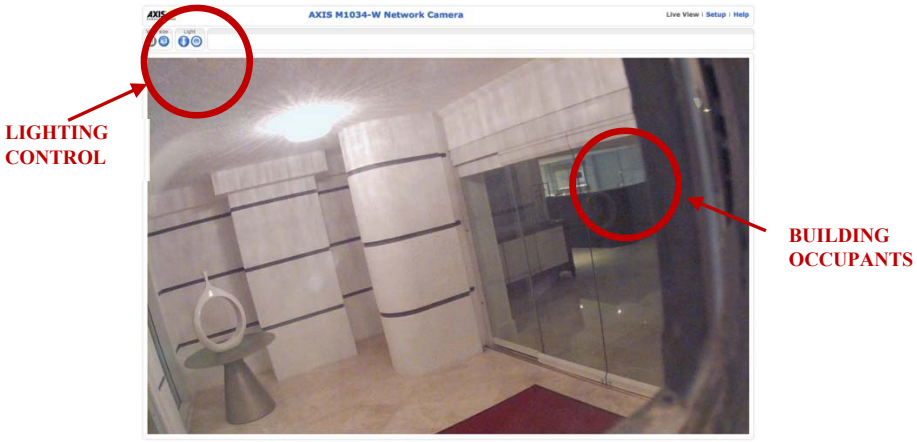

Figure 5 - Access to Building Camera and Lighting without user authentication

For the purpose of this paper, we limit our attention to the function of producing chilled water to demonstrate the STPA-Sec method. The system boundary therefore, includes all equipment and personnel of the MIT CUP that enable the production of chilled water using centrifugal chillers. The motivation for down-selecting the centrifugal chiller for this initial analysis is as follows:

- Centrifugal chillers are some of the more modern pieces of equipment utilized by the central utilities plant. It has a relatively complex system architecture where several different processes have to be

synchronized in order to deliver the *primary-value function* of producing chilled water. And therefore, it uses complex, software-based control algorithms (via PLCs) for process control. Information about chillers is readily available from manufacturers. It also represents a small enough system to demonstrate the application and utility of the STPA-Sec method.

- Chillers are often overlooked as a non-critical pieces of equipments. In engaging with plant personnel, it was discovered that the general mental model about chiller loss of function was that, it would not result in catastrophic damage, rather it would be more of an inconvenience for building occupants. As noted by Angle [1], unlike traditional equipment failures which probabilistically impact one component at a time, cyberattacks have the unique capacity to target several components simultaneously. Under worst-case environmental conditions, a cyberattack on a hot summer day, would prevent the Central Utilities plant from providing any chilled water to the MIT campus. More than occupant inconvenience and loss of productivity, such an attack could cause the computer servers to crash if alternate methods are not utilized to reject heat from the server room.

  Although not cyber-related, a chiller accident at the Los Alamos National Labs in 1997 resulted in $3.2 million in damage to the facility and to the equipment used for nonproliferation and International Security Operations [11]. In the post-accident report, one of the causal factors for the accident was the '*incorrect maintenance categorization of equipment'*. Since the equipment involved in the propagation of the accident (including sump pumps and chiller) were not considered mission-critical, they were not rigorously maintained. This lead to component failures, which under worst-case environmental conditions (i.e. subfreezing outside temperatures and human error leading to improper reservoir temperature set-point setting), resulted in freezing and subsequent bursting of the chiller coils, causing widespread flooding. The point is that chillers are often considered as non-critical pieces of equipment, where as they pose a significant risk in terms of damage potential.

- Chillers use Variable Frequency Drives (VFD). VFD is a type of adjustable-speed drive used in electro-mechanical drive systems to control AC motor speed and torque by varying motor input frequency and voltage. It is widely used in CUP chiller operation for driving chilled water pump motors as well as for chiller compressor capacity control. As noted by Angle [1], VFDs are large energy storage devices where software control is used to maintain proper range of DC bus voltage; with minor modifications to the VFD firmware, it is demonstrated that the capacitors can be exploded one by one. In a separate indictment of VFDs, Zetter [12] noted that VFDs from 4 different manufacturers all provided read/write capability to reset the speed without authentication; VFD could be queried for the critical speed of the attached load and then commanded to drive the load at the dangerous speed, leading to permanent damage of the motor as will be shown later in the analysis.

We will now briefly describe the STPA-Sec method.

IV. THE METHOD

STPA-Sec is an analytical method that finds its theoretical basis in an accident causality model based on *Systems Theory* called STAMP (System-Theoretic Accident Model and Processes). STAMP treats *safety* (and by extension, *security*) as a 'dynamic control problem' rather than a 'failure prevention problem'. Traditional causality models used for safety analysis, attribute accidents to an initial component failure or human error that cascades through a set of other components. Such models are adequate for systems with limited complexity, or systems that exhibit linear interactions and simple cause-and-effect linkages [2, 7].

More complex, software-intensive systems, that are increasingly becoming commonplace in industrial settings, present new challenges in the form of losses caused not only by component failure, but also unsafe interactions among components (none of which may have failed), system requirements and design errors and indirect sociotechnical interactions resulting in unidentified common-cause failures of barriers [2, 7]. For such complex systems, STAMP offers a more robust and comprehensive accident causality model because of the following reasons [2]:

- It works top-down, rather than bottom up i.e. instead of using *external threats*, it uses *outcomes* to derive security requirements
- It includes software, humans, organizations, safety culture, etc. as causal factors in accidents and other types of losses without having to treat them differently or separately.
- It allows creating more powerful tools, such as STPA, accident analysis (CAST), identification and management of leading indicators of increasing risk, organizational risk analysis, etc.

However, it is important to understand, that STAMP is not an analysis method; rather it is a model or set of assumptions about how accidents occur [2]. The two most widely used STAMP-based tools that provide an analysis method are STPA (System-Theoretic Process Analysis) and CAST (Causal Analysis based on STAMP). STPA is *forward-looking* (i.e. a tool for hazard analysis) while CAST is backward-looking (i.e. a tool for analyzing loss events that have already occurred). The basic steps in STPA are shown in Figure 6.

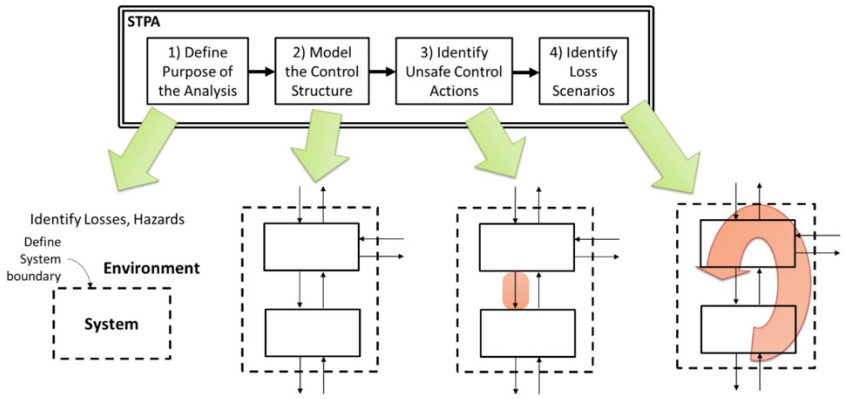

Figure 6 - Overview of the Basic STPA Method [2]

STPA starts with defining the purpose of the analysis by defining system-level losses that the analysis aims to prevent (*Step 1* in Figure 6). The next step (*Step 2* in Figure 6) is to build a hierarchical functional control structure that captures the functional relationships and interactions by modeling the system as a set of feedback control loops. A general form of a control loop is shown in Figure 7. In the control structure, each level of the structure enforces the required constraints on the behavior of the components at the next lower level. Missing or lack of enforcement of relevant constraints can lead to elevated risks, which may result in a loss event(s) under worst-case environmental conditions.

The third step (*Step 3* in Figure 6) is to analyze control actions in the control structure to examine how they could lead to unacceptable losses identified in the first step. These *unsafe control actions* are used to create functional requirements and constraints for the system. Finally, the last step (*Step 4* in Figure 6) identifies reasons why unsafe control actions might occur. Scenarios are created to explain:

1. How incorrect feedback, inadequate requirements, design errors, component failures, and other factors could cause unsafe control actions and ultimately lead to losses.
2. How safe control actions might be provided but not followed or executed properly, leading to a loss.

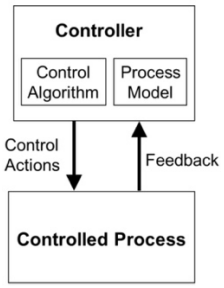

Figure 7 - Generic Control Loop [2]

STPA-Sec is an extension to STPA to include security analysis. The initial steps in the analysis are identical to those for safety: identifying the losses to be considered, identifying system hazards or security vulnerabilities, drawing the system functional control structure, and identifying unsafe, or in this case, insecure, control actions. The only difference is the addition of intentional actions in the generation of the causal scenarios, in the last step in the process [7].

According to Young [7], use of a systems-theoretic approach to security, *"requires a reframing of the usual security problem…into one of strategy rather than tactics. In practice, this reframing involves shifting the majority of security analysis away from guarding against attacks (tactics) and more toward design of the broader socio-technical system (strategy)"*. This means, that instead of focusing on threats from adversaries which are outside the control of the system, security efforts should be focused on controlling system vulnerabilities. This would prevent not only disruptions from known threats, but also disruptions introduced by unknown threats, such as insiders. In other words, in STPA-Sec, the source of the disruption does not matter; what matters is identifying and controlling the inherent vulnerabilities [7].

According to Young [7], the STPA-Sec method does not circumvent a formal threat analysis but proposes to perform the threat analysis only after developing a deeper systemic understanding of the context under which the threats may operate and the disruptions that could actually lead to critical loss events.

We will now provide a high-level description of the system under analysis (i.e. the centrifugal chiller).

V. CHILLER OPERATION

A chilled water system consists of a centrifugal chiller or a combination of chillers, air-handling units (AHU), cooling towers as well as the auxiliary equipment including pumps, water purification system and piping as shown schematically in Figure 8. The centrifugal chiller is a machine that removes heat from a liquid via a vapor-compression cycle. Figure 9 shows the basic refrigeration circuit which consists of the following four main components:

Evaporator – The evaporator in a centrifugal water-cooled chiller is usually a shell-and-tube heat exchanger that removes heat from the chilled water return line (from the building AHU) lowering its temperature in the process. The heat is used to boil the refrigerant, changing its state from liquid to vapor. The evaporator is typically flooded i.e. the chilled water is placed in the tubes while the refrigerant is placed in the shell, completely submerging the tubes.

Compressor – The compressor assembly is made up of a prime mover, driven by an electric motor and a centrifugal compressor. The centrifugal compressor is a dynamic device similar to a centrifugal water pump. It raises the pressure and temperature of the refrigerant by converting kinetic energy into pressure.

Condenser – Similar to the evaporator, the condenser is usually a shell-and-tube heat exchanger. In this case, it removes heat from the refrigerant gas causing it to condense to a liquid. The heat raises the temperature of the cooling water, referred to as the condenser water. The condenser water carries the heat to the cooling tower where the heat is rejected to the atmosphere.

Expansion Device – After the refrigerant condenses to a liquid (in the condenser), it passes through a pressure reducing device (generically known as a metering device). This can be as simple as an orifice plate or as complicated as an electronic modulating expansion valve. As the pressure is reduced by allowing small amounts of condensed refrigerant to pass through the valve, it's temperature decreases, cooling the refrigerant.

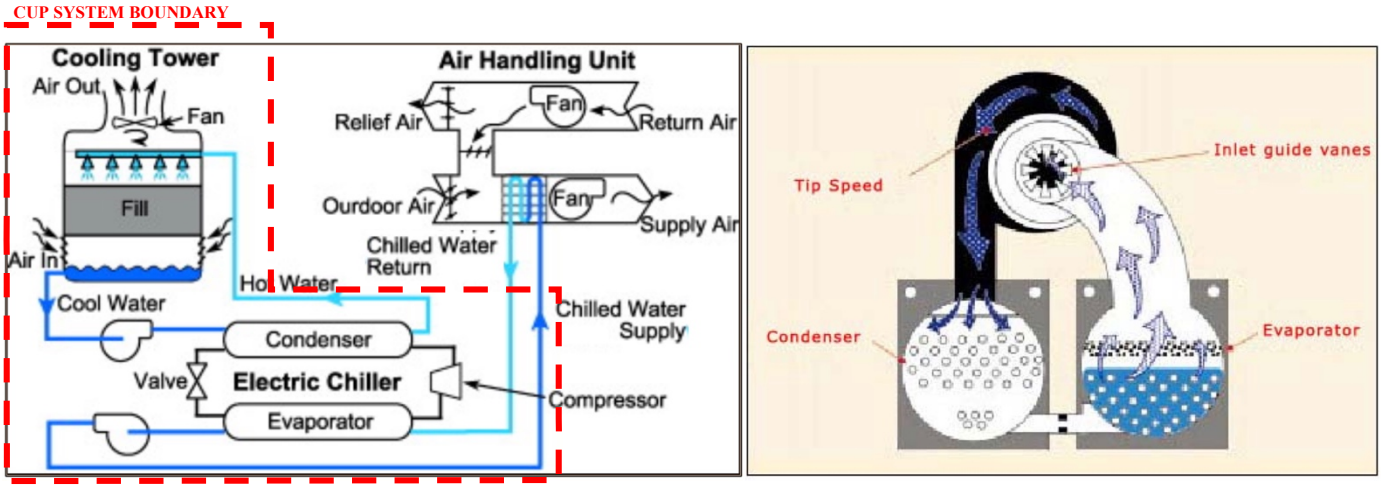

Figure 8 - Schematic of a typical chilled water system

As may be evident, there are three independent fluid loops which function together to enable delivery of chilled water to the campus; 1) a closed water circuit that runs chilled water between the building Air Handling Unit (AHU) and the evaporator, 2) a closed refrigerant loop, which through the change of its states, enables transfer of heat from the chilled water loop to the condenser water loop, and 3) an open water loop, absorbing heat from the refrigerant and rejecting it to the atmosphere through cooling towers. Make-up water

from the main water distribution line ensures the water in the cooling tower reservoir is maintained at the desired level.

As shown in Figure 8, for MIT CUP, the operation of the air-handling units is outside the purview of CUP control system. Instead, the chilled water differential temperature and pressure between the supply and return lines is monitored and maintained, while each building's air-handling units are managed by third-party Building Management Systems (BMS).

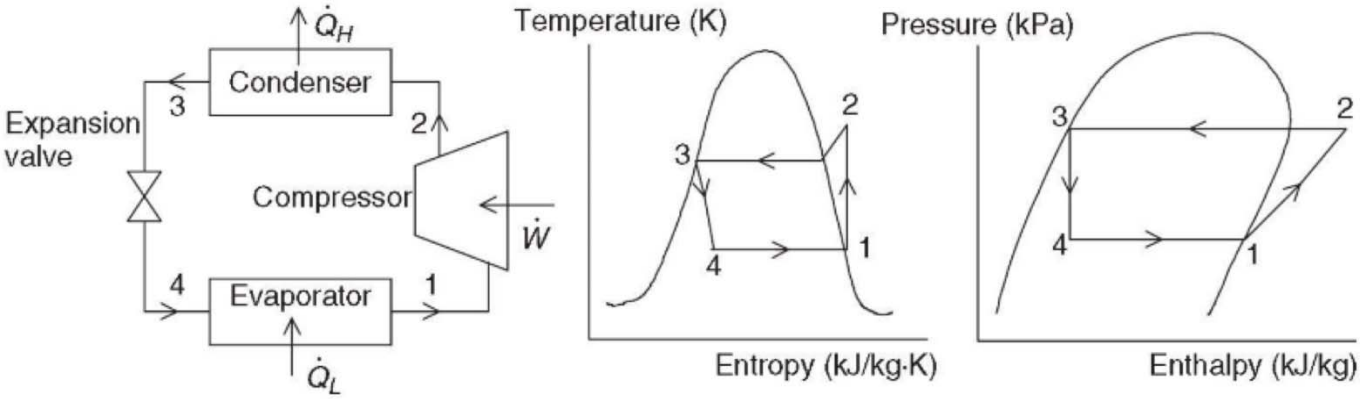

Figure 9 - Chiller vapor curves

VI. ANALYSIS

In this section, we demonstrate the application of STPA-Sec to the MIT CUP chiller use case. This section is divided into subsections where each subsection represents one step in the basic STPA-Sec diagram as presented in Figure 6. The hazard scenarios generated in this section are used to guide an in-depth security analysis and derive mitigation strategies presented later.

Step 1 – Define Purpose of Analysis

We define the system under analysis as the *control of the centrifugal chiller in its different system states*. The STPA-Sec method starts with an identification of system losses and hazards based on expert knowledge. It is emphasized that as part of STPA method's *top-down* approach, high-level system losses are considered rather than component losses. System-level losses are defined as any loss that would be unacceptable to the stakeholders. Table 1 provides a prioritized list of system losses. Here, (**L-1***) loss of equipment*, is given higher preference over (**L-2**) *loss of cooling*, which is the *primary-value function* of the system.

The justification for this prioritization is based on impact severity; an equipment loss would potentially result in a longer-term loss of function of the system. Downstream losses are not considered since they are outside the control of the system. For instance, in the case of the Los Alamos National Lab's chiller accident [11], the presence of radiological sources in the basement, significantly complicated the reclamation efforts and cost, after the freezing and bursting of the chiller coils that subsequently flooded the basement. Even though quite significant, such losses are outside the system boundary and hence not considered.

Based on system losses, system hazards are defined as shown in Table 2. System hazards are conditions or system states that will result in a system loss under worst-case environmental conditions. Table 3 maps system losses to hazards, showing which losses can potentially be caused by each hazard.

Table 1 - System-Level Losses

L-1: Loss of equipment (financial/operational)
L-2: Loss of cooling (environmental control)
L-3: Death, dismemberment or injury to plant personnel

Table 2 - System-level Hazards

Hazards	**Related Losses**
H-1: Chiller is operating beyond normal operational limits	L-1, L-2, L-3
H-2: Chiller violates correct sequence of operations	L-1, L-2, L-3
H-3: Chiller is unable to provide accurate feedback about status	L-1, L-2, L-3
H-4: Chiller releases asphyxiate gasses	L-2, L-3
H-5: Chiller does not respond to local (chilled water) demand	L-2

Table 3 - Mapping of Hazards to Losses

Hazard	L-1	L-2	L-3
H-1	x	x	x
H-2	x	x	x
H-3	x	x	x
H-4		x	x
H-5		x	

Based on the system hazards, system-level constraints are derived by essentially inverting the system hazards as presented in Table 4. For instance, if a system-level hazard is defined as 'Chiller is operating beyond normal operational limits', a system-level constraint can be derived as 'the chiller must not operate beyond normal operational limits'; the key idea is to specify system conditions or behaviors that need to be satisfied in order to prevent hazards (and ultimately prevent losses). In the current use case, no specific security-related system constraints are defined. Instead, overall system operation is considered from a top-down perspective without specifying a singular approach or solution to prevent hazards. This enables a broader exploration of solution space further down the analysis. These constraints are ultimately refined during the analysis to comprehensively encapsulate safety and security needs.

Table 4 - System-Level Constraints

Hazards	**Related Losses**	**Constraints**
H-1: Chiller is operating beyond normal operational limits	L-1, L-2, L-3	SC-1.1: Chiller must not operate beyond normal operational limits. SC-1.2: If chiller is operating outside limits, then the violation must be detected and measures taken to prevent operation.
H-2: Chiller violates correct sequence of operations	L-1, L-2, L-3	SC-2: Chiller must not violate correct sequence of operations.
H-3: Chiller is unable to provide accurate feedback about status	L-1, L-2, L-3	SC-3: Chiller must provide accurate feedback about status (of operations) at all times.
H-4: Chiller releases asphyxiate gasses	L-2, L-3	SC-4.1: Chiller must prevent release of asphyxiate gasses. SC-4.2: If chiller releases asphyxiate gasses, measures must be taken to alert proximate workers.
H-5: Chiller does not respond to local (chilled water) demand	L-2	SC-5.1: Chiller must be sized adequately to meet local demand. SC-5.2: If chiller is unable to meet local demand, measures must be taken to prevent damage to critical loads due to loss of cooling.

Step 2 – Model and Control Structure

The next step in the STPA-Sec method is to model the hierarchical control structure. At its most fundamental level, the control structure models *control loops* composed of *controlled processes* and *controllers.* The *controllers* receive *feedback* about the *controlled process,* and then based on some *control algorithm* or logic provide *control actions* to adjust the process. The STPA control structure enables modeling not only physical processes, but also human operations, including control of procedures, information and policies. This makes STPA extremely versatile for cybersecurity applications because it provides visibility of complex functional interactions between the physical systems as well as information flows between control systems, operators, management and even government and regulatory bodies as a whole. Due to the complexity of the system, not every controlled process is modeled in the control structure; instead, abstraction of the physical model is used to convey the essence of the flow of *control actions* and *feedback*.

Figure 10 illustrates the high-level functional control structure for the MIT CUP. In this view, the system under analysis i.e. the chiller and associated equipment, is abstracted as the *cooling system* under the boundary of the central utilities plant. The figure also illustrates how the operators have the ability to control the CUP through both the DCS as well as the equipment's local *Human-Machine Interface* (HMI) screen. The operator actions, in turn, are controlled via operating procedures and instructions by Plant Engineers. Both Plant Engineers and operators report to CUP's Operations Management which enforces MIT leadership's enterprise level goals and vision through policies and standards. The MIT's leadership, in turn, is controlled by municipal, state and federal regulations enforced via certificates and licensure for operating the plant.

Figure 11 expands the *cooling system* to reveal more details of the underlying physical processes. As noted earlier, the *cooling system* is not only composed of the *chiller controller*, but also a *chilled water loop* for absorption of heat from the buildings, a *condenser water loop* for rejection of heat to the outside environment as well as a *water purification system* for controlling water chemistry to ensure longer service life of the equipment.

For the purpose of this paper, we focus our attention on a single control loop – the *Chiller Cooling Capacity Controller* – and its interaction with the different logical components in the system for analysis. An overview of the chiller PLC architecture is shown in Figure 12, while a functional control structure derived from this architecture is shown in Figure 13. For ease of reference in the text, each *node* and *connection* in the control loop is labeled with prefixes *N-* and *C-*, respectively.

The chiller PLC controls the cooling capacity of the chiller in response to chilled water temperature deviation from the set-point by adjusting the speed of the compressor motor. The PLC receives feedback from several sensors monitoring various physical processes, including refrigerant discharge and suction temperatures and pressures, condenser and evaporator water temperature, pressure and flow, compressor oil temperature and pressure, guide vane position etc. Based on a computation of this information in combination with operator defined set-points (through the DCS), the chiller PLC decides if the compressor motor needs to be activated and if so, defines a target speed for the compressor motor via the VFD control unit. Once the chiller achieves the desired chilled water temperature, the PLC commands the VFD control unit to ramp down the motor to a stop. The PLC also monitors motor voltage, current and temperature in order to keep it within safe operating limits.

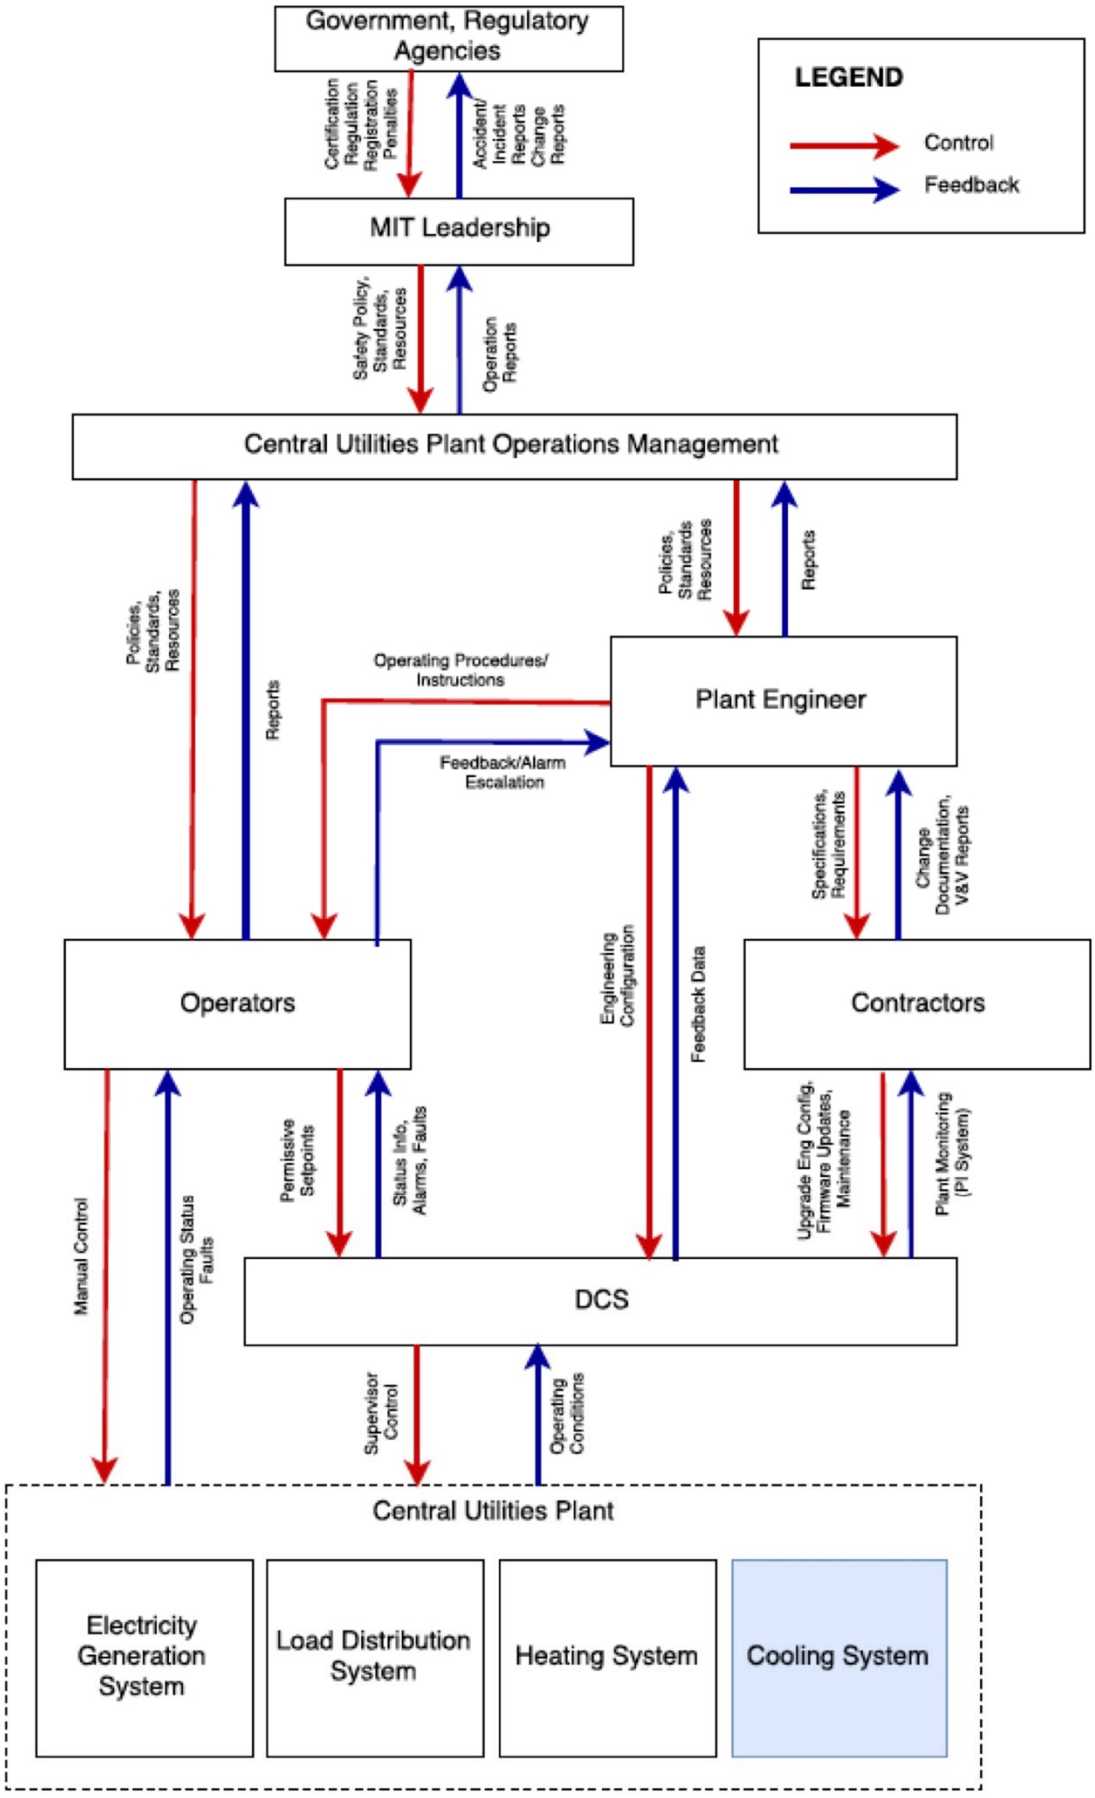

Figure 10 – Overall CUP functional control structur

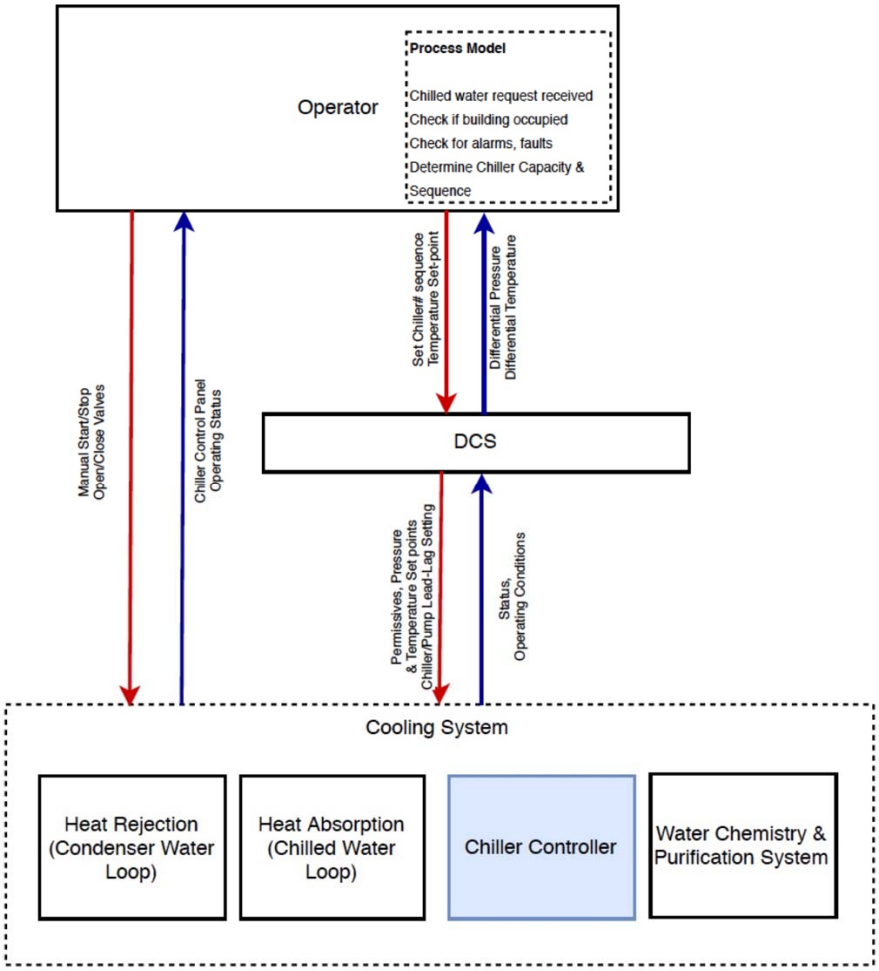

Figure 11 – Cooling System functional control structure

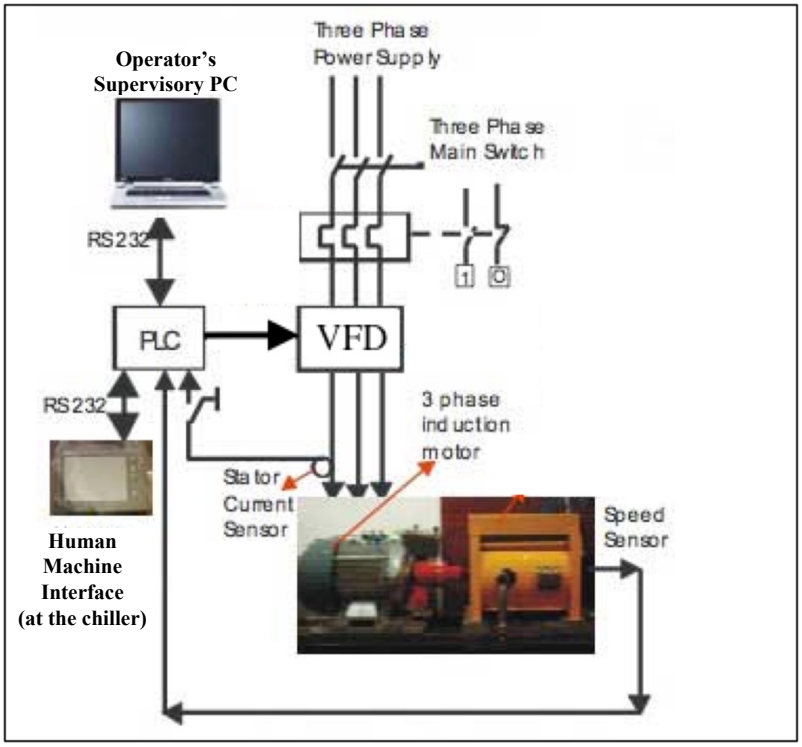

Figure 12 - Overview of the Chiller PLC Controlling Compressor Speed via VFD

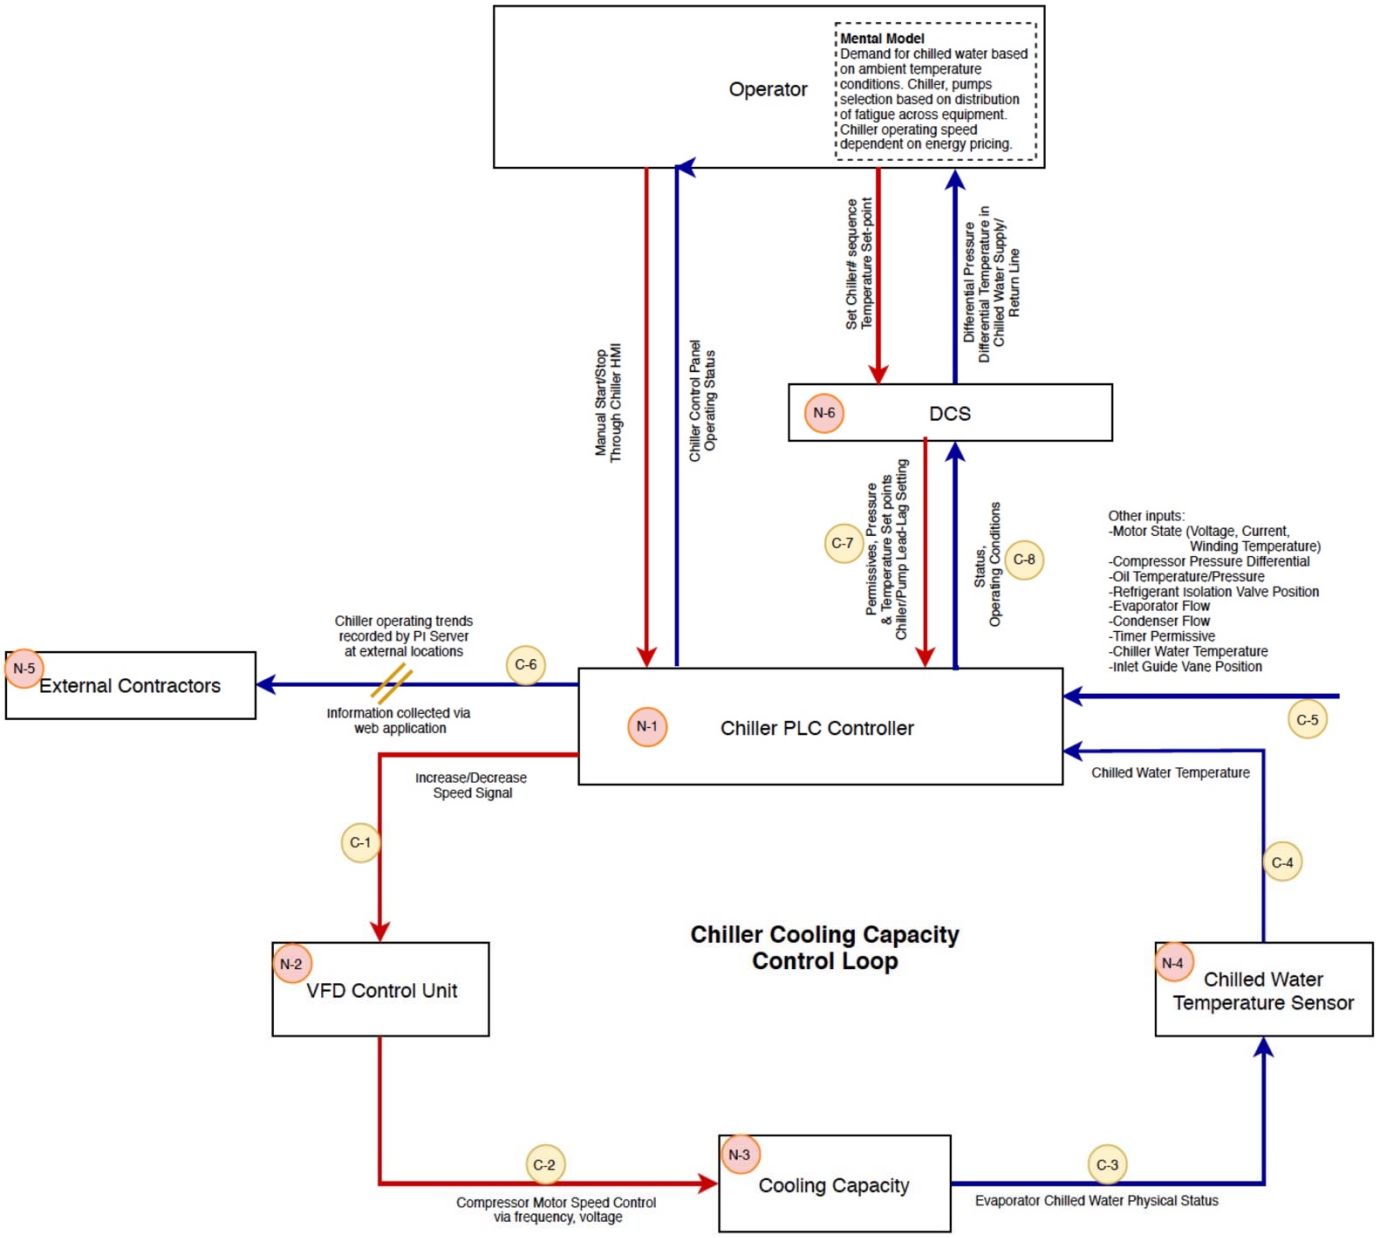

Figure 13 – Zooming into Chiller Controller

Step 3 – Identify Unsafe Control Actions

The next step in the STPA-Sec method is to identify *Unsafe Control Actions.* We begin by defining system states that are relevant for the correct operation of the chiller. Based on the chiller PLC control loop, Table 5 lists the relevant system variables along with values that each variable can take at any given time. One interesting insight that can be derived from Table 5 is that while for traditional safety analysis, system states can be either *'Within Limits'* or *'Outside Limits',* when considering malicious actions, *unconventional* system states also need to be considered.

For instance, Compressor Motor Speed in Table 5 is shown to have four system states; *Within Limits, Critical, Reverse, Unstable.* Since the compressor motor has several resonant frequencies which occur as the motor ramps up to its operating speed (i.e. there is no fixed range of values where the motor can safely operate), we abstract the *'Outside Limits'* system state under *'Critical Speed'* system state to encompass all compressor motor values where the operation of the motor would be unsafe. Similarly, whereas a traditional safety analysis would not consider *Reverse rotation* or *unstable (oscillatory) rotation* since these system states do not occur naturally, a security analysis needs to consider all unconventional system states that are possible by malicious actions.

Table 5
System Variables that are relevant for the correct operation of the Chiller Controller and their possible values

#	Name	Values
1	Compressor Motor State (Voltage, Current, Temperature)	Within Limits; Outside Limits
2	Evaporator Refrigerant Temperature	Within Limits; Outside Limits
3	Compressor Pressure differential (Discharge vs. Suction)	Within Limits; Outside Limits
4	Oil Temperature/Pressure	Within Limits; Outside Limits
6	Evaporator Flow	Within Limits; Outside Limits
7	Condenser Flow	Within Limits; Outside Limits
8	Timer permissive	Available; Not Available
9	Chilled Water Deadband Temperature Range	Within Limits; Outside Limits
10	Inlet Guide Vane Position	Within Limits; Outside Limits
11	Compressor Motor Speed	Within Limits; Critical; Reverse; Unstable

In STPA, for each possible control action, a decision has to be made whether the control action is hazardous in a given system state or not. A control action can be hazardous if (i) it is not provided, (ii) it is applied at all, (iii) it is applied too early, too late or out of order or (iv) if it is stopped too soon or too late in a given system state [6]. For our analysis, we define *Increasing Compressor Motor Speed* as the *control action*, while considering the command to start/stop the compressor motor as implicit to this control action. Table 6 shows a list of the hazardous control actions identified for the chiller controller. A hazardous control action is present for a specific system state and can cause a set of hazards [6]. We will now discuss the selected control action in the context of these system states and hazards.

Table 6

Control Actions by different system states. A '-' indicates that the status of the variable is irrelevant for the hazardous behavior of the highlighted control action. A control action can be unsafe at any time it is performed in a given state, or only if provided too early or too late or not at all. '1' indicates control action is unsafe in the given state and '0' indicates vice versa.

	System Variables	#1	#2	#3	#4	#6	#7	#8	#9	#10	#11	Providing Causes Hazard	Not Providing Causes Hazard	Too Early, Too Late, or Out-of-Order	Applied too long, Stopped too soon	Hazards
Increase Compressor Motor Speed	CA-1	Out	-	-	-	-	-	-	-	-	-	1	0	1	1	H-1, H-2, H-5
	CA-2	-	Out	-	-	Out	-	-	-	-	-	1	0	1	1	H-1, H-2, H-4, H-5
	CA-3	-	-	Out	-	-	Out	-	-	-	-	1	0	1	1	H-1, H-4, H-5
	CA-4	-	-	-	Out	-	-	-	-	-	-	1	0	1	1	H-1, H-2, H-5
	CA-6	-	-	-	-	Out	-	-	-	-	-	1	0	1	1	H-1, H-2, H-5
	CA-7	-	-	-	-	-	Out	-	-	-	-	1	0	1	1	H-1, H-2, H-5
	CA-8	-	-	-	-	-	-	Not Avail	-	-	-	1	0	1	1	H-2
	CA-9	-	-	-	-	-	-	-	Out	-	-	1	0	1	1	H-2
	CA-10	-	-	-	-	-	-	-	-	Out	-	1	0	1	1	H-1, H-5
	CA-11	-	-	-	-	-	-	-	-	-	Critical	1	0	0	1	H-1, H-5
	CA-12	-	-	-	-	-	-	-	-	-	Reverse	1	0	0	1	H-1, H-5
	CA-13	In	In	In	In	In	In	Avail	Out	In	In	0	1	0	0	H-5
	CA-14	In	In	In	In	In	In	Avail	In	In	Unst.	1	0	0	1	H-1, H-5

Compressor Motor State

As noted earlier, the chiller PLC controls the chiller capacity by changing the compressor speed via the VFD control unit in response to chilled water temperature deviation away from the set-point. In some circumstances, however, a request by the PLC to increase compressor speed can prove to be hazardous. For instance, if the compressor motor is operating at its temperature or current limit, a command by the PLC to increase motor speed would result in overheating; excessive heat can lead to premature loss of motor winding insulation, resulting in the motor burning itself out [13]. Some motors have thermal protection relays that shut off the motor when subject to unsafe operating temperatures. At this stage, we are not taking credit for any protective devices installed in the plant; the point is to enumerate all vulnerabilities so that a comprehensive understanding of the attack surface and impact severity can be developed.

Evaporator Refrigerant Temperature

Similarly, if the compressor speed is increased when the refrigerant temperature in the evaporator is at its lower limit and this condition is coupled with an inadequate water flow in the evaporator, the evaporator coils can potentially freeze (CA-2). This is because, increasing the compressor speed reduces the suction pressure at the compressor inlet which lowers the boiling point of the refrigerant, hence increasing the evaporation rate in the evaporator. Coupled with a low/no-flow condition in the chilled water loop, this condition can cause the water in the evaporator tubes to freeze, subsequently leading to tube rupture.

Surging

Another characteristic hazardous condition for centrifugal chillers is *surging*. This can occur when the compressor inlet flow is reduced such that the head developed by the compressor is insufficient to overcome the pressure at the discharge of the compressor. Once surge occurs, the output pressure of the compressor is drastically reduced, resulting in flow reversal within the compressor [4]. The flow reversal applies significant dynamic forces on the impeller which subjects the compressor components (such as thrust bearings, bearings, casing) to large axial force changes due to the rotor rocking back and forth. The flow reversal within the compressor also results in hot compressed vapor returning to the compressor inlet. If not controlled, the temperature at compressor inlet can cause tight-tolerance compressor internals to expand at different rates, leading to damage from friction. Intermittent operation in surge is not normally detrimental

to the machine. However, if prolonged operation in surge is not arrested, it can cause permanent damage to the compressor.

Figure 14 illustrates generic compressor performance curves, superimposed with *surging* phenomenon through two different operating conditions; (1) increase in compressor speed and, (2) change in suction/discharge flow or pressure characteristics. In order to understand the phenomenon, consider the compressor operating at *Point A* (in the figure on the left (Figure 14 (a))); an increase in compressor speed would move the operating point of the compressor to Point B, which is at the surge limit (CA-3). Similarly, as shown in the figure on the right (Figure 14 (b)), a reduction in water flow in either the condenser or evaporator water loops would cause the differential pressure across the compressor to increase (CA-5, CA-6) i.e. move the compressor operating point from Point A to Point B in the figure on the right (Figure 14 (b)). In both system states, increasing the compressor speed would cause the compressor to surge.

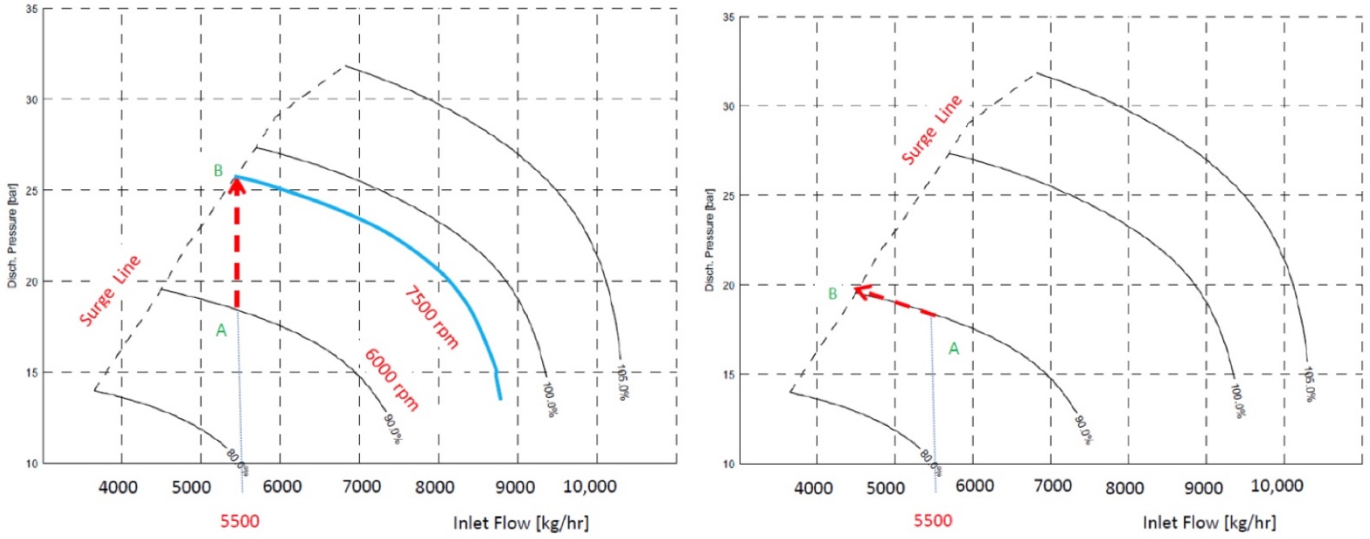

Figure 14 – Surge due to (a) Increase in Speed, (b) Change in Suction or Discharge Flow or Pressure [4]

Inlet Guide Vane Position

Some chillers use a combination of Inlet Guide Vanes (IGV) and speed control to optimize capacity control of the chiller. The inlet guide vanes are stationary blades with variable pitch that provide a mechanism to alter the swirl pattern on the inlet flow to the compressor. The performance of the compressor is optimized by adjusting the position of the guide vanes in combination with the speed of the compressor. An incorrect guide vane position, coupled with an increase in compressor speed has the potential to push the operating point of the compressor into the surge region (CA-10).

Lubrication Oil

Centrifugal compressors need oil forced around some of their internal components (such as gears, thrust bearings etc.) to provide lubrication and remove heat caused by friction. The lubrication oil has to be at the correct temperature and pressure for it to perform its intended function; it must be thin enough to lubricate properly at the high speeds of rotation of the compressor but also thick enough to handle the heat and refrigerant contamination that can occur. If the lubrication oil conditions are beyond design limits (in terms

of temperature and pressure), and the compressor is commanded to start, the compressor would essentially destroy itself because of the excessive heat build-up through friction in the internal components (CA-4).

Timer Permissive (Anti-Recycle or Anti-Short Cycle Timer)

The chiller PLC also limits the number of compressor startups in a given period of time. This is to allow the motor sufficient time to cool down between startups. Forcing the motor to violate this permissive could potentially burn the motor (CA-8).

Chilled Water Set-point

During normal operation, once the chilled water set-point is achieved, the chiller PLC ramps down the compressor. If instead of ramping down upon achieving the set-point, the compressor motor speed is increased beyond the set-point, there is potential for the evaporator coils to freeze (CA-9). Alternatively, if the chiller PLC sends a *Start/Run* permissive to the compressor via the VFD control unit, but the compressor motor does not start, it also presents a hazardous condition since the chiller plant would not be able to supply chilled water to the campus (CA-13) which could result in downstream losses including loss of cooling for campus critical loads such as the IT server room.

Motor Critical Speed & Reverse Rotation

The compressor motor, like all motors, has a critical speed; a speed at which mechanical resonance occurs. Causing the motor to run at the critical speed, can cause considerable damage to the motor, drive and compressor components. When driven via VFD, a motor may have several critical speeds; the traditional approach is to program the VFD to skip over the critical speeds. According to Zetter [12], if a motor is run at its critical speed the vibrations can 'destroy the bearings and (the) motor shaft' (CA-10).

Another unsafe state for the compressor motor is reverse rotation (CA-12). Some VFDs allow reverse rotation by changing some parameters on the PLC controlling the drive (via the input/output (I/O) card of the drive). Reversing the direction of rotation of the impeller, would still continue to deliver the fluid from the low-pressure suction side to the higher-pressure discharge side in the positive direction of flow, but the flow characteristics would be adversely impacted. This would significantly reduce compressor efficiency; the typical response of the chiller PLC would be to increase compressor speed to achieve the cooling capacity. As shown in Figure 14 (b), an increase in speed would push the operating point beyond the compressor surge limit.

Motor Instability – Stuxnet Condition

Lastly, inspired by the Stuxnet attack, *unstable motor speed* is considered as a system state. In the Stuxnet case, a *Siemens S7-300 PLC*, controlling a motor via a *Vacon VFD* was attacked; when certain criteria were met, the code periodically modified the frequency (from 1,410 Hz to 2 Hz to 1,064 Hz), thus affecting the operation of the connected motors by changing their rotational speed and ultimately damaging the centrifuges [14]. Even though, the electric-driven compressors at CUP do not operate at the same supersonic speed as the centrifuges, the control architecture of the compressor motor is almost identical – the motor speed is controlled by a PLC via a VFD. This means that a similar throttling of motor frequencies as executed by the Stuxnet code, would apply significant thermal stresses and dynamic loads to the motor and cause permanent damage to the motor and by extension to the compressor and chiller.

Based on the enumeration of hazardous system states in the context table, we condensed the list into a set of unsafe control actions as shown in Table 7. In order to keep this list to a manageable level, some system state variables have been abstracted under a common term referred to as the 'permissive'; if any of the variables is out of bound, it is considered that the chiller PLC does not have the permissive to issue the control action and vice versa.

Table 7
Unsafe Control Actions

Action By	Control Action	Not Providing Causes Hazard	Providing Causes Hazard	Too soon, Too late, Out of order	Stopped too soon, Applied too long
Chiller Controller	Increase Compressor Speed via VFD	**UCA-1:** Chiller Controller does not increase compressor speed when refrigerant temperature is below setpoint --> [H-5]	**UCA-2:** Chiller controller increases compressor speed when permissives for this action are unavailable --> [H-1, H-2, H-4]	**UCA-3:** Chiller controller increases compressor speed before permissives for this action are available --> [H-1, H-2, H-5]	**UCA-7:** Chiller controller continues to increase compressor speed when permissives for this action become unavailable --> [H-1, H-2, H-5]
		N/A	**UCA-4:** Chiller controller increases compressor speed but in the reverse direction --> [H-1, H-5]	N/A	N/A
		N/A	**UCA-5:** Chiller controller increases compressor speed but to a value that is different than the one requested. --> [H-1, H-5]	N/A	N/A
		N/A	**UCA-6:** Chiller controller sends the signal to increase speed but it is executed incorrectly (successive ramp-up and ramp downs of compressor speed at an unsafe rate via VFD (Stuxnet case)) --> [H-1]	N/A	N/A

Step 4 – Identify Loss Scenarios

In the previous subsection, we contextualized how a given control action would become hazardous under various system states. To finally connect all the artifacts generated in the analysis, a hierarchical list of loss scenarios is required. A loss scenario is a textual representation of causal factors that can lead to unsafe control actions resulting in hazardous system states that can potentially culminate into system losses. According to Levenson [2], two types of loss scenarios must be considered:

A. Scenarios that lead to *unsafe control actions*
B. Scenarios in which control actions are improperly executed or not executed altogether

Scenarios leading to *unsafe control actions* could be a result of *Unsafe controller behavior* or *inadequate feedback;* alternatively, scenarios leading to improperly executed or altogether ignored *control actions* could be a result of the *control path* or the *controlled process* itself as illustrated schematically in Figure 15 and detailed in Table 8 [2] below:

Table 8 - Generation of Loss Scenarios

A. Identifying Scenarios that lead to Unsafe Control Actions

1. *Unsafe Controller Behavior*

- **a.** Failure involving controller
- **b.** Inadequate Control Algorithm
  - **i.** Flawed implementation of the specified control algorithm
  - **ii.** Specified control algorithm is flawed
  - **iii.** Specified control algorithm becomes flawed over time due to changes/degradation
- **c.** Inadequate process model
  - **i.** Controller receives incorrect feedback/information
  - **ii.** Controller receives correct feedback/information but interprets it incorrectly or ignores it
  - **iii.** Controller does not receive feedback/information when needed (Delayed or never received)
  - **iv.** Necessary controller feedback/information does not exist
- **d.** Unsafe Control input (from another controller)

2. *Inadequate Feedback and information*

- **a.** Feedback or information not received
  - **i.** Feedback/info sent by sensor but not received by controller
  - **ii.** Feedback/info is not sent by sensor but is received by controller
  - **iii.** Feedback/info is not received or applied to sensor
  - **iv.** Feedback/info does not exist in control structure or sensor does not exist
- **b.** Inadequate feedback is received
  - **i.** Sensor responds adequately but controller receives inadequate feedback/info
  - **ii.** Sensor responds inadequately to feedback/info that is received or applied to sensor
  - **iii.** Sensor is not capable or not designed to provide necessary feedback/info

B. Identifying Scenarios in which control actions are improperly executed or not executed

3. *Scenarios involving the Control Path*

- **a.** Control Action not executed
  - **i.** Control action is sent by controller but not received by actuator
  - **ii.** Control action is received by actuator but actuator does not respond
  - **iii.** Actuator responds but the control action is not applied to or received by the controlled process
- **b.** Control Action improperly executed
  - **i.** Control action is sent by controller but received improperly by actuators
  - **ii.** Control action is received correctly by actuator but actuator responds inadequately
  - **iii.** Actuator responds adequately, but the control action is applied improperly at the controlled process
  - **iv.** Control action is not sent by controller, but actuators or other elements respond as if it had been sent

4. *Scenarios related to the Controlled Process*

- **a.** Control action not executed
  - **i.** Control action is applied or received by the controlled process but the controlled process does not respond
- **b.** Control action improperly executed
  - **i.** Control action is applied or received by the controlled process but the controlled process responds improperly
  - **ii.** Control action is not applied or received by the controlled process but the process responds as if the control action had been applied or received

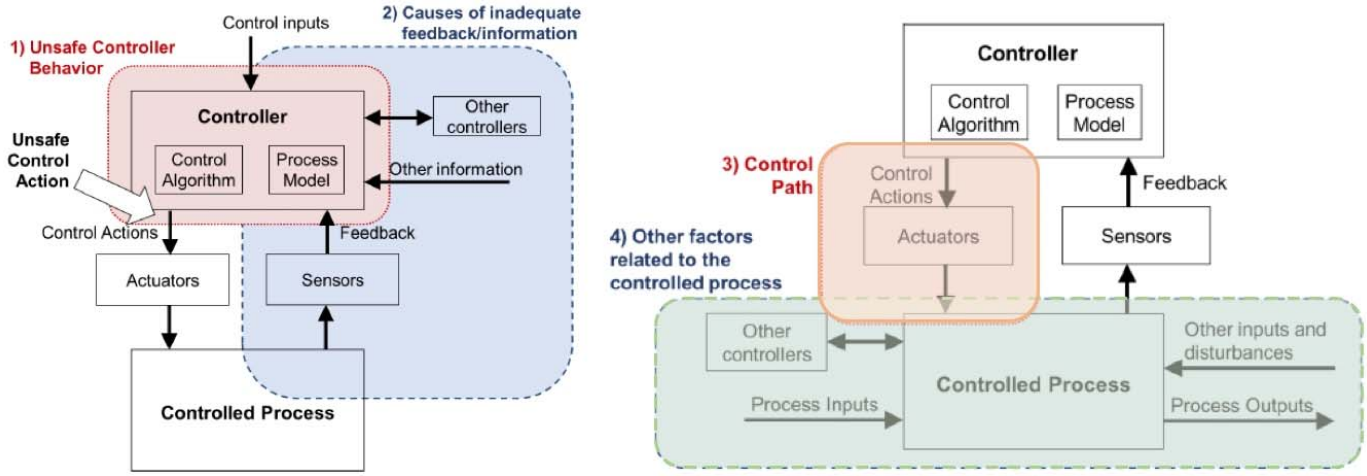

Figure 15 – Generic control look illustrating factors that can result in a) unsafe control actions b) control actions not or improperly executed

Considering the *Loss Scenario* causal factors in the context of cybersecurity, the chiller PLC control loop (shown in Figure 13) was mapped against traditional availability and integrity threats as shown in Table 9, based on the approach presented by Friedberg [6]. This threat-map helps to quickly identify the nodes and connections that provide the widest attack surface as well as pose the greatest risk in terms of impact severity and can help in prioritizing cyber-defense strategies for the various nodes and connections in the system.

Table 9
Mapping between the Availability/Integrity Threats and Nodes and Connections in the Control Loop

Threats		Nodes and Connections at Control Layer													
		N-1	N-2	N-3	N-4	N-5	N-6	C-1	C-2	C-3	C-4	C-5	C-6	C-7	C-8
Integrity Threats	Command Injection	x	x				x	x	x				x	x	
	Command Drop	x	x				x	x	x					x	
	Command Manipulation	x	x				x	x	x					x	
	Command Delay	x	x				x	x	x					x	
	Feedback Injection	x			x		x				x	x			x
	Feedback Drop	x			x		x			x	x	x	x		x
	Feedback Manipulation	x			x		x			x	x	x	x		x
	Feedback Delay	x			x		x				x	x			x
Availability Threats	Communication Delay	x	x		x	x	x	x	x		x	x	x	x	x
	Communication Drop	x	x		x	x	x	x	x		x	x	x	x	x
	Node Overloaded (delay)							x	x		x	x		x	x
	Node Overloaded (drop)							x	x		x	x		x	x
Hazards	Operation Beyond Limits (H-1)	x	x	x	x		x	x	x		x	x		x	x
	Operation Sequence Violation (H-2)	x					x				x	x		x	x
	Inaccurate Feedback (H-3)	x			x		x			x	x	x	x		x
	Release of Asphyxiate Gases (H-4)	x					x				x	x		x	x
	Inability to meet local demand (H-5)	x	x	x	x		x	x	x	x	x	x		x	x

Next, we review each unsafe control action in the context of the causal factors list (provided in Table 8) and the availability/integrity threat-map to generate a structured list of *Loss Scenarios* as presented in Table 10.

Table 10 - Loss Scenarios for Chiller PLC Control Loop

UCA-2, -3, -7: Chiller Controller increases compressor speed when permissives for this action are not available **OR** before the permissives for this action are available **OR** controller continues to increase speed when permissives have become unavailable.

	Loss Scenario	Associated Causal Factors	Rationale	
1.0	Chiller controller incorrectly believes that it has the permissive to increase compressor motor speed i.e. controller does the right thing but the information it basis its decision on, is corrupted.	1. Inadequate/malformed process model: - Controller receives incorrect feedback/ information (LS-1.c.i)	-Malicious feedback injection to controller (N-1) about physical status from sensors (C-5)	Unsafe Controller Behaviour
		- Controller does not receive feedback/ information when needed (Delayed or never received) (LS-1.c.iii)	-Malicious command injections spoof controller node (N-1) so that it is overloaded and becomes unavailable; controller assumes previous incorrect state	
		2. Feedback or information not received: - Feedback/info is not sent by sensor but received by controller (LS-2.a.ii)	-Malicious feedback injection to controller (N-1) from connections (C-5 or C-4)	Inadequate Feedback and Information
		- Feedback/info sent by sensor but not received by controller (LS-2.a.i)	-Malicious feedback drop at connections (C-5 or C-4); controller assumes previous state where it had the permissive to increase speed	
		3. Inadequate feedback is received: - Sensors responds adequately but controller receives inadequate feedback/info (LS-2.b.i)	-Communication drop or delay at C-4 e.g. communication that evaporator flow is low is not communicated to the controller; controller assumes it has permissive to increase speed	
1.1	Chiller controller interprets correct feedback incorrectly and increases speed when it does not have the permissive to do so.	1. Inadequate/malformed process model: - Controller receives correct feedback/ information but interprets it incorrectly or ignores it (LS-1.c.ii)	-Malicious feedback manipulation at controller (N-1) from sensors (C-5 or C-4) causes the controller to assume incorrect state	Unsafe Controller Behaviour
		2. Inadequate/malformed control algorithm: - Flawed implementation of the specified control algorithm (LS-1.b.i)	-Malicious command manipulation on Chiller controller (N-1) causes the controller to undertake incorrect actions such as increase compressor speed when the opposite is required	
1.2	Chiller controller increases compressor speed when permissives for this action are not available after equipment addition or plant configuration changes	1. Inadequate control algorithm: - Specified control algorithm becomes flawed over time due to equipment configuration changes (LS-1.b.iii)	-Replacement of components or change in system configuration makes previously applied security measures obsolete. For instance, update of firmware on VFD inadvertently removes previously applied security measures	Unsafe Controller Behaviour

System Loss: Either of these scenarios would result in irreversible damage to the chiller compressor. For instance, if the oil temperature and pressure is not at the required level, running the compressor even for short periods of time would result in gears and bearings to wear out (it can be thought of as equivalent to driving a car without lubrication oil in the engine).

UCA-4: Chiller Controller increases compressor speed when permissives for this action **are** available, but the VFD actuates the motor in the reverse direction

	Scenario	Associated Causal Factors	Rationale	
2.0	Chiller controller commands the compressor to increase speed, but the compressor motor begins to run in the reverse direction. Compressor speed increases but in the reverse direction	1. Control action improperly executed: - Control action is sent by controller but received incorrectly by VFD(LS-3.b.i)	-VFD drive allows reverse rotation; Malicious command manipulation at connection (C-1) or node (N-1) sends reverse rotation signal to VFD	Scenarios involving Control Path
		- Control action is received correctly by VFD but VFD responds inadequately (LS-3.b.ii)	-VFD's (actuator N-2) control algorithm maliciously manipulated to send reverse rotation command to motor	
		2. Feedback or information not received: - Feedback/info does not exist in control structure to inform the controller if the command from the controller has been executed correctly (LS-2.a.iv)	-Feedback from VFD to controller considered redundant and hence not provided - instead the controller bases its decision on temperature readings to determine if an increase in speed is required; reverse rotation is not detected.	Inadequate Feedback and Info

System Loss: Reversing the rotation still enables forward-direction fluid flow, but prevents the compressor from pumping the fluid to the correct discharge pressure. This results in the chiller surging i.e. the compressor is unable to raise the head sufficiently to lift the refrigerant from the evaporator to the condenser, leading to reverse flow through the compressor. Refrigerant level and compressor motor current fluctuate drastically (several times a minute). Such a conditions can result in the compressor thrust assembly and bearings to be damaged in addition to causing further damage to the gearbox.

UCA-5, 6: Chiller Controller increases compressor speed when permissives for this action **are** available, but the VFD actuates the motor to a value that is different than the one requested.

	Scenario	Associated Causal Factors	Rationale	
3.0	Chiller controller sends a start/run permissive to the VFD control unit with a target value, but the VFD sets the compressor speed to a value other than the target value	1. Control action improperly executed: - Control action (i.e. Target Speed) is sent by controller but received incorrectly by VFD (LS-3.b.i)	-VFD is independently energized i.e. independent control path from network to VFD exists which enables its control parameters to be tampered with. Malicious command manipulation on connection (C-1) tampers signal to VFD	Scenarios involving Control Path
		- Control action is received correctly by VFD but VFD responds inadequately (LS-3.b.ii)	-VFD is compromised; command manipulation at VFD results in output voltage/frequency to be incorrectly set (such as by multiplying or dividing by a large number)	
		2. Feedback or information not received: - Feedback/info does not exist in control structure to inform the controller if the command from the controller has been executed correctly (LS-2.a.iv)	-Feedback from VFD to controller is considered redundant and hence not provided - instead the controller bases its decision on temperature readings to determine if an increase in speed is required. Therefore, the controller does not know if the speed has been correctly set or not	Inadequate Feedback and Info

UCA-5, 6: Chiller Controller increases compressor speed when permissives for this action **are** available, but the VFD actuates the motor to a value that is different than the one requested.

	Scenario	Associated Causal Factors	Rationale	
3.1	Chiller controller sets the target speed to the attached load's critical speed(s) **OR** throttles compressor speed between upper and lower operating limits	1. Inadequate/malformed control algorithm: - Flawed implementation of the specified control algorithm (LS-1.b.i)	-VFD allows reading of critical speeds and programming to skip critical speeds via network. Malicious command manipulation on Chiller controller (N-1) causes the controller to undertake incorrect actions such as set compressor speed to critical speed of the attached load or throttles compressor speed between extreme values (as in the Stuxnet case)	Unsafe Controller Behaviour
		2. Inadequate feedback is received: - Sensors respond adequately but controller receives (and by extension transmits to DCS and operator) incorrect feedback/info (LS-2.b.i)	-Feedback from sensors is scattered by malicious actor (Communication drop or delay); instead normal operating values are maliciously injected to controller and operator; controller or operators do not know correct operating conditions	Inadequate Feedback and Info

System Loss: Setting the speed incorrectly could result in loss of cooling function for the chiller i.e. chiller unable to meet local cooling demand. Throttling the compressor speed between extreme values or running the compressor at its critical speed(s) could result in permanent damage to compressor thrust assembly and bearings as well as gearbox due to fatigue.

UCA-1: Chiller controller sends the signal to increase speed but the command is not executed

	Scenario	Associated Causal Factors	Rationale	
4.0	Chiller controller sets a new target speed for the VFD control unit, but the command is not executed	1. Control action not executed: -Control action (Target speed) is sent by controller but not received by VFD (LS-3.a.i)	-Malicious command drop or delay; VFD spoofed with malicious data and unavailable to respond to controller command	Scenarios involving Control Path
		-Control action (Target speed) is received by VFD but VFD does not respond (LS-3.i)	-VFD is incapacitated; as shown by Angle [1], malicious logic may have blown capacitors on the VFD, thus incapacitating it	
		2. Feedback or information not received: - Feedback/info does not exist in control structure to inform the controller if the command from the controller has been executed correctly (LS-2.a.iv)	-Feedback from VFD to controller is considered redundant and hence not provided - instead the controller bases its decision on temperature readings to determine if an increase in speed is required. Therefore, the controller does not know if the command has been received by the VFD	Inadequate Feedback and Info

System Loss: This scenario would prevent the chiller from performing its primary function of chilled water at the desired set-point. Such a condition would result in loss of cooling for the campus.

VII. DISCUSSION

Up until now, it has been shown how control actions under certain system states can become hazardous and potentially result in system-level losses. We will now analyze how new component, procedural and managerial constraints can be defined and enforced at various levels of the hierarchical control structure to prevent unsafe control actions from materializing into hazardous system states that have the potential to cause system-level losses. Figure 16 shows a summary of the proposed constraints superimposed on the hierarchical control structure of the cooling capacity control loop; the proposed constraints are discussed below.

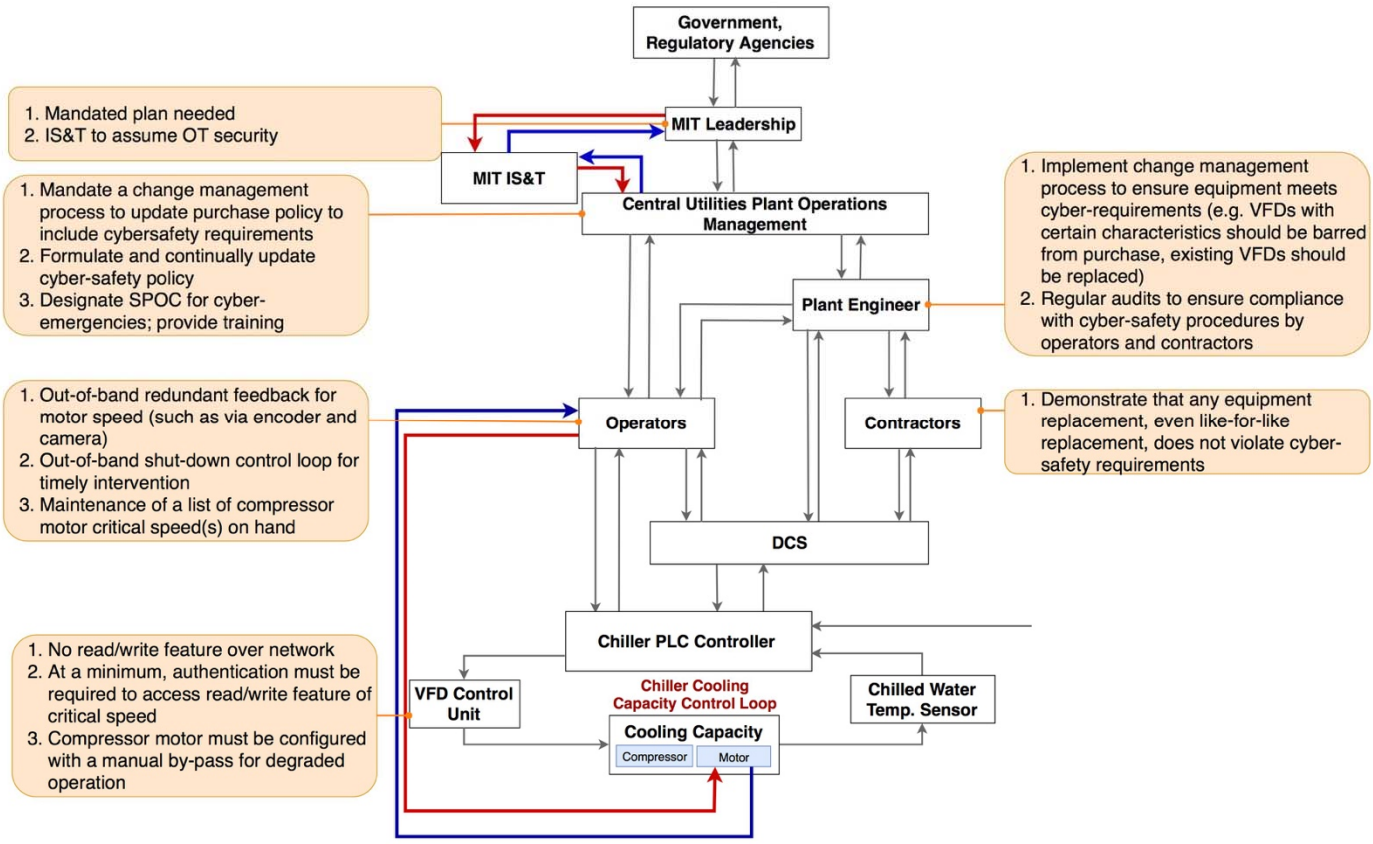

Figure 16 - Component, Procedural and Managerial Constraints defined throughout the hierarchical control structure

VFD Selection – Reverse Rotation via PLC

A review of the loss scenarios indicates that scenarios #2 through #5 involve VFDs. Regulating the selection of VFDs to have certain characteristics would limit some of these loss scenarios. For instance, scenario #2 is caused by the VFD's ability to enable reverse rotation via digital input from PLC. This scenario could be prevented by constraining the selection of VFDs to only those VFD types that do not allow reverse rotation via the PLC; instead require physical reversal of cables to enable motor reverse rotation (as shown by component-level constraints CC-1.0 and CC-1.1 in Table 11).

VFD Energization Source

Similarly, loss scenario 3.0 is possible because the VFD drive has an additional source of energization other than the PLC, making it vulnerable to attack; sometimes this architecture is implemented to have the flexibility to control the VFD from the SCADA/Building Management System etc. The logical constraint is that the VFD must be protected from being energized from anywhere other than the chiller PLC (CC-2.0).

Furthermore, this protection should be enforced at the VFD level rather than the PLC or DCS level so that changes/updates to the DCS or Chiller PLC firmware do not inadvertently remove this protection (CC-2.1).

Out-of-Band Control Loops

These constraints would not eliminate the hazard; however, they would limit the attack surface for a malicious actor. The impact of this scenario could be limited further by defining yet another constraint to measure and display the compressor motor speed to the operator through a secondary redundant mechanism, such as with an encoder attached to a display unit that is independent of the primary control system. A further constraint could be defined to enable compressor motor shutdown by operator via a redundant network that is independent of the primary control system; the logic being that if discrepancies between motor speed readouts are noticed, the operator is empowered to intervene and take effective evasive measures that are independent of the control network under attack. Such an intervention would also require the operator to have motor critical speeds on-hand so that in the event of a cyberattack, the operator can identify anomalous behavior and take remedial control actions.

VFD Selection – Read/Write feature over network

Certain VFD manufacturers allow the critical speed(s) to be *readable and writeable without authentication* [12], which can enable occurrence of loss scenario 3.1. The critical speed setting is not a setting which needs to be changed during operation; it only ever needs to be changed if the motor is swapped out or replaced. This scenario could be prevented by defining a constraint to not allow use of frequency drives that have a writeable feature (over a network protocol) (CC-3.0).

VFD Speed Feedback to PLC

The occurrence and impact severity of loss scenario 3.1 could be limited by the already defined constraints CC-2.2, 2.3 and 3.0 i.e. redundant display of compressor speed with availability of operator intervention via independent network and use of VFDs without read/write feature over network protocol. Loss scenario 4.0 could be mitigated by defining a constraint to configure the VFD to provide the output speed feedback to the PLC; in some control configurations this feedback feature from the VFD is considered redundant and hence not implemented. In the context of cybersecurity, however, this feature could prove useful in confirming that a control command, as sent by the PLC, has been received by the VFD without disruption in the connection path or at the VFD.

Mechanical By-Pass for VFD

The impact of scenario 4.0 could be mitigated by mandating the use of a traditional mechanical bypass for the compressor motor (CC-4.1) i.e. a separate motor starter that is mechanically interlocked with its companion VFD output contactor in a way that allows only the VFD or the bypass to operate the motor at any given time and requires manual activation to engage the bypass [15]. Such an implementation would enable function of the compressor to deliver the primary-value function of providing chilled water to the MIT campus, albeit in a degraded mode i.e. at the expense of speed control and efficiency.

Mechanical Safety Devices

Scenario 1.0 and 1.1 could be prevented by instituting use of mechanical safety devices such as thermal protection relays and freeze stats to protect against motor overheating and evaporator coils freezing conditions, respectively. Timers that depend on manual inputs to change settings could be implemented to lock out the compressor rather than implementing compressor startup delays via the PLC (CC-5.0).

It is emphasized that some of the constraints defined above, such as use of redundant networks, backup controls, mechanical safety devices etc., to ensure availability are not novel ideas for industrial control system safety but instead are quite common in the industry. The point is that these redundant/backup systems have traditionally been architected with equipment reliability and availability in mind, not cybersecurity. Therefore, a rethinking of the control architecture is required in the context of cybersecurity.

Table 11 - Component Level Constraints

Component Level Constraints	**Loss Scenarios**
CC-1.0: VFD drive selected for use at CUP must not allow reverse rotation via digital command from PLC.	LS-2.0
CC-1.1: VFD drive must be wired such that reverse rotation of the drive is not possible without physically reversing the wire connections.	LS-2.0
CC-2.0: VFD must be protected from being energized from anywhere other than the chiller PLC.	LS-3.0
CC-2.1: VFD energization protection must be enforced at the VFD level so that changes/updates to the Building Automation, DCS or Chiller PLC firmware do not inadvertently remove this protection.	LS-3.0
CC-2.2: Compressor motor must display motor speed via a independent, redundant mechanism to the operator.	LS-3.0, LS-3.1, LS-4.0
CC-2.3: Compressor motor must provide a feature to be shutdown via methods that are independent of the primary control system	LS-3.0, LS-3.1
CC-3.0: VFD must not have writeable feature over network protocol; at a minimum authentication must be required to access read/write feature on VFD	LS-3.1
CC-4.0: VFD must be configured with a feedback to PLC to confirm command receipt	LS-4.0
CC-4.1: Compressor motors controlled by VFD must be configured with a manual bypass	LS-4.0
CC-5.0: Thermal protection relays must be used on all motors to protect against overcurrent.	LS-1.0, LS-1.1

Formal Change Management

Whereas traditional hazard or vulnerability analysis methods focus almost exclusively on technical or equipment constraints, the unique aspect of STPA-Sec analysis is that it additionally considers constraints for the broader socio-organizational control environment as well. For instance, the specific types of VFDs that may be installed at the CUP may be specified by typical hazard analysis methods, but this constraint

would be largely futile if it is not coupled with a formal change management process that is mandated by the CUP management and ensures that all existing VFDs are accounted for and replaced with the correct type of drive. We now analyze some proposed management constraints.

Equipment Configuration Changes & Contractor Management

Loss Scenario 1.3 is possible because addition or removal of equipment from the plant's *secured* control system, has the potential to change the security architecture of the plant. For instance, a technician may swap out a faulty frequency drive with an equivalent drive from a different manufacturer. The new drive may violate component constraints CC-1.0 and CC-2.0 (i.e. allows reverse rotation or read/write capability) and in doing so compromise the security architecture. Hence, management level controls are required to mandate review of cyber-safety impact for any equipment changes at the CUP, including any equipment introduced by contractors.

This means that rather than having a passive approach where maintenance of large equipment, such as chillers, boilers etc., is outsourced to contractors, MIT CUP personnel need to actively scrutinize all contractor actions from a cyber-safety lens (LS-1.3). Any new equipment brought into the plant or any change in configuration needs to be expressly authorized by the CUP management or plant engineer who must ensure that the addition of the new equipment or change in configuration would not adversely impact any component or system level constraints. To expand on this further, the contractor's ability to monitor operation of the equipment in real-time via Plant Information (PI) servers or push firmware updates remotely must also be constrained as doing so limits CUP personnel's ability to control the change process.

Mandated Cyber-safety Plan

Furthermore, in performing this analysis it was discovered that MIT CUP does not have a cybersecurity policy; in the event of a cyberattack, no single person is designated as the Single-Point-of-Contact (SPOC) to manage the emergent situation. By extension, the plant does not have any standard operating procedures to deal with a cyberattack; the plant personnel, including both engineers and operators have not had any training in cybersecurity.

This lack of cybersecurity policy stems from the fact that the regulatory bodies, licensing the operation of the plant, do not require CUP to have a cybersecurity policy in place as a pre-condition to licensure. This is because MIT, being a private institution and CUP exclusively generating electricity for the MIT campus, is not considered a contributor to the Bulk Energy System (BES) and hence is exempt from showing compliance to the cybersecurity standards. A management constraint is therefore required to ensure that CUP develops and maintains a robust cybersecurity policy (MC-2.0) of its own.

Mandated Role for IS&T

A further constraint is proposed at the MIT leadership level to mandate MIT Information Services and Technology (IS&T) or an equivalent body with domain knowledge of operational technology (OT) and to have hierarchical control and authority over cybersecurity policy. By separating the function of cybersecurity from CUP Operational Management's domain, it could be ensured that cybersecurity constraints are not compromised for equipment availability for operations. For instance, most manufacturing departments have quality departments which report directly to the CEO and are outside the sphere of influence of plant production or operations; quality departments are not concerned with plant production, their mandate is to ensure product quality even at the expense of production. A similar arrangement for cybersecurity would

introduce checks and balances that would ensure that cybersecurity policies are not compromised by production/operation considerations.

All the management constraints discussed above are itemized in Table 12. To conclude, we have shown how the impacts of the attack scenarios can be mitigated by eliminating or reducing system vulnerabilities by defining component as well as management-level constraints on the system.

Table 12 - Management Level Constraints

Management Level Constraints	**Loss Scenarios**
MC-1.0: Any change to equipment configuration must not be performed without reanalyzing cybersafety architecture	LS-1.3
MC-1.1: Like-for-like equipment must also undergo a complete cybersecurity assessment to ensure the cyber-safety architecture is not compromised.	LS-1.3
MC-1.2: Any equipment brought into the plant by contractors must also be subject to the same rigorous cyber-safety analysis	LS-1.3
MC-1.3: Contractors must not be able to issue firmware updates to chiller PLC remotely.	LS-1.0, LS-1.1
MC-2.0: MIT CUP must have a comprehensive cybersecurity policy in place that provides guidance and training to plant personnel and engineers for dealing with cyberattacks.	All
MC-3.0: MIT leadership must mandate MIT IS&T or equivalent to develop and implement cybersecurity policy	All

VIII. CONCLUSION

In this paper, we analyzed a single representative control loop (the compressor motor speed control) of an archetypal industrial control system (i.e. the centrifugal chiller) at a small-sized power plant in the context of cybersecurity using a vulnerability analysis method based on *Systems Thinking*. Starting with system-level losses and hazards we traced the functional control structure of the plant and abstracted out the compressor capacity control loop for a detailed analysis of a single control action under various system states. We then generated loss scenarios under which the unsafe control actions would result in system-level losses. Finally, we proposed new constraints at various layers of the functional control structure (starting at the process layer and going all the way back up to the enterprise and regulatory level) to prevent the system from entering unsafe system states.

In the process of performing this analysis, we uncovered several insights about the system (i.e. the centrifugal chiller) which were not obvious at the onset of the analysis; such as, the selection of a component at the process layer, is ultimately linked to a policy level decision at the enterprise level. For instance, the selection of VFDs of certain types (such as with read/write capability over network), though improves convenience through increased functionality, simultaneously introduces new vulnerabilities for the system. Successful elimination of this vulnerability requires an organizational change management process that not only ensures that existing VFDs meet the cyber-safety specification, but that new purchases are systematically vetted out for such vulnerabilities. In this way, the success of the proposed constraint requires the support and cooperation of additional office functions, such as *Purchasing* and *Quality* departments. Typical hazard or vulnerability analysis methods focus primarily on the technical aspect of the system, rather than taking a broader view of the system as a whole.

In addition, the analysis highlighted missing feedback loops both for components (e.g. independent compressor motor speed feedback to operator via encoder) as well as for processes (e.g. re-validation of the security architecture of the plant by the engineer due to addition/replacement of equipment from the plant). In both these cases, *insights emerged almost naturally* due to the use of the functional control structure. The functional control structure also provided a bird's eye view of the entire system by combining organizational, human and technical controllers in a single diagram, enabling a broader view of the system along with key leverage points for enforcement of the proposed constraints.

Even though some of the component-level recommendations may appear obvious to operational technology personnel since these may have been implemented in the past to ensure equipment *reliability and availability*, it is important to restate these in the context of cybersecurity. A vendor or contractor may not see the implications of replacing a piece of equipment with an equivalent type of equipment, but such an action would violate the cyber-safety constraint that could result in a loss.

The STPA-Sec analysis uncovers such vulnerabilities and then proposes constraints at the procedural level (e.g. requisition of new equipment requires engineer to demonstrate the purchase does not violate cyber-safety specification) and further through policy by higher-level controllers, to more effectively manage vulnerabilities. Ultimately, the method reimagines the security problem as a dynamic control problem where the focus is on identifying and controlling vulnerabilities within the system rather than capabilities of external threat actors that are beyond the control of the system.

Using a top-down approach and starting with system-level losses, the analysis always maintains focus on the bottom-line i.e. what constraints, if violated, would result in the system entering an unsafe state that could propagate into a system-level loss. This enables the STPA-Sec method to be more strategic in

identifying the most critical vulnerabilities. Furthermore, since each step of the analysis is always tied back to system-level losses, the method provides traceability between recommendations and losses. This helps to communicate with policy and decision-makers who can more clearly see how the recommended changes can help mitigate vulnerabilities in the system and how those vulnerabilities are linked to system-level losses.

In conclusion, the STPA-Sec method provides a well-guided and structured analysis process to identify vulnerabilities in complex socio-technical systems. It ties system-level losses to violation of constraints at both the component-level as well as the process level and provides recommendations to make the system more resilient by controlling vulnerabilities by defining additional constraints.

IX. REFERENCES

[1] M. Angle, "Identifying and Mitigating Cyber Attacks that could cause Physical Damage to Industrial Control Systems", August 2017, [Online]. Available: http://web.mit.edu/smadnick/www/wp/2017-14.pdf (Last Accessed: July 30, 2018)

[2] N. Levenson, STPA Handbook, March 2018

[3] "MIT Central Utilities Plant Second Century Project", Single Environment Impact Report, Epsilon Associates, EEA #15453, [Online]. Available: http://powering.mit.edu/sites/default/files/documents/05122016_SEIR_MIT_Submittal_rev%20reduced2.pdf (Last Accessed: July 30, 2018)

[4] "Centrifugal Compressor Surge-Basics, Mechanism", April 2017, [Online]. Available. http://www.mechanicalengineeringsite.com/centrifugal-compressor-surge-basics-mechanism/ (Last Accessed: July 30, 2018)

[5] E. Colbert, "Security of Cyber-Physical Systems", Journal of Cyber Security and Information Systems, Volume 5, Number: 1 – Cyber Science & Technology at the Army Research Laboratory (ARL), January 2017. [Online]. Available: https://www.csiac.org/journal-article/security-of-cyber-physical-systems/ (Last Accessed: July 30, 2018)

[6] I. Friedberg, "STPA-SafeSec: Safety and security analysis for cyber-physical systems", Journal of Information Security and Applications, June 2016, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214212616300850 (Last Accessed: July 30, 2018)

[7] W. Young, NG. Levenson, "An integrated approach to safe and security based on systems theory", Commn ACM 57 (2), pp. 31-35, Feb 2014, [Online]. Available: https://cacm.acm.org/magazines/2014/2/171683-an-integrated-approach-to-safety-and-security-based-on-systems-theory/abstract (Last Accessed: July 30, 2018)

[8] K. Stouffer et. Al., "Guide to Industrial Control Systems (ICS) Security", Rev. 2, NIST Special Publication 800-82, May 2015, [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-82r2

[9] "MIT Cogeneration Plant – Homepage", [Online]. Available: http://cogen.mit.edu (Last Accessed: Feb 10, 2018)

[10] "Chiller Training Video", Solano College, [Online]. Available: https://www.youtube.com/watch?v=NqkjiBn0Qto (Last Accessed: July 30, 2018)

[11] B. Twining, "Type B Accident Investigation Board Report Chiller Line Rupture at Technical Area 35, Building 27, Los Alamos National Laboratory", Feb 1998, [Online]. Available: https://www.energy.gov/sites/prod/files/2014/04/f15/9711lanl.pdf (Last Accessed: July 30, 2018)

[12] K. Zetter, "An easy way for hackers to remotely burn industrial motors", Dec 2016, [Online]. Available: https://www.wired.com/2016/01/an-easy-way-for-hackers-to-remotely-burn-industrial-motors/(Last Accessed: July 30, 2018)

[13] "What causes electric motors to overheat?", Industrial Motors and Machining Inc, [Online]. Available: http://www.electricmotorrepairsales.com/blog/electric-motor-overheat-reasons.php (Last Accessed: July 30, 2018)

[14] E. Chein, "Stuxnet: A Breakthrough", Symantec Official Blog, November 2010, [Online]. Available: https://www.symantec.com/connect/blogs/stuxnet-breakthrough (Last Accessed: July 30, 2018)

[15] T. Trullinger, "Web Exclusive: VFD bypasses and backups: Which should you use?", Consulting-Specifying Engineer, Dec 2012, [Online]. Available: https://www.csemag.com/single-article/web-exclusive-vfd-bypasses-and-backups-which-should-you-use/1c977d00d32608b95f670db50970889e.html (Last Accessed: July 30, 2018)