INL/CON-10-19561
PREPRINT

Notional Examples and Benchmark Aspects of a Resilient Control System

3rd International Symposium on Resilient Control Systems

Craig G. Rieger

August 2010

Notional Examples and Benchmark Aspects Of a Resilient Control System

Craig G. Rieger, *Senior Member, IEEE*
Idaho National Laboratory, Idaho Falls, Idaho, USA

Abstract — **Digital control system technology has pervaded most industries, leading to improvements in the efficiency and reliability of the associated operations. However, the ease of distributing and connecting related control systems for the purposes of increasing performance has resulted in interdependencies that can lead to unexpected conditions. Even with less complex designs, operators and engineers alike are often left with competing goals that are difficult to resolve. A fundamental reason for this dichotomy is that responsibilities lie with different disciplines, and operations are hosted on separate control systems. In addition, with the rising awareness of cyber security and diverse human interactions with control systems, an understanding of human actions from a malicious and benevolent standpoint is necessary. Resilience considers the multiple facets of requirements that drive the performance of control systems in a holistic fashion, whether they are security or stability, stability or efficiency, human interactions or complex interdependencies. As will be shown by example, current research philosophies lack the depth or the focus on the control system application to satisfy these requirements, such as graceful degradation of hierarchical control while under cyber attack. A resilient control system promises to purposefully consider these diverse requirements, developing an adaptive capacity to complex events that can lead to failure of traditional control system designs.**

Keywords: resilient control; cyber awareness; human systems; complex networked control systems; data fusion; and cyber physical systems.

I. INTRODUCTION

To be considered a benchmark, an ad hoc or defacto standard is established that forms a basis for performance comparisons. In the case of this paper, notional industrial examples are used to illustrate limitations of current control system research philosophies in achieving resilience. Those limiting aspects are categorized and form a qualitative benchmark by which a new research philosophy for resilient control systems is based. Before discussing the philosophy of resilient control systems, however, one must first consider its definition. A resilient control system is defined as one that maintains state awareness and an accepted level of operational normalcy in response to disturbances, including threats of an unexpected and malicious nature [1]. The notional architecture of such a system is defined in Figure 1.

Figure 1 is laid out to parallel the functional aspects of a traditional control system, with interfaces to the plant and operation on the left, control philosophy and data analysis in the center, and human interfaces on the right. While at a fundamental level these aspects will remain in a resilient control system, the details of the control philosophy, data analysis, and presentation of information have changed. At the left, Operational Data and Controls provide interfaces not only to the operation, but also to other indicators of plant performance. The security of the system must be considered alongside the stability, as it is one indicator of control system integrity. The efficiency of the operations is important to the economics of operation, and even if the operation is stable, the efficiency may impact the financial margin for the associated asset owner. Moving to the center of the figure, a Data Fusion approach is employed that can process this diverse data to proactively recognize threats within each performance measure and prioritize response. The Mixed Initiative Control framework provides mechanisms to integrate automation and human response in an optimized manner, benefitting from the inherent resilience in both. The Hierarchical, Multi-agent Control Design provides an adaptive mechanism for optimizing control system performance to measures of normalcy [2]. To the right, information is targeted to the consumer of the information, tailoring what is presented and how to ensure a reproducible response.

While one might say that resilient designs are built upon dependable computing research, these research philosophies do not characterize the aspects of the design that involve the implementation of advanced control theory for feedback control functionality [1]. Dependable computing research considers the malicious faults as a source of failure but does not consider these faults or the associated consequences in terms of impact on the unique design considerations of a control system. Even so, cyber research is very immature in its development of widely accepted solutions.

While fault tolerant or reconfigurable control technologies have been under research for some time, neither of these provides the comprehensive plan to maintain performance in the face of threats. While little research has been done to link the aspects of control action with the fault detection and diagnosis (FDD) [4], which is necessary to determine state awareness, even current detection understanding does not consider malicious action to undermine normal system behavior. In addition, the development of a higher level hierarchical methodology that fuses and prioritizes incoming information to ensure state awareness and optimize response is lacking. Furthermore, the consideration of both the human and automation as working partners in the control algorithm is disconnected and separate, with little research to measure and base the level of automation on the resilience of the human [5].

In the sections that follow, a review of limitations in current research philosophy will be discussed, notional examples provided, and resilience improvements in current technology suggested.

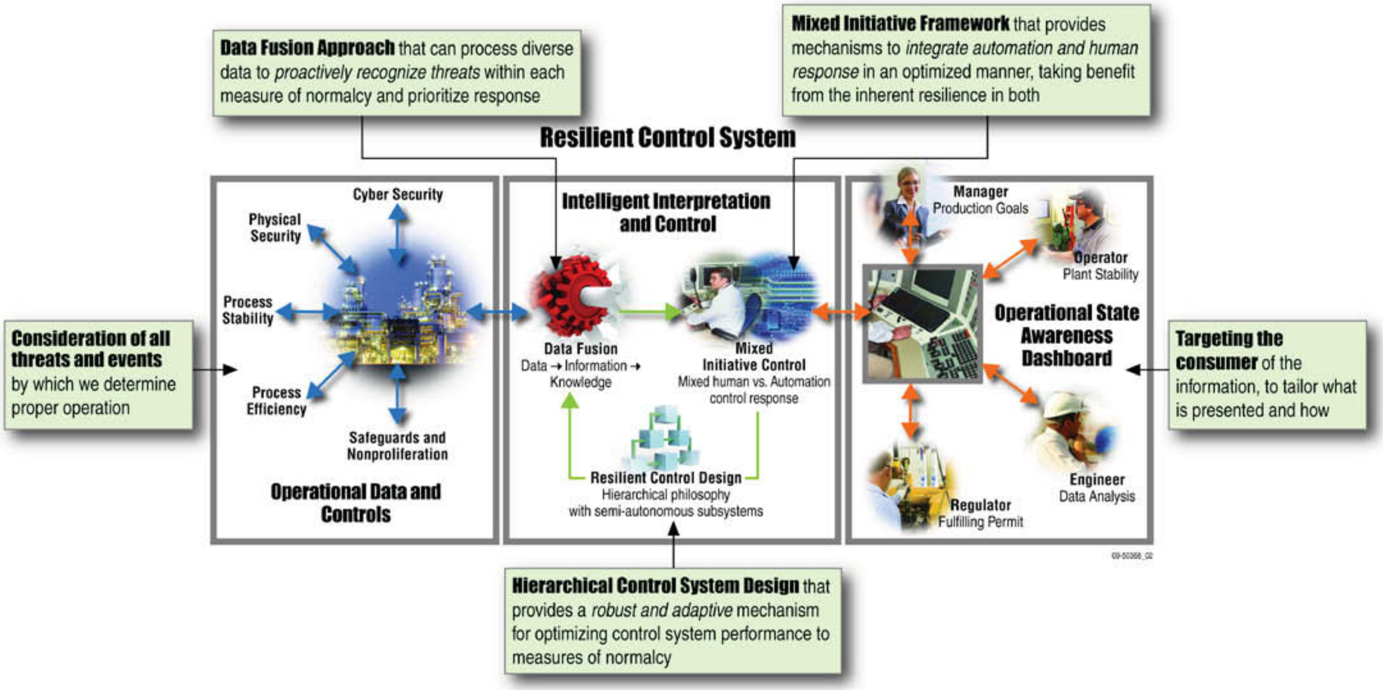
Data Fusion Approach that can process diverse data to proactively recognize threats within each measure of normalcy and prioritize response
Mixed Initiative Framework that provides mechanisms to integrate automation and human response in an optimized manner, taking benefit from the inherent resilience in both
Resilient Control System
Cyber Security
Physical Security
Process Stability
Process Efficiency
Safeguards and Nonproliferation
Operational Data and Controls
Intelligent Interpretation and Control
Data Fusion
Data → Information → Knowledge
Mixed Initiative Control
Mixed human vs. Automation control response
Resilient Control Design
Hierarchical philosophy with semi-autonomous subsystems
Manager
Production Goals
Operator
Plant Stability
Operational State Awareness Dashboard
Engineer
Data Analysis
Regulator
Fulfilling Permit
Consideration of all threats and events by which we determine proper operation
Targeting the consumer of the information, to tailor what is presented and how
Hierarchical Control System Design that provides a robust and adaptive mechanism for optimizing control system performance to measures of normalcy
09-50368_02

Figure 1. Resilient Control System Framework.

II. Current Research Philosophy Limitations

A. *Unexpected Condition Adaptation*

1) Achievable hierarchy with semi-autonomous echelons

The concept of cascading or supervisory control is not a new one. However, control systems that implement these concepts have often been designed for coupling few control loops, or have been designed to provide parameter and set point modification specific to known operating conditions. Logic interlocking of control system actions has also been a standard practice for regimenting the responses to defined states. While these methodologies have been successfully implemented in very reliable applications, they have also been an underlying cause of failure [6], [7]. Furthermore, many examples exist that illustrate the failure of complex automation due to the loss of required sensors or data for the decision process, or because of an incomplete understanding of the operating modes for an application [8]. The basic limitation of these designs is intrinsic brittleness, in that while there may be a well-defined path for degradation, unexpected or untested disturbances may cause the system to extend outside of the desired boundaries of operation. Whether the control system designers lacked full understanding of the interdependencies or made simplifying assumptions in regard to the loss of information, graceful degradation did not occur. The defined path for degradation depends on the engineer to understand the fault paths in order to develop a corresponding action, which can include redundant pathways.

2) Complex interdependencies and latency

While the implementation of advanced control theory into control systems of today has been only gradual, the acceptance of technology has lead to the migration of industrial operations to some form of digital control system. These systems are often designed and implemented in a local fashion, assuming facility stability will be achieved by the stabilization at the local level. However, as more and more systems become interconnected, even across corporate and industrial sectors, less is understood regarding the couplings that may exist. Concerns over the impacts of latencies on complex control algorithms have generated basic research involving the mitigation of these latencies. Simplistic methods that determine the sensitivity of individual control techniques provide a limited view of how individual or a small number of feedback loops might be made more robust. However, these methods do not speak to the resilience of the system when multiple, often unknown, latencies may exist due to the complexity of the system. To ensure that graceful degradation of control system designs are provided in light of resilience, the overall architectures of these systems needs to reflect a hierarchy that resolves itself to threats.

B. *Human Interaction Challenges*

1) Human performance prediction

Human performance has long been an issue in regard to control system design and operation. While traditional concerns were often specific to ensuring correct judgment through the appropriate presentation of information, which remains a concern today, a more comprehensive interest involves measuring operator effectiveness for multiple control system interactions. Not unlike a feedback loop in a control system design, effectiveness can be measured by placing sensors on the operator that can be analyzed to prognosticate the risk in whether the right judgment will be made or to supplement with automation. Unlike traditional design, however, the human contribution to resilience can be beneficial or detrimental. In addition, research to measure human effectiveness is useful, but far from conclusive. Furthermore, design of control systems involve a number of disciplines, and like any measure of effectiveness, will be no better than the competency and proficiency of the least capable individual.

2) *Cyber awareness and the intelligent adversary*

Cyber security is a concern by a large group of security professionals. However, the effort to focus specifically on those that understand security overlooks the fact that other work disciplines are also responsible for security, or rather the security of the control system operation. In addition, designs that require more passwords to enhance layered security protection, has diminishing benefit because of individual performance to implement consistently. Research and development to date has investigated better mechanisms to detect attacks, understanding attack vectors, and developing threat models. However, while some existing near term improvements might be realized through these research paths, they do little to improve the inherent resilience of the overall design. To provide measureable improvement in this design, one must reduce the overall complexity of implementing and maintaining the security protections. To accomplish this task, a combination of passive and active techniques will be required. Those techniques that have an inherent nature to deflect the attacker, such as decoys and randomization, will be passive. Active techniques not covered here include traffic rerouting based upon perceived threat.

C. *Goal Conflicts*

1) *Multiple performance measures*

Performance of control systems is based upon several measures, not just process stability, but also physical and cyber security, process efficiency, and process compliancy. Fulfillment of these goals has been the responsibility of a diverse group of organizations, with little in the form of standardized philosophy. Such a philosophy, or measurement of the capacity remaining in the system before failure, is required and must be based on each measure of performance. As these measures are also not considered in a holistic manner, fulfillment of overall performance or determination of priority can elude both the designers and the operators of the control system. Without holistic view, or frame of reference, little in the way of normalizing comparisons to ensure the proper weighting of each is achieved. This situation can create a greater risk of responding to a lower priority situation or creating conflicting priorities.

Considering energy efficiency, for instance, heat sinks and sources exist throughout processes and processing facilities. Without a mechanism to consider all of these points, it is unrealistic to assume that the optimum efficiency is reached, or even known. It is the latter which is most important, specifically from a state awareness standpoint. To prioritize a response based upon a metric, an understanding of the efficiency is required for comparison. Therefore, while the optimum efficiency may be sacrificed to ensure safety, prioritization is based on complete understanding.

2) *Lack of state awareness*

Adding to multiple measures of performance, maintaining observability on conditions that characterize a shift from normal must also be achieved. While research in the areas of prognostics, diagnostics, condition-based maintenance, and online monitoring have been a research field for over two decades, very little of this technology has engendered itself into a control system design. The development of an observer for an unmeasured variable provides a useful analogy, but does not encompass what is effectively a separate "online monitoring" discipline that uses multiple techniques to assess system condition. When considering measures such as cyber security, these technologies have not been applied, and in general, the field of cyber security metrics is in its infancy.

III. Examples

Several notional examples are developed below to illustrate the limitations of current research philosophies in addressing the need. These include a power transmission substation, chemical facility reactor, and a heating, ventilation, and air conditioning (HVAC) system for a hazardous facility. In discussing current research philosophies, it is assumed that existing control system technologies are included, and will therefore not be identified separately. In identifying the limitations and resulting impact, the term adaptive capacity will be used. Adaptive capacity, in this context, may be used to describe the allowable loss in system functionality before a loss of acceptable performance is recognized. Finally, references to operations refer to individual, identifiable aspects of an industrial operation, often referred to as unit operations within the process industries. In like fashion, an operator (dispatcher for power operations) is considered the individual with direct responsibility for monitoring plant condition and making appropriate changes to maintain normal operation. A designer is considered the individual who develops the control system theory of operation.

A. *Power Transmission Substation Scenario*

A transmission substation is currently interfaced to a supervisory control and data acquisition system (SCADA), down to the relay level for monitoring and control. This substation is part of a large transmission system, and has state-of-the-art gear with IEC-61850 interfaces. The communications to each relay is over an Internet Protocol (IP) based network, with standard Information Technology (IT) routing and segmenting equipment. The substation switchgear is protected from the environment by a building, and for physical security, has a locked door and surrounding fence with a locked gate. Inspection visits to the substation occur infrequently and normally only for maintenance purposes.

During a backshift, a dispatcher receives an indication that one of the relays in a substation has caused a breaker to isolate a critical line, which would reduce capacity and potentially cause a blackout in a section of a neighboring large city. However, the monitoring data from a downstream phasor measurement unit (PMU) indicates that power is still flowing to the city. Within one scan cycle of the substation SCADA system, the breaker status returns to the normal closed position. To be cautious, the dispatcher sends a crew to investigate the associated substation to confirm the status of the breaker. In the course of investigation, the crew finds that the lock on the fence is missing and the door to the substation is unlocked. Had the door been opened, a status alarm should have been indicated back at the control center. Investigation of the substation indicates that the breaker with the suspect condition is still closed and operating correctly. As the substation appears undisturbed, other than the missing fence lock, it is concluded that further investigation is unnecessary.

In a week following this incident, a different crew is on the backshift when numerous calls are received that the neighboring city has lost power. The indication in the control center seems to reflect an absence of problems. Again, a crew is dispatched to the substation that had been inspected the week before. While the physical security of the substation does not appear to be compromised, it is found the breaker previously investigated is now open, as are several others, creating an overload condition on the transmission lines to the city that eventually tripped power. A foreign wireless communications device was found with an investigation by security, indicating a back door was created by an individual entering the substation.

B. Chemical Facility Reactor Scenario

A chemical reactor unit operation is automated with a state-of-the-art distributed control system (DCS). The DCS provides multivariable control of the reactor, which is provided via an optimal control methodology. The sensors that provide the data for this multivariable design are interfaced to multiple redundant controllers based on proximity to the process equipment, requiring exchange of data to the controller hosting the optimal control design. The communications system is an IP based design, which interconnects all of the controllers. The DCS system is isolated from the business systems via standard information technology (IT) devices, namely firewalls, segmentation, and demilitarized zone (DMZ) protections.

During the operation of this system, a failure of a group of sensors occurs. Depending on the type of failure, this event may or may not be recognized and responded to by current research philosophies. If they fail to normally accepted high or low levels and generate an alarm, they will be easily addressed. However, if they fail in a known good state or normal range, limitations in current research philosophies become apparent and do not address the issues addressed by the situation in a holistic manner. This failure could be due to cyber attack specific to an OLE for Process Control (OPC) server or a wireless access point, or it could be due to software failing in an undesirable or unexpected manner.

C. Hazardous Facility HVAC Scenario

A facility that is producing hazardous substances has an advanced HVAC system for regulating pressures within the building. By maintenance of pressures with the most hazardous areas at the lowest pressure and normally occupied spaces at the highest, the migration of hazardous substances can be prevented. The system design also uses supervisory control, in that a neural network design implements night-time setbacks increases the air conditioning set points to reduce overall energy usage. The primary temperature and differential pressure control on the system are through some form of PID algorithm, with each hazardous zone having its own controller and separate temperature controllers for the hazardous and occupied zones. Intake and discharge ducting and blowers are common for the hazardous and occupied spaces, but each area has an individual header. In the case of the hazardous areas, high efficiency filters are used to remove pollutants.

During the morning before the workers arrive, the exhaust airflow from the facility gets largely blocked due to an abnormal failure of a damper. This creates a back pressure on both the hazardous and occupied zones of the facility. In response to the reduced airflow, the inlet damper of each hazardous zone closes to maintain the required differential pressure. However, minimum facility flows are not maintained and the damper controls are not able to equalize consistently, allowing periods where potential migration of hazardous species may occur. In addition, the drop in airflow prevents cooling and allows the temperature to increase in both the occupied and hazardous areas. As the airflow through the air conditioning coils has dropped, the PID controller continues to increase the amount of coolant to the coils until they freeze—which the freeze protection switch, or freeze stat, fails to prevent due to improper positioning. The regime that the facility is now operating within has also gone outside of the training for the neural network, but as the occupied period is reached, the neural network decreases the temperature set points without regard to the abnormal situation.

IV. Power Transmission Substation Scenario

A. Unexpected Condition Adaptation

Because of a cyber event, intrinsic redundancy within the substation to prevent overload has been compromised. The attacker can use readily available vendor software for configuring the relays, which are often left with the original vendor default passwords, to configure relay settings to abnormally fault or not to bypass any intended equipment and reliability protections. Cyber security is the responsibility of the security group, not the dispatcher, so any available data on the security would have to be recognized by the security engineer and reported to operations. However, in this case, port security is not turned on in the substation to prevent connection of the foreign wireless device, or access point, to the communications system. As intrusion detection systems (IDS) are outwardly focused, no anomalies were detected.

The hierarchy of current power designs still requires a centralized control, in spite of the assets being distributed. Although some form of decentralized control is present (for example, the dependence on relay devices to protect the transmission lines from abnormal circumstances) energy management system (EMS) applications provide a centralized monitoring and control of the transmission system. During a loss of communications to the energy management system, the individual devices continue to operate at the last settings. As "coordination" of the system is dependent upon the EMS and dispatcher interaction, the cyber compromise provided in this example created a situation where overloading of the system is possible without the awareness of the central authority.

B. *Human Interaction Challenges*

A physical security alarm should have been recognized by the dispatcher at the control console when the intruder entered the substation. As it turns out, these alarms can become a nuisance and are often low priority to operation alarms. In the case of this substation, the alarm had been turned off due to recent maintenance operations and had never been reactivated. The attacker bypassed the IDS systems by attacking at the endpoint devices, in this case a network, and one connection allowed access to multiple devices. Had port protections been implemented, the attacker would still have acquired direct access to the network switch and network connectivity.

C. *Goal Conflicts*

Operation of the power system and the cyber security of the system is a multidisciplinary responsibility. However, the health of the system is equally as important if critical assets are compromised. As the dispatcher has no authority in this area or data, no attention to this consideration will be made. A status alarm from the substation door being opened would have provided some evidence of an impending compromise, but these alarms are considered low priority even when they are active.

V. Chemical Facility Reactor Scenario

A. *Unexpected Condition Adaptation*

The reconfigurable control philosophy would consider known failures and provide a corrective action to maintain a level of normalcy desired based on the failure and available actions. However, if the sensors failed in an acceptable range, meaning the information appears good but is not accurate, the FDD may not address the problem. This limitation exists because of both the state of failure and the cause, which can also include cyber compromise. As the mechanisms of cyber attack are not characterized in current FDD design and can exhibit themselves in terms of corrupt data or latencies, this type of failure would be missed entirely. In addition, the methodology to characterize multiple sensor failures and develop an appropriate control response will still be necessary.

By its nature, an optimal control algorithm is developed for an operation, implying a multi-input, multi-output (MIMO) control design. The implementation of the design may be using a supervisory hierarchy, where the outputs of the optimal control algorithm feeds set points to individual Proportional-Integral-Derivative (PID) controllers, or directly connected to field devices through interface hardware. With current optimal control philosophies, whether the design is H_2 or H_∞, either incorrect data in the transfer of information can cause the resulting algorithm to break down. Timings specifically are crucial to an optimal control algorithm [9], and the failure of the sensors is from a cyber attack that injects latency, the control system may behave unexpectedly. While mechanisms to switch the optimal control algorithm to single loop control can be conceived, these would also depend upon knowledge that the sensors have failed.

B. *Human Interaction Challenges*

Given the failure, an experienced operator may detect that the sensor is not reading the correct value. However, as no alarms are expected, it would require a conscientious individual to determine it in a timely fashion. Depending on the timeliness of the response needed, action on this item may not occur until other operations in the facility are affected, generating another alarm. Depending on the number of operations now affected, the burden on the operator has increased, which will also affect the ability of the operator to respond. In the case of a cyber compromise, current research philosophy does not consider the operator as a player in the cyber security response. However, without some understanding of the fact that the sensor has failed and the circumstances, there is a risk of undesirable responses being taken.

C. *Goal Conflicts*

The failure of the sensors can be due to many events; in this case, a software failure or cyber attack that causes the sensor to fail in a normal range. While the consequences of such a failure can vary, this in itself emphasizes the need for a mechanism to measure the resilience or adaptive capacity remaining so that a response can be prioritized. With the failures unknown to the operator, reaction to other events that are perceived as critical will result. As the responsibility of the operator to recognize unexpected events such as this can vary between industries, it is known that training and expectations can often be very prescriptive. Abnormal event scenarios tend to be defined into operator training, and events outside these scenarios then relegated to the engineer to respond.

VI. Hazardous Facility HVAC Scenario

A. *Unexpected Condition Adaptation*

It is clear that the PID controllers for pressure and temperature control will attempt to maintain a set point as long as they are enabled. As this is normally an automatic operation during even occupied periods, operator intervention is not expected. The failure of the damper created a back pressure disturbance and the PID controllers, acting independently in response, established pressure gradients that were conducive to migration of hazardous species. Independent operation, which worked well during normal operation and expected disturbances, failed when an unexpected event occurred. The introduction of hierarchy would have benefited this situation. However, implementation of traditional supervisory designs has limitations, as can be seen with the nighttime setback neural network. This neural network is designed and trained for particular conditions, and can provide unpredictable performance when outside of these conditions. A method of detecting the condition often may not be satisfactory, specifically if this condition is unlikely or unknown. In traditional supervisory designs, knowledge of disturbances that may impact the control design is necessary to ensure adequate response to conditions.

B. Human Interaction Challenges

While the discussion in the previous two examples focused on the individual interacting with the control system during the scenario (i.e., a dispatcher or operator), the designer clearly plays a significant role in the adequacy of the control system to respond to abnormal events. A designer bases implementation of control logic and controller algorithms on individual experience, or on that of a team. Any response by the control design will be based upon an awareness of potential failures during the design process. The direct implementation of an advanced control method will be limited, or possibly create a worst possible situation, when unrecognized failures occur. The driving motivation to fully automate can lie in the desire to provide a reproducible response. However, as is the case in HVAC designs, even if the designer allows for manual control, the system must be attended. In general, however, locking out operators from full manual operation can create more problems than the automation was intended to remove [10].

C. Goal Conflicts

The HVAC system design has two apparent operating goals: temperature and differential pressure. In this case, however, there is no overarching mechanism to ensure that one goal is maintained over the other. During normal circumstances, this may not be an issue, but as an unexpected event occurs, this design characteristic may create one. A clear mechanism to consider the two goals would allow for prioritization of the differential pressure versus the temperature goal. In addition, a proper pathway to degradation is also needed to prevent undesirable changes, such as that of the neural network controller continuing to reduce the temperature set point when coil freezing was occurring. To enable the prioritization, however, an awareness of the physical failure through direct measurement or analysis would be needed. As there is no way to ensure that all unexpected conditions can be determined directly, a mechanism where established conditions can be monitored and anomalies detected provides a more resilient solution.

VII. Resilient Control System Enhancements

A. Power Transmission Substation

In recent years, several researchers have proposed multi-agent designs as a way to increase performance of these systems [11], [12]. While all of these have essentially been academic exercises, their introduction has merit in the fact that they consider multi-level autonomy and peer negotiation for resources. Dependence on a centralized control room and automation, such as an EMS, which creates a dependence upon communication links that can be destroyed during natural disasters or cyber attack. What multi-agent design can prevent are rapid shifts in load as resources can be negotiated at the substation, depending on available assets. What this scenario also portrays is the need to have cyber awareness, which is not an independent aspect of control systems. While a substation break-in was the entry point for the malicious device for this scenario, other avenues (such as an employee's compromised thumb drive) can be used to develop holes in current security protections. It is clearly important to ensure designed protections and controls are not overridden by malicious changes. All information that is pertinent to the condition of the system is required. Without the cyber information, decisions are made based upon an inaccurate understanding of the situation. The ability to respond quickly based on state awareness of the conditions provides adaptive capacity to the control system.

B. Chemical Facility Reactor

The root cause of the failure in this scenario is loss of state awareness. While traditional redundancy, even triple modular redundancy, can maintain sensor information, unexpected failures lead to difficulties in prediction. This failure, while the root cause, indicates many of the aspects of resilience that are needed to fulfill the ultimate goal of maintaining an acceptable level of operability. A carefully defined data fusion framework is appropriate to confirm full state awareness of the sensor loss [13]. The analysis and prioritization of the data will provide a basis for response, and is necessary to discriminate between failures that are software/hardware related and cyber related, as the appropriate response will change. The appropriate response to the situation will be a blend of automation and human interaction. If a cyber event occurs, there is little in the way of current automation measures that would restore this situation. Therefore, primarily human response is expected, and differs dependent on the player. For the operator, the knowledge that the failure is cyber related would prevent an inappropriate action, such as making set point adjustments for the operation or call instrument technicians to repair the failure. For the security engineer and network technicians, modifications to isolate the network path for the cyber attack will be part of an appropriate response. If the data fusion system recognizes the software failure, an automated "resilience" response may be used to complement the traditional fault tolerance design and provide, as a minimum, an indication of a module to be removed by a technician. The associated optimal control algorithm may be operated in a degraded fashion, where feasible to use remaining inputs, or require the direct interaction of an operator to regulate.

C. Hazardous Facility HVAC

This system design would have benefited from a hierarchical control methodology and operator interaction. The disconnection between individual differential pressure controllers and an overall operating philosophy, which could respond to degradation, prevented an adaptive capacity to be maintained while the controllers were reacting to the failure. That is, an optimal pathway was needed to maintain the necessary pressure differential gradient. This same overall operating philosophy would allow prioritization of the differential pressure versus temperature controls, and preventing reaction by the neural controller until the differential pressure had stabilized and the failure corrected. In this circumstance, the availability of a trained operator could have mitigated this event somewhat by disengaging ineffective automation and giving similar preference to ensuring differential pressure on the most hazardous cells. While human intervention may have helped, however, facility owners may be unlikely to hire round-the-clock operators unless the risk is great. Therefore, a control design with more adaptive capacity for this event would be a preferred option. This design can not only improve the efficiency of the system, but maintain the prioritization of response.

VIII. CONCLUSIONS

Current research philosophies consider aspects of control system reliability, including reconfigurable control theory and cyber security. However, limitations within these philosophies are exhibited in the area of goal conflicts, unexpected condition adaptation and human performance. As the complexity of the control systems has increased with distribution and interdependencies, the desire to characterize a response has evolved [14]. However, this response has not considered the multi-faceted, multidisciplinary nature of the problem. A need exists for resiliency in control systems [15].

In the case of the reconfigurable control, for instance, some key assumptions are made, specifically that the failure is known. To have information that can detect and predict failures can be complex, but often requires analysis of data sets to determine anomalies. The methods of performing this type of analysis are varied, using empirical, statistical, first principles and intelligent systems techniques. The area of online condition monitoring might, in fact, be considered a separate discipline from control engineering. Without a full-state awareness of system condition, a reconfigurable control design may not have the appropriate information necessary for the control system to react. With state awareness, a resilient control design maintains an adaptive capacity to respond to unexpected events.

Cyber security protections on existing control systems have been inherited primarily from those developed for the IT world. The next generation of cyber protections can be expected to follow the same course. However, there are fundamental differences between the design requirements of control systems as compared to IT applications. Often, control systems are expected to operate 24/7 and at 99.99% plus reliability. IT systems, while expected to be reliable, do not normally have such stringent requirements, notably because they are not depended upon to “keep the lights on” or “prevent releases of toxic substances.” As was suggested in one example, a control system has several stakeholders in the implementation, maintenance and response to cyber events. Without considering cyber security as a performance parameter in the design, there is little hope of ensuring an adequate response to cyber events. As the consequence of such a failure can be great, a resilient control system framework considers cyber security as one of the performance parameters for design.

Human performance provides an aspect of control system performance that can be beneficial or detrimental. While not necessarily statistical in nature, the discipline of human factors has demonstrated the need for understanding the interaction of human performance so that predictive and desirable responses can be achieved. In considering resilience, however, the potential benefits of the human are also considered [16]. The automation mechanism and human form a team in response to disturbances and implementing the operating philosophy. The knowledge, experience, and questioning attitude of an operator can provide the ability to adapt to abnormal circumstances. This adaptive capacity can be recognized within all the areas of performance. For example, the operator is not currently considered a stakeholder in cyber security. However, in a resilient control system paradigm, any disturbance to the operation is considered one that requires response. When a system has a cyber security compromise, this is no longer the case. The operator’s choice to delay response, in itself, can prevent the worst circumstances from being propagated.

Additional benefits of the resilient control system depicted in Figure 1 could be discussed in greater detail. For example, consider the global nature of facility energy efficiency. To achieve optimal energy performance for numerous unit operations, a controlled optimization of these in context to other performance measures is necessary. Note that the combination of performance measures is another important aspect of resilience. As in the HVAC example provided, while adding a nighttime setback for temperature control provides benefit, the action of the neural network supervisor was totally decoupled from the performance of the individual PID differential pressure controls. As a result, instead of aiding in the overall operation of the system against the priority of the moment, it acted to diminish the response.

The concept of resilient control systems cover many discipline areas, which implies that one of the benefits of resilient controls research is indeed a bridging of talents to solve a higher need. The reader is encouraged to develop a gap analysis of their own example of interest, and apply this example against the introduced benchmark(s) summarized in Table 1.

IX. ACKNOWLEDGEMENT

This manuscript has been authored by Battelle Energy Alliance, LLC under Contract No. DE-AC07-05ID14517 with the U.S. Department of Energy, performed as part of the Instrumentation, Control, and Intelligent Systems Distinctive Signature (ICIS) of Idaho National Laboratory.

REFERENCES

Table 1. Benchmark Aspect Summary

Current Research Philosophy Limitations	**Power Transmission Substation**	**Chemical Facility Reactor**	**Hazardous Facility HVAC**
Unexpected Condition Adaptation	Brittle Hierarchy	Brittle Hierarchy & Latency	Brittle Hierarchy
Human Interaction Challenges	Operator & Malicious Attacker	Operator, Security Engineer & Malicious Attacker	Operator & Design Engineer
Goal Conflicts	Physical & Cyber Security	Process Stability & Cyber Security	Process Stability & Efficiency

[1] C. G. Rieger, D. I. Gertman, and M. A. McQueen, “Resilient Control Systems: Next Generation Design Research”, 2nd Conference on Human System Interactions, Catania, Italy, pp. 632 – 636, May 2009.

[2] I. Terzic, A. Zoitl, M. Rooker, T. Strasser, P. Vrba, and V. Marik, “Usability of Multi-agent Based Control Systems in Industrial Automation,” *4th International Conference on Industrial Applications of Holonic and Multi-agent Systems*, Vol. 5696, pp. 25-36, 2009.

[3] A. Avizienis, J. Laprie, B. Randell, and C. Landwehr, “Basic Concepts and Taxonomy of Dependable and Secure Computing, *IEEE Transactions on Dependable and Secure Computing*, Vol. 1, No. 1, pp. 11-33, 2004.

[4] Y. Zhanga and J. Jiangb, “Bibliographical Review on Reconfigurable Fault-Tolerant Control Systems,” *Annual Reviews in Control*, Vol. 32, Issue 2, pp. 229-252, December 2008.

[5] R. L. Boring, “Reconciling Resilience with Reliability: The Complementary Nature of Resilience Engineering and Human Reliability Analysis,” *Proceedings of the Human Factors and Ergonomics Society 53rd Annual Meeting*,” pp. 1589-1593, October 2009.

[6] P. Dongbo, L. Feng, Z. Xuelian, and L. Tao **“**Functional safety in Building Automation and Control Systems,” *3rd IEEE Conference on Industrial Electronics and Applications*, pp. 467-470, June, 2008.

[7] M. Zhivich and R. K. Cunningham, "The Real Cost of Software Errors," IEEE Security and Privacy, vol. 7, no. 2, pp. 87-90, March/April, 2009.

[8] K.R. Rohloff, "Sensor Failure Tolerant Supervisory Control," 44th IEEE Conference on Decision and Control, pp. 3493-3498, December, 2005.

[9] F.-Y. Wang and D. Liu, *Networked Control Systems: Theory and Applications*, Springer-Verlag, London, UK, 2008.

[10] G. Traufetter, "The Computer vs. the Captain: Will Increasing Automation Make Jets Less Safe?", *Der Spiegel*, July, 2009.

[11] C. Rehtanz, *Autonomous Systems and Intelligent Agents in Power System Control and Operation*, Springer-Verlag, Berlin, Germany, 2003.

[12] D. P. Buse and Q. H. Wu, IP Network-based Multi-agent Systems for Industrial Control, Springer-Verlag, London, UK, 2007.

[13] H. B. Mitchell, *Multi-Sensor Data Fusion*, Springer-Verlag, Berlin, 2007.

[14] S.P. Meyn, *Control Techniques for Complex Networks*, Cambridge University Press, New York, NY, 2008.

[15] S. M. Mitchell and M. S., Mannan, "Designing Resilient Engineered Systems," Chemical Engineering Progress, Vol. 102, No. 4, pp. 39-45, April 2006.

[16] E. Hollnagel, D. D. Woods, and N. Leveson, *Resilience Engineering: Concepts and Precepts*, Ashgate Publishing, Aldershot Hampshire, UK, 2006.