PERSPECTIVES ON DESIGN CONSIDERATIONS INSPIRED BY SECURITY AND QUANTUM TECHNOLOGY IN CYBERPHYSICAL SYSTEMS FOR PROCESS ENGINEERING

Helen Durand [a1], Jihan Abou Halloun [a,], Kip Nieman [a], and Keshav Kasturi Rangan [a]

[a] Department of Chemical Engineering and Materials Science, Wayne State University, Detroit, MI 48202

Abstract

Advances in computer science have been a driving force for change in process systems engineering for decades. Faster computers, expanded computing resources, simulation software, and improved optimization algorithms have all changed chemical engineers' abilities to predict, control, and optimize process systems. Two newer areas relevant to computer science that are impacting process systems engineering are cybersecurity and quantum computing. This work reviews some of our group's recent work in control-theoretic approaches to control system cybersecurity and touches upon the use of quantum computers, with perspectives on the relationships between process design and control when cybersecurity and quantum technologies are of interest.

Keywords

model predictive control, cybersecurity, quantum computing, quantum control.

Introduction

Digitalization is one of the major themes of process systems engineering today, and has been made possible by advances in computing and algorithms. Despite the great advances posed by new capabilities in optimization, autonomy, and data analytics, challenges remain with use of data and simulation for process engineering. One of these challenges is cybersecurity risks that can be introduced through greater reliance on computation and communication in system operation. One of the avenues for addressing cybersecurity of control systems that has been pursued in process systems engineering is to develop strategies for detecting attacks through process data that appears anomalous (e.g., Wu et al. (2018)) or through the development of strategies which adjust system operation to locate attacks (e.g., Narasimhan et al. (2022); Oyama and Durand (2020)). It has also been suggested that cybersecurity should play a role in hazard analysis Cormier and Ng (2020). An important question for the topic of cybersecurity of control systems is clarifying the role that process systems engineers could play in developing cybersecurity strategies, and attempting to outline what is likely to be impactful in this field traditionally dominated by computer scientists. On the quantum computing front, quantum computation has been a technology of interest in recent years in process systems engineering for topics such as optimization and machine learning (e.g., Ajagekar et al. (2022); Ajagekar and You (2020); Harwood et al. (2021)). In our recent work Nieman et al. (2022), we have performed an initial exploration of control actions computed by quantum computers. A challenge for process systems engineering as a discipline is to understand and locate all use cases for these new devices. We argue that the study of quantum computation holds benefits for process systems engineering beyond the algorithms themselves in facilitating an understanding within the discipline of principles of control of quantum systems and the consideration of how this control might be used to achieve larger-scale benefits. This could provide insights into new future process design directions.

Cybersecurity and Process Operation

The Stuxnet worm is one of the most famous instances of a cyberattack on an industrial control system. This incident, which took place in Iran, involved an attack that broke centrifuges at a uranium enrichment plant while falsifying data to mask its presence Karnouskos (2011). Considering that this was a successful attack on an industrial control system, an important question is how cybersecurity of control systems should be relevant to chemical engineers (rather than information technology and computer science professionals) beyond understanding cybersecurity risks so that engineers can be aware of times that they may need to seek guidance from colleagues more experienced in cybersecurity before installing or upgrading components of automated systems. This is important both to guide research directions toward those that will be impactful to the field (given the critical-

[1] Corresponding author: Helen Durand (E-mail: helen.durand@wayne.edu).

ity of preventing security breaches) and for determining how to update chemical engineering education to make students aware of cybersecurity practice in a manner appropriate for their workplace roles.

At first, the need for chemical engineers to be involved in this discipline may appear somewhat limited. One of the reasons is implied by the suggestions in Cormier and Ng (2020) regarding including cybersecurity considerations in risk analysis. As is highlighted in Cormier and Ng (2020), chemical engineers already design processes in a way that considers possible failure scenarios. This is part of standard safety protocols in areas such as HAZOP analysis and layers of protection analysis. Cormier and Ng (2020) argues that the risk analysis framework should be altered when considering the possibility of attacks, so that failure events that might seem highly unlikely in the absence of an attack become more likely when it is considered that an attacker could exploit vulnerabilities. In addition, it is considered that part of the risk management should include the information technology level (rather than just planning to put traditional safety systems in place to avoid serious accidents when the attackers get through). Furthermore, mathematically, cyberattacks have characteristics similar to faults in that they result in unintended actuator actions impacting a process. The key difference between faults and attacks, as highlighted in Rangan et al. (2022), is that attacks are deliberate actions to analyze process vulnerabilities and exploit them while covering it up. The aspect of covering up the faulty behavior is not typically a feature of a fault created in the absence of an attack. Though faults do have this difference from an attack, it appears to be solid to consider attacks at the risk analysis level with potentially different "blocking" measures put in place to prevent the attacks than might be made to only prevent accidents, and then to consider some minimal level of risk.

One might consider, however, that as attackers increase their capabilities, it is inconvenient to need to consistently attempt to keep up-to-date in security practices to prevent attacks from breaking through. It would be preferable to set up systems where it was not possible for an attacker to break in or, perhaps more ideally, where their presence in the system is considered of no consequence. Posing the problem in this fashion makes it clear that there is a true research challenge in this direction for chemical engineers, which is the question of what it would take to achieve process designs, combined with operating strategies, that could permit this level of flexibility. Traditional process design procedures do not seem to be enough for achieving this concept, which is a mixture of inherent resilience in design and inherent safety in design.

Though it would be exciting to have designs for processes in the traditional process design methodology that can handle cyberattacks on the control systems in a manner that is agnostic to attacks (e.g., the profits and safety are approximately the same when someone is attacking the system compared to when they are not), it would be hard to imagine a traditional process system today having these capabilities. These requirements seemingly conflict from a traditional process design and control viewpoint. For the design to be agnostic to attacks, it seems it would need to be somewhat unresponsive to actuator actions. If it is unresponsive to actuator actions, it seems that it is unlikely to be very profitable and then would essentially be able to be operated in open-loop. This implies that addressing the question of whether there are any possible combined process design and control strategies that might achieve the cybersecurity goal will require fundamental advances in process design and control.

At first, it may appear to be an impossible concept. However, it should also be recognized that there are many dynamic considerations for processes that are often neglected. For example, process designs would not typically involve setting up interesting transport fields that give the same area-weighted average value of a desired property (e.g., temperature, concentration, or velocity) but have different radial variations (or event different time-varying radial variations). This is one direction that might be explored, as it could attempt to use properties of process systems that are usually not exploited in seeking to deal with attacks (e.g., potentially to detect them by failing to detect required patterns in these auxiliary dynamics) without impacting overall product requirements. Perhaps the physics of computing devices, sensors, and actuators could also be explored in greater detail. Despite the challenge of attempting to understand how to break into this area, we can summarize learnings from our work focusing on cybersecurity of control systems to provide some insights into how process dynamics and control system design interact, with the goal that this may provide insights for future works at trying to understand how processes might be designed in a fashion that is resilient to attacks (or to understand why this is not possible if it turns out not to be).

Cybersecurity of Control Systems: An Overview of Results on Control-Theoretic Cyberattack Detection Using Lyapunov-Based Economic Model Predictive Control

In Durand (2018), our group introduced a notion of cyberattack-resilient control design in which a process design is considered resilient against attacks if there is no possible input policy starting from t_0 that, regardless of the disturbance trajectory, is able to drive the closed-loop state out of a set of safe states if it is initialized within a set of allowable initial states. Despite that reasonable system designs meeting this definition are challenging to conceive, we can expect that we would be able to gain insights into such designs by considering other notions related to detection of cyberattacks and control in the presence of cyberattacks, and then checking whether it is possible under these detection and control policies to cause the closed-loop state to enter unsafe operating regions without the attack being detected. In Oyama and Durand (2020), we introduced three detection policies for sensor attacks when the controller of a process is an optimization-based control law known as Lyapunov-based economic model predictive control (LEMPC) Heidarinejad et al. (2012), and analyzed the extent to which they guarantee safety. These were subsequently modified to suit cases with actuator attacks Rangan et al. (2022) and sensor and actuator attacks that could occur at the same time Oyama et al. (2022). The major characteristics of these strategies are as follows:

- Detection Strategy 1 is based on driving the closed-loop state over time through different regions of state-space by using properties of the Lyapunov-based stability constraints in LEMPC to ensure that the closed-loop state can be driven along a desired path and to ensure that it is known that if it follows that path, the Lyapunov function should be decreasing between two sampling periods. This enables attacks to be flagged if it is not.

- Detection Strategy 2 is based on using state predictions to cross-check the state measurements to verify that they are consistent with expectations coming from the case of no attack (but with potentially noise and disturbances). The LEMPC is designed to be sufficiently conservative to ensure that at least a sampling period of time is available after an attack is not detected at a sampling time before the closed-loop state leaves a safe operating region.

- Detection Strategy 3 is based on using redundant state estimates to cross-check one another. The LEMPC is designed to be sufficiently conservative to ensure that at least a sampling period of time is available after an attack is not detected at a sampling time before the closed-loop state leaves a safe operating region.

Table 1 highlights major characteristics of these strategies.

Table 1: Three detection strategies summary.

Property	**Strategy 1**	**Strategy 2**	**Strategy 3**
Type	Active	Passive	Passive
Safety Guarantees with Undetected Sensor Attacks	None	Sampling period of safety	Safety guaranteed until detection
Safety Guarantees with Undetected Actuator Attacks	Safety guaranteed until detection	Safety guaranteed until detection	None
Limitations with Sensor Attacks	Poor safety guarantees	State predictions can be corrupted by sensor attacks without detection	Requires some state measurements to not be attacked
Limitations with Actuator Attacks	Creates non-standard operating conditions	Requires a redundant control law calculation	Poor safety guarantees

In Oyama et al. (2022), Detection Strategies 1 and 3, and 2 and 3, were combined to form strategies that guaranteed safety until attack detection even if both actuators and sensors could be attacked at the same time (as long as not all sensors were attacked so that at least one of the redundant state estimators was still intact). The fact that these combinations provide safety guarantees when attacks occur on both the actuators and sensors at the same time is not surprising when considering the strengths of each technique with respect to sensor-only and actuator-only attacks in Table 1. What is somewhat surprising is that combining two passive detection strategies in the case of the combination of strategies 2 and 3 is capable of enabling the detection of sensor and actuator attacks simultaneously with strong safety guarantees. The reasons for the different strategies working or not working was highlighted in Oyama et al. (2022) through the definition of cyberattack discoverability. Specifically, a "cyberattack discoverable" system is defined, colloquially, to be a system in which the state trajectories under different measurement noise and disturbance profiles that are within the expected bounds and distributions of these quantities are able to be distinguished from those under the actual noise and disturbance profiles. This definition can be used also to inspire detection and control frameworks for noting cyberattacks in a system and to benchmark whether a method would be expected to work before an in-depth exploration is performed. Specifically, it shows that the design principles for detection and control policies needs to be that the control design creates an expectation for what the closed-loop profile should look like in a way that would be difficult for an attacker to replicate without getting noticed as an attacker. Any case in which they cannot be noticed but can provide the attack would be a situation in which the cyberattack is not discoverable. The goal of the integration of the control and detection policies is to attempt to create situations where the lack of discoverability does not mean that safety is lost by seeking to make the undetected cases essentially equivalent to a form of bounded measurement noise that can be protected against.

We can use this notion to understand how a combination of two passive attack detection policies could also create a situation where all attacks could be discovered before creating safety issues. Specifically, the two passive policies that are combined (Strategies 2 and 3) guarantee that attacks will be distinguishable from non-attacked cases for the actuator and the sensor cases individually if the attack is detected (Strategy 2 makes this guarantee for the actuator case, and Strategy 3 makes this guarantee for the sensor case). The combination of the two policies therefore enables the simultaneous attacks to be discovered. Notably, this means that passivity versus activity of an attack detection method is not the determining factor in whether it has the ability to make control policies under normal compared to attacked operation noticeable. However, it is reasonable to note that a passive strategy does not enable on-line testing of hypotheses that a system is or is not under attack.

From the process design perspective that we have been utilizing in this paper to consider the future of cybersecurity research for chemical engineers, some of the interesting strategies discussed above for cyberattack-handling are those, such as Detection Strategy 3 in the presence of attacks on the sen-

sors or Detection Strategy 1 in the presence of attacks on the actuators, which are capable of continuing to operate a process in the presence of undetected attacks. These operate a system in such a way that they either force the attacker to continue to stabilize the process or to be noticed as an attacker. These therefore provide an interesting framework for the consideration of interactions between process control and process design under attacks that might be exploited in future work at the design/control interface for seeking to develop strategies for cyberattack-handling that are simultaneously flexible and resilient. In particular, we stated a vision above in which it would be most desirable to have processes where they were agnostic to attacks on the system. These two strategies which are able to stabilize a process even in the presence of attacks provide an indication that there could exist conditions under which the process design and detection policies can work together to make the system "insensitive," in a sense, to certain types of attacks. The goal of the future of the design/control interface from a cybersecurity perspective is to locate more of these conditions, to better understand their interactions with process profits, flexibility, and safety, and to eventually develop systematic protocols for achieving such types of resilient systems with the goal of maximizing process flexibility for incorporating advances in computer science and electrical engineering.

From these discussions, we can see that a goal of designing future control-theoretic integrated cyberattack detection and control policies should be to locate new strategies that minimize the set of attacks that cannot be discovered and that also continue to enable undetected cases to be mathematically equivalent to measurement noise. It would seem that one of the most effective ways for handling discoverability would be to design policies by which the attacker's trajectory is distinguishable from the typical operating strategy because the attacker cannot predict what the typical operating strategy will be and therefore will "mess up" and show themself. However, before further commenting on this, it is necessary to clarify why a control policy in Durand (2018) that incorporated randomization was not successful at preventing the success of an attacker targeting the sensors. This specific policy was designed as follows: n_p LEMPC's, plus an auxiliary Lyapunov-based controller, were assumed to be available. All supported operation within subsets of a larger "safe" set. To attempt to prevent the attacker from knowing which control law would be applied at a given time (under the hope that the attacker will not be able to figure out what control law will be used next to exploit it), one of the set of $n_p + 1$ controllers is selected at every sampling time. However, to ensure closed-loop stability under normal operation (i.e., in the absence of any cyberattack on the sensors), the implementation strategy for this randomized control law policy required that at every sampling time, one of the $n_p + 1$ controllers could be suggested by the random number generator, but that this control law could only actually be used if the closed-loop state measurement at the sampling time was within a subset of the "safe" set supported by the selected control law. If not, a different control law would need to be selected. This enabled rigorous guarantees of closed-loop stability of a nonlinear process operated under such a control policy in the absence of an attack.

This strategy was not being used in Durand (2018) for attack detection; rather, it was being explored for its ability to slow an attacker down at achieving their goal. Therefore, it is not fully tied to the notion of discoverability. However, the reasons that it does not succeed at significantly slowing an attacker can provide insights on discoverability of attacks. The specific reasons that this policy did not work are as follows: 1) though there is an element of "randomness" in the selection of the control laws at a sampling time, it is not true randomness. For example, if the state measurement is in a given region of state-space, some of the controllers may not be able to be selected due to the state measurement being outside of the region in state-space which they support. This reduces the number of controllers that are actually potential controllers for a given state measurement (in some regions lowering it to perhaps only 2 controllers, which decreases the random element of this strategy significantly); 2) there was not an attempt made to ensure that the controllers available at a given point in state-space would compute significantly different control actions from one another. Therefore, an attacker could end up with cases where even if there were multiple control laws that might be selected, multiple might compute similar or identical inputs (e.g., all at the upper bound), so that the attacker might again know what the expected state measurement would be at the next sampling time; 3) the number of controllers available was dictated by the state measurement, but the state measurement had been falsified. Therefore, the strategy, in some sense, "plays into the attacker's hands" by putting the randomness of the strategy at the mercy of the attacker. One could imagine ways of re-designing the controller to attempt to overcome the challenges in each direction, and in general, the ability to look at potential ideas and analyze why they do or do not work and what may be allowable ways to change them so that they come closer to causing more attacks to be discoverable is a major benefit of the control-theoretic studies in cybersecurity of process control systems.

In contrast to this failed randomization strategy, we can recognize that there are other cases where random number generation is quite important to cybersecurity, in particular in the case of cryptography. In information security, an important concept is the notion of how hard it would be for someone to figure out what the random number was. We can use this thinking to inspire another strategy for cyberattack detection which we call "Directed Randomization" in Oyama, Messina, Kasturi Rangan, Nieman, Tyrrell, Leonard, Hinzman, Williamson, and Durand (Oyama et al.). This strategy involves pre-computing two potential control actions at every point in state-space. These control actions should be designed to maintain the closed-loop state in a safe operating region over the following sampling period, but also to create a set of possible states at the end of the sampling period that cannot overlap with one another (even in the presence of noise and disturbances). Then, at every sampling time, one of the two control actions is randomly selected based on the state measurement. The attacker may know which of

the two control actions *could* be selected, but does not know which actually *was* until the next measurement is taken at the beginning of the next sampling period. Because the control actions were designed to ensure that there is no overlap between the potential states at the beginning of the next sampling period under either of the control actions, the next state measurement puts the attacker in a bit of a pickle; the attacker will be caught if they provide a state measurement outside the expected region, but even if they knew which of the two regions is expected, they still could not know with certainty whether to provide a state measurement in one compared to the other. They have a 50/50 chance of being right. Intuition indicates that this could become a somewhat challenging game for them to keep playing to evade detection, and it fits within the concept of designing a controller that makes certain attack policies discoverable by designing the detection policy in such a way that it is "hard" (in a loose sense) for the attacker to figure out how to always provide state measurements consistent with expectations. This does have some disadvantages of needing to determine the bias policies throughout state-space *a priori* and then to somehow store these, which could get to be a large space to search when a value is needed and use a significant amount of memory. However, it does aid in providing some clarity to the concept of how one might try to design control laws that put the closed-loop state into conditions where an attacker could find it difficult to replicate those states reliably.

Cybersecurity, Randomness in Quantum Computation, and Future Manufacturing

Above, it was noted that randomness could play a role in the future of cybersecurity for process control systems. Quantum systems can create randomness through nondeterminism in the physics of their measurement (and through other issues such as unknown interactions with the environment) that could make quantum computation an interesting element to consider when considering cybersecurity of control systems. However, the benefits of considering quantum computation when exploring the future of process systems engineering go beyond whatever role it might play in realizing randomness in a cybersecurity context or even in computation itself. Quantum computers are examples of systems in which the quantum properties of matter (superposition, entanglement, and interference) are already being exploited in devices. One of the exciting aspects of quantum computation as a research area today is that the full potential of these computing devices are not fully known.

Quantum computers should not be considered to be "faster" computers; rather, they are "different" computers. Algorithms for them therefore must be formed differently. Two common types are quantum annealers (which can solve a type of optimization problem) and gate-based quantum computers (which can perform more universal computations). Algorithm design for gate-based quantum computers occurs today at the circuit level (meaning deciding on specific manipulations of quantum states that must be performed in a row to attempt to get answers from these devices that are both useful and efficient). Everything a quantum computer does is not "better" than what a classical computer does. For example, one thing that can be done on a quantum computer is to manipulate the state of a quantum system (considered to initially be denoted by $|0\rangle$, or $[1\ 0]^T$ in vector form) to a superposition of $|0\rangle$ and $|1\rangle = [1\ 0]^T$ ($\frac{|0\rangle+|1\rangle}{2}$). The gate that would achieve this when representing a quantum circuit is known as the Hadamard gate (represented in matrix form as $H = \frac{1}{\sqrt{2}}[1\ 1;\ 1\ -1]$). The fact that a superposition can be created is not necessarily useful. For example, measuring the quantum state after it is in a superposition (with respect to the computational basis of $|0\rangle$ and $|1\rangle$) will transform the state back to either $|0\rangle$ or $|1\rangle$ with a 50/50 probability. This is no better than randomly guessing whether the state would end up in $|0\rangle$ or $|1\rangle$ Yanofsky and Mannucci (2008).

The benefits of quantum computation come from clever consideration of the mathematics of the original problem and of the impacts of a series of gates on multiple (often interacting) quantum states. Algorithms may need to grapple with the probabilistic nature of quantum mechanics (e.g., some algorithms, such as addition based on the Quantum Fourier Transform (QFT) Ruiz-Perez and Garcia-Escartin (2017), are deterministic in theory, whereas others give an answer with a certain probability), as well as with the lack of fault-tolerance of these devices today (a feature typically referred to as "noise"). In Nieman et al. (2022), our group provided investigations into the intersection of nondeterminism in computing results or control actions with control (including a simulation of a simple control law on a quantum simulator using a quantum addition code drawn largely from Anagolum (2018)), suggesting that there may be potential for both unique algorithms on quantum computers and devices without fault-tolerance to be stabilizing for processes.

Though the exploratory nature of investigating quantum computing for process systems engineering makes it interesting, we argue that there is more that might be learned from studies of quantum computation for future manufacturing in chemical engineering beyond algorithm development. Specifically, for the exploitation of quantum phenomena in manufacturing (e.g., the translation of quantum control into workable manufacturing devices through, for example, direct manipulation of quantum states), it might be asked whether, as in the case of quantum computation, low-level quantum phenomena might be harnessed and scaled up for any next-generation process design purposes. The concept that quantum state manipulation might be scaled up in some sense for useful purposes has not been restricted to computation; for example, quantum control has been explored for various purposes, including manipulating the length of a hydrogen fluoride (HF) bond Magann et al. (2021). These manipulations can be modeled using the time-dependent Schrodinger equation ($\hat{H}\psi(x,t) = i\hbar\frac{\partial\psi(x,t)}{\partial t}$, where $\hbar$ is the reduced Planck constant, $\hat{H}$ represents the Hamiltonian operator, and the wavefunction is $\psi(x,t)$, explicitly dependent upon space and time coordinates) with consideration of the impacts of an electromagnetic field on the potential term in the Hamiltonian. Quantum control involves many physics principles that go beyond a typical process systems engineering education, suggesting that there may be many unexplored concepts in this

domain for chemical engineers. For example, studies involving quantum control of HF have considered representing the potential in the Hamiltonian as the sum of an unperturbed and perturbed term, where the perturbed term accounts for the influence of the electromagnetic field; an understanding of the modeling efforts in this case requires studies of rigorous quantum mechanics.

Conclusion

This work reviewed recent work from our group on cybersecurity of control systems and highlighted several features of quantum computation and our view of its relevance as a new direction not only for process systems engineering computations, but also for drawing inspiration for exploring new design possibilities. The vision that we cast throughout for process design research might re-evaluate prior advances in simultaneous design and control of process systems (e.g., Sandoval et al. (2008)) for these unestablished areas.

Acknowledgement

Financial support from the Air Force Office of Scientific Research (award number FA9550-19-1-0059), National Science Foundation CNS-1932026 and CBET-1839675, and Wayne State University is gratefully acknowledged.

References

Ajagekar, A., K. Al Hamoud, and F. You (2022). Hybrid classical-quantum optimization techniques for solving mixed-integer programming problems in production scheduling. *IEEE Transactions on Quantum Engineering 3*, 1–16.

Ajagekar, A. and F. You (2020). Quantum computing assisted deep learning for fault detection and diagnosis in industrial process systems. *Computers & Chemical Engineering 143*, 107119.

Anagolum, S. (2018). Donew, https://github.com/sashwatanagolum/donew.

Cormier, A. and C. Ng (2020). Integrating cybersecurity in hazard and risk analyses. *Journal of Loss Prevention in the Process Industries 64*, 104044.

Durand, H. (2018). A nonlinear systems framework for cyberattack prevention for chemical process control systems. *Mathematics 6*(9), 169.

Harwood, S., C. Gambella, D. Trenev, A. Simonetto, D. Bernal, and D. Greenberg (2021). Formulating and solving routing problems on quantum computers. *IEEE Transactions on Quantum Engineering 2*, 1–17.

Heidarinejad, M., J. Liu, and P. D. Christofides (2012). Economic model predictive control of nonlinear process systems using Lyapunov techniques. *AIChE Journal 58*(3), 855–870.

Karnouskos, S. (2011). Stuxnet worm impact on industrial cyber-physical system security. In *Conference of the IEEE Industrial Electronics Society*, pp. 4490–4494.

Magann, A. B., M. D. Grace, H. A. Rabitz, and M. Sarovar (2021). Digital quantum simulation of molecular dynamics and control. *Physical Review Research 3*, 023165.

Narasimhan, S., N. H. El-Farra, and M. J. Ellis (2022). Active multiplicative cyberattack detection utilizing controller switching for process systems. *Journal of Process Control 116*, 64–79.

Nieman, K., K. Kasturi Rangan, and H. Durand (2022). Control implemented on quantum computers: Effects of noise, nondeterminism, and entanglement. *Industrial & Engineering Chemistry Research 61*(28), 10133–10155.

Oyama, H. and H. Durand (2020). Integrated cyberattack detection and resilient control strategies using lyapunov-based economic model predictive control. *AIChE Journal 66*(12), e17084.

Oyama, H., D. Messina, K. Kasturi Rangan, K. Nieman, K. Tyrrell, A. F. Leonard, K. Hinzman, M. Williamson, and H. Durand. Cyberattack detection with image-based control, directed randomization, and distributed lyapunov-based economic model predictive control. *Digital Chemical Engineering submitted*.

Oyama, H., D. Messina, K. K. Rangan, and H. Durand (2022). Lyapunov-based economic model predictive control for detecting and handling actuator and simultaneous sensor/actuator cyberattacks on process control systems. *Frontiers in Chemical Engineering 4*, 810129.

Rangan, K. K., J. Abou Halloun, H. Oyama, S. Cherney, I. A. Assoumani, N. Jairazbhoy, H. Durand, and S. K. Ng (2022). Quantum computing and resilient design perspectives for cybersecurity of feedback systems. *IFAC-PapersOnLine 55*(7), 703–708.

Rangan, K. K., H. Oyama, and H. Durand (2022). Actuator cyberattack handling using Lyapunov-based economic model predictive control. *IFAC-PapersOnLine 55*(7), 489–494.

Ruiz-Perez, L. and J. C. Garcia-Escartin (2017). Quantum arithmetic with the quantum fourier transform. *Quantum Information Processing 16*(6), 1–14.

Sandoval, L. R., H. Budman, and P. Douglas (2008). Simultaneous design and control of processes under uncertainty: A robust modelling approach. *Journal of Process Control 18*, 735–752.

Wu, Z., F. Albalawi, J. Zhang, Z. Zhang, H. Durand, and P. D. Christofides (2018). Detecting and handling cyberattacks in model predictive control of chemical processes. *Mathematics 6*(10), 173.

Yanofsky, N. S. and M. A. Mannucci (2008). *Quantum Computing for Computer Scientists*. Cambridge University Press.