Military Cyber Affairs

Volume 7
Issue 1 *HammerCon 2024 Issue*

Article 7

May 2024

Using Digital Twins to Protect Biomanufacturing from Cyberattacks

Brenden Fraser-Hevlin
Washington State University, b.fraser-hevlin@wsu.edu

Alec W. Schuler
Washington State University, alec.schuler@wsu.edu

B. Arda Gozen
Washington State University, arda.gozen@wsu.edu

Bernard J. Van Wie
Washington State University, bvanwie@wsu.edu

Follow this and additional works at: https://digitalcommons.usf.edu/mca

 Part of the Cognitive Psychology Commons, Cognitive Science Commons, Computer and Systems Architecture Commons, Computer Law Commons, Digital Communications and Networking Commons, Intellectual Property Law Commons, International Relations Commons, Military, War, and Peace Commons, National Security Law Commons, Other Computer Engineering Commons, and the Systems Science Commons

Recommended Citation

Fraser-Hevlin, Brenden; Schuler, Alec W.; Gozen, B. Arda; and Van Wie, Bernard J. (2024) "Using Digital Twins to Protect Biomanufacturing from Cyberattacks," *Military Cyber Affairs*: Vol. 7 : Iss. 1 , Article 7.
Available at: https://digitalcommons.usf.edu/mca/vol7/iss1/7

This Article is brought to you for free and open access by the Open Access Journals at Digital Commons @ University of South Florida. It has been accepted for inclusion in Military Cyber Affairs by an authorized editor of Digital Commons @ University of South Florida. For more information, please contact digitalcommons@usf.edu.

Using Digital Twins to Protect Biomanufacturing from Cyberattacks

Cover Page Footnote

The authors would like to acknowledge funding from the Griffiss Institute VICEROY program through award No. SA10012021MM0336 for the Northwest Virtual Institute for CyberSecurity Education & Research (CySER). Undergraduate students assisting with the research include Mikayden Weise and Adam Caudle, both supported through stipends from the VICEROY program.

This article is available in Military Cyber Affairs: https://digitalcommons.usf.edu/mca/vol7/iss1/7

Using Digital Twins to Protect Biomanufacturing from Cyberattacks

Brenden Fraser-Hevlin, Alec W. Schuler, B. Arda Gozen,
and Bernard J. Van Wie

1. Introduction

Cybersecurity is critical to the national defense of the United States. This security interest also includes the biological domain in which medical information and biological systems are at severe risk. Attacks in both the cyber and biological domains can severely compromise national security and have shared characteristics: they are complex, fast paced and covert in nature, and difficult to predict. They have relatively low costs compared to traditional weapons, and they can significantly impact the effect of conventional attacks if used in combination.[1] A major intersection of these domains lies in the field of biomanufacturing: specifically, the production of biological medicines and devices for making them. Recent developments have resulted in increased focus on cybersecurity in this field. In 2019, the U.S. Bipartisan Commission on Biodefense held a panel to assess cybersecurity threats in the biotechnology industry, specifically examining the security vulnerabilities that would arise from the combination of cyber and biological domains.[2] Dixon (2021) raised the issue of cybersecurity in biological systems as a potential new grey area in warfare. He defines "cyberbiosecurity" as a form of defense that covers the biological and medical information that comes at risk when living and nonliving systems are combined.[2] New devices like biosensors are one example of potential risks. They can be used to monitor biological parameters in real time; while they are intended to observe the growth and health of living materials like cells, they could also be used in monitoring of biological weapons, as a surveillance tool. Another example arises with companies using bioinformatics algorithms to synthesize Deoxyribonucleic Acid (DNA) and proteins. The field of synthetic biology has seen the application of semiconductors, computer aided design and Artificial Intelligence (AI) being used to design and build biological material.[2] The possibilities for scientific innovation in this field are considerable but introduce dangerous security risks. Biological research data could be exposed, or sequences of proteins could be altered to make faulty or even toxic molecules rather than the intended biopharmaceuticals.[2] Thus, to fully assess these risks, the commercial biotechnology industry must be examined further.

Biomanufacturing is on the verge of a fourth industrial revolution. The creation of artificial intelligence, machine learning, and cloud data storage has led to the rise of this Industry 4.0, which will allow for smarter, leaner, and faster

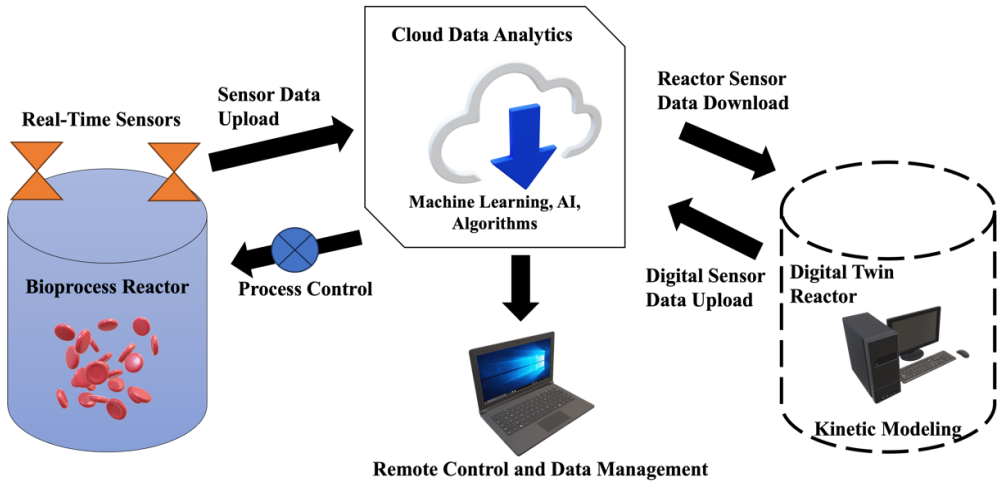

Figure 1: Digital Twin Reactor Technology. The system includes a simulated primary reactor, with connection to a digital twin and real-time collection and transfer of cloud data.

manufacturing in comparison to the past.[3] Thanks to the continued development of the Internet of Things (IoT), it is now possible to fully monitor production in unprecedented ways, with a dynamic connection between physical systems and digital cloud analytics.[3] Biotechnology has recently been expanding beyond the production of primary and secondary metabolites to new products such as cells and biological tissues.[4] The recent rise of these new technologies has been a response to the increasing complexity of producing medicines. With these products, real-time sensing and control is a necessity to ensure that the medicines retain their quality and can be manufactured efficiently and economically. Integration of cloud data analytics allows for process updates as soon as any change occurs, and integration of dynamic models to predict product outcomes. Even the Food & Drug Administration (FDA) has begun development on studies of models that can be used in regulation of real-time processes using cloud data, helping address the lack of standardization of this new technology.[5] One newer concept in this field is digital twins: virtual systems that match the dynamic behavior of physical machines.[5] As shown in Figure 1, a digital twin generally involves a physical product transmitting data to a cloud or internet-connected device, which then sends data to a virtual machine. Digital twins often incorporate process models, including first principles models (based on process physical properties and kinetic models), as well as compartmental models and hybrid models that incorporate AI and machine learning.[3] With the high sensitivity of biological products, having a digital twin can be extremely valuable. Recent surveys found that 75% of companies that were using Internet of Things (IoT)-connected devices were also using digital twins or plan to start using them.[6] Digital twins generally use either mechanistic, (mathematical model-based) or data-driven (machine learning) approaches to simulate the production of their counterpart systems. Park et al.[7] identified kinetic modeling as an ideal mechanistic method to represent the growth of cells and integrate a process with a digital twin. They show many bioprocessing operations are based on fundamental kinetic models, and addition of digital twins allows for

real-time integration of a mechanistic model with the actual process, updating the bioprocess as needed based on predictions from the kinetic model. They argue digital twins also provide value given their ability to combine inputs from numerous sources and compare these inputs to key process indicators in real time. Digital twins, along with machine learning and data analytics, may also have utility in predicting and detecting cyberattacks of complex IoT-enabled manufacturing processes.

Digital twins have many useful advantages, including lower machine downtime, enhancement of planning and scheduling in production, and performance of virtual commissioning (i.e. providing a "commission twin" that simulates a process before investing real resources and expenses).[6] The ability of digital twins to collect and analyze data from every stage of a product life cycle, and connect to the IoT in real time, is significant.[3] Nonetheless, these innovations introduce new vulnerabilities that could be exploited by cyber criminals. The biomanufacturing industry is highly susceptible, with current cyber defenses considered by experts to be at an inadequate level. Any disruption to the global production chain can be massive, with the recent COVID-19 pandemic as a significant example.[8] The healthcare sector has already been targeted frequently by hackers in the past. From 2010-2019, the number of data breaches in healthcare in the U.S. went up by 10% each year.[9] In 2017, Merck & Co was targeted by ransomware hackers, causing the shutdown of their computer systems used to control manufacturing causing a major delay in vaccine distribution, and total expenses of $1 billion.[9] With the frequent advances and increased sophistication in cyber attacker technology, these events are only expected to become worse.

There is a need for standardized methods to track the collection of real-time data in networked operations and design with cyber protection in mind. Mylrea et al.[8] proposed a framework for use of digital twins in this industry, which would be expected to improve preparedness for these attacks over time. The current gaps in digital twins are tied directly to cybersecurity, with technology, policy, and the workforce all identified as lacking effectiveness to detect and prevent cyberattacks. In general, there is a need for better cyber-physical monitoring, enhanced security control, and improved audit trails.[8] The future of process control development will be key in providing both improvements to digital twin performance, and overall improvement in cyber protection for Digital Twin (DT) systems. Already, the chemical industries have integrated model predictive control and self-optimizing control, which use algorithms based on historical process data.[10] Now, these advances must be applied in the new, but rapidly growing biomanufacturing industry. Previous work has demonstrated successful implementation of process control for cybersecurity in physical systems. Balta et al. created a digital twin architecture for an existing commercial 3D printer system, collecting data and introducing attacks that interrupted the system's performance, with anomaly detection models designed to communicate with a novel network controller and to provide counteract attacks.[11] Yet, examples of physical system implementation specifically for the biomanufacturing sector are limited. Currently, this field is

exposed to cyber-intrusion and would benefit from the development of process-control based systems to prevent attacks.

In the present study, we propose a novel model system in which process control principles are used in tandem with a digital twin to detect cyberattacks and faults on a simulated small-scale biotechnology operation. The system provides a baseline to study the behavior of process control-based cybersecurity integrated with digital twins in a highly sensitive biologic process. This allows for efficient control of the process, but also prediction of the impact of sending cloud data to a digital twin wherein there are associated vulnerabilities. We lay the groundwork for further development of this critical technology, demonstrating the importance of being able to deliver real-time data using predictive models to increase efficiency and ultimately lower the cost of the process. The system is based on a novel Centrifugal Bioreactor (CBR) with for cancer immunotherapy.[12] Expanding on the past development of the CBR, we now introduce a real-time process control framework to automatically regulate cell growth associated metabolites and simultaneously prevent cyberattacks. We have introduced a Digital Twin Reactor (DTR) counterpart which can collect data via Wi-Fi from the CBR in real time, and a simulated hack method that sends false data to the digital twin. In our system, we grow a culture of bovine killer T cells in a Simulated Physical Reactor (SPR) standing in for the CBR, with hacked signals sent to the DTR over a two-week period. In this study we seek to demonstrate successful integration of a real-time control system, to model, predict, and mitigate impacts of potential cyberattacks on biomanufacturing. The goal is to provide efficient process control while compensating for the associated vulnerabilities of delivering valuable cloud data. Our control system uses a DTR to increase efficiency and connect the capabilities of cloud analytics and machine learning, while also anticipating the possibility of cyberattacks that this introduces.

2. Methods

2.1 Centrifugal Bioreactor Design

2.1.1 Physical Reactor Development

Originally developed by Van Wie et al. for growth of mammalian and microbial cells,[13] the CBR may be used for high density hybridomas,[14] scaffold free tissue culture[15, 16, 17] and most recently, for expansion of killer T cells for use in cancer immunotherapy.[12] The recent developments of the system are a promising step forward for next generation biomanufacturing. Within this industry, immunotherapy is an area that will benefit from future innovations, and integration of digital twins and real time control. In this new form of medical treatment, a patient's own immune system is used to fight cancers and other diseases; immune cells from the blood are isolated from the patient and expanded outside of their body.[18] This treatment can be applied to many cell types, including Cytotoxic T Lymphocytes (CTLs) and genetically engineered Chimeric Antigen Receptor T cells.[18] Currently, the process has only been made available to a select number of

patients, and mostly on a small scale, due to limitations in density and efficiency of the bioreactor systems used to expand the cells.[12] The prototype system mentioned above and developed by Kaiphanliam et al.[12] resolves these obstacles and uses a balance of centrifugal and fluid forces, as well as perfusion with continuous addition of fresh medium and removal of waste, to culture T cells faster than existing systems currently used in clinical and industrial settings. As shown in Figure 2, there are centrifugal and fluid forces acting on each cell in the chamber, due to the rotating motion of the reactor and the fluid flow from fresh medium entering the chamber, respectively. The force balance allows for achievement of higher cell densities and processing in a shorter time compared to existing reactor systems used for the same treatments. Past work has optimized the performance of the CBR using kinetic models, based on measured growth data from various cell lines.

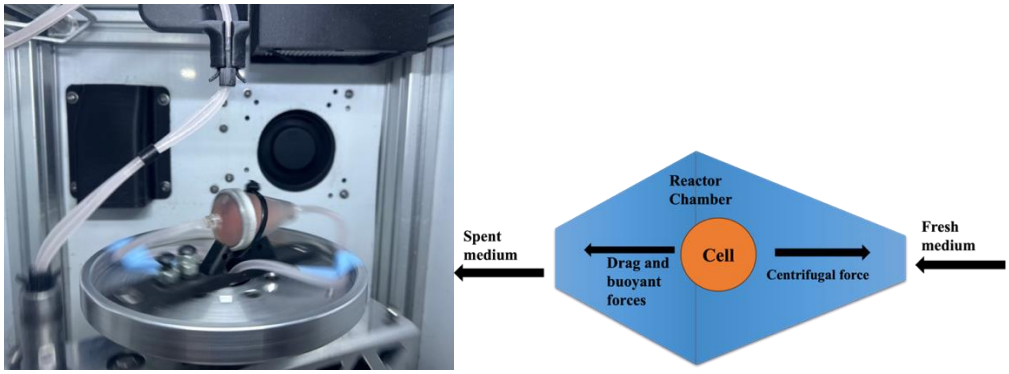

Figure 2: Centrifugal bioreactor design. At left, conical cell chamber in the physical system prototype, mounted on a centrifuge disc with anti-twister tubing above. Right, balance of forces within a conical chamber.

2.1.2 Cell Growth Modeling

The kinetic growth model used to represent the CBR system is based on a generalized set of Monod equations first published by Han et al.[19] The model was most recently used in modified form by Kaiphanliam et al. (2023),[12] to optimize the growth of bovine cytotoxic T cells (CTLs) in the CBR. The model accounts for the effects of the primary nutrients required by the cells to grow, glucose and oxygen, using these factors to determine the growth rate of the CTLs.[12] Three of the primary model equations are provided here. Equation (1) represents a simplified equation for the cell growth rate in terms of the maximum specific growth rate, μ_{max}, and the concentration of cells, C_{Cell}. Equation (2) provides the rate of change in glucose concentration inside the CBR chamber as a function of the difference between the glucose concentrations in the fresh feed stream entering ($C_{G,FF}$) and the spent medium exiting the process ($C_{G,out}$) as well as the volume of the system (V_{Rxtr}). It also includes a yield coefficient, Y_{GC}, which represents the amount of glucose used as cells expand. Finally, Equation (3) represents the concentration of dissolved oxygen inside the chamber, $C_{O,out}$, based on the entrance oxygen

concentration ($C_{O,0}$) and the oxygen consumption rate (*OCR*), a measure of how much oxygen cells use as they grow[12] and the total oxygenated fresh feed (Q_{FF}) plus the reoxygenated recycled medium rate.

$$\frac{dC_{cell}}{dt} = \mu_{max} \cdot C_{Cell} \quad (1)$$

$$\frac{dC_G}{dt} = \frac{Q_{FF}}{V_{Rxtr}}\left(C_{G,FF} - C_{G,out}\right) - Y_{GC} \cdot \left(\frac{dC_{Cell}}{\partial t}\right) \quad (2)$$

$$C_{0,out} = C_{O,0} - \frac{OCR \cdot N_{cell}}{(Q_{FF} + Recycle\ Rate)} \quad (3)$$

2.1.3 Reactor Control Simulations

In the present study, the physical CBR system is replicated via simulation in MATLAB Simulink, hereon referred to as the SPR. We have created a matching DTR in another installation of Simulink on a separate desktop computer and connected it to the SPR computer using MathWorks ThingSpeak, a cloud based IoT interface for MATLAB. Both simulations are based on the mathematical models introduced earlier; the model allows for calculation of live cell count based on the concentration of oxygen in the system. Sensor readings for glucose and oxygen are sent from the SPR computer to the DTR computer through ThingSpeak, integrating process control algorithms and allowing for real-time updating of fresh medium flowrates. The simulation is designed to be integrated with digital controllers, making it possible to transmit data from sensors and command flow rate changes in the actual physical CBR, and allowing for simple replacement of the virtual reactor with the real physical CBR system in future work.

a.

b.

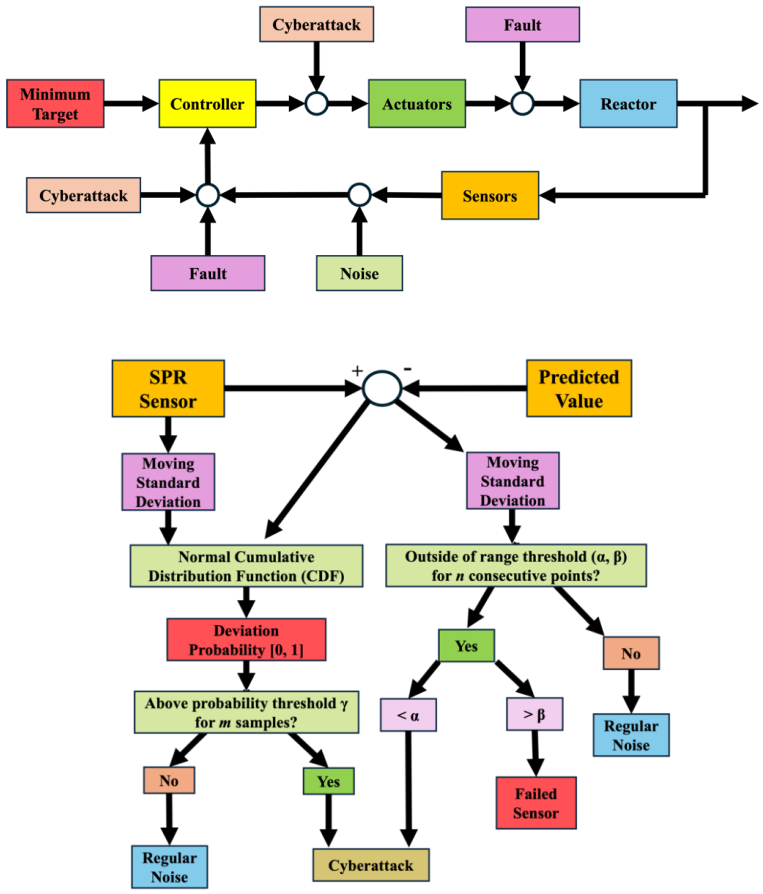

Figure 3: Process control simulation design in MATLAB Simulink. (a) Design of Simulink model. including inputs of sensor data, response from controllers and actuators, and effects of cyberattack. (b) Disturbance detection and response model, showing the two possible responses based on whether the sensor reading falls within the normal CDF or outside of threshold, distinguishing attacks from faults.

The SPR simulation contains the CBR itself, sensors and pumps for glucose and oxygen, its controller, and a series of subsystems to emulate attacks and system disturbances. As shown in Figure 3 the controller receives inputs in the form of simulated glucose and oxygen sensor data, produced by the kinetic growth model, and adjusts outputs of fresh and recycled medium pumps as needed. The overall objective of the controller is to keep the glucose concentration above a specified minimum amount of 50 mg/dL while keeping the net flowrate, the sum of the fresh feed and recycle rates, constant until the target cell count is reached. The cell count is determined via calculation using the known dissolved oxygen concentration coupled with the oxygen consumption rate per cell per hour. This basic process control cycle consists of reading a moving average of glucose sensor data that has been contaminated with emulated Gaussian noise and increasing fresh feed rates when the average drops below a specified threshold, causing a spike in glucose concentration. Once a specified cooldown window is passed, the cycle can repeat,

increasing the fresh feed by larger amounts as the cell count increases until a maximum fresh feed rate is reached.

Here, we have simulated a culture of CTLs in the SPR in MATLAB Simulink. The simulation utilizes the forward Euler method via Simulink's built-in equation solvers to build the growth model discussed earlier and inputs the previously measured kinetic parameters for the CTLs[12] in the model. In this work we set the simulation to culture cells for 300 hours at a sampling rate of 1 sample per sim hour, starting with an initial seeding of 1 x 10^6 viable cells. The process is modulated with a custom control framework and cyberattacks are introduced as discussed in the following sections. Cell growth, glucose and oxygen data have been collected and stored in Simulink in real-time.

2.2 Digital Twin Reactor

The objective in using the DTR is to provide fault and disturbance detection to the SPR through analysis of sensor concentrations and pump flowrates sent via ThingSpeak. There are eight signals (referred to as fault codes) the DTR can send to the SPR: no faults detected, oxygen sensor disturbance, oxygen sensor failure, glucose sensor disturbance, glucose sensor failure, recycle pump failure, fresh feed pump failure, and cyberattack detected. In general, these fault codes are triggered by finding unexpected differences in SPR sensor data to the DTR simulated sensor data, with the manner and method of the difference determining which type of fault is found. The fault codes are generated and sent to both the operator and the SPR, but no corrective actions are taken in response.

2.3 Fault and Attack Simulations

Although the types of faults a sensor and an actuator can experience are vast and varied, for the purposes of this simulation, only a single common case for each is considered. Based on recent experiments in which we performed initial tests on cell culture oxygen sensors, as well as past studies in literature,[20] it is evident that oxygen sensor failure degrades its value to a low fixed reading that is experimentally determined, with small amounts of electrical noise present in the system. Conversely, real-time glucose sensors for cell culture, being a relatively recent development, have no predetermined failure condition. However, our initial experimental observations with a commercially purchased glucose sensor indicate a large increase in noise far beyond that seen during regular operation. Each of the actuators, the fresh medium and recycled medium peristaltic pumps, most likely fail by a motor, gearbox, or electrical fault, thereby shutting off their flow to the cell chamber.

Our simulation framework features two subsystems that are used to emulate various system faults or cyberattacks. In the fault subsystem, each failure is

individually toggleable, and can be triggered at any predetermined time. The pump failures simulations are the simplest, and operate by reducing their respective flowrates to zero, regardless of controller's commands. Known types of sensor faults can be categorized either as incipient failure, where the sensor readings are inaccurate due to insecure process conditions (even though the sensor itself has no actual fault), or abrupt failure, where the systems are damaged or suddenly shut down, which is clearly identified by major drops or increases in the sensor reading at the moment of the damage.[21] Types of incipient failure include sensor bias, a very common occurrence where the sensor output is switched with a continual reading at an incorrect value, staying constant. Sensor drift is another type of incipient fault where the readings stray from the sensor's calibration setting over time, with any subsequent readings going up or down as a function of time. Sudden failures can include damage to the system, but also noise from outside sources, usually due to physical impacts on the sensor.[21]

To emulate oxygen sensor failure, assumed to be a stuck sensor fault, the sensed value is replaced with a fixed low non-zero value with added Gaussian noise at a higher level than what is typically added to the sensor, similar to methods discussed in previous literature.[22] Glucose sensor failure, assumed to be a bias fault, is simulated by adding additional Gaussian noise to the sensor's reading, with a range about four-fold larger than typically added by the standard noise subsystem.

The attack subsystem is designed to falsify sensor data to deceive the SPR into making erroneous control decisions. This may be in the form of falsifying a system fault or may be done to convince the controller that a reactor run is completed despite being far from it. Faults and cyberattacks can be the most difficult to differentiate between if the attack is designed to simulate a sensor fault.[23] In these attack simulations, a variety of parameters can be implemented: initial time of attack, attack onset time (e.g., instantly or over several timesteps), constant value to add or subtract, and scaling factor of sensor reading. Similar models have been used in literature to represent attacks of a constant linear magnitude[24]. Depending on the value of each parameter, different types of faults can be simulated, such as bias, drift, offset, and scaling, as well as combinations thereof.[25,26] For example, an attack simulating a combination drift-scaling fault would, at a certain time and over the course of tens of samples, scale the sensor reading up by a scaling factor α and add a constant value δ. Conversely, a falsified steady-state reactor condition would consist of replacing and exponentially decreasing an oxygen sensor reading with a stagnant value.

2.3.1 Failure and Cyberattack Detection

The detection subsystem of the proposed framework is designed to differentiate between the three disturbances considered: noise, faults, and cyberattacks. This is accomplished by comparing SPR sensor data to DTR simulations in tandem with various analysis tools in real-time. There is wide variety of both cyberattacks and

faults to which a process control system may be susceptible, but in this study the DTR is designed to recognize only the single fault condition prescribed for each sensor and actuator and treats all other non-noise disturbances as cyberattacks. Figure 3(b) shows two of the many components of the detection subsystem that is used to determine the type of disturbance detected. There are several subsystems involved that produce unique flags as different disturbances are detected. Depending on the types, order, and hierarchy of flags triggered, fault source or attack detection flags are generated. To inform an operator that a disturbance of some kind has been detected, regardless of source, two precursor flags are used, one each for the glucose and oxygen sensor. These are triggered preemptively when any subsystem detects a problem and waits for additional faults to occur before deciding on the source and cause of the issue. Although these flags are generic and can be triggered in response to a wide variety of issues, they serve to give an operator time to prepare a response to any disturbance pre-emptively.

The oxygen sensor failure detection subsystem compares sensor readings against a known fault range and is triggered when they stay within this range for an abnormally long period of time. The glucose sensor failure subsystem performs similarly except that it looks at the standard deviation of the data instead; an approximate known noise level of a regularly performing glucose sensor is considered, and if the standard deviation of a 10-sample moving average exceeds a threshold for several timesteps in a row, a glucose sensor fault code is generated. The fresh medium pump failure subsystem is used to observe the glucose sensor in response to an SPR-controller increase in fresh feed-recycle ratio. If a moving average of the glucose concentration does not increase several times within a window after the fresh medium pump flowrate has been increased, then the fresh feed pump is found to be unresponsive to commands and therefore has likely failed. The recycled medium pump failure subsystem utilizes oxygen sensor data instead, as the recycle pump provides most of the net flowrate and therefore contributes more oxygen to the reactor than the fresh medium pump. Recycle pump failure is evidenced by a sharp and sudden drop in oxygen concentration, followed by a more gradual but still abnormally rapid drop in oxygen thereafter. The detection subsystem is primed by this large drop and is formally triggered by consistent drops in oxygen concentration for several data points after.

The cyberattack detection system chiefly revolves around the comparison of DTR data to SPR outputs. Both the SPR and DTR simulations are a series of differential equations based on the CTL kinetic model, and when solved using the same methods, the same timesteps, and the same initial conditions, they agree with each other exactly. Therefore, the values from the DTR and SPR readings should vary from each other only by a factor of the noise present from the sensors. This noise is simulated within the SPR by adding Gaussian noise to the sensor data, parameters of which are experimentally determined from actual glucose and oxygen sensors. The first detection subsystem is used to compute the magnitude of this noise by determining the standard deviation of the difference between the ideal DTR value and the true SPR reading on a moving window and used to verify that

the output of this function is within the acceptable range. If the value is too low, for both sensors an attack flag is triggered, as this would only be caused by an attack that scales down the sensor values considerably or outright replaces the sensor reading with a fixed value. For an oxygen sensor with a high standard deviation reading, the cause can certainly be attributed to falsified high noise or an upward scaling attack. For the glucose sensor, a high standard deviation may be due to a similar attack method but could also be attributed to a simple sensor fault, requiring additional detection methods to deduce the true cause. This involves the second detection method which utilizes a normal cumulative distribution function (CDF) to compute a deviation probability between the sensor reading and simulated value. A rolling window data set is used in conjunction with the difference between the two values to determine a probability between 0 and 1 that the sensor readings have deviated from their expected values, with higher probabilities implying a higher chance of deviation. If the probability exceeds a threshold for a specified period of samples, this flag is triggered. If a deviation flag is thrown in tandem with a low or high noise flag for either sensor, the ultimate result is determined to be a cyberattack. Conversely, if a deviation flag is triggered in tandem with an oxygen fault flag or a pump failure flag, then the DTR determines that sensor or actuator fault is the true cause of the deviation.

3. Results

3.1 Simulated Physical Reactor Culture

Modeling of operation in a SPR culture under conditions of no anomalies is important to understand how key variables vary as the cells grow and how the process may be controlled to keep the cells growing with enough O_2 and glucose. Results of the SPR bioreactor culture model are shown in Figure 4 below. The culture expands as expected based on the kinetic model, increasing from 1×10^6 to 12×10^8 viable cells in 300 hours. Glucose drops from 185 mg/dL down to 20 mg/dL after the cells reach exponential growth and consume the glucose, as anticipated, and dissolved oxygen concentration drops from an initial value of 0.245 mM down to 0.15 mM, decreasing as the cells grow.

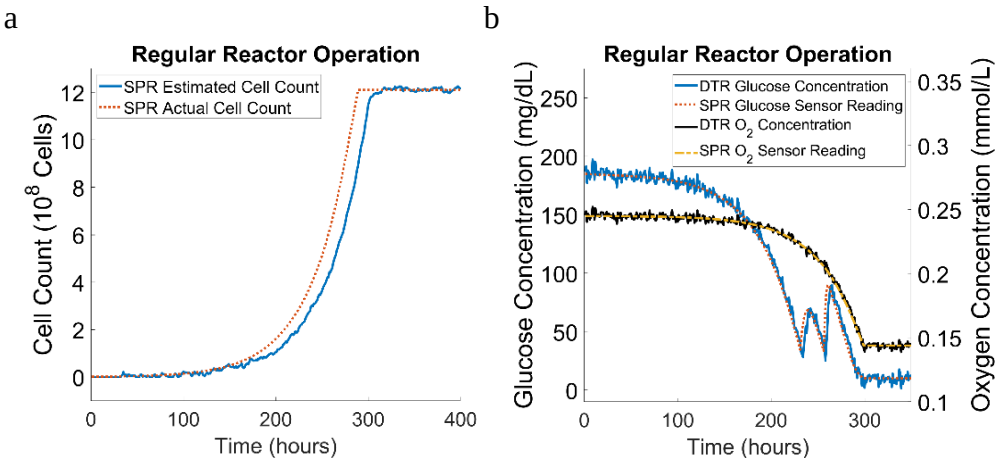

Figure 4: Results of SPR/DTR culture under regular operation. Plot shows the resulting increase in cell count and decreases in glucose and O_2 concentrations. At left, (a) cell count in SPR/DTR. At right, (b) glucose concentrations and dissolved O_2 concentrations in SPR/DTR.

3.2 Digital Twin Reactor Culture

As discussed earlier, the DTR is designed to collect data from the simulated SPR sensors and subsequently predict anticipated results using the kinetic model to determine if the process is running normally or if there is a sensor attack or failure. The DTR and SPR, when given the same flowrate parameters follow the same paths for glucose concentration, oxygen concentration, and cell count in the absence of system faults or cyberattacks. The minor discrepancies present can be explained by communication delays between the two simulations, because although every attempt is made to ensure that both reactors receive the same inputs simultaneously, an occasional delay leads to one sensor reading leading the other by a single timestep margin. Such differences are rarely more than 2%. When noise is added to the SPR's sensors before being overlayed upon DTR data, the difference is not visually noticeable as the noise is larger in magnitude than the difference between the simulations. As such, the DTR only generates false positives for sensor faults and cyberattacks entirely due to communication delays from using ThingSpeak. The most common occurrence is when the communication delay coincides with a change in flowrate. While the SPR will respond to its controller's command, the DTR will continue running with the old flowrate until the new command finally reaches it. The DTR will then read consistently lower glucose values than the SPR, triggering a false positive for a disturbance, typically in the form of a cyberattack. If communication delays are not present – or if they do not occur during a critical control action time – the simulations match well, as shown in Figure 4. This is because the disturbance detection is designed to explicitly ignore sections of data lost due to connection delays.

3.3 Fault and Cyberattack Simulations

Fresh feed pump failure detection, shown in Figure 5(a), is found to be as slow as 145 samples on one occasion and is rarely ever faster than 90 samples. This is because a flowrate increase needs to be initiated for the detection algorithm to run. If the failure occurs long before the reactor glucose concentration reaches the minimum threshold, then it can take tens or hundreds of samples to reach the concentration criteria needed to require a flowrate increase, such as in the case of a fresh feed pump failure at 80 hours. Additionally, if the failure occurs after the fresh medium flowrate has been maximized, there will be no further increases to command, and therefore no trigger for the fresh feed pump failure detection will occur. Recycle pump failure, shown in Figure 5(b), typically takes 10 to 30 samples to detect, although the software may be used to accurately determine an issue that has occurred.

As shown in Figure 5(c) and 5(d), and 5(e) and 5(f), respectively, glucose sensor and oxygen sensor faults are detected most reliably of any disturbance. On average, the glucose disturbance flag is generated within 20 samples of the fault first occurring, and the oxygen disturbance flag within 15 samples. In both cases, the digital twin consistently and correctly identifies the issue for each respective fault within 10 samples thereafter independent of the time the fault occurs. Actuator failure detection, particularly the fresh medium pump, is generally less reliable in terms of accuracy and timing than sensor failure detection. For the oxygen and glucose sensor attacks that do not emulate the predetermined fault condition, a disturbance flag is consistently generated between 15 and 25 samples, with an additional 10 samples required to accurately determine that the source of the disturbance is an attack. Attacks that onset instantaneously are typically detected a few samples sooner than those that onset over a period of tens of samples. For both sensors, both attack detection methods – deviation and noise detection – are utilized to make the correct decision about the source of the disturbance. When the attacks are designed to perfectly falsify sensor failures, the attack detection does not correctly identify the source. While the detection system correctly generates a disturbance flag, it attributes the disturbance to the incorrect source. Glucose sensor attacks of this variety reliably trigger a glucose sensor fault flag, while oxygen sensor attacks generate a mix of oxygen sensor faults and recycle pump faults.

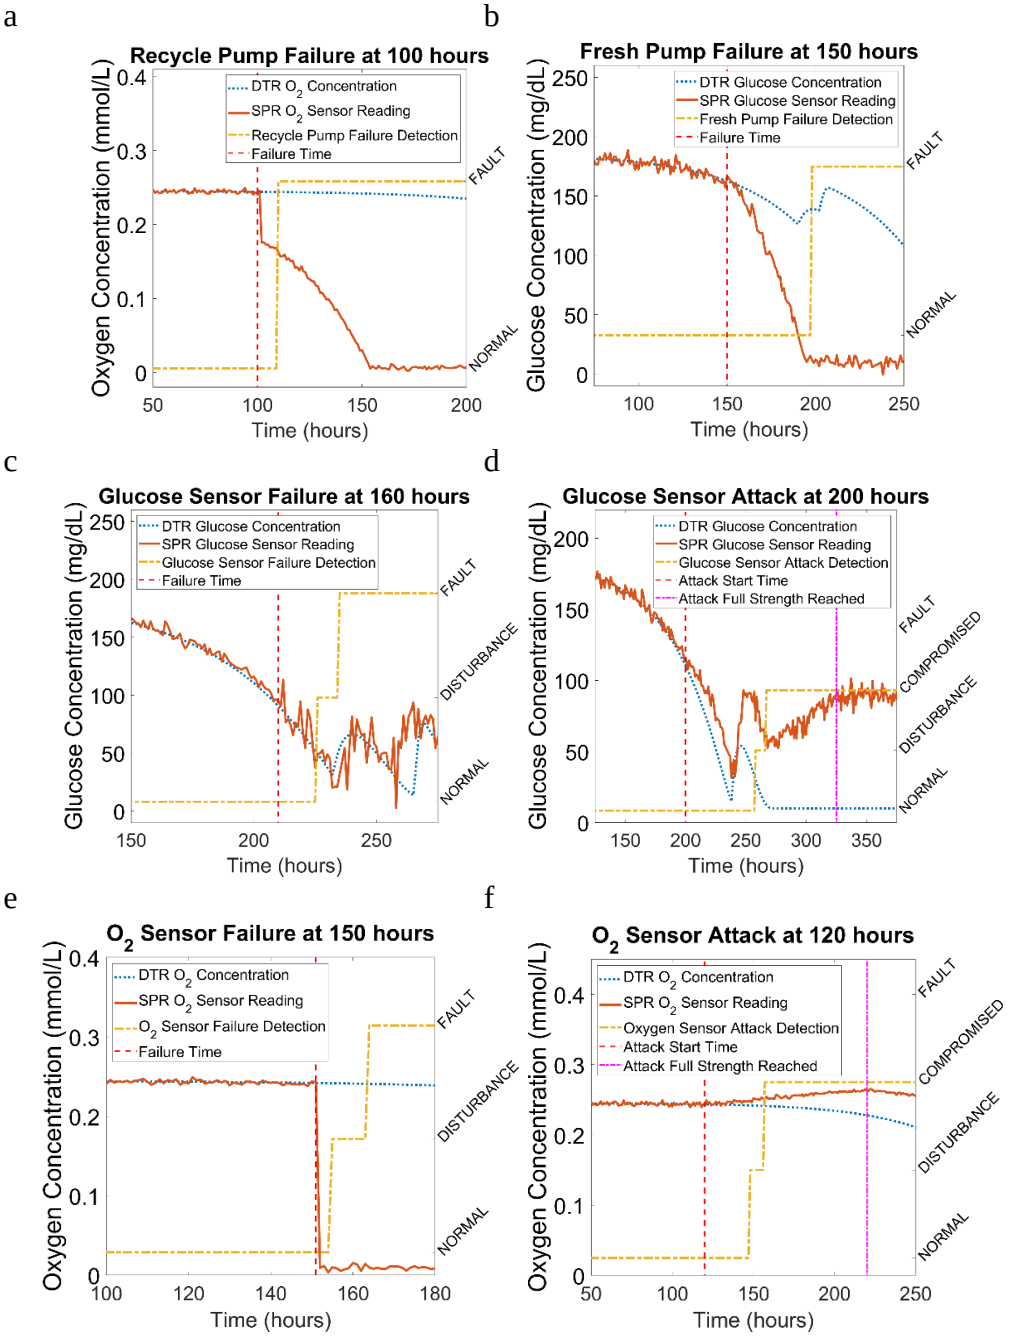

Figure 5: Data from process simulation, comparing equipment faults, cyberattacks, effects on SPR/DTR operation, and response at various time points. (a), failure of recycle pump (b), failure of fresh feed pump, (c), failure of glucose sensor, (d), cyberattack on glucose sensor, (e), failure of O_2 sensor, (f), cyberattack on O_2 sensor.

4. Discussion

4.1 Comparison of Simulated Physical Reactor to Digital Twin Reactor

Under normal operation, both the SPR and the DTR perform as expected. The outcomes found in our study match the results achieved in the real physical CTL CBR culture performed by Kaiphanliam et al.[12] Thus, the simulations are well-suited to provide a baseline process control framework to the real-world CBR process in future work. However, introduction of attacks and faults presents further complexity for the process, especially when a system is to be brought online with the function of cloud data transfer. For a complicated mechatronic system with network connection, determining the appropriate administrative and technical countermeasures in response to a disturbance requires an accurate assessment of the source. Falsely attributing a component failure to a cyberthreat can waste resources securing a network that is already secure or investigating something other than the physical system. Conversely, conflating a cyberattack for a component failure not only squanders replacement parts and diagnostics time, but more importantly allows a cyberthreat to spread throughout a network completely undetected, potentially intruding into attached databases or other network-connected machines. Therefore, simulation detection methods place an emphasis on properly assessing the source of the disturbance and require additional processing power and time to do so. Under certain circumstances, such as a falsified sensor failure, the DTR may incorrectly attribute a cyberattack as a physical fault or vice-versa, but an operator would nonetheless be made well-aware of an issue with the system. These fault codes are designed to be utilized as a guide in tandem with rigid troubleshooting procedures such as physical testing of sensors purported to have failed rather than a standalone detection and response mechanism. The delays to decreased glucose readings being recorded by the DTR, which can lead to false detection of an attack.

Upcoming work with the DTR/SPR combination will focus on full integration of the DTR with the physical prototype CBR system, including connection to a Wi-Fi integrated microcontroller. The demonstrated application of the DTR here allows for plug-and-play replacement in the real physical reactor. With applications being expanded beyond the simulated systems, this allows for anticipation of new types of equipment failures. For example, the real system will allow for consideration of sensor and pump replacement, in the instance of a failure detection or the occurrence of an attack that permanently damages a pump. An automated system that re-directs flow to a different pump in the event of a detected failure would be a valuable addition to the system.

4.2 Fault and Cyberattack Simulations

The digital twin is best suited to detect oxygen sensor and glucose sensor failure as evidenced by the consistent detection of such faults. Nearly all simulations involving a timed sensor fault resulted in DTR detection within a few samples independent of the failure time. When detecting fresh feed or recycle pump failures, the DTR is useful in detecting both, but the conditions under which they can be detected in a timely manner are limited. For example, the fresh medium pump fault detection can only be tracked after the controller has commanded a flowrate increase to the pump, limiting the detection only to times when such an increase is required. When the fresh medium pump failure occurs early in the run, at 80 hours, the fault is initially attributed to a glucose sensor attack, before the flowrate increase command at 220 hours finally causes the detection subsystem to accurately assess a pump failure. This can be solved by having the DTR poll the SPR for a fresh medium flowrate increase if the glucose sensor is reading lower than anticipated to confirm that the glucose concentration is still responsive to flowrate changes.

Many of the shortcomings of the digital twin can be addressed with additional sensors elsewhere in the system. Redundant glucose and oxygen sensors, or sensors placed at different points in the fluid path, e.g., both upstream and downstream of the cell chamber, would be useful to verify sensor operation as well as to better validate the inner mechanics of the cell chamber.[27] Additionally, visual verification, either with video processing or by having an operator physically check on the system, would greatly augment the digital twin's troubleshooting process. In the case of pump failure, confirmation of whether a peristaltic pump is rotating would be invaluable information to the DTR. Performing such physical checks, which may or may not be achievable without human intervention, would allow for more accurate fault and cyberattack detection, including a wider variety of types of such disturbances. Machine learning may also be a useful tool in training a digital twin to detect both faults and cyberattacks.[3] Hypothetically, every single failure mode for each sensor and actuator can be simulated at every single possible time during the reactor run, with the final data set being used to train a machine learning algorithm. Such algorithms can also be configured to learn past events, such as false positive flags to continually improve their performance. Furthermore, such training data can be shared across different bioreactor systems across secure cloud integrations through federated learning approaches.[28] Through such methods, the DTR can more accurately predict the source of the fault should one occur within the SPR, and it can further be used to determine if the system is being attacked if no pattern is matched. This will require an attacker to have an extremely in-depth understanding of the operation of the system to falsify a sensor failure, as more minor differences between the machine learning dataset and the actual data will be more likely to be attributed to a cyberattack than with the current methodology. Coupled with the addition of more sensors and verification tools, a digital twin with machine learning capabilities can be a very powerful and resilient process control tool.

In general, the larger the magnitude of the attack, the more of an effect it has on the cell growth rate, while also making them easier to detect. Therefore, most attacks that slip past the detection system are typically not of enough magnitude to have a significant effect on the growth rate or final cell count. Perfectly emulated falsified failure attacks, which would require intimate knowledge of the sensor's specification and extensive data logging to design, cannot be properly detected, and use of our system will reveal it to be a fault rather than an attack. This sort of attack is difficult to properly detect and counter using purely process control methods and highlights the need for additional layers of security. These can be more technical methods such as more typical cybersecurity measures but can also include administrative controls like restricting access to system specifications and materials lists to approved personnel. The more knowledge a would-be attacker has about a system, the more it could be tuned to be undetectable, meaning preventing an attacker from gaining that information can be just as important as being able to detect it. For a device that has publicly available component lists or specifications, this type of attack may not be detectable at the process level, thereby reinforcing the need for more traditional cybersecurity measures deployed elsewhere. Implementing cybersecurity measures on a local process controller is a final defense method and should be utilized as such.

4.3 Scale-up of the CBR/DTR Combination for Commercial Biomanufacturing

The integration of DTR systems in biomanufacturing is an important step forward to increase product yields and refine efficiency for biological processes. In comparison to other chemical industries, biomanufacturing is still on a smaller scale, particularly for immunotherapy production, with many current bioreactors limited by their maximum cell densities and overall efficiency.[12] As these processes increase in complexity, and more capacity for higher production levels is introduced, the number of DTR-based systems will need to go up. The current level of security for the industry opens many risks that need to be addressed for safe implementation of this technology on a large scale. As presented here, a framework for process-focused cybersecurity exists, but for a single, small-scale reactor. Our current work focuses on a single CBR system with an 11-mL chamber, suitable for production of T cells for a single immunotherapy treatment. Future developments will focus on expansion to larger scales, with several chambers and centrifuge disks present, to achieve greater cell yields for multiple patients. In this case, for every subsequent reactor connected to the system, an additional digital twin should be added. This will introduce further potential for cyberattacks, and thus, the control system should be re-examined and analyzed to expand upon possible attack types. Additionally, as biosensor design expands and allows for better reading of critical process parameters including glucose and dissolved oxygen, for larger scale-processes, the anticipated sensor readings and failures will need to be modified. Lastly, one aspect that our control system does not address is regulation of the centrifuge speed. While this is a constant parameter throughout a standard culture process, it may be subject to dynamic changes at higher cell densities or differing

cell types, depending on the possible need for high centrifuge speeds at greater cell densities. At any point, if the centrifuge rotor were to be remotely operable as part of the control loop, it would become vulnerable to attacks. This would require the design of new centrifuge failure scenarios akin to those designed for the pump failure modes. Nonetheless, our system is prepared for future large-scale expansion.

4.4 Application of the CBR/DTR Combination in Cyber and Biological Defense

With the continued development of biotechnologies, the analysis of potential cyber threats in this domain will need to be expanded. Our system provides a key small-scale example that can be applied to some of the most dangerous threats to national security in the biological domain: high-containment laboratories or HCLs. These labs contain the deadliest and most dangerous infectious diseases, vaccines, and pathogens.[29] Many of these facilities are already being connected to networks to increase process efficiency. Data for toxic pathogens are being stored in cloud databases, making this a prime target for opportunistic attackers. In these labs, automated bioreactors are already in use. While traditional cybersecurity is a critical point of defense for these systems, it is also necessary to implement security principles in the design of bioreactor automation. The development of new means to synthesize proteins and DNA may also result in greater threats of biological warfare. Already, many companies are commercially synthesizing DNA, an industry that makes millions every year.

Current efforts to integrate cybersecurity in this field have focused on data storage and analysis, using computer algorithms to compare all sequences to a database of known dangerous sequences.[30] But, the security systems have not addressed the actual manufacturing systems used to produce the sequences, which is a key area of weakness. Our case provides a basis for the application of security-focused process control. Many of the process control systems in the biotech industry are designed to anticipate various failure modes including those due to human error and are thus reducing the impact of human interaction by increasing automation.[31] These are also designed to predict equipment failure, as in our case. However, as discussed earlier, there are limited examples of an automated control system for biotechnology processes that may be used to anticipate cyberattacks in real time. For biomanufacturing process control systems to be fully implemented in a safe manner, there is a need for dynamic modeling to predict and identify threats as they happen, which we have provided in the present work. High-security laboratories, both in industry and in government settings, are an ideal starting point for real-world implementation of our system. These facilities employ bench-scale operations that produce dangerous products and would benefit from the integration of dynamic failure and attack prediction. Once a successful framework has been established on this scale for the most critical of products, it can then be expanded to large commercial scale processes, providing dynamic prevention of cyberattacks.

5. Conclusions

In closing, we have developed a novel model system to study the impact of cyberattacks in next-generation smart biomanufacturing. Our automated process system controls the concentrations of glucose and oxygen, as well as feed and recycle flowrates, for a novel centrifugal bioreactor for T cell therapies, using process control simulations. The system may be used to reliably detect both attacks and sensor failures using process control and a digital twin reactor that runs simultaneously in parallel to the real system, ensuring that the instance of an attack is immediately detected and distinguished from an equipment failure. The framework we designed is well-suited for implementation in laboratories that handle sensitive biological research and for bioreactors that produce dangerous pathogens and viruses. The system is designed to allow for simple, plug and play integration of controller hardware in addition to physical bioreactors. In future work, we will be implementing the process control system in our physical centrifugal bioreactor prototype. As the field of biosecurity expands, and new potential threats of biological warfare are introduced, it is critical that the full scope of cybersecurity in this industry is examined, and every process step is protected. In our work, we demonstrate that process-control-based cybersecurity is a feasible first step towards preventing cyberattacks within the massive biotechnology industry.

Acknowledgement

The authors would like to acknowledge funding from the Griffiss Institute VICEROY program through award No. SA10012021MM0336 for the Northwest Virtual Institute for CyberSecurity Education & Research (CySER). Undergraduate students assisting with the research include Mikayden Weise and Adam Caudle, both supported through stipends from the VICEROY program.

References

1. Koblentz GD, Mazanec BM. Viral warfare: The security implications of cyber and biological weapons. *Comparative Strategy*. 2013;32(5):418-434. doi:10.1080/01495933.2013.821845
2. Dixon T. The grey zone of cyber-biological security. *International Affairs*. 2021;97(3):685-702. doi:10.1093/ia/iiab041
3. Gargalo CL, de las Heras SC, Jones MN, Udugama I, Mansouri SS, Krühne U, Gernaey KV. Towards the development of digital twins for the bio-manufacturing industry. In: Herwig C, Pörtner R, Möller J, eds. *Digital Twins: Tools and Concepts for Smart Biomanufacturing*. Springer International Publishing; 2021:1-34. doi:10.1007/10_2020_142
4. Zhang YHP, Sun J, Ma Y. Biomanufacturing: History and perspective. *Journal of Industrial Microbiology and Biotechnology*. 2017;44(4-5):773-784. doi:10.1007/s10295-016-1863-2

5. Chen Y, Yang O, Sampat C, Bhalode P, Ramachandran R, Ierapetritou M. Digital twins in pharmaceutical and biopharmaceutical manufacturing: A literature review. *Processes*. 2020;8(9). doi:10.3390/pr8091088
6. Shao G, Helu M. Framework for a digital twin in manufacturing: Scope and requirements. *Manufacturing Letters*. 2020;24:105-107. doi:10.1016/j.mfglet.2020.04.004
7. Park SY, Park CH, Choi DH, Hong JK, Lee DY. Bioprocess digital twins of mammalian cell culture for advanced biomanufacturing. *Current Opinion in Chemical Engineering*. 2021;33:100702. doi:10.1016/j.coche.2021.100702
8. Mylrea M, Fracchia C, Grimes H, Austad W, Shannon G, Reid B, Case N, Lawless WF, Llinas J, Sofge D, Mittu R. BioSecure digital twin: Manufacturing innovation and cybersecurity resilience. In: Lawless WF, Llinas J, Sofge DA, Mittu R, eds. *Engineering Artificially Intelligent Systems: A Systems Engineering Approach to Realizing Synergistic Capabilities*. Springer International Publishing; 2021:53-72. doi:10.1007/978-3-030-89385-9_4
9. Guttieres D, Stewart S, Wolfrum J, Springs SL. Cyberbiosecurity in advanced manufacturing models. *Front Bioeng Biotechnol*. 2019;7:210. doi:10.3389/fbioe.2019.00210
10. Udugama IA, Lopez PC, Gargalo CL, Li X, Bayer C, Gernaey KV. Digital twin in biomanufacturing: Challenges and opportunities towards its implementation. *Systems Microbiology and Biomanufacturing*. 2021;1(3):257-274. doi:10.1007/s43393-021-00024-0
11. Balta EC, Pease M, Moyne J, Barton K, Tilbury DM. Digital twin-based cyber-attack detection framework for cyber-physical manufacturing systems. *IEEE Transactions on Automation Science and Engineering*. 2023:1-18. doi:10.1109/TASE.2023.3243147
12. Kaiphanliam KM, Fraser-Hevlin B, Barrow ES, Davis WC, Van Wie BJ. Development of a centrifugal bioreactor for rapid expansion of CD8 cytotoxic T cells for use in cancer immunotherapy. *Biotechnology Progress*. 2023;n/a(n/a):e3388. doi:10.1002/btpr.3388
13. Van Wie BJ, Brouns TM, Elliot ML, Davis WC. A novel continuous centrifugal bioreactor for high-density cultivation of mammalian and microbial cells. *Biotechnology and Bioengineering*. 1991;38(10):1190-1202. doi:10.1002/bit.260381011
14. Detzel CJ, Mason DJ, Davis WC, Van Wie BJ. Kinetic simulation of a centrifugal bioreactor for high population density hybridoma culture. *Biotechnology Progress*. 2009;25(6):1650-1659. doi:10.1002/btpr.240
15. Detzel CJ, Van Wie BJ, Ivory CF. Fluid flow through a high cell density fluidized-bed during centrifugal bioreactor culture. *Biotechnology Progress*. 2010;26(4):1014-1023. doi:10.1002/btpr.395
16. Detzel CJ, Van Wie BJ. Use of a centrifugal bioreactor for cartilaginous tissue formation from isolated chondrocytes. *Biotechnology Progress*. 2011;27(2):451-459. doi:10.1002/btpr.551
17. Nazempour A, Quisenberry CR, Van Wie BJ, Abu-Lail NI. Nanomechanics of engineered articular cartilage: Synergistic influences of transforming growth

factor-β3 and oscillating pressure. *J Nanosci Nanotechnol*. 2016;16(3):3136-3145. doi:10.1166/jnn.2016.12564
18. Garcia-Aponte OF, Herwig C, Kozma B. Lymphocyte expansion in bioreactors: Upgrading adoptive cell therapy. *Journal of Biological Engineering*. 2021;15(1):13. doi:10.1186/s13036-021-00264-7
19. Han K, Levenspiel O. Extended Monod kinetics for substrate, product, and cell inhibition. *Biotechnology and Bioengineering*. 1988;32(4):430-447. doi:10.1002/bit.260320404
20. Xu X, Du Z, Bai Z, Wang S, Wang C, Li D. Fault diagnosis method of dissolved oxygen sensor electrolyte loss based on impedance measurement. *Computers and Electronics in Agriculture*. 2023;212:108123. doi:10.1016/j.compag.2023.108123
21. Li D, Wang Y, Wang J, Wang C, Duan Y. Recent advances in sensor fault diagnosis: A review. *Sensors and Actuators A: Physical*. 2020;309:111990. doi:10.1016/j.sna.2020.111990
22. Huang P, Chen X, Chai Y, Ma L. A unified framework of fault detection and diagnosis based on fractional-order chaos system. *Aerospace Science and Technology*. 2022;130:107871. doi:10.1016/j.ast.2022.107871
23. Tertytchny G, Nicolaou N, Michael MK. Classifying network abnormalities into faults and attacks in IoT-based cyber physical systems using machine learning. *Microprocessors and Microsystems*. 2020;77:103121. doi:10.1016/j.micpro.2020.103121
24. Saxena S, Bhatia S, Gupta R. Cybersecurity analysis of load frequency control in power systems: A Survey. *Designs*. 2021;5(3). doi:10.3390/designs5030052
25. Sinha A., Das D. SNRepair: Systematically addressing sensor faults and self-calibration in IoT networks. *IEEE Sensors Journal*. 2023;23(13):14915-14922. doi:10.1109/JSEN.2023.3277493
26. Tao H, Peng T, Yang C, Gao J, Yang C, Gui W. Voltage and current sensor fault diagnosis method for traction converter with two stator current sensors. *Sensors (Basel)*. 2022;22(6). doi:10.3390/s22062355
27. Kajmakovic A, Diwold K, Römer K, Pestana J, Kajtazovic N. Degradation detection in a redundant sensor architecture. *Sensors (Basel)*. 2022;22(12). doi:10.3390/s22124649
28. Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y. A survey on federated learning. *Knowledge-Based Systems*. 2021;216:106775. doi:10.1016/j.knosys.2021.106775
29. Crawford E, Bobrow A, Sun L, Joshi S, Vijayan V, Blacksell S, Venugopalan G, Tensmeyer N. Cyberbiosecurity in high-containment laboratories. *Frontiers in Bioengineering and Biotechnology*. 2023;11. https://www.frontiersin.org/articles/10.3389/fbioe.2023.1240281
30. Puzis R, Farbiash D, Brodt O, Elovici Y, Greenbaum D. Increased cyber-biosecurity for DNA synthesis. *Nature Biotechnology*. 2020;38(12):1379-1381. doi:10.1038/s41587-020-00761-y
31. Mantle JL, Rammohan J, Romantseva EF, Welch JT, Kauffman LR, McCarthy J, Schiel J, Baker J, Strychalski EA, Rogers KC, Lee KH.

Cyberbiosecurity for biopharmaceutical products. *Frontiers in Bioengineering and Biotechnology*. 2019;7. https://www.frontiersin.org/articles/10.3389/fbioe.2019.00116

Appendix

All simulations were designed in MATLAB 2023a. As of submission, copies of the files and programs used for this publication can be found at: https://github.com/alecschulerWSU/CySER_digitalTwinReactor

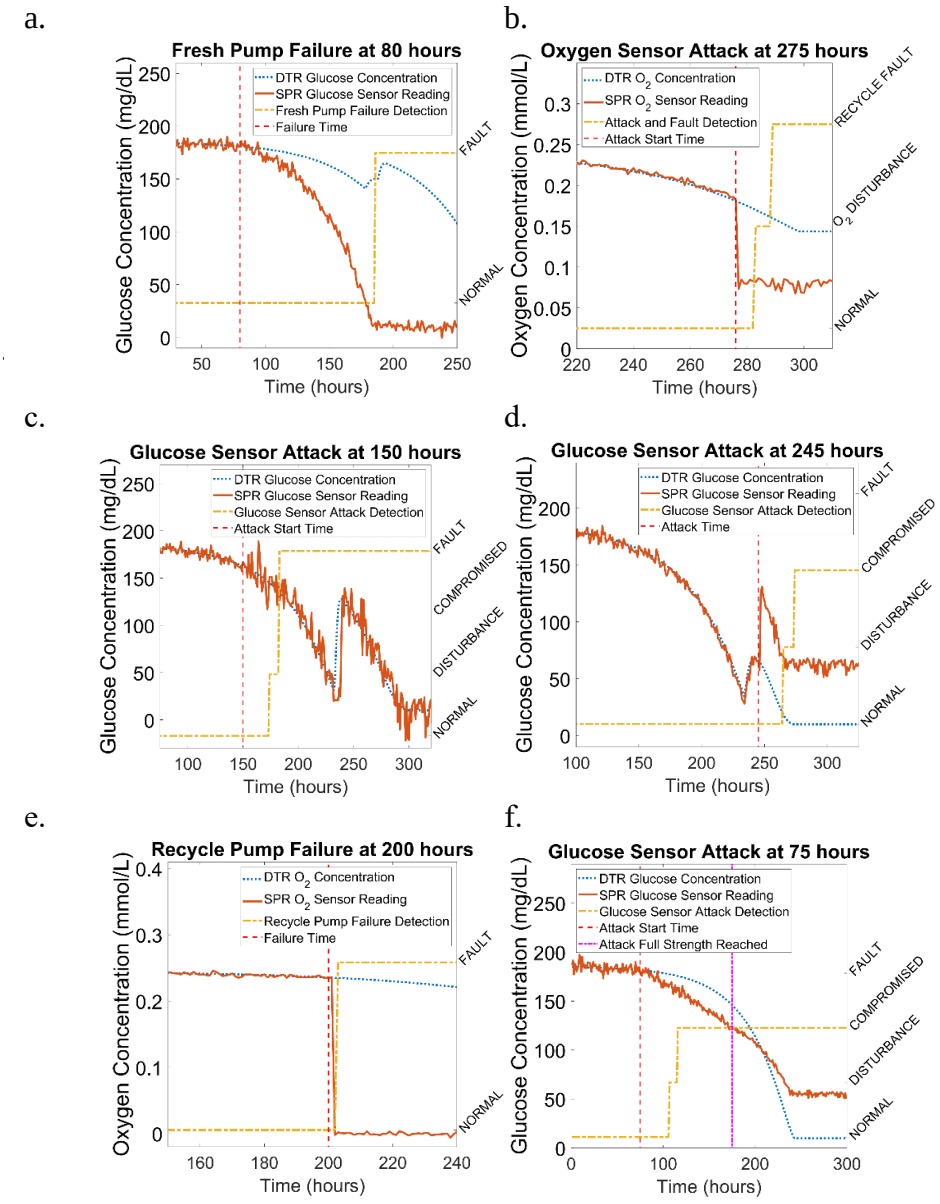

Additional data from faults and cyberattack simulations and their effects on SPR/DTR operation, along with detection. (a) Second fresh pump fault simulation with over 100 sample detection time. (b) Falsified oxygen fault attack with incorrect detection as recycle pump failure. (c) Second constant value glucose sensor attack (d) Falsified glucose fault attack with incorrect detection as fault. (e) Second recycle pump fault simulation with 10 sample detection time. (f) Second glucose sensor attack with correct detection within 30 samples.