Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

SIMULATING ENERGY AND SECURITY INTERACTIONS IN SEMICONDUCTOR MANUFACTURING: INSIGHTS FROM THE INTEL MINIFAB MODEL

Gabriel A. Weaver

Infrastructure Analysis Division
Idaho National Laboratory
Idaho Falls, ID 83415, USA

Jacob Shusko, John J. Hasenbein, and Erhan Kutanoglu

Operations Research and Industrial Engineering
The University of Texas at Austin
Austin, TX 78712, USA

Gonzalo Martinez-Medina and Krystel K. Castillo-Villar

Mechanical Engineering Department
The University of Texas at San Antonio
San Antonio, TX 78249, USA

Paulo C.G. Costa

Cyber Security Engineering Department
George Mason University
Fairfax, VA 22030, USA

ABSTRACT

Semiconductor manufacturing, particularly wafer fabrication, is a highly complex system of processes and workflows. Fabrication facilities must deal with re-entrant flows to support multiple types of wafers being produced simultaneously, each with their own deadlines and specifications. The manufacturing process itself depends upon the ability to control and programmatically adjust a variety of environmental conditions. In addition, wafer fabrication consumes large amounts of energy, particularly electricity. Emerging technologies including networked devices may help reduce the energy footprint but can introduce cybersecurity risks. Therefore, this paper presents its modeling and simulation framework to quantify tradeoffs between operational measures of performance, energy consumption, and cybersecurity risks. We augment the Intel Minifab model with an *Industrial Control Systems (ICS)* reference model based on the *Purdue Enterprise Reference Architecture (PERA)* as well as tool-level energy consumption data from a semiconductor manufacturing testbed.

1 INTRODUCTION

Semiconductor manufacturing is a highly complex process from both theoretical and practical standpoints. Unlike other types of workflows, semiconductor manufacturing processes have re-entrant flows in which wafers at varying stages may be processed at the same machine. In addition, wafer fabrication plants have to deal with multiple types of wafers being produced simultaneously, each with their own deadlines and specifications. The overall manufacturing process, particularly wafer fabrication, takes several months to complete and as a result, the same disruption—whether an accidental fault or intentional cyberattack—may have different impacts in terms of resources lost. The impact and resultant cost of its lost wafers may be more significant given recent semiconductor supply chain shortages, and this is reflected in recent federal initiatives to address supply chain issues.

The semiconductor manufacturing process depends upon the ability to control and programmatically adjust a variety of conditions. For example, *Heating, Ventilation, and Air Conditioning (HVAC)* systems must sense and regulate temperature, humidity, and particles in the air to ensure the quality of wafers produced. Machines within these rooms depend upon advanced automation systems, backed by communications networks, in order to orchestrate scheduling and routing, and provide *Human-Machine Interfaces (HMIs)*

978-1-6654-7661-4/22/$31.00 ©2022 IEEE

for operators, and access databases by which to configure and operate machines. Furthermore, the *Industrial Internet of Things (IIoT)* within the manufacturing environment promises more efficient, safe, and customized manufacturing workflows by collecting and analyzing data from vast sensor and device networks. Services such as predictive maintenance, informed by data such as temperature, pressure, and equipment usage history, promise to prevent unexpected and costly equipment failures. Furthermore, the introduction of IIoT devices can provide stakeholders with raw energy consumption at the tool level or an energy signature, this data can be utilized by multi-dimensional anomaly detection methods to highlight deviations from nominal values.

Although the IIoT promises more efficient use of resources, resulting in increased productivity and lower manufacturing costs, such technologies also have the potential to introduce risks. Specifically, increased dependencies of manufacturing workflows on emerging technologies may make fabrication plants (fabs) more susceptible to cyberattacks, including but not limited to ransomware. For example, the *Taiwan Semiconductor Manufacturing Company (TSMC)* suffered $170M in losses after having to suspend production for approximately 3 days due to the WannaCry ransomware (Dignan 2018; Ting-Fang and Li 2018). More recently, Tower Semiconductor paid $250,000 in order to retrieve key information encrypted during a ransomware attack so as to quickly resume production (Orbach 2020). Given the importance of semiconductor manufacturing globally, nation-state campaigns such as Operation Skeleton Key (Osborne 2020) against Taiwanese manufacturers underscore the need to quantify the risk of targeted, cyber-originating disruptions to semiconductor manufacturing. This aligns with recent efforts, including the SEMI Cybersecurity Standard, Specification for Malware Free Equipment Integration (SEMI 2022), as well as NIST 800-82, Guide to *Operational Technology (OT)* Security (Stouffer et al. 2022).

This paper presents a modeling and simulation framework to help stakeholders quantify and explore tradeoffs among a fab's operational measures of performance, energy consumption, and security risks. Specifically, this paper consists of three key contributions. First, there is a need for an approach to augment existing manufacturing reference models with critical infrastructure dependencies to capture dependency-based risks within the current semiconductor manufacturing ecosystem. Therefore, we present a reference framework that integrates the Intel Minifab model (Kempf 1994) with concepts based on the *Purdue Enterprise Reference Architecture (PERA)* (Williams 1994; ENISA 2018), capturing dependencies between the manufacturing and OT environments due to the IIoT. Second, using the reference model, we present a data analysis and simulation pipeline, that integrates data from actual semiconductor testbeds into the model to simulate the energy consumption and production of the Minifab model under the baseline and cyberattack-induced disrupted conditions. The energy baselining considers dependencies between electrical power usage (kWh) and machine state to estimate energy consumption over a planning horizon with high granularity. Third, we present a use case, based on a *Nearly Orthogonal Latin Hypercube (NOLH)* experimental design, to calibrate tool and fab-level optimization models using production and power consumption response surfaces computed using our pipeline. As such, this work sets the stage for being able to explore and optimize tradeoffs between energy consumption, cybersecurity, and production.

This paper is structured as follows: Section 2 surveys the state of the art and practice relative to our contributions. Section 3 provides background on IIoT for semiconductor manufacturing as well as a threat catalog given these emerging technologies. In Section 4, we use a multilayered network formalism to integrate the Intel Minifab model with an IIoT reference architecture based on PERA, enabling a simulation pipeline to model cyber-originating disruptions to semiconductor manufacturing workflows. Section 5 presents a ransomware-inspired use case, based on our pipeline, to construct response curves to calibrate optimization models at the tool and fab level as well as compute an energy baseline. Finally, Section 6 suggests future research and Section 7 presents our concluding remarks.

2 RELATED WORK

This section reviews the state of the art in the literature to understand the context and novelty of our contributions. For each of our contributions, we discuss relevant prior work. First, the analysis pipeline described in this paper augments representations of existing semiconductor manufacturing reference models with critical infrastructure dependencies, such as communications/IT, to enable dependency-based risk analysis within the emerging semiconductor manufacturing ecosystem. Within the academic literature, there is an acknowledged need to be able to represent and analyze semiconductor manufacturing control systems at different hierarchical levels (Shao et al. 2019). Previous work done by Fowler describes a model for semiconductor manufacturing consisting of four levels of simulation including individual tools, the manufacturing site, internal supply chain, and the end-to-end supply chain (Fowler, Mönch, and Ponsignon 2015). An overview of reference models for semiconductor manufacturing sites is provided by Kopp et al. (2020) who compare the scale of different models, including the Intel Minifab, MIMAC, Harris, and SEMATECH 300 models.

Our work focuses on cyber-originating disruption models for semiconductor manufacturing and as such, augments the well-known reference model for the Intel Minifab (Kempf 1994) with concepts based on the *Purdue Enterprise Reference Architecture (PERA)* for *Operational Technology (OT)* (Williams 1994; ENISA 2018; Langill, Hegrat, and Peterson 2019). There are a variety of emerging cybersecurity standards within the industry including the SEMI Cybersecurity Standard, specification of malware-free equipment integration (SEMI 2022); the NISTIR 8183A, Cybersecurity Framework Manufacturing Profile (Stouffer et al. 2019); and more generally the Guide to Operational Technology Security (NIST 800-82r3) (Stouffer et al. 2022). Texas Instruments recently motivated the need for a Cyber-Physical IT assessment tool for semiconductor manufacturing companies (Cayetano et al. 2018). Within the digital twin literature, however, some claim that "less is known or published on the topic of usability of a digital twin for cybersecurity" (Pokhrel, Katta, and Colomo-Palacios 2020). Our intent is to address this gap by augmenting a reference model for ICS networks and thereby enabling analysis of upstream infrastructure disruptions that impact manufacturing site operations and energy consumption.

Second, we provide a use case within the semiconductor manufacturing space that sets the stage for stakeholders to consider tradeoffs among operational measures of performance, energy consumption, and cybersecurity. Research on digital twins for manufacturing spans decades and includes the Virtual Factory (Fisher 1986), work by Jain et al. (2015) on the concept, and an entire panel held on the topic more recently (Shao et al. 2019). Furthermore, extensive research has also been conducted on digital twins for smart manufacturing system design (Leng et al. 2021; Haddod and Dingli 2021; Wang et al. 2021). As mentioned earlier, however, less work has been done on digital twins for cybersecurity (Pokhrel, Katta, and Colomo-Palacios 2020). Nonetheless, we claim that little research has been done into modeling the risks introduced by the adoption of these emerging IIoT technologies. The intent of our second contribution is to provide an approach to consider the impact of cybersecurity controls, such as those described by NISTIR 8183A, and their impact on operational measures of performance and energy consumption.

A variety of disruption modeling approaches have been considered within the simulation literature. In terms of short-term planning (less than 3 months), researchers have looked at disruptions due to changes to product mixes and machine breakdowns due to maintenance events (Scholl 2008). In considering stochastic modeling of machine downtimes and times between failures, Rose demonstrated (Rose 2004) the importance of the shape of the distribution in addition to the mean and variance of failure data. Moreover, Rose concluded that using exponential distributions can lead to misleading results. Other work by on long-term (longer than 6 months) disruptions briefly mentioned the impact of IT problems, but did not discuss this in detail (Scholl et al. 2012).

One intent of our framework is to start to understand whether properties of cyber-originating disruptions caused by an intelligent adversary can be represented by existing disruption models in the literature or whether models based on new assumptions need to be developed. For example, traditional stochastic disruption models may be well suited to model the timing of accidental infection by ransomware, such as appears to be the case with the WannaCry incident at TSMC described in Section 3 (Dignan 2018; Ting-Fang and Li 2018). Given that multiple tools may depend on a given cyber asset, however, other incidents will likely simultaneously affect more machines than an unexpected machine failure. Alternatively, an intelligent adversary targeting a semiconductor manufacturing plant may seek to explicitly time a disruption in order to maximize losses, motivating the need for other measures of impact beyond yield such as the embedded energy of products lost. Finally, we should note that traditional disruption models focused on machine availability *may* be applicable to some threat models but attacks affecting other security properties such as data integrity and confidentiality do not align with such an approach.

Finally, as our third contribution, we compute several simulation-driven response curves to be used to calibrate optimization models which will consider energy consumption, security, and production. In order to explore the space of disruption parameters, we adopt an NOLH approach to understand how different types of disruption models impact production as well as energy consumption. Such an approach employs experimental design principles, as advocated by Sanchez (Sanchez, Sanchez, and Wan 2020), and integrates data analytics with simulation (data farming). One key analysis resulting from this effort is the ability to consider energy consumption within a manufacturing site as well as at the level of an individual tool. Our energy baselining approach builds on work done by Kannaian to predict energy consumption within semiconductor manufacturing (Kannaian 2018), and others in the literature (Gopalakrishnan, Mardikar, and Korakakis 2010; Hu and Chuah 2003).

3 BACKGROUND

Industrial Control Systems (ICS) have been used to monitor and control a variety of semiconductor manufacturing processes, such as HVAC systems, to regulate temperature and filter air particles, as well as to schedule and route wafer lots through a facility. These processes are supported by *Operational Technology (OT)*, the computer systems used to manage industrial operations. Traditionally, these include *Programmable Logic Controllers (PLCs)*, *Remote Terminal Units (RTUs)*, *Human-Machine Interfaces (HMIs)*, and *Supervisory Control and Data Acquisition (SCADA)* software systems (McGinley 2021). The *Purdue Enterprise Reference Architecture (PERA)*, provides a framework to classify IT and OT devices into five levels and a basis to understand communication paths between them (Williams 1994; Fabrio, Gorski, Spiers, Diedrich, and Kuipers 2016). The notional layers of PERA also provide a means by which to implement and secure IT/OT networks through network segmentation (Fabrio et al. 2016).

Emerging *Industrial Internet of Things (IIoT)* applications include predictive maintenance, optimal routing of wafer lots, real-time monitoring, and machine-learning enhanced metrology (Technavio 2020). The integration of sensors and cloud-based analytics within the OT environment has motivated some to develop new reference models for network architectures that more closely reflect the current and emerging state of the practice (Langill, Hegrat, and Peterson 2019). For example, the *European Union Agency for Cybersecurity (ENISA)* published a high-level reference model for smart manufacturing that extends PERA to include sensors and cloud-based analytics (ENISA 2018). The ENISA model classifies IT, OT, and IoT devices into six levels and provides an overview of high-level communication paths that support manufacturing processes.

Traditional information security focused on the properties of *Confidentiality, Integrity, and Availability (CIA)* (Barker 2003). Although these categories are critiqued within the academic literature (Schneider 2011), we employ the CIA security model to categorize general threats within the cybersecurity industry that impact semiconductor manufacturing. More recently, practitioners have also used the PERA model to describe cyberattacks on ICS, even developing the ICS kill chain to describe how an adversary gains initial access, often via enterprise IT systems, and subsequently propagates to OT systems to affect industrial processes (Assante and Lee 2015).

Confidentiality ensures that data or information is not disclosed in an unauthorized manner (Barker 2003). Loss of confidentiality within semiconductor manufacturing ICS enables several adversarial behaviors defined within the MITRE *Adversarial Tactics, Techniques, and Common Knowledge (ATTCK)* for ICS taxonomy. Specifically, Theft of Operational Information for Impact (Alexander, Belisle, and Steele 2020) may result in loss of manufacturer's trade secrets or IP by exfiltration of sensitive data (e.g., information stored within a data historian). For example, Operation Chimera conducted during 2018 and 2019 targeted

Taiwanese semiconductor manufacturers for IP theft and data exfiltration (Osborne 2020). Moreover, with the adoption of IIoT-based applications, data may be gathered and processed by sensors and sent directly to the cloud (Langill, Hegrat, and Peterson 2019). Beyond operational information, other sensitive information such as valid usernames and passwords (e.g., Valid Accounts for Lateral Movement (T0859)) may be harvested and subsequently used for lateral movement through a manufacturing ICS network.

Data integrity prevents improper modification or destruction of data and includes non-repudiation and authenticity (Barker 2003). Loss of data integrity within a semiconductor manufacturing environment could result in Damage to Property (T0879) as well as Loss of Productivity and Revenue (T0828) through a mask tampering attack. Mask tampering attacks can introduce flaws into wafers produced or introduce malicious logic into a system. For example, in 2018, companies such as Amazon and Apple found a hidden chip on their server motherboards that could provide remote access to servers with the vulnerability (Singh and Szczys 2018). In addition, given increased reliance on sensors for more precise control of the manufacturing environment, the integrity of data generated by such sensors is another aspect of control systems cybersecurity.

Availability ensures timely and reliable access to information (Barker 2003). Availability of key services within an ICS environment is critical to manufacturing networks in which downtime of an asset may result in Loss of Productivity and Revenue (T0828) and Damage to Property (T0879). Over the past several years, ransomware attacks have affected semiconductor manufacturing plants. In August 2018, the TSMC had to shut down its plants for approximately 3 days after being infected with a variant of the WannaCry ransomware. As a result, TSMC had to shut down production at its ‘most advanced’ iPhone chip manufacturing site, resulting in approximately $170M USD in losses that affected quarterly earnings (Dignan 2018; Ting-Fang and Li 2018). Other known ransomware victims within the semiconductor manufacturing industry include MaxLinear, infected by the Maze ransomware in 2020 (Brennan 2020); X-Fab, which had to halt six production facilities in July 2020 for multiple weeks (Stegall 2020); and Tower Semiconductor, which suspended operations in some facilities on September 6, 2020 but resumed operations by September 10 after having paid the ransom (Orbach 2020). According to a report by Dragos, the Manufacturing sector dominated the ICS Sectors most impacted by ransomware in 2021 with 211 known incidents, Food and Beverage was next with 35 incidents (Dragos 2021).

4 MODEL

Although the adoption of IIoT may result in more efficient semiconductor manufacturing workflows, it also may increase the attack surface for the manufacturing system due to networked components with potential exposure to the cyberattacks. As a result, stakeholders need to be able to consider potential disruptions within a manufacturing workflow, relative to its dependencies on communications and computer networks. One key consideration for ICS security is to consider the risks of using specific technologies to implement or improve an operational process (Langill, Hegrat, and Peterson 2019). By augmenting the Intel Minifab model with a notional reference architecture such as PERA or related model, we hope to provide an integrated approach to discuss cybersecurity relative to business process (via graph-theory) and associated measures of performance (via simulation). As a result, stakeholders can consider the impact of adopting new technologies as well as how newly-discovered vulnerabilities relate to their operational processes.

4.1 Manufacturing Facility Model

Our analysis pipeline validates and parses a graph-based representation of the manufacturing workflow. In such a representation, vertices correspond to resources (machines, physical routing locations), while edges define possible routes between these machines. Queueing network parameters, such as service time and server capacity, are encoded as network attributes. Table 1 provides a more detailed overview of the manufacturing model parameters and measures of performance.

Table 1: Model Inputs and Outputs.

Simulation Model Parameters and Outputs	
Lot Parameters	
Variable	Description
w_l	Lot l wafer type (e.g. P_a, P_b, TW)
S	Number of steps in workflow, same for all wafer types
$s_l \in S$	Last completed step in workflow
A_w	Arrival rate for lots of a given type
Resource Parameters	
M	Workflow resources, including machines and physical routing locations
R_s	Adjacency matrix of size $\lvert M \rvert \times \lvert M \rvert$ defined for each workflow step S. Routes wafer lots based on their last completed step
$\mu_m^{s,w}$	Service time at machine m relative to last completed workflow step and wafer type
c_m	Number of lots that can be simultaneously serviced at machine m
K_m	Number of lots that can wait to be serviced
D_m	Queueing discipline at resource m
Outputs	
u_m	Mean utilization of machine m
$\widehat{WIP}$	Average work in progress
$\widehat{CT}$	Average cycle time

The manufacturing network is used as an intermediate representation to instantiate the simulation as well as provide a future capability for graph-theoretic analyses on the workflow. To implement the manufacturing network within our data processing and analysis pipeline, we encode workflow and simulation model parameters using the *Yet Another Workflow Language (YAWL)*. YAWL is a markup language for workflows, based on Petri Nets, that was extended to accommodate assumptions of real-world workflow constructs (Van Der Aalst and Ter Hofstede 2005). The simulation module parses the YAWL manufacturing workflow into an intermediate graph representation using the NetworkX library (Hagberg, Swart, and S Chult 2008). This intermediate representation is subsequently used to instantiate the simulation via the Ciw library, a discrete event simulation library for queuing networks (Palmer et al. 2019).

Within our simulation, lots of 25 wafers (Kempf 1994) move through a manufacturing workflow. By assumption, a lot (l) contains wafers of a single type (w_l). For example, the Intel Minifab model has three wafer types, P_a, P_b and test wafer TW. Depending upon their type, lots arrive at different rates (A_w) to the start of the workflow network and are routed to various workflow resources according to (R_s) depending upon their last completed step (s_l) and wafer type (w_l). Having arrived at a resource $m \in M$, lots wait in a queue to be serviced according to the FIFO discipline. Operational measures of performance output by the simulation include average *Work in Progress (WIP)* ($\widehat{WIP}$), average cycle time ($\widehat{CT}$), and utilization of machines, (u_m). These performance measures are averaged over multiple replications of each scenario.

Verification and Validation (V&V) was conducted to ensure correct implementation and accuracy of the model. Verification of the simulation model ensures that it is properly implemented with respect to the specification. As part of this effort, development of the simulation and the entire supporting pipeline included extensive unit testing to ensure proper code behavior. Moreover, model parameter units are explicitly specified and processed using the `pint` library. Validation of the simulation model ensures the accuracy of the model against the specification or real-world system. Validation efforts focused on the operational measures of performance.

4.2 Multilayered Networks: An ICS Disruption Model

We employ a multilayered network formalism (Boccaletti et al. 2014), as a natural approach to represent the communication flows represented by PERA-like reference models of ICS systems. At a high level, we define a multilayered network whose layers correspond to a subset of Levels as defined by the ENISA model (ENISA 2018). Cross-layer edges within the multilayered network correspond to known information flow dependencies within the reference model. For example, the Intel Minifab workflow network, encoded by YAWL, defines a network layer that corresponds to a Level 0 manufacturing process. Additional Levels may be explicitly encoded within the multilayered network representation depending upon the scope of the analysis of business function or threat model. More information about this approach may be found in (Weaver 2021).

Specific business functions enabled by the Level 0 workflow—such as defect inspection following lithography—may have different implementations and, as a result, different direct and indirect dependencies on IT/OT devices and Levels. Emerging deep learning algorithms for wafer metrology and defect detection depend on real-time sensor readings and in-process signals (Rambo and Sperling 2021). In contrast, manual inspection, while potentially less efficient as machine learning technologies improve, has fewer dependencies on real-time sensor data. Sensor data and other in-process signals can be stored in data historians such as that provided by the *Open Automation Software (OAS)* platform. In fact, OAS was selected by Intel as the standard to archive data at an archive rate of 10 million tags per second.

In the case of vulnerability analysis, recently-disclosed vulnerabilities associated with the OAS platform include the ability to obtain directory listings (CVE-2022-27169), obtain usernames and passwords (CVE-2022-26077), and craft packets that lead to a denial of service and loss of communication (CVE-2022-26026) (Vijayan 2022). While the PERA and associated reference models allow stakeholders to put such vulnerabilities in the context of upstream and downstream data flows at a high level, the intent behind a machine-actionable representation of such reference models is to model and simulate their potential impact on operational workflows. In addition, such an ICS representation may help bridge the gap between complex systems analysis and applied systems security (Weaver, Yardley, and Emmerich 2021).

5 USE CASE

The operational impact of a cyber-originating disruption on a manufacturing workflow depends on different factors that we want to explicitly model and explore. We consider the impact of varying security investment on production and energy consumption within a manufacturing fab accidentally infected by ransomware. We employ a *Nearly Orthogonal Latin Hypercube (NOLH)* experimental design with 33 design points (Sanchez 2011; Sanchez, Sanchez, and Wan 2020). For each design point, we simulate a model of a fab for 32 months with 6 months of warm-up based on the Intel Minifab model. These simulation results allow us to develop functional representations that consider the impact of the level of cybersecurity investment and product arrival rates on measures of production, specifically WIP, as well as energy consumption.

5.1 Experimental Design Factors

Security Investment: The first disruption factor considered is the security investment, which affects the *mean time to repair (MTTR)*. Within our experimental design, there are three factor levels for security investment that correspond to three different triangular distributions for MTTR, based on real-world ransomware incidents. When the level of security investment is *low*, the MTTR ranges from 3 to 21 days with a mode of 10, when there is a *medium* level of investment, the MTTR ranges from 0.5 to 14 days with a mode of 7, and with a high level of investment, the MTTR ranges from 0.5 to 5 with a mode of 3. This assumes that investment in security will affect an organization's ability to quickly respond to and recover from a cyber-originating disruption, which may not always be the case. Nonetheless, this approach allows companies to compare the impact of different response times and their overall effect on production and energy consumption.

Disruption Location: The location of machines indirectly affected by a cyber-originating event is also an important factor. Cybersecurity controls to mitigate the propagation of ransomware through a victim network include network segmentation and workload-specific micro-segmentation. The same ransomware may propagate at different rates following initial access to a victim network. In this experiment, we control the disruption location to machine 6 (lithography). We can, however, extend the model to allow security investments to affect the quantity and type of machine affected based on how different security controls impact malware propagation.

Disruption Start Time: Another disruption factor within our experimental design is the start time. The disruption may occur anytime immediately after the warmup period through the end of the 18-month simulation time. Our design has two factor levels for the disruption start time, both of which correspond to triangular distributions. The first level has the disruption start in the first half of the year following the warmup period while the second level occurs in the latter half of the year. Given that the model assumes a homogeneous arrival rate of product mixes, this model is not well suited to understand the impact of timing of targeted ransomware attacks on a fab.

Wafer Arrival Rates: The final factor within our experimental design is the wafer lot arrival rate. Higher wafer lot arrival rates may result in greater productivity. On the other hand, with less slack in the system, the impact of disruptions may be greater. Therefore, this experimental factor ranges between 70% and 95% maximum utilization across all machines in the facility. This maximum utilization is then translated into a wafer lot arrival rate for each type in the model.

5.2 Manufacturing Site WIP Response Curves

Response curves and simulation WIP outputs for the manufacturing site as a whole were fitted and plotted over 33 design points; see Figure 1. The x-axis is the maximum utilization (which is proportional to the arrival rates). Note that recovery time (not shown) is a function of the overall security investment. These results, based on 10 independent replications per design point, provide an initial sketch of the relationship between the security investment and maximum utilization (therefore arrival rates). Overall, the results demonstrate that longer recovery times associated with a lower security budget result in larger mean WIP values (averaged again over 10 replications), namely low, medium and high investments have average WIPs of 19.2, 12.4, and 11.7 respectively. To fit the WIP outputs, we assume the following functional form: $WIP(\rho) = c\left(\frac{\rho}{1-\rho}\right)^k$; ρ is the maximum utilization. This functional form is suggested by the general theoretical relationship between traffic intensity and average queue size in single-station, multiple-server models. Each security level's data points were fit to parameters c and k which are displayed in the legend of Figure 1. We fit our functional form to the data, using the `scipy` implementation of the Levenberg-Marquardt algorithm (Moré 1978) that solves unconstrained non-linear least square problems.

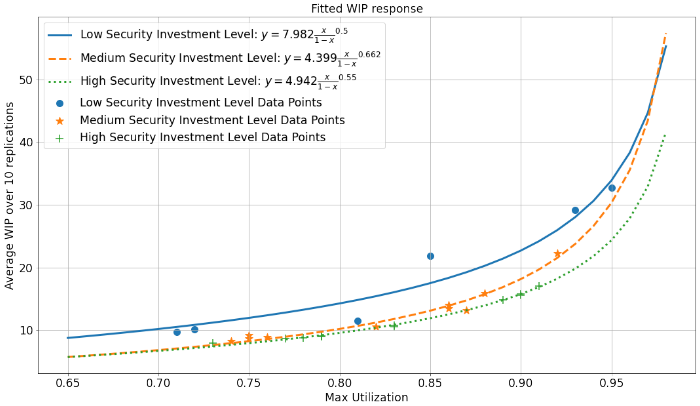

Figure 1: Average WIP responses to arrival rates (max utilization) over the 10 replications for each level of security investment. The legend states the fitted function to each security level's set of data points.

5.3 Tool-Level Energy Baselining

In addition to response plots for the fab as a whole, we also consider energy consumption of individual machines in the fab for which raw power consumption is available. Given a raw active power signal from a physical testbed, we estimate an energy baseline for a machine over the simulation time. Figure 2 shows data from the *Integrated Nanotechnology Research Facility (INRF)* at UCI: a raw power signal for the PlasmaTherm machine, sampled for 3 hours at a frequency of 1 Hz. The PlasmaTherm is used in processes immediately preceding and following lithography: *Plasma Enhanced Chemical Vapor Deposition (PECVD)* and *Reactive Ion Etching (RIE)*. Although a signal from a lithography machine is ideal, we consider these two processes, provided by the PlasmaTherm, as part of the function of the Minifab's lithography machine. In order to estimate power consumption over the simulation, the machine utilization u_m is used to interpolate between time spent in UP and IDLE states. Therefore, the raw power signal is partitioned into sets of time intervals corresponding to these states and power consumption values within each partition are averaged to estimate these values. Figure 3 shows that the energy used in the process increases as the maximum utilization rate increases, as expected. On the other hand, Figure 4 illustrates that recovery time does not show a clear relationship with energy usage or with the level of security investment. It may be possible for a medium level investment sample to have a longer recovery time than a high or low level investment.

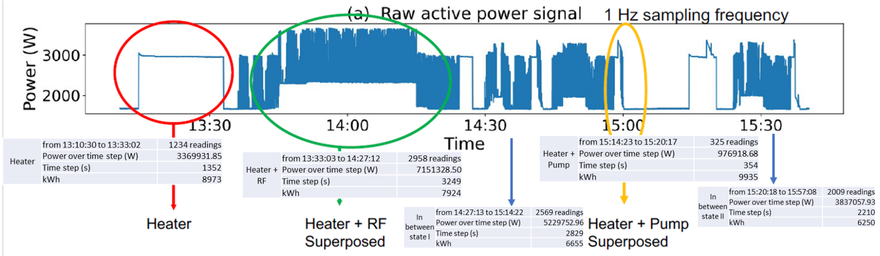

Heater	from 13:10:30 to 13:33:02	1234 readings
	Power over time step (W)	3369931.85
	Time step (s)	1352
	kWh	8973

Heater + RF	from 13:33:03 to 14:27:12	2958 readings
	Power over time step (W)	7151328.50
	Time step (s)	3249
	kWh	7924

In between state I	from 14:27:13 to 15:14:22	2569 readings
	Power over time step (W)	5229752.96
	Time step (s)	2829
	kWh	6655

Heater + Pump	from 15:14:23 to 15:20:17	325 readings
	Power over time step (W)	976918.68
	Time step (s)	354
	kWh	9935

In between state II	from 15:20:18 to 15:57:08	2009 readings
	Power over time step (W)	3837057.93
	Time step (s)	2210
	kWh	6250

Figure 2: Processing the raw power signal of a PlasmaTherm sensor to estimate energy consumption.

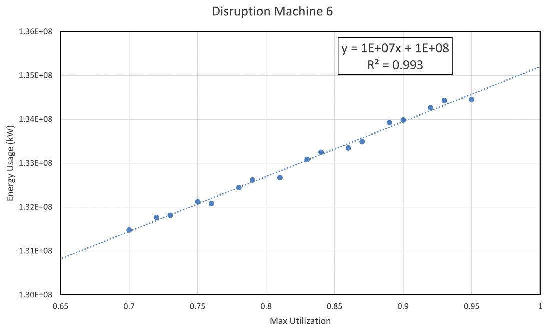

Figure 3: Energy usage and max utilization.

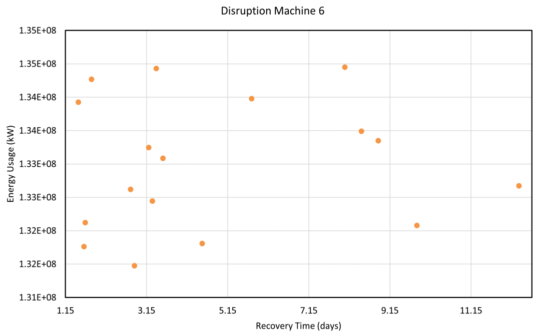

Figure 4: Energy usage and recovery time.

6 FUTURE WORK

Scalable approaches to support analysis of business impact and energy consumption relative to cybersecurity is important to consider, especially with larger workflow models. For example, the MIMAC and SEMATECH 300 reference models have 260/275 machines at 85/103 work stations respectively (Cayetano et al. 2018). Again, we note that modeling data-integrity attacks or confidentiality breaches are less clearly related to traditional stochastic disruption models and are left for future research. Finally, other approaches to analyze power consumption, such as performance measures employed by Hu (Hu and Chuah 2003), may provide further insight into the tradeoffs between cybersecurity and energy consumption.

7 CONCLUSION

Emerging technologies may help reduce the energy footprint of manufacturing systems such as wafer fabs but can introduce cybersecurity risks. Therefore, this paper presents a modeling and simulation framework to quantify tradeoffs between operational measures of performance, energy consumption, and cybersecurity controls. By augmenting the Intel Minifab model with interdependencies in the Communications/IT and Energy Sectors, we have demonstrated an approach, based on experimental design, to explore these interacting factors. Such work sets the stage for optimization at the facility and tool-level relative to cybersecurity and energy consumption.

ACKNOWLEDGMENTS

Research was sponsored by the *Cybersecurity Manufacturing Innovation Institute (CyManII)* and was accomplished under Grant Number DE-EE0009046. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

Alexander, O., M. Belisle, and J. Steele. 2020. "MITRE ATT&CK for Industrial Control Systems: Design and Philosophy". Technical Report MP01055863, The MITRE Corporation, Bedford, MA.

Assante, M. J., and R. M. Lee. 2015. "The Industrial Control System Cyber Kill Chain". Technical report, SANS Institute.

Barker, W. C. 2003. "Guideline for Identifying an Information System as a National Security System". Technical Report 800-59, *National Institute of Standards and Technology (NIST)*, Gaithersburg, Maryland.

Boccaletti, S., G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. 2014. "The structure and dynamics of multilayer networks". *Physics reports* 544(1):1–122.

Brennan, S. 2020. "MaxLinear targeted by Maze ransomware attack". *ITPro*. https://www.itpro.com/security/ransomware/356186/maxlinear-targeted-by-maze-ransomware-attack, accessed 6th May 2022.

Cayetano, T. A., A. Dogao, C. Guipoc, and T. Palaoag. 2018. "Cyber-physical IT assessment tool and vulnerability assessment for semiconductor companies". In *Proceedings of the 2nd International Conference on Cryptography, Security and Privacy*, 67–71. Wuhan, China: Association for Computing Machinery.

Dignan, L. 2018. "TSMC says variant of WannaCry virus brought down its plants". *ZDNet*. https://www.zdnet.com/article/tsmc-says-variant-of-wannacry-virus-brought-down-its-plants/, accessed 6th May 2022.

Dragos 2021. "ICS/OT Cybersecurity Year in Review 2021". Technical report, Dragos. https://www.dragos.com/year-in-review/, accessed 9th July 2022.

ENISA 2018. "Good Practices for Security of Internet of Things in the context of Smart Manufacturing". Technical report, *European Union Agency For Network and Information Security (ENISA)*. https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot, accessed 11th October 2022.

Fabrio, M., E. Gorski, N. Spiers, J. Diedrich, and D. Kuipers. 2016. "Recommended Practice: Improving Industrial Control System Cybersecurity with Defense-in-Depth Strategies". https://www.cisa.gov/uscert/sites/default/files/recommended_practices/NCCIC_ICS-CERT_Defense_in_Depth_2016_S508C.pdf, accessed 6th May 2022.

Fisher, E. L. 1986. "An AI-based methodology for factory design". *AI Magazine* 7(4):72–72.

Fowler, J. W., L. Mönch, and T. Ponsignon. 2015. "DISCRETE-EVENT SIMULATION FOR SEMICONDUCTOR WAFER FABRICATION FACILITIES: A TUTORIAL". *International Journal of Industrial Engineering* 22(5):661–682.

Gopalakrishnan, B., Y. Mardikar, and D. Korakakis. 2010. "Energy analysis in semiconductor manufacturing". *Energy Engineering* 107(2):6–40.

Haddod, F., and A. Dingli. 2021. "Intelligent Digital Twin System in the Semiconductors Manufacturing Industry". In *Artificial Intelligence in Industry 4.0*, edited by A. Dingli, F. Haddod, and C. Klüver, 99–113. Cham, Switzerland: Springer.

Hagberg, A., P. Swart, and D. S Chult. 2008. "Exploring network structure, dynamics, and function using NetworkX". Technical Report LA-UR-08-05495, Los Alamos National Lab, Los Alamos, New Mexico. https://www.osti.gov/biblio/960616, accessed 11th October 2022.

Hu, S.-C., and Y. Chuah. 2003. "Power consumption of semiconductor fabs in Taiwan". *Energy* 28(8):895–907.

Jain, S., D. Lechevalier, J. Woo, and S.-J. Shin. 2015. "Towards a virtual factory prototype". In *Proceedings of the 2015 Winter Simulation Conference (WSC)*, edited by L. Yilmaz, W. K. V. Chan, I.-C. Moon, T. Roeder, C. M. Macal, and M. D. Rossetti, 2207–2218. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Kannaian, T. K. 2018. “Capacity model with sustainability scope to predict the semiconductor manufacturing’s energy consumption and carbon dioxide emissions”. Technical report, Texas State University, San Marcos, Texas. https://digital.library.txstate.edu/handle/10877/7888, accessed 11th October 2022.

Kempf, K. 1994. “Intel five-machine six step mini-fab description”. Technical report, Arizona State University, Tempe, Arizona.

Kopp, D., M. Hassoun, A. Kalir, and L. Mönch. 2020. “SMT2020—A semiconductor manufacturing testbed”. *IEEE Transactions on Semiconductor Manufacturing* 33(4):522–531.

Langill, J., B. Hegrat, and D. Peterson. 2019. “Is The Purdue Model Dead?”. In *Proceedings of the 2019 S4 Conference*. Digital Bond. https://www.youtube.com/watch?v=KfxPF9xjFrE, accessed 9th July.

Leng, J., D. Wang, W. Shen, X. Li, Q. Liu, and X. Chen. 2021. “Digital twins-based smart manufacturing system design in Industry 4.0: A review”. *Journal of Manufacturing Systems* 60:119–137.

McGinley, M. 2021. “Parting Ways with Purdue? The Effect of Industry 4.0 on ICS Security Architectures”. In *Security BSides Athens*. Security BSides. https://youtu.be/W1df3NWOUjo, accessed 9 July 2022.

Moré, J. J. 1978. “The Levenberg-Marquardt algorithm: Implementation and theory”. *Numerical Analysis* 630:105–116.

Orbach, M. 2020. “Israeli chipmaker Tower confirms cyberattack forced it to shut down systems”. *CTech*. https://www.calcalistech.com/ctech/articles/0,7340,L-3848290,00.html, accessed 11th October 2022.

Osborne, C. 2020. “Black Hat: Hackers are using skeleton keys to target chip vendors”. *ZDNet*. https://www.zdnet.com/article/black-hat-hackers-are-now-using-cobalt-strike-and-skeleton-keys-to-target-semiconductor-firms/, accessed 6th May 2022.

Palmer, G. I., V. A. Knight, P. R. Harper, and A. L. Hawa. 2019. “Ciw: An open-source discrete event simulation library”. *Journal of Simulation* 13(1):68–82.

Pokhrel, A., V. Katta, and R. Colomo-Palacios. 2020. “Digital twin for cybersecurity incident prediction: A multivocal literature review”. In *Proceedings of the 42nd International Conference on Software Engineering Workshops*, 671–678. Institute of Electrical and Electronics Engineers, Inc. and Association for Computing Machinery.

Rambo, S., and E. Sperling. 2021. “AI In Inspection, Metrology, and Test”. *Semiconductor Engineering*. https://semiengineering.com/ai-in-inspection-metrology-and-test/, accessed 9th July, 2022.

Rose, O. 2004. “Modeling tool failures in semiconductor fab simulation”. In *Proceedings of the 2004 Winter Simulation Conference (WSC)*, edited by J. Smith, B. Peters, R. G. Ingalls, and M. D. Rossetti, 1910–1914. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Sanchez, Susan M. 2011. “NOLHdesigns spreadsheet”. http://harvest.nps.edu/, accessed 11th October 2022.

Sanchez, S. M., P. J. Sanchez, and H. Wan. 2020. “Work smarter, not harder: A tutorial on designing and conducting simulation experiments”. In *Proceedings of the 2020 Winter Simulation Conference (WSC)*, edited by R. Thiesing, T. Roeder, K.-H. G. Bae, S. Lazarova-Molnar, Z. Zheng, B. Feng, and S. Kim, 1128–1142. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Schneider, F. B. 2011. “Blueprint for a Science of Cybersecurity”. Technical report, Department of Computer Science, Cornell University, Ithaca, New York. https://ecommons.cornell.edu/bitstream/handle/1813/22943/SoS%20blueprint.pdf, accessed 11 October 2022.

Scholl, W. 2008. “Coping with typical unpredictable incidents in a logic fab”. In *Proceedings of the 2008 Winter Simulation Conference (WSC)*, edited by T. Jefferson, J. Fowler, S. Mason, R. Hill, L. Moench, and O. Rose, 2030–2034. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Scholl, W., M. Mosinski, B. P. Gan, P. Lendermann, P. Preuss, and D. Noack. 2012. “A multi-stage discrete event simulation approach for scheduling of maintenance activities in a semiconductor manufacturing line”. In *Proceedings of the 2012 Winter Simulation Conference (WSC)*, edited by O. Rose, A. Uhrmacher, M. Rabe, C. Laroque, R. Rasupathy, and J. Himmelspach, 1–10. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

SEMI 2022. “Specification for Cybersecurity of Fab Equipment”. Technical Report SEMI E187, *Semiconductor Equipment and Materials International (SEMI)*. https://store-us.semi.org/products/e18700-semi-e187-specification-for-cybersecurity-of-fab-equipment, accessed 6th May 2022.

Shao, G., S. Jain, C. Laroque, L. H. Lee, P. Lendermann, and O. Rose. 2019. “Digital twin for smart manufacturing: The simulation aspect”. In *Proceedings of the 2019 Winter Simulation Conference (WSC)*, edited by Y.-J. Son, P. Haas, N. Mustafee, M. Rabe, H.-K. G. Bae, C. Szabo, and S. Lazarova-Molnar, 2085–2098. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Singh, I., and M. Szczys. 2018. “Malicious Component Found on Server Motherboards Supplied to Numerous Companies”. *Hackaday*. https://hackaday.com/2018/10/04/malicious-component-found-on-server-motherboards-supplied-to-numerous-companies/, accessed 9th July 2022.

Stegall, A. 2020. “Ransomware attack halts X-FAB production in Lubbock, worldwide”. *KCBD News*. https://www.kcbd.com/2020/07/15/ransomware-attack-halts-x-fab-production-lubbock-worldwide/, accessed 6th May 2022.

Stouffer, K., M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, and S. Lightman. 2022. “Guide to Operational Technology (OT) Security”. Technical Report SP 800-82r3, *National Institute of Standards and Technology (NIST)*, Gaithersburg, Maryland. https://csrc.nist.gov/News/2022/guide-to-operational-technology-ot-security, accessed 6th May 2022.

Stouffer, K., T. Zimmerman, C. Tang, J. Cichonski, N. Shah, and W. Downard. 2019. "Cybersecurity Framework Manufacturing Profile Low Security Level Example Implementations Guide: Volume 1–General Implementation Guidance". Technical Report NISTIR 8183A, *National Institute of Standards and Technology (NIST)*, Gaithersburg, Maryland.

Technavio 2020. "Research Report: Semiconductor Wafer Inspection Equipment Market (2020-2024) — Growing Demand for IoT Devices to boost the Market Growth". *BusinessWire*. https://www.businesswire.com/news/home/20200901005137/en/, accessed 9th July.

Ting-Fang, C., and L. Li. 2018. "TSMC malware attack hit 'most advanced' iPhone chip site". *Nikkei Asia*. https://asia.nikkei.com/Asia300/TSMC-malware-attack-hit-most-advanced-iPhone-chip-site2, accessed 6th May 2022.

Van Der Aalst, W. M., and A. H. Ter Hofstede. 2005. "YAWL: Yet Another Workflow Language". *Information Systems* 30(4):245–275.

Vijayan, J. 2022. "Critical OAS Bugs Open Industrial Systems to Takeover". *Dark Reading*. https://www.darkreading.com/application-security/critical-oas-bugs-industrial-takeover, accessed 9th July 2022.

Wang, H., S. Chen, M. S. U. I. Sami, F. Rahman, and M. Tehranipoor. 2021. "Digital Twin with a Perspective from Manufacturing Industry". In *Emerging Topics in Hardware Security*, edited by M. Tehranipoor, 27–59. Cham, Switzerland: Springer.

Weaver, G. A. 2021. "Scientific Data Management for Interconnected Critical Infrastructures". In *Proceedings of the 2020 Joint Conference for Digital Libraries (JCDL 2021)*, 192–201: Institute of Electrical and Electronics Engineers, Inc.

Weaver, G. A., T. Yardley, and D. P. Emmerich. 2021. "Continuous Infrastructure Assessment for Key Business Functions in Changing Environments". In *Proceedings of Resilience Week 2021*, 1–8: Institute of Electrical and Electronics Engineers, Inc.

Williams, T. J. 1994. "The Purdue enterprise reference architecture". *Computers in industry* 24(2-3):141–158.

AUTHOR BIOGRAPHIES

GABRIEL A. WEAVER is a Senior Critical Infrastructure Analyst at Idaho National Laboratory. Weaver holds a Ph.D in Computer Science from Dartmouth College. His research interests include modeling and simulation of critical infrastructure systems and applied systems security. His email address is gabriel.weaver@cymanii.org.

JOHN J. HASENBEIN is a professor at The University of Texas at Austin in the graduate program in Operations Research and Industrial Engineering within the Department of Mechanical Engineering. His research interests include scheduling and analysis of semiconductor wafer fabs, stochastic optimization, and resilience models for natural disasters. His email address is has@me.utexas.edu.

ERHAN KUTANOGLU is an associate professor of Operations Research and Industrial Engineering within the Department of Mechanical Engineering at the University of Texas at Austin. Dr. Kutanoglu specializes in applied operations research regarding manufacturing and service logistics and supply chain management. His email address is erhank@austin.utexas.edu.

GONZALO MARTINEZ-MEDINA is a Ph.D student at the University of Texas at San Antonio, in the Texas Sustainable Energy Research Institute working with Dr. Krystel Castillo-Villar. His research focuses on energy baselining and the design of a optimization model to reduce possible cyber-attacks in the manufacturing industry. His email address is gonzalo.martinezmedina@cymanii.org.

JACOB WILLIAM SHUSKO is a graduate student at The University of Texas at Austin in the Operations Research and Industrial Engineering program within the Department of Mechanical Engineering working with Dr. John J. Hasenbein. His research interests include statistical modeling, machine learning and simulation methods for IoT systems. His email address is jacob.shusko@utexas.edu.

KRYSTEL K. CASTILLO-VILLAR is the Vice President of Energy Efficiency at the Cybersecurity Manufacturing Innovation Institute (CyManII), director of the Texas Sustainable Energy Research Institute and Professor at The University of Texas at San Antonio. Her research interests are mathematical modeling of supply chains; logistics; optimization of large-scale scenarios; and statistical quality control and reliability. Her email address is krystel.castillo@cymanii.org.

PAULO CESAR COSTA is Vice President for Securing Manufacturing Automation at CyManII and Supply Chain Security Professor at George Mason University. His research focuses on probabilistic reasoning and ontologies, with applications in cyber and transportation security, data fusion, and decision support. His email address is paulo.costa@cymanii.org.