Advanced Reactor Safeguards & Security

Artificial Intelligence for Advanced Reactor Cybersecurity Protection

Small Modular Reactor – Threat Hunting Representations for Embedded Anomaly Tracking (SMR-THREAT)

Final Report

Glenn A. Fink, PhD, CISSP
Michael R. Moore, PhD
J. Michael Vann, PhD
September 2025

SRNL-TR-2025-00651, Revision 1

DISCLAIMER

Printed in the United States of America

Prepared for
U.S. Department of Energy

SRNL-STI-2025-00651
Revision 1

Keywords: Cybersecurity, Nuclear Materials Downblending, machine learning, safety

Retention: *Varies (Track number is applicable)*

Small Modular Reactor – Threat Hunting Representations for Embedded Anomaly Tracking (SMR-THREAT) FY25 Final Report

Glenn A. Fink, PhD, CISSP
Michael R. Moore, PhD
J. Michael Vann, PhD

September 2025

Savannah River National Laboratory is operated by Battelle Savannah River Alliance for the U.S. Department of Energy under Contract No. 89303321CEM000080.

PREFACE OR ACKNOWLEDGEMENTS

This report summarizes FY25 accomplishments on the Small Modular Reactor Threat Hunting Representations for Embedded Systems Anomaly Tracking (SMR-THREAT) project, which was funded April 2024 through September 2025. SMR-THREAT is an application of machine learning techniques to cyber data found in SMRs related to safety, security, and safeguards (3S) conditions. SMR-THREAT investigates ways that cybersecurity data found in network flows can be used to identify 3S threats as early as possible, especially in unattended, remote reactors. The report documents the test system used as a surrogate for SMRs, the data collected from it, and the initial analysis showing the predictive and diagnostic value of cyber data for 3S applications. The report concludes with recommendations for future work including making better predictions from the data and processes documented herein.

The author would like to thank the staff of the Melter project who willingly answered questions, collected data, and even modified their data collection system to accommodate the needs of this project. Persons who were especially helpful included Zach Sechrist, the project manager, Michael Lowrey, the computer network engineer, and Justin Dobey, the programmable logic controller (PLC) design authority. Brian Gerkin, our cyber engineer, conducted the cyber data collection and advised on its use. Ian Webb, power engineer, designed the out-of-band sensor system and collected that data. This project would not have been possible without their help, and the authors look forward to working together on follow-on work coming from this initial effort.

EXECUTIVE SUMMARY

The SMR-THREAT project is intended to provide threat hunters with information about safety and security conditions of concern arising especially in remote and unattended small modular reactors (SMRs) by examining cybersecurity data. At this writing there are no SMRs actively deployed in the US power grid. Thus, we use a surrogate, the Melter system, which, although it is not an SMR, is a mobile nuclear system that shares many safety and security characteristics with SMRs and is thus a suitable surrogate for analysis. The Melter is installed at remote sites and used to down-blend nuclear materials of concern to safe radioactivity levels. We discuss this process, the robotic and remote-controlled systems involved, and the computer network that controls them. We also describe the data that was collected from the Melter and our initial analysis of it.

Data of interest include:

1. *Safety* data such as sensor readings (e.g., furnace temperatures and coolant flow rates), control logic signals (e.g., set points and valve positions), and performance metrics (e.g., efficiency measures).
2. *Security* data such as cyber communication data including network data packet captures, computer and system event logs, door, lock, and occupancy information, and intrusion detection alerts.

This report does not discuss safeguards. The key thesis of this work is that most of the safety and security information is present in and may be extracted from the computer network communications (cyber) data generated during system operation. The key finding of this report is that cyber data contains information that may correlate to and predict safety and security events although the correlation requires some *a priori* knowledge of the system at this point. We believe this finding will also apply to SMRs, assuming an effective digital communication architecture, an appropriate collection approach, and the ability to extract application data from the network communications data. We arrive at this finding by examining the cyber data from the Melter, showing how the roles and activities of devices within the network may be discovered, how devices are clearly separable into their types and functions, and how communication patterns in network traffic predict or correlate with out-of-band or operator-observed events.

This project seeks to process all available data as text using language modeling machine learning techniques. These models can ingest any digital data for processing. By harvesting the control-systems data from the cyber data we may obtain a consistent view of the overall system with all events sharing a common timing regimen. This gives a total ordering of all events and enables discussion of causality. However, this report covers the collection and initial analysis of cyber and physical data using more conventional methods. Not all information desired is available as part of the cyber data. Some of the non-cyber data is available in digital form and some must be digitized from other formats such as images from video surveillance and separately recorded operator logs. The operator logs give a timestamped, view of the operation of the system from the operator's point of view. Similarly, to make use of handwritten data, spoken information, or data that is wholly kept private to some part of the system, it must be collected and processed separately and may, with effort be integrated with the cyber data. Future directions include finding cost effective ways to make more of this data collectable in the cyber data stream and a discussion of simulating off-normal events in a safety-critical system like an SMR.

TABLE OF CONTENTS

1.0 Introduction 1
1.1 Test System Description 2
2.0 Data Description 5
2.1 Network Data Description 5
2.2 Out-of-band Data Description 7
2.3 Out-of-Band (External) Sensing 10
2.3.1 Measurement Setup 11
2.3.2 Data Collection Details 13
2.3.3 Operational Parameters 14
2.4 Data Limitations 17
3.0 Algorithm Development 18
3.1 In-band OT Algorithm Development 19
3.1.1 Packet Bitmaps 19
3.1.2 Packet Interarrival times and Jitter 22
3.2 Out-of-Band Sensing Algorithm Development 26
3.3 State Based AI – Enabling the Integration of In-Band and Out-of-Band Techniques 27
3.3.1 Traffic mapping using Association Rules Mining 28
3.3.2 Alternative Packet Pattern Mapping Techniques 30
3.4 Identifying sensor measurement trends and inferring process states 30
4.0 Conclusions 30
5.0 Future Work 30
6.0 References 32

LIST OF TABLES

Table 1: Operator Log Excerpt 9
Table 2 – 23 June 2025 Operational Logs for a test event 14
Table 3 - 23 June 2025 Operational Logs for Test Event 15
Table 4 - Stages of A Priori Knowledge Versus Techniques for In-Band OT Analysis 18
Table 5. Acceptable tolerances for IPAT and Jitter 22
Table 6. Packet timing statistics 23
Table 7. Connection identifier assignments 28
Table 8 – Subset of Results from Association Rules Mining 29

LIST OF FIGURES

Figure 1. Melter Facility Layout 3
Figure 2. Melter Facility Process Modules 4
Figure 3. Simplified diagram of the data collection points 5
Figure 4. Melter Network block diagram 6
Figure 5. Detected network flow connectivity 7
Figure 6. Three-phase current, temperature, and vibration during the 6/23/2025 melt operation. 8
Figure 7. Thermocouple data (deg C). 8
Figure 8. Close-up of events between about 13:35 and 14:15 on 6/23/2025 9
Figure 9. Input power sensor placement 11
Figure 10. Current Probe Details 12
Figure 11. Vibration and Temperature Probe Details 12
Figure 12. Temperature Sensing Placement 13
Figure 13. Data Acquisition System 13
Figure 14 – Overview of Cyber Mitigation Capabilities Versus Availability of Documentation on the Process Control System 19
Figure 15. Bitmap representation of all packets colored by connection 20
Figure 16. Bitmap representation of various types of motors in the Melter → PLC 21
Figure 17. Bitmap representation of device ←→ PLC communication 21

Figure 18. IPAT patterns distinguish classes of devices. 23

Figure 19. High-Jitter apparent correlation to operator events. 24

Figure 20. HEPA Jitter falls after furnace restart events. 24

Figure 21: Furnace HMI and PLC communication does not correlate to furnace restart events. 25

Figure 22- An example of a Code layer scatter plot using only reconstruction loss for training. NOTE: The cluster colors and labeling were manually assigned by a human analyst. 26

Figure 23 –Analysis of a blind dataset from a test facility is shown as a scatter plot of human labeled (i.e., "classes") cluster points versus time. 27

Figure 24 – Example of generation of 2D images from simple packet labeling schemes that lend themselves to multiple levels of AI/ML 30

LIST OF ABBREVIATIONS

3S	Safety, Security, and Safeguards
BERT	Bidirectional Encoding Representations from Transformers
DASP	Dimensionally Aligned Signal Projection
GPU	Graphics Processing Unit
HC	Host Country
HEPA	High-Efficiency Particulate Air
HMI	Human-Machine Interface
HMM	Hidden Markov Model
ICMP	Internet Control Message Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
LHM	Lower HEPA Module
LLM	Large Language Model
LMC	Lower Melt Cell
LTM	Lower Transfer Module
Melter	Mobile Melt-Consolidate
NMOC	Nuclear Materials of Concern
PLC	Programmable Logic Controller
RTU	Remote Terminal Unit
SMR	Small Modular Reactor
SMR-THREAT	Small Modular Reactor Threat Hunting Representations for Embedded Anomaly Tracking
SRNL	Savannah River National Laboratory
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
UDP	User Datagram Protocol
UMC	Upper Melt Cell
UPS	Uninterruptible Power Supply
UTM	Upper Transfer Module
WUNM	Weapons-Usable Nuclear Material

1.0 Introduction

The increasing reliance on digital systems and remote operations for future Small Modular Reactors (SMRs) presents a paradigm shift in nuclear plant monitoring and control especially in deployment situations where they could be unattended or remotely operated. This heightened digitalization, while offering significant operational efficiencies, exponentially expands the cyber-physical attack surface, introducing complex new risks to the essential pillars of safety, security, and safeguards (3S). While the current nuclear fleet relies on network isolation and significant on-site staffing, SMRs produce less power and require a new, less costly approach to mitigating 3S risks—one where cybersecurity is not merely a defensive layer but an integrated, predictive sensor for 3S integrity.

This research explores the ambitious aspiration of engineering SMRs to make the widest possible range of 3S events, both unintentional (safety) and intentional (security and safeguards), detectable and predictable through the OT system's cyber data. Leveraging an expanded digital architecture—including advanced sensors, cameras, and access controls—we aim to transform the OT network from a system of control into a pervasive monitoring and early warning system.

However, this aspiration must contend with present-day realities and challenges. Traditional 3S events, such as physical sabotage of analog components or human error with no digital footprint, will remain difficult for cyber data alone to detect without advanced and redundant cyber-aware monitoring. Furthermore, insider threats that bypass or corrupt digital systems, and sophisticated material diversions designed to evade digital safeguards, persist as critical risks in current systems. This research explores the which 3S events are detectable and which are not in the cyber data of a surrogate SMR system (the Melter) to help identify the necessary design constraints and features that will be needed to maximize cyber-based monitoring, ultimately providing a blueprint for making unattended and remote A/SMRs demonstrably safer, more secure, and proliferation-resistant than current systems that rely purely on physical measures for 3S protection.

The SMR-THREAT project will create new cyber analysis tools to assist cybersecurity personnel in their hunt for malicious activity of fused cyber and operational data found in SMR control systems. The resulting tools will ingest data from system communication logs, process reports, and performance metrics.

A/SMRs are cyber-physical systems that use programmable logic controllers (PLCs) and remote terminal units (RTUs) to control critical processes that require exact controls. These systems generate extensive data flows from digital instrumentation and control systems to support remote monitoring, automation, and grid integration. Cybersecurity data from components and systems covering safety and security must be correlated to obtain a holistic view of reactor state. But correlating multimodal data from physical processes, video, cyber sensors, and grid transactions is challenging because of the variety of formats, volumes, and velocities of data. SMR-THREAT employs statistical modeling and neural networks to fuse timestamped cyber flows and physical transactions from various components by aligning timescales and entity behaviors. This will characterize normal relationships between digital and physical systems. Interactive visual interfaces of the changes in normal state over time will then help threat hunters identify anomalies, signaling potential issues. By integrating cyber and physical data, SMR-THREAT will enhance monitoring and protection to provide a unified perspective on the health and operational security of the plant in a way not possible with current monitoring tools.

The unified data foundation and automated learning of expected behaviors will provide a technical basis for remote and unattended operation proposed for some SMRs. Intuitive visualizations can be created from this fused data to reduce the difficulty plant staff may have in continuously monitoring safety and

security status. Cost savings will result from partial automation of behavioral modeling for cyber compliance activities. SMR-THREAT will create advanced analytics to strengthen cyber defense for modern nuclear plants.

Safety data streams include (1) Time-series sensor measurements (temperature, pressure, flow rates) from reactor protection systems and plant control instruments monitoring physical processes, (2) Control logic signals and actuator outputs from distributed control systems: setpoints, valve positions, pump speeds, etc., and (3) Plant performance metrics (heat/power generation rates, efficiency metrics) indicating overall plant production health.

Security data streams include (1) Packet captures from TCP/IP traffic between digital components of the plant control network and connections to external systems (ERP, grid operators), (2) Syslog and event log data from computers, servers, firewalls, and access control systems, (3) Authentication traffic and user account changes, and (4) Intrusion detection/prevention system event records and alerts.

By ingesting and correlating multimodal data into the THREAT analytics engine, the toolset will create an integrated perspective on SMR platform activities not feasible with current analytic techniques. This fusion of divergent data sources is necessary to provide comprehensive monitoring and threat evaluation capabilities for advanced nuclear plants.

1.1 Test System Description

Because there are no operational SMRs in the US, we utilize a surrogate system: the Melter, developed by SRNL and its partners. This project will demonstrate concept feasibility using cyber security and physical process logs from the prototype Melter being developed by SRNL and its partners as a proxy for data from an actual SMR. The Melter melts and down-blends excess weapons-usable nuclear material (WUNM) to yield a more proliferation-resistant and reduced attractiveness metallic or cermet ingot. The Melter mirrors key components and functionalities of SMRs and exhibits similar safety, security and safeguards (3S) dimensions as an SMR, providing a valuable testing ground for our approach. We have collected an initial set of data from this system including security data from cyber traffic among communicating programmable logic controllers (PLCs), and safety and performance data from the many digital and analog sensors.

The demonstration setup emulates key components and missions common to SMRs without the complexity, expense, and risk of implementing the data collection system in an experimental SMR. The demonstration environment generates realistic cyber and physical data flows on a small scale to provide an integrated 3S perspective for modeling normal behaviors. This will result in a low-risk proof of concept demonstrating viability for follow-on research. Typical melt operation last approximately 18 hours. With cyber data accumulating at approximately 4,500 packets per second we expect about two gigabytes of cyber data per hour. Data gathered from out-of-band sensors is in addition to this and the amount collected depends on the type and frequency of collection.

Real data from our experiments will allow operators to compare their physical activities and observations (as noted in the operator log file) to the analysis results, enriching both the analytic methods and the operators' ability to observe the system. Later this analytic pipeline can be fed by online processes to create near-real-time analysis of the data. The combination of real data from actual events, sensor noise, operator inputs, etc. will provide the diversity and fidelity needed to demonstrate the THREAT analytics on anomalous 3S patterns that are viable for analytics of SMR systems under test and in operation. The remainder of the introduction will briefly describe the unique Melter facility, its layout and function.

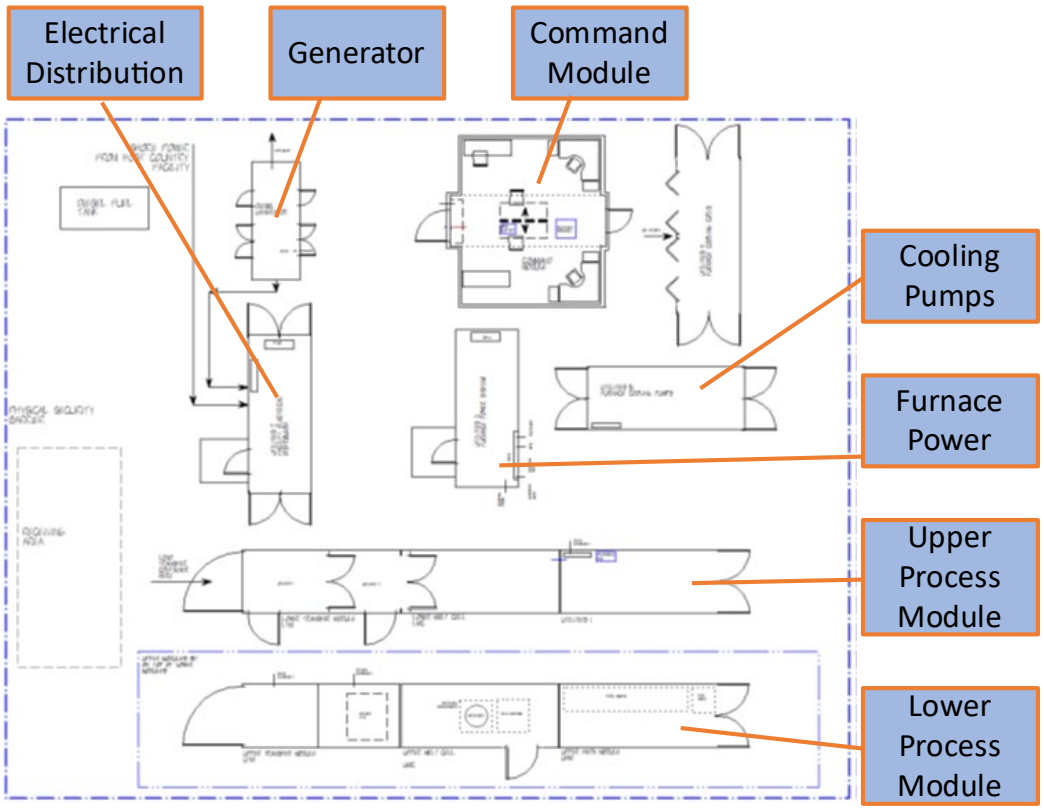

Figure 1. Melter Facility Layout

The Melter ships in a set of shipping containers, each containing machinery and controls for a part of the process (see Figure 1). The Command Module contains the operators' area with the computers and PLC controls for the cameras, actuators, and sensors. The Furnace subsystem is in several containers, including electrical power, cooling pumps, and the furnace proper. The Melter is designed to connect to shore power, but also has generators for backup power. There is a power distribution container that connects generation to each container. The Upper and Lower Process Modules (see Figure 2) have mechanisms to transport the NMOC to and from the furnace. The furnace is where the core function of the Melter is performed.

The Melter system is shipped in six ISO sea/land containers (stacked two high) or modules that house equipment, systems, and personnel. A water tanker or host facility water line provides emergency cooling water to the closed loop furnace cooling system, along with a primary line power feed and backup diesel generator and accompanying diesel fuel tanker in case power to the coolant pump fails. Additionally, a command module interfaces with the utility and treatment modules and allows for full furnace control. The Melter treatment modules house a furnace shielding and containment box, furnace, trolley, dolly, crane hoist, HEPA filter system, off-gas system, and regulated exhaust system. The adjacent utility modules house the utility systems including the furnace control panel, UPS, electrical distribution panel, etc.

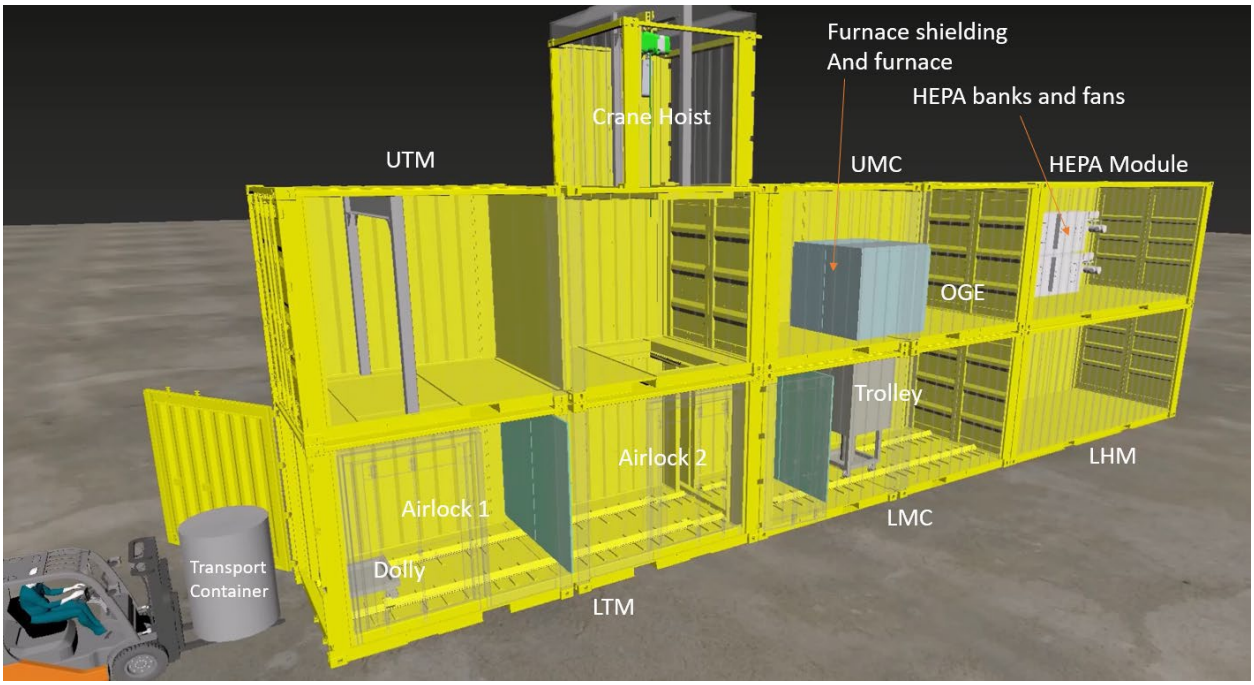

Figure 2. Melter Facility Process Modules

The materials to be melted are placed into a three-layer crucible surrounded by a metallic transport container which is then moved to the exterior door to airlock 1 of the Lower Transfer Module (LTM) of the Process Modules via fork truck. The Melter dolly is positioned inside of airlock 1 and moves the transport container under the crane hoist in the Upper Transfer Module (UTM) in airlock 2. The crane removes the crucible from the transport and the dolly moves the emptied transport container back to the first airlock. The trolley then moves under the raised target materials and the crane lowers the crucible onto it. Next the trolley moves the crucible to the Lower Melt Cell (LMC) and raises the crucible into the furnace in the Upper Melt Cell (UMC).

During the melt process, the expected temperature in the crucible (~1,600°C) must be monitored to remain well below the boiling point for uranium (4,131°C) or plutonium (3,230°C) so there will not be significant fissile material accumulation on the HEPA filters. Upon completion of the treatment process and ingot cooling, the lift is unlocked, the crucible with its final product ingot will be lowered back into the trolley, and the loading process reversed.

All positions of the dolly and trolley, crucible, and transport container are monitored by video cameras and moved via robotic controllers. Sensors along the track record their location and report it to the operator. The “airlocks” are not airtight, but they operate as barriers that funnel the air flowing between the chambers with sensors to detect improper airflow that might produce contamination. The HEPA filters in the filtering module ensure that any radioactive contaminants are scrubbed from the air. The crane hoist, door, trolley, and dolly positions, temperatures, air pressures, furnace operations, and more are all reported to the operators in the command module and controlled by PLCs and transferred over the TCP/IP (computer network protocol) network where they can be copied and collected at the management port of the main switch. Thus, safety data can in large part be subsumed by the same collection method as the security data represented by the cyber data.

Data that does not appear in the cyber data includes the safeguards data and some operator-entered safety or operations status data. The latter data is stored in a FileMaker Pro database on the Monitoring Laptop

and must be collected separately because it does not cross the network. A few other data sources are not part of the information system because they are not controlled centrally. For example, the backup generators may have electronic components, but they are operated manually and do not transmit status or control information. Another example includes the PLC used to control the operation of the transfer container. This PLC is not operated on the same network where the monitoring laptop is. Finally, the electronics used to automate the release of cooling water for the furnace if the facility loses power are not yet represented in the collected cyber data. Doubtless this will be the case with some subsystems in an SMR as well. Collection of data from these sources is not covered in this project and will need to be addressed separately. While the Melter is arguably not as complex as a true SMR, it certainly has enough complexity to be a realistic stand-in for 3S synthesis and cyber threat detection.

2.0 Data Description

The operational goal of this project is to develop algorithms and/or methods that can be used to protect remotely operated nuclear processes from cyber-attacks. Our thesis is that most of the safety and security data can be obtained or inferred by collecting the TCP/IP data from a central point. Barring encryption, the status information for all the devices can be recovered from the TCP/IP data. We refer to the TCP/IP communications data as cyber data or in-band data. In-band includes OT traffic between the PLC and various sensors and actuators. The other main type of data we collected we call "out-of-band" data, which is collected from independent sensors strategically placed to measure physical phenomena that can be attributed to the physical processes of the Melter (or similar system). Another important out-of-band data set we collected is the operator log data, which record system actions and observations from a human perspective and timescale. These data are collected completely independently of the in-band data, and events found in them may be used as ground truth to determine whether the in-band cyber data reflects the same events. Figure 3 shows a diagram of where, what kind, and the capture rate of all the kinds of data we collected.

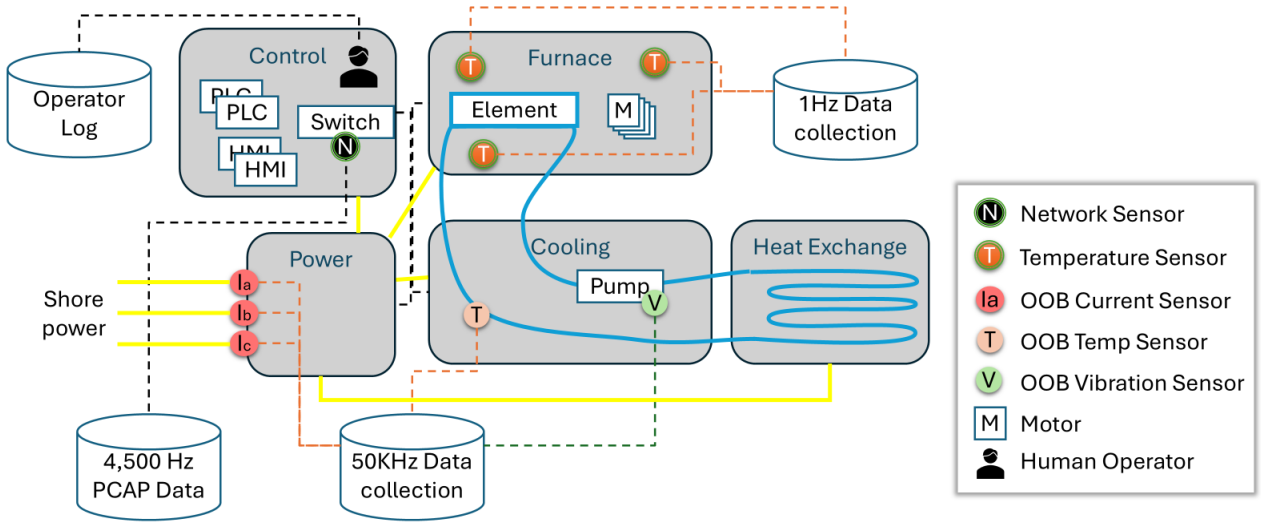

Figure 3. Simplified diagram of the data collection points

2.1 Network Data Description

The Melter network is comprised of two main networks, the management network (192.168.0.0/24) and the operations network (192.168.1.0/24). See the network block diagram (Figure 4) for the allocation of components to these networks.

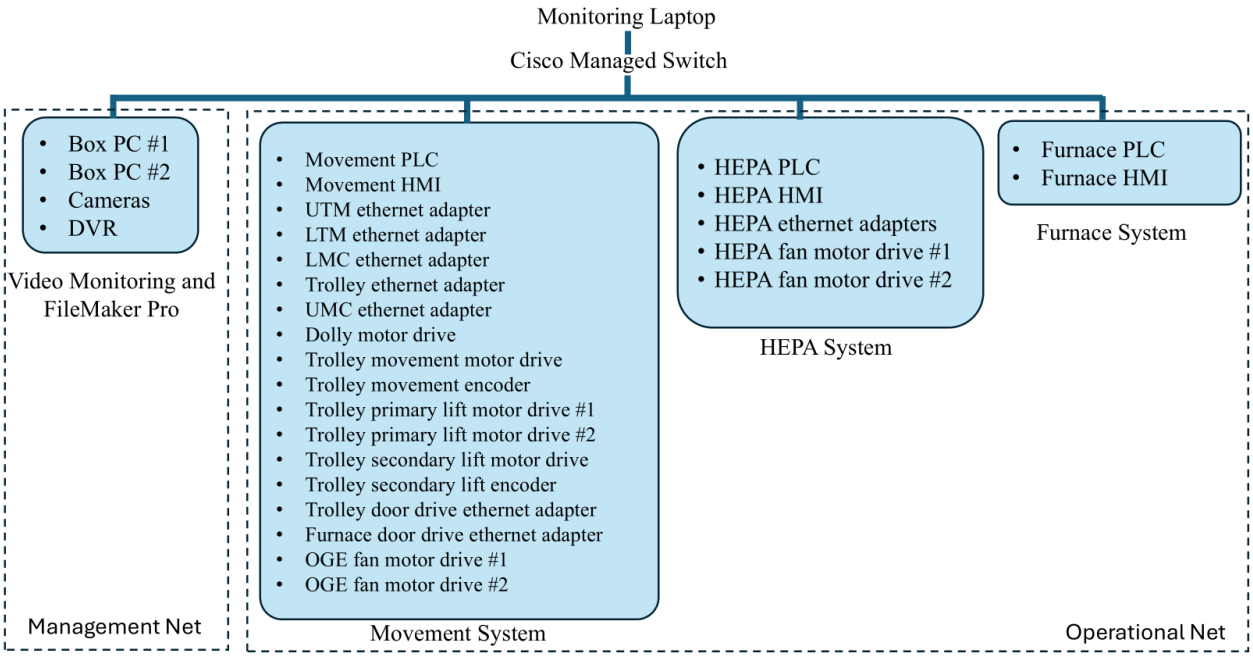

Figure 4. Melter Network block diagram

While the Melter is not a large system, it has complexity comparable to some of the smallest known SMRs and many of the components have 3S dimensions. There are four computers (two laptops and two workstations), ten cameras (and plans for up to 24), a digital video recorder, networking equipment, two human-machine interface (HMI) machines and their corresponding PLCs, and 22 device Ethernet adaptors (essentially Remote Terminal Units (RTUs) or field busses). The Melter is expected to have multiple generations as versions of it are planned for deployment to many locations.

Figure 5 shows the network connectivity of the Melter during the Bump Test on 6/23/2025 during the period from 1350 to 1420 (early startup operations). In this demonstration, the decision was to exclude the camera flows from the management network, so the management network does not appear. All three PLCs from the network block diagram appear and another network component contains the management PC, which communicates with all the PLCs using multicast/broadcast. The largest component contains both the movement and HEPA PLCs which communicate with each other but whose subordinates do not communicate with the other PLC's. The greatest amount of traffic (by an order of magnitude) is between the movement PLC and the trolley lift motors (primary #1 and #2 and secondary), which in turn are communicating over multicast, forming the only cycles in the diagram.

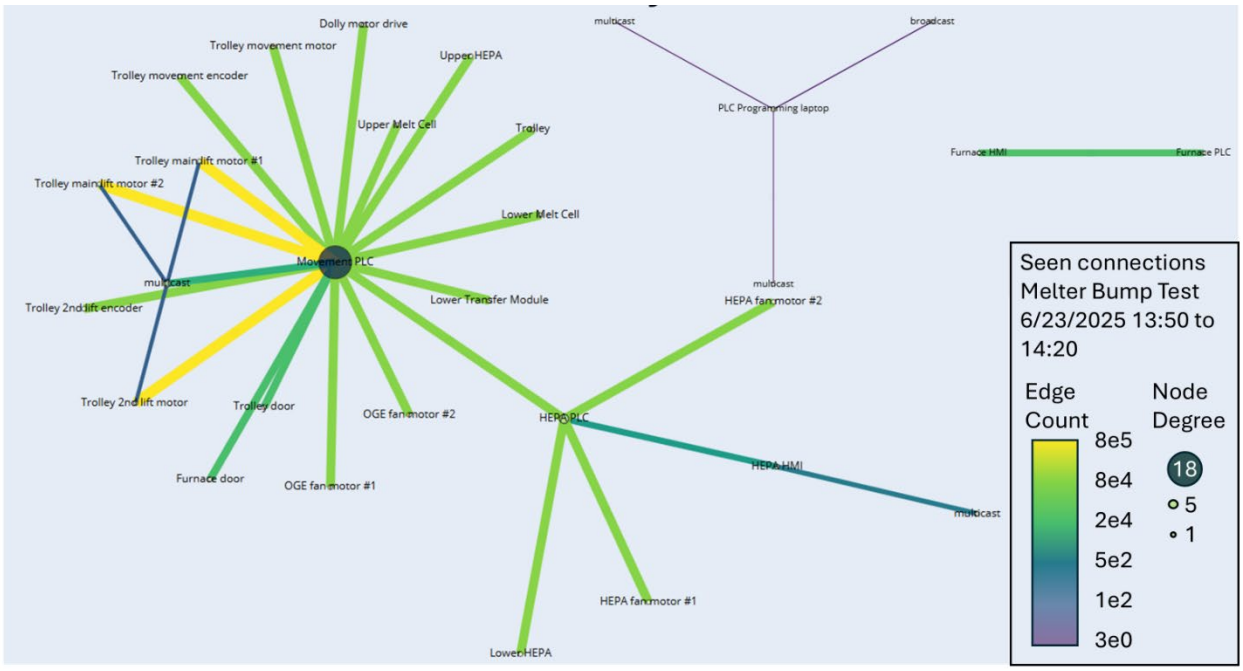

Figure 5. Detected network flow connectivity

The dominant communication protocol is IPv4 User Datagram Protocol (UDP), and only a tiny fraction is Transmission Control Protocol (less than 0.01%). While most of the nodes of the diagrams are host IP addresses, multicast and broadcast addresses represent connections to multiple machines and should be seen as hyper-edges. In our previous report we examined some network anomalies by looking more closely at the cameras on the management network. For this report we are focusing on identifying mostly safety and status related events on the movement, HEPA, and furnace PLCs by comparing their traffic with physical events related by the out-of-band data.

Additionally, some in-band temperature sensing data was collected from the temperature sensors on the furnace. This data was collected at 1 Hz from several locations within the furnace predesignated by the manufacturer. Although this data is reported in-band, we connected to it separately and collected it out-of-band.

2.2 Out-of-band Data Description

The physical events tell the story of the system from the operational perspective. There are two main sources: the out-of-band sensors and the operator logs. Figure 6 shows data from the out-of-band sensors: the three-phase current and temperature and vibration. Overlaid on the data are pink vertical bands representing operator events (see Table 1) that occurred at various times during the melt operation. Figure 7 shows the temperature values reported by the furnace, in-band, for comparison. Note that the temperatures in Figure 6 are being measured on the return coolant line from the furnace while these are from sensors reported by the furnace itself.

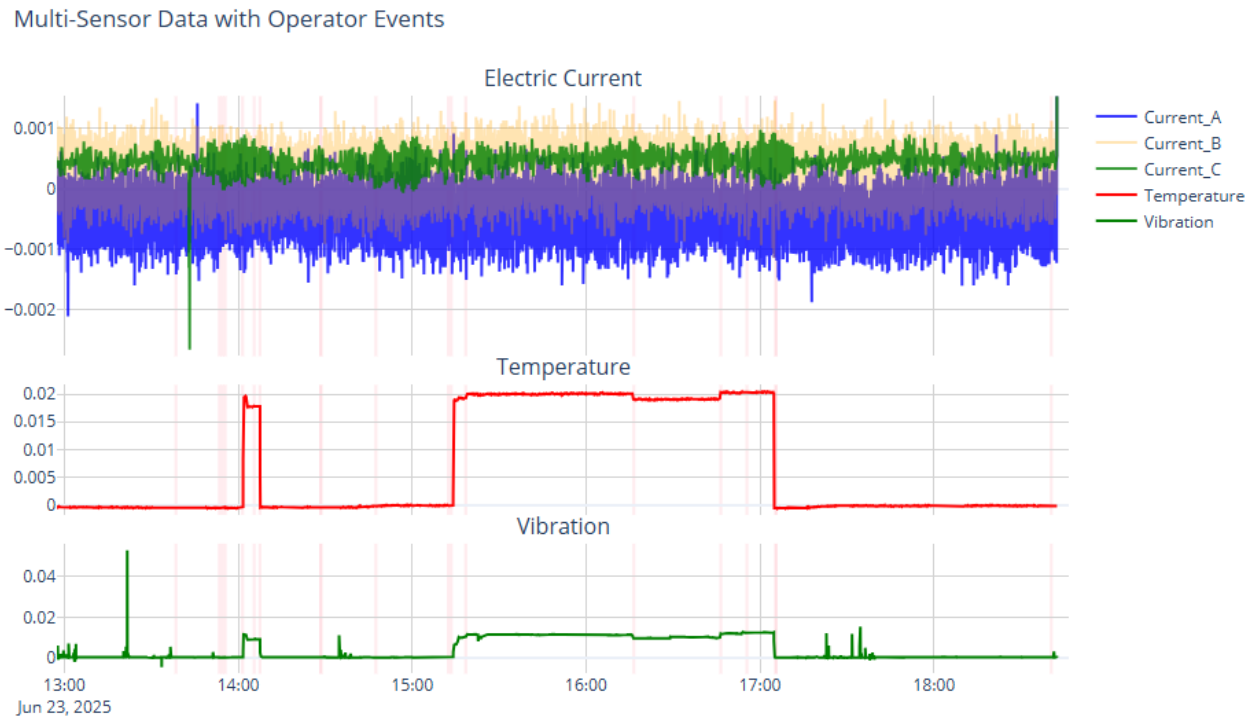

Figure 6. Three-phase current, temperature, and vibration during the 6/23/2025 melt operation.

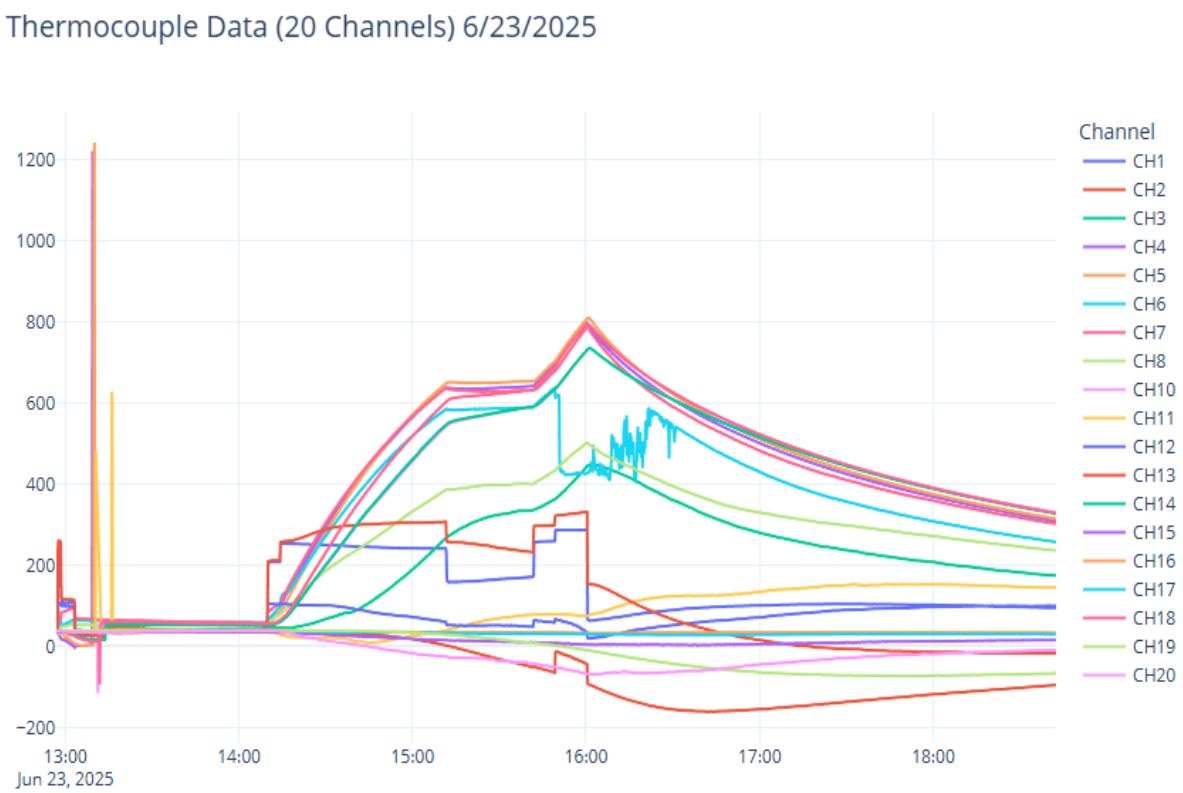

Figure 7. Thermocouple data (deg C).

Of note is how there is a significant excursion in Current_C around 14:00 hours. This section is expanded in Figure 8 to make these events clear. At about 13:38, the operator records flipping a breaker. Within a minute or so we see the excursion first in Current_C and then in Current_A. At 13:53 the operators report pressing the reset button and then at 13:55 they press the furnace start button. At 14:05 they take the furnace to 15%, then at 14:07 they notice a temperature spike and reduce to 10%. At 14:07 they power off the furnace to investigate further. There is some visible correlation between the operator actions and the out-of-band values sensed, although the times reported by the operators are not necessarily accurate compared to those reported by the sensors.

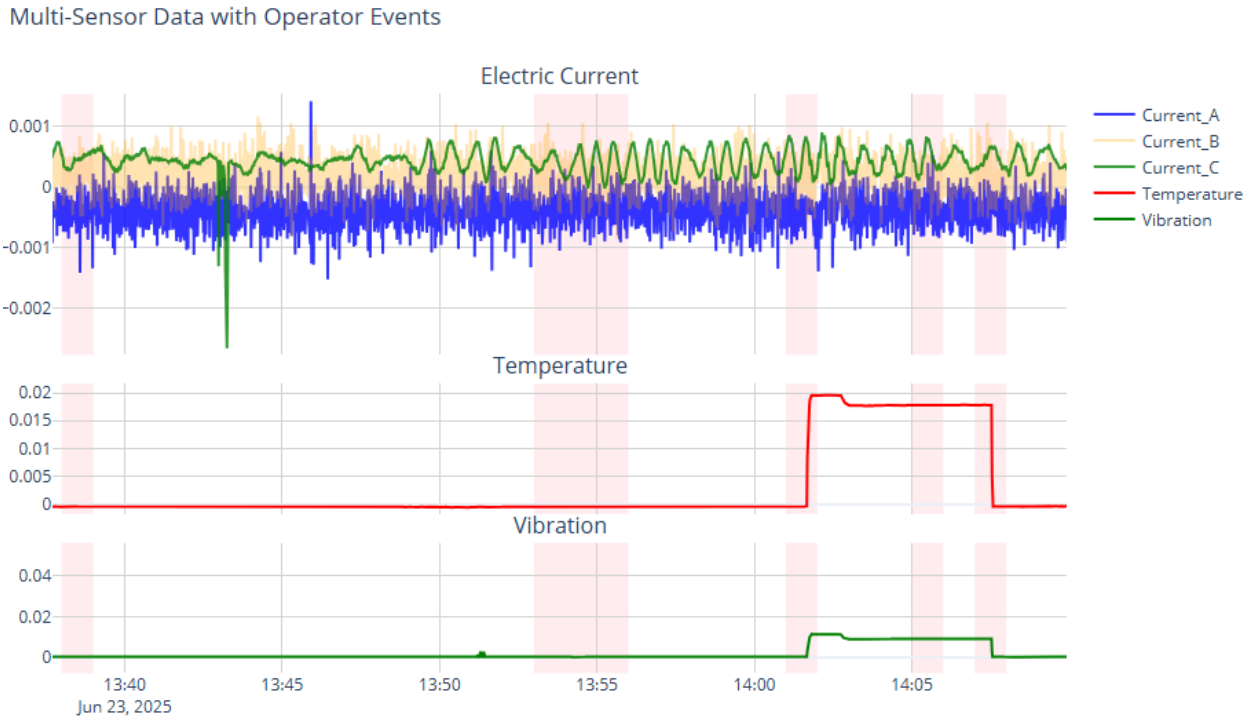

Figure 8. Close-up of events between about 13:35 and 14:15 on 6/23/2025

The Melter operators use a FileMaker Pro database to log important status changes and events. This data covers many of the areas that are not covered directly by the cyber data, although prior to our analysis, it was unclear whether any of the anomalies reported in this log are detectable from the cyber data alone. A successful outcome would be to show that the in-band data predicts or reflects these anomalies. Most 3S incidents will have a physical manifestation, and early detection of these conditions through timely analysis of the cyber data is the goal of this project.

Table 1 shows an excerpt of the operator log during the early start-up for the Bump Test conducted 23 June 2025. This section of time has been the focus of our first investigation into causality between network and physical events.

Table 1: Operator Log Excerpt

Start Time	Event/Notes
13:38	Flipped breaker to give furnace supply full 480V power (turned on furnace control power about 5 minutes earlier)
13:53	Pressed reset button. Now no alarms currently present

13:54	Pressed the GLD alarm button and then reset. No alarms currently present
13:55	Pressed the furnace start button and got a green light as expected
14:01	Started furnace at 15%
14:05	Furnace bumped down to 10% due to a quick spike in temperature on two of the thermocouples. Jumped up 100 degrees Celsius suddenly.
14:07	Furnace turned to 0%. Holding to see about a situation concerning a thermocouple. A couple of thermocouples leads were flipped around by accident as temperatures were going down as the melt was heated up. Taking a break to switch the leads around.
14:28	Pressed and held the GLD stop to ensure working properly. Pressed reset button immediately after.

The table shows a raw stream of events typed in by the operators at the one-minute level of precision (far less precise than the 4500 Hz of the network data or the 50 KHz of the physical sensor data). There are far fewer of these events (about 20 in a six-hour test, 18 hours if you include overnight operations). Things that went right and wrong have been noted producing an event stream that gives a high-level view of the first few minutes of the melt test operation. It also contains many indicators of problems that might have been more readily understood if operators were able to see errors in the cyber data in real time.

Several kinds of information are noticeable in the log:

1. Operator-initiated events, e.g., “Flipped breaker” and “Pressed Reset”
2. Operator notes for future action, e.g., “set up a meeting to review procedure.”
3. Operator hypotheses, e.g., “Since the charge is lower in the furnace, maybe all reading the same temperature.”
4. Qualitative observations, e.g., “No alarms currently present.”
5. Semi-quantitative observations, e.g., “Jumped up to 100 degrees Celsius suddenly.”
6. Quantitative observations, “Water temperature was high at 114F.”

Of these types, quantitative observations alone contain automatically extractable information, and even these vary in format and ordering, making it more difficult to use them for automated analysis. Operator log data provides a much clearer overview of the state of the system but does not include details that may only be visible in the cyber data. The Melter is not designed for fully autonomous operation, and the design of autonomous SMRs will need to consider the degree to which human observation can be replaced by automated means and how those means may be protected from intentional deception and unintentional occlusion. The other kind of out-of-band data collected in this exercise came from external sensors.

2.3 Out-of-Band (External) Sensing

Although the thesis of this work is that the cyber data can predict and diagnose the safety and security of ASMRs, the need for independent sensing has been well established especially by stakeholders involved in protecting the electric grid infrastructure [Darknet] as well as other cyber-physical systems. One approach uses LED links for remotely connecting measurement devices and analytical capabilities [Bogg2020], which states that, “*Out-of-band communication channels* (emphasis added) are needed to

provide a final layer of defense that is resilient in the case of attackers compromising devices and the entire network infrastructure."

During FY25 measurements were performed as a first step in developing the ability to detect anomalous behavior based on any inconsistencies between measurements and events reported via the OT traffic and actual physical conditions as measured and transmitted via a separate set of sensors and communication links. For instance, if an OT devices or component has been compromised, then although the network traffic may show that a command to increase temperature had been sent, and increasing temperatures are reported, the malware may be spoofing these readings. This possibility implies the operator would not be able to detect this without some out-of-band indication. [NSA/CISA] Therefore, some effort toward developing independent verification of physical measurements, states is needed. Still for fully remote operations, even this out-of-band reporting must be brought in-band to communicate back to the system owner.

The approach involves collecting out-of-band sensing data along with sufficient operational logs to make consistent deception very difficult. AI-based algorithms for cyber (or other anomalous) activities could be developed to find small inconsistencies from falsified reporting between in-band and out-of-band reports.

2.3.1 Measurement Setup

A key data collection was performed 23-25 June, 2025 that involved collecting strategically placed vibration, temperature, power line current sensors along with "in-band" sensors such as network traffic and a temperature sensor that is part of the process control system. Sufficient sensor data and operational log data were collected to support the desired capability development. The three sensor modalities collected as "out-of-band" sources are:

1. Vibration data (sensor attached to a frame that was near a motor that controlled the cooling loop for the induction furnace cabling)
2. Temperature data (Thermocouple that was attached to the outside of a metal piece of piping that held the induction furnace cooling liquid)
3. All three phases of current going into the entire Melter (at the main feed point)

Refer to the simplified block diagram showing the notional placement of the out of band sensors in Figure 3. The following graphics show exact placement of the out-of-band sensors within the Melter system. Figure 9 shows the three phase shore power that the current sensors monitor.

Figure 9. Input power sensor placement

Figure 10 shows the current probe employed and the location inside the Melter where the current sensors were placed.

Rogowski Coil model used:
https://www.pemuk.com/products/cwt-range/cwt

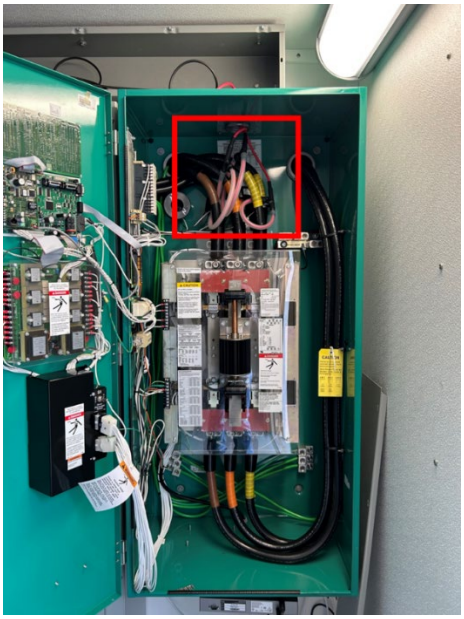

Figure 10. Current Probe Details

Figure 11 shows the placements of the vibration and temperature sensors on the trolley carriage.

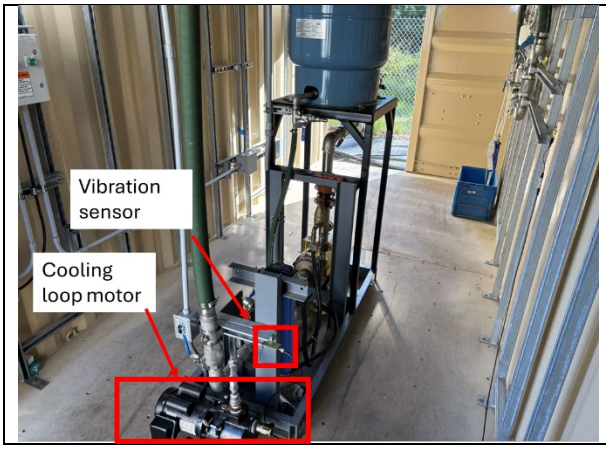

Vibration sensor model used:
https://www.te.com/en/product-1006015-1.html

Figure 11. Vibration and Temperature Probe Details

Figure 12 shows additional thermocouple temperature sensor placement. These sensors were placed by the Melter staff for their own use and they collect furnace temperature data at 1 Hz.

Figure 12. Temperature Sensing Placement

Figure 13 shows the data collection system, which consists of a laptop connected to a data acquisition system and the cables connected to the sensors. The cables to the three phases of current sensor were kept as close as possible to the same length so that the timing of the component signals would not be confounded.

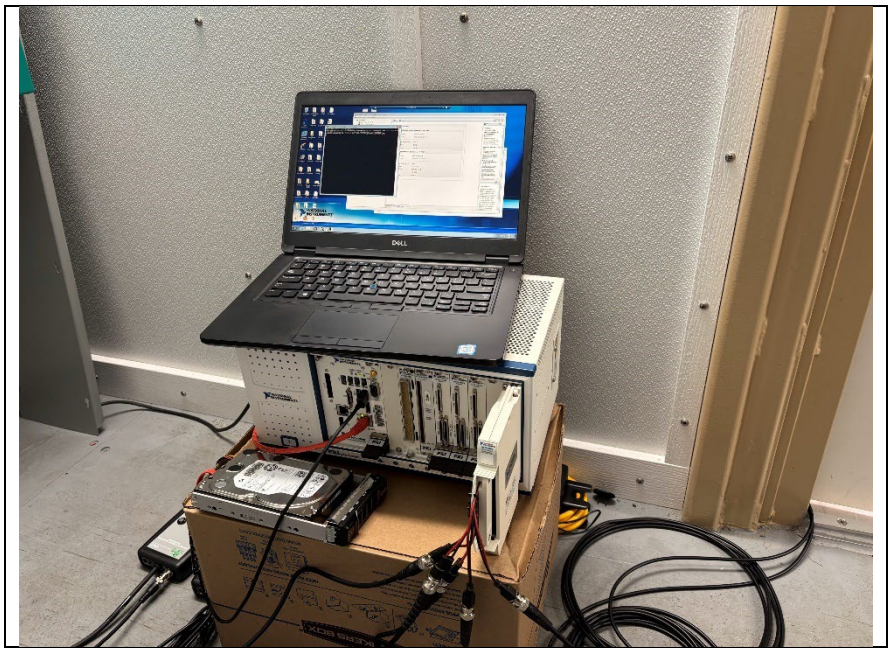

Figure 13. Data Acquisition System

2.3.2 Data Collection Details

Several events of interest appear to have happened during the first melt, including a breaker-trip from a VFD/motor, along with associated troubleshooting. The data appears to include many physical events, which should provide a rich feature-searching challenge. A third short capture (33 mins) was performed before the first melt on 6/23 to provide a baseline for any model training. This short capture was with the facility on standby (facility back-up UPS charging, lights, air conditioners, PLCs, cameras, computers, monitors, etc. were idling, not actively being used).

There were melt operations on both 6/23 and 6/25. The full melt process was captured on 6/23, however the operational team started earlier on 6/25 than scheduled, so the out-of-band sensors for the first hour of operations that day was not captured. The data collection team recorded the sequence of events that occurred on both melt days as shown in the logs (Table 2 and Table 3) below.

2.3.3 Operational Parameters

For the 23 June 2025 IFE Test Melt operation, the temperature was recorded as 35 degrees Celsius and the Relative Humidity as 56%.

Table 2 – 23 June 2025 Operational Logs for a test event

#	Start Time	Event/Notes
1	13:38	Flipped breaker to give furnace supply full 480V power (turned on furnace control power about 5 minutes earlier)
2	13:53	Pressed reset button. Now no alarms currently present
3	13:54	Pressed the GLD alarm button and then reset. No alarms currently present
4	13:55	Pressed the furnace start button and got a green light as expected
5	14:01	Started furnace at 15%
6	14:05	Furnace bumped down to 10% due to a quick spike in temperature on two of the thermocouples. Jumped up 100 degrees Celsius suddenly.
7	14:07	Furnace turned to 0%. Holding to see about a situation concerning a thermocouple. A couple of thermocouples leads were flipped around by accident as temperatures were going down as the melt was heated up. Taking a break to switch the leads around.
8	14:28	Pressed and held the GLD stop to ensure working properly. Pressed reset button immediately after.
9	14:28	Paused operations after noticing the cooling fan for the heat exchanger wasn't operating.
10	14:47	Troubleshooting the heat exchanger problem and noted the breaker feeding the VFD for the cooling fan had tripped. As a result the water to the furnace was running a little warmer (110 F) than nominal (90 - 95 F). Reset the breaker and the cooling fan activated.
11	15:12	Pressed green start button for furnace
12	15:13	Turned furnace power to 12%
13	15:18	Turned furnace power to 15%
14	16:16	Turned furnace power to 11%. Began hold at 600C.
15	16:46	Turned furnace power to 16%
16	16:55	Started noticing a red glow from the furnace while material reading around 700C
17	17:05	Turned furnace down to 0%.
18	17:05	Alarm on heat exchanger went off. The cooling fan was found stopped and the water temperature was high at 114F. Breaker tripped again. Reset breaker but killed power to the pump temporarily. Pump stopped working and pressure started dropping. Turned back on the pump but emergency auto water did not kick in. Lesson learned: emergency water doesn't run properly when the main pump is running. Also believe that VFD breaker may be tripping due to VFD overtemperature. VFD rated for 122F and when hitting the VFD with an IR thermometer, we saw parts of it were 128F.

19	18:40	Returned back from a scheduled discussion to ensure the system was prepared for its overnight running. Because of the high temperatures in the cooling and cooling pump modules, we placed fans in them. Note: When a fan was plugged into the receptacle of the cooling pump module, the two 20A breakers in slots 2 & 4 as well as the 50A breaker in slot 1/3/5 had all tripped. Think these may have tripped after the fan was plugged into the receptacle instead of before, otherwise the pumps in the cooling pump module should have had their power feed cut earlier when we were troubleshooting at 17:05.
20	19:20	System deemed in a safe state and everyone left at the MNF heads home.

Table 3 - 23 June 2025 Operational Logs for Test Event

Start Time	Event/Notes
7:22	Glenn throws the breaker to give full 480V power to the furnace power supply
7:30	Press Stop/Reset and GLD alarm doesn't go off right away. Goes off after Roy clears the alarms.
7:35	Furnace power turned up to 15%. Thermocouple #13 showing strange numbers, believe it could be possible coupling with the furnace electromagnetic field.
7:36	Thermocouple #6 showing strange numbers as well, believe it could also be due to coupling. And thermocouple #11's terminals may have been flipped as it's climbing down in temperature.
8:15	Started signal data collection
8:16	Increase furnace power to 17%
8:17	Walked over to correct the vibration sensor placement
8:30	Started collecting more data in 30 minute increments. Water Temp: 85 F Furnace Temp: 586 C (Channel 7) Outdoor temp: 90 F Humidity: 67%
8:37	Began 30 minute hold.
8:38	Decrease furnace power to 14%
9:00	Water Temp: 87 F Furnace Temp: 616 C (Channel 7) Outdoor temp: 88 F Humidity: 70%
9:07	Increase furnace power to 17%
9:18	Increase furnace power to 20%
9:30	Water Temp: 90 F Furnace Temp: 795 C (Channel 7) Outdoor temp: 91 F Humidity: 65%
9:32	Increase furnace power to 22%

9:34	Lost thermocouple #6 due to passing its temperature limit. Furnace temp (Channel 7) at the time: 842 C
9:38	Increase furnace power to 25%
9:47	Increase furnace power to 28%
9:59	Increase furnace power to 31%
10:00	Water Temp: 93 F Furnace Temp: 1046 C (Channel 7) Outdoor temp: 94 F Humidity: 61%
10:08	Increase furnace power to 34%
10:14	Increase furnace power to 37%
10:15	Thermocouple #7 stalled at 1185 C
10:25	Thermocouple #7 is lost due to passing its temperature limits
10:30	Water Temp: 95 F Furnace Temp: 1209 C (Channel 10) Outdoor temp: 95 F Humidity: 60%
10:41	Furnace power turned down to 0% to allow cooling. Melting operations on hold due to thermal concerns around the radiation-hardened camera. The pyrometer failed as it started outputting a current of 2mA (around a temperature of 90 C at its tip), likely due to heat and don't want to risk the $70k radiation-hardened camera that is sitting right beside it. Pyrometer has a rated temperature of 65 C, while the radiation-hardened camera has a rated temperature of 55 C. Never been an issue before, but thinking that because the ambient temperature is so high, the air between the furnace cover and radiation-hardened camera/pyrometer is not being cooled off as we're accustomed to.
11:00	Water Temp: 97 F Furnace Temp: 1246 C (Channel 10) Outdoor temp: 95 F Humidity: 60%
11:20	E-stop pressed on remote Furnace pedestal in order to trip the breaker to the furnace to disable the 480V power provided.
11:30	Water Temp: 100 F Furnace Temp: 1372 C (Channel 14) Outdoor temp: 98 F Humidity: 53%
12:00	Water Temp: 100 F Furnace Temp: 1080 C (Channel 18) Outdoor temp: 101 F Humidity: 50%
2:30	Water Temp: 105 F Furnace Temp: Unknown Outdoor temp: 102 F Humidity: 47%
4:00	Water Temp: 105 F Furnace Temp: Unknown

	Outdoor temp: 101 F Humidity: 40%
7:21 (6/26/2025)	Cut off furnace control power
7:35	Stopped network data collection

The multi-channel data files are:

- "pxie4305_stream_20250623_084428" – 33 minute baseline capture
- "pxie4305_stream_20250623_125654" – Full melt that occurred on 6/23
- "pxie4305_stream_20250625_081515" – Melt that occurred on 6/25, minus this first hour or so

The channels are designated as follows:

- Channel 0: Vibration
- Channel 1: Temperature
- Channel 2: Phase A current
- Channel 3: Phase B current
- Channel 4: Phase C current

NOTE: on 6/25, the vibration and temperature leads were inadvertently swapped when the sensors were rewired, so for the capture that occurred on 6/25, the channels are:

- Channel 0: Temperature
- Channel 1: Vibration
- Channel 2: Phase A current
- Channel 3: Phase B current
- Channel 4: Phase C current

2.4 Data Limitations

Because this demonstration was a limited first attempt at setting up and running an external data collection system on a production system we only used five sensors (temperature, vibration, and 3x current). This small number can never give a holistic view of the physical system. Next year, we plan more sensors and collection from a system that was designed with data collection in mind, Argonne's Mechanisms Engineering Test Loop (METL) that emulates a sodium fast reactor's sodium cooling loop.

Limitations of the in-band network data include devices that do not share their status information across the network producing data that cannot be collected centrally. The generators, coolant pumps, backup water supply, and air conditioning units are self-contained and do not report status or accept commands from the PLCs. Some subsystems (e.g., emergency water turn-on and water pump pressure valves), are designed to be autonomic and self-correcting (e.g., triggered on power loss or low pressure). Still, other systems produce a status that is directly read by the operators without the intervention of a PLC like the pyrometer for the furnace that enables operators to read the furnace current and temperature by looking at wall-mounted dials.

Human observation is required to detect failures in these devices, a situation that will not suffice for an autonomous SMR, which must be designed to be self-sufficient. These requirements would be mitigated

in such a system, unlike for the Melter. Thus, we hypothesize in this work that failures in all systems will be at least partially observable in the cyber data.

3.0 Algorithm Development

As mentioned above, this report discusses two broad approaches to achieve mitigation of cyber-attacks on physical systems – in-band and out-of-band. This section addresses both approaches in different subsections.

Table 4 gives a categorical overview of possible levels of a priori knowledge and some example techniques for developing optimal stages of knowledge about the state of the system.

Table 4 - Stages of A Priori Knowledge Versus Techniques for In-Band OT Analysis

Stage 1	Stage 2	Stage 3	Stage 4
Length and timing values only, with No packet info readable	Only source and destination are decipherable	Source, destination, and packet contents readable but unparsed data	Fully decipherable packets including meaning of each field
Use statistics of packet arrival times to detect anomalous activity	Map the network and try to infer some Stage 3 values using Image-based ML	Use physical knowledge to infer contents of certain packets – enables some Stage 4 techniques	Complete process model and states can be inferred
e.g., [Moor2017]	e.g., Vann	e.g., [Moor2018]	

Increasing Levels of A Priori Knowledge

Following are some possible components of the algorithm development process. This is not a complete list, but merely examples. And, even if the currently envisioned Melter process already has a well-mapped state mapping and OT topology, these methods can be used to independently verify if remote changes are being made to the process and adapt the cyber-attack mitigation methods to any intentional changes in the Melter process.

Figure 14 gives a more detailed flowchart / information / algorithm development view of both the FY25 efforts and the proposed FY26 efforts. As shown in the middle column of database sources, as the amount of prior knowledge is increased, more advanced cyber mitigation algorithms can be developed. Note that the out-of-band algorithms (Vann IP) is the only one that works with no network (OT) traffic or knowledge. This method can be used stand-alone or used to augment other methods. For this reason, we strongly suggest that future approaches include at least some independent means of communicating sensor data back to wherever remote monitoring algorithms can be executed.

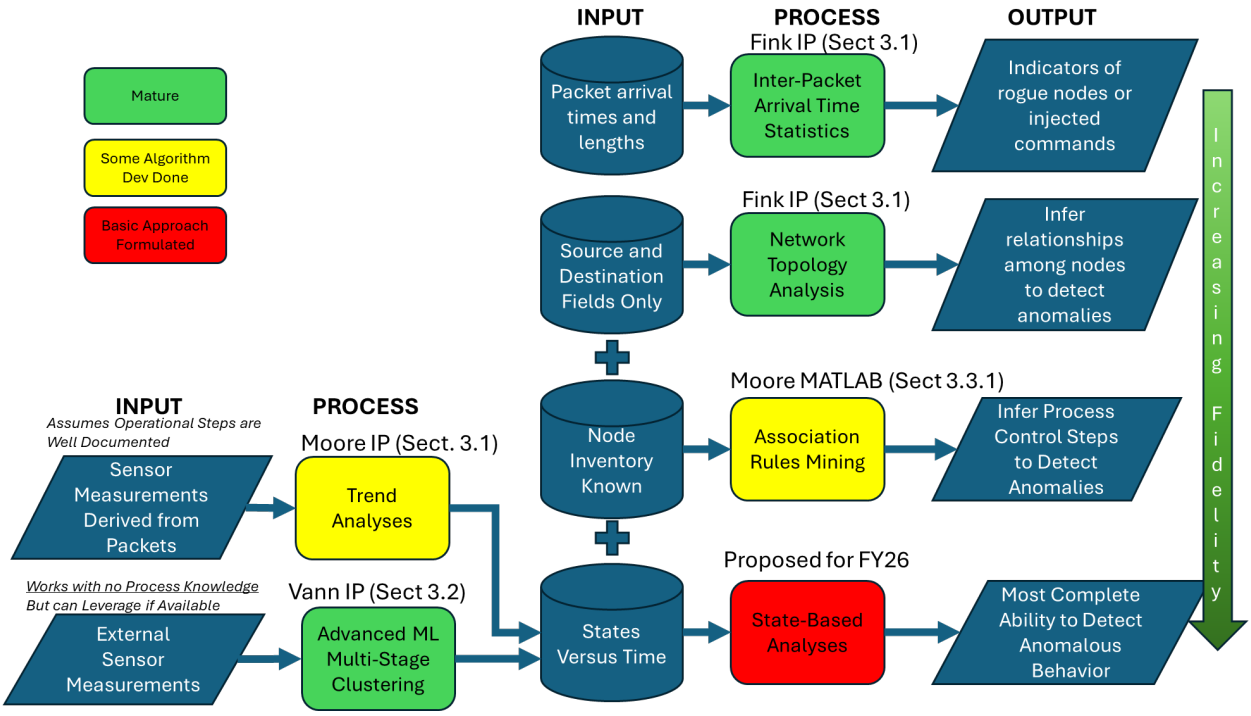

Figure 14 – Overview of Cyber Mitigation Capabilities Versus Availability of Documentation on the Process Control System

3.1 In-band OT Algorithm Development

The initial approach of this project was to use large language model embeddings on network packet contents to determine similarity between devices and to estimate normality, but this approach has been consistently thwarted by the site's operational security that is designed for plant operations and not research. As we are standing up an open science network, we expect these limitations to become obsolete early in fiscal year 2026. Some work-arounds using autoencoders were done, but they could be completed in time for this report. Thus, we focus on our statistical analysis of network traffic data.

OT traffic is very different from information technology (IT) network traffic. OT traffic is by necessity extremely regular and sparsely filled with information. Most of the communication is in star networks with a PLC at the center sending commands and receiving status from its subordinate controllers. Most OT devices continually send status messages as regularly as heartbeats. This keeps the whole system continually aware of what is on the network.

3.1.1 Packet Bitmaps

To understand the kinds of traffic on the network our second approach was to examine the packets as a bitmap to find patterns. Figure 15, Figure 16, and Figure 17 show bitmap representations of various packets on the OT network, one packet per line. When a bit is on, a colored pixel is shown. When it is off the pixel is black. The color of each packet's pixels is determined by the connection identifier to distinguish the connections from one another. Figure 15 shows 600 packets (about 0.13 seconds worth) as they appear in temporal order. All the conversations are interleaved so a variety of colors and patterns appear. This display also shows the relative frequencies of each type of communication.

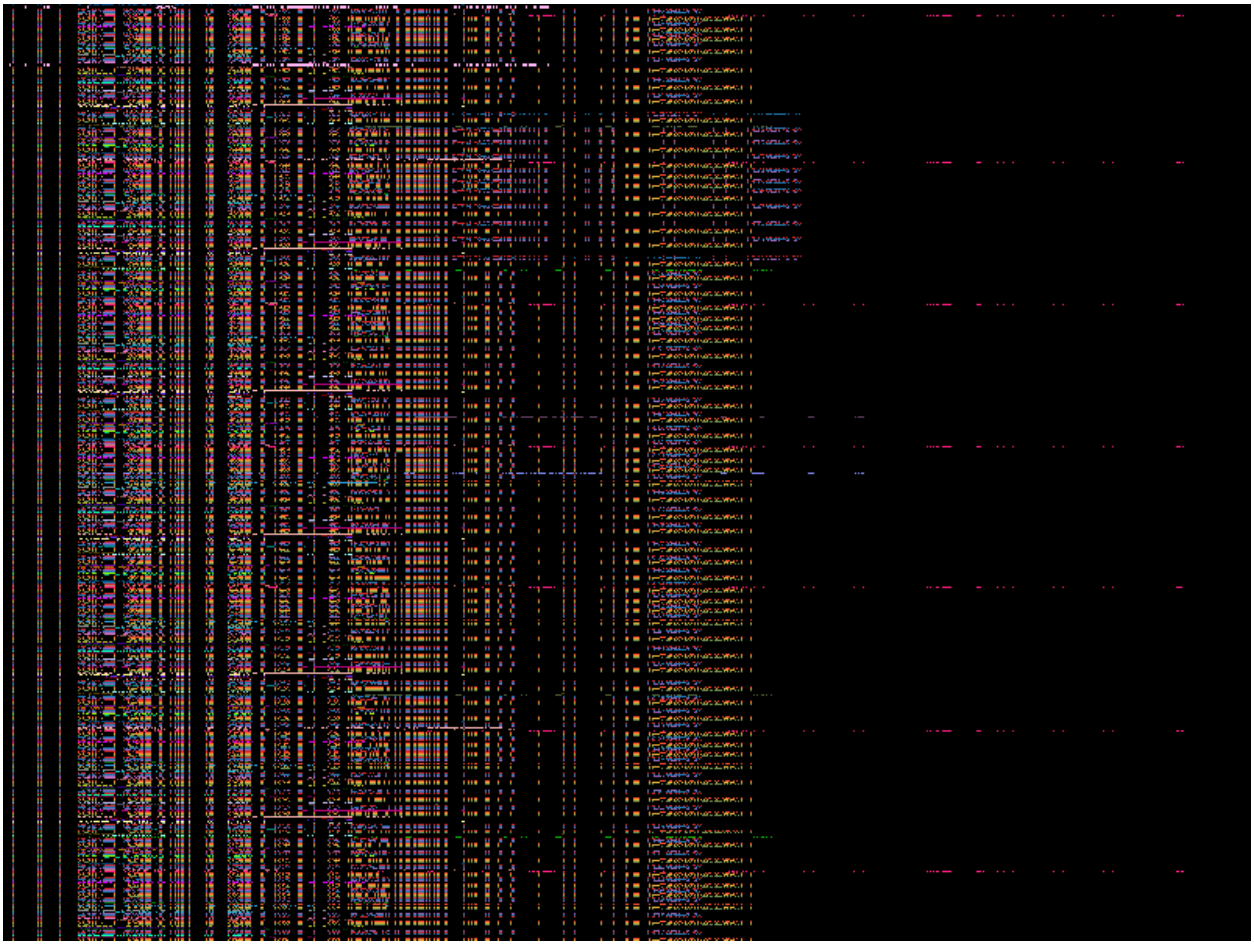

Figure 15. Bitmap representation of all packets colored by connection

Figure 16 shows isolated streams of packets from each of the motors in the Melter talking to the Movement PLC. Although they are all motors and their communications all have similarities, the pictures illustrate visual differences in the communication patterns of different types of motor: lift motors, fan motors, movement motors, and door motors. Each pattern is very distinct, so that if another lift motor showed up, its fingerprint could easily be recognized.

Figure 17 contrasts the communication patterns of device-to-PLC with PLC-to-device messages. Not only are the communication patterns of the devices to the PLCs distinct, the PLC communicates with separate patterns to each type of device. Clearly, although PLCs are distinct from other devices, they also communicate with many different patterns, making it hard to identify a single PLC from its communications with its devices.

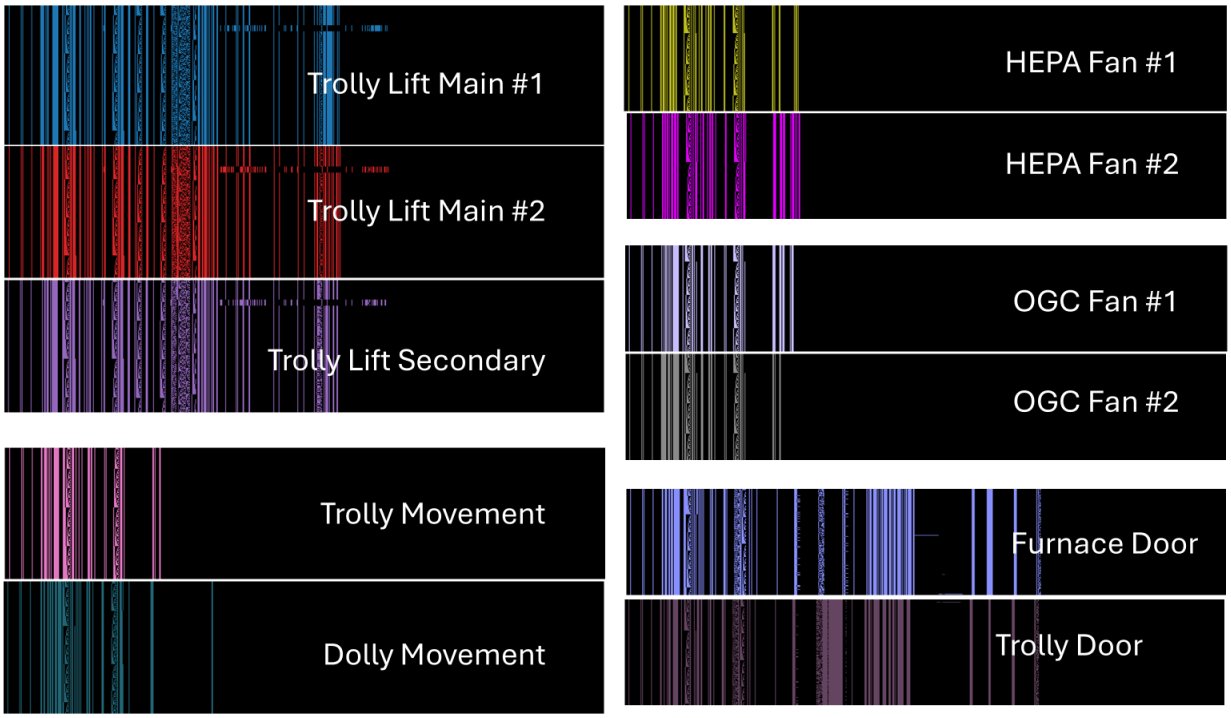

Figure 16. Bitmap representation of various types of motors in the Melter → PLC

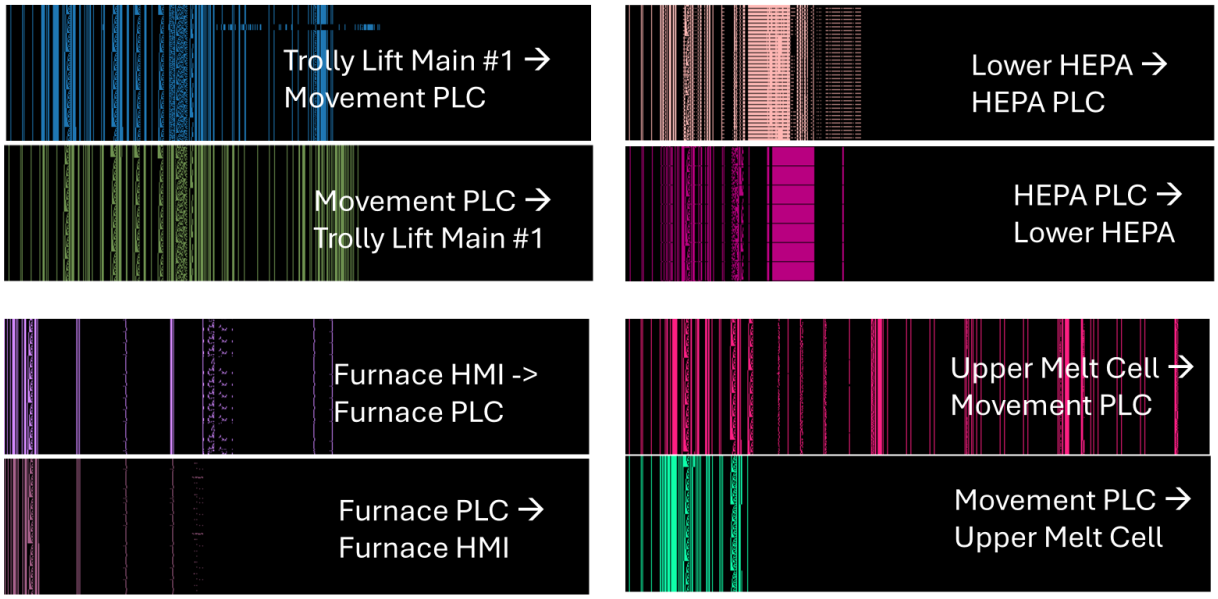

Figure 17. Bitmap representation of device ←→ PLC communication

Having established the patterns of communications in the OT network, it was clear that devices could be distinguished to some degree from their individual utterances without knowing what they meant. This agnostic approach is important because new device types are being added all the time, and we desire our approaches to be as widely applicable as possible. Thus, our third approach was to look at packet interarrival times (IPAT) and jitter.

3.1.2 Packet Interarrival times and Jitter

Network traffic whose content is unknown (stage 3 from Table 4) may still be characterized by the number of packets that are sent in a time window, the inter-packet arrival times (IPAT) and the packet delay variation (PDV or jitter). IPAT measures the time difference between consecutive packets arriving in a stream. This metric proves essential for analyzing traffic patterns and detecting anomalies. PDV quantifies the variation in IPAT across a packet stream. Network engineers commonly refer to PDV as jitter, which directly impacts real-time application performance.

In an SMR's OT network, the difference between IPAT and jitter is as important as the difference between a consistent heartbeat and an irregular one. The network's function is to maintain absolute predictability, where any deviation, however small, is considered a potential failure that could lead to catastrophic events. The IPAT is not just a statistical average but an engineered requirement, such as a process sensor transmitting data every 10 milliseconds. Any IPAT measurement outside this narrow, predetermined window is a malfunction because each data packet is a discrete control or safety signal. A missed packet or one that arrives late can put the entire system out of its programmed time cycle, potentially causing a safety system to enter a failed state.

Jitter is even more significant. Jitter represents the variation in the IPAT. For a hard real-time nuclear system, the acceptable jitter is effectively zero. The tolerance is so tight that the network must operate deterministically, meaning every action happens precisely when it is supposed to, every time. Any measurable jitter could indicate a severe underlying problem, such as network congestion, a compromised device, or an imminent hardware failure. In a safety-critical system, jitter is not just a nuisance; it is a diagnostic signal of a potentially disastrous condition. Jitter may be measured as the second difference of packet arrival times or as a rolling standard deviation of IPAT. The former approach gives a noisier measure of jitter that is effective for considering the limits of its magnitude for hard real-time applications. The latter approach gives a more smoothed measurement that is good for streaming media, etc. We use the former method for the most part in our analysis. Table 5 shows example expected times for IPAT and Jitter in high consequence applications such as SMRs.

Table 5. Acceptable tolerances for IPAT and Jitter

Metric	Acceptable SMR tolerance	Context and implication
IPAT	Fixed and deterministic. A packet rate of milliseconds or microseconds, with virtually no variation.	A sudden or sustained change in IPAT is a system fault. The safety system must be designed to either ignore or respond to an unexpected change in timing.
Jitter	Sub-microsecond or effectively zero.	Any measurable jitter is considered a fault. For systems controlled by protocols like IEEE 1588 (PTP), the synchronization accuracy is measured in microseconds to ensure perfect timing across all networked components.

In a one gigabyte network data sample with 8,114,333 packets (representing about 30 minutes of traffic), the IPAT, the jitter (second difference method), and the jitter (rolling standard deviation IPAT method) were found to be as shown in Table 6. Mean IPAT appears to be sub-millisecond while jitter is much higher than preferred for real time and nuclear applications.

Table 6. Packet timing statistics

	IPAT	PDV (2nd difference)	PDV (rolling stdev)
mean	2.176503e-04	-2.242945e-11	2.734594e-04
std	2.736907e-04	4.287013e-04	1.631409e-05
min	0.000000e+00	-2.584000e-03	2.178249e-04
max	2.613000e-03	2.610000e-03	3.377424e-04

We investigated both IPAT and jitter to determine what could be learned from them. We found it most useful to recalculate both metrics on a per-connection basis. We refer to a connection simply as a conversation between two endpoints with network addresses. Connections may be refined by considering the protocol being used as well. Connections may be considered as two-way (A ←→ B) or one-way (A → B or B → A). Then IPAT and PDV (jitter) tell us more about the communicators than the overall system.

We found that different patterns of two-way, protocol-agnostic connections were very distinct for different kinds of devices (Figure 18). The figure shows three minutes of IPAT data for seven different connections. The left subfigure shows IPAT for connections between the main PLC and three motors. The right subfigure shows IPAT for connections between the main PLC and four device controllers. The motors behave so similarly over time that their IPAT distributions completely overlap. The controllers' IPAT measurements show a pattern of sequential connections from the PLC to each of its controllers in turn, causing a very regular zigzag pattern, offset by a few seconds each.

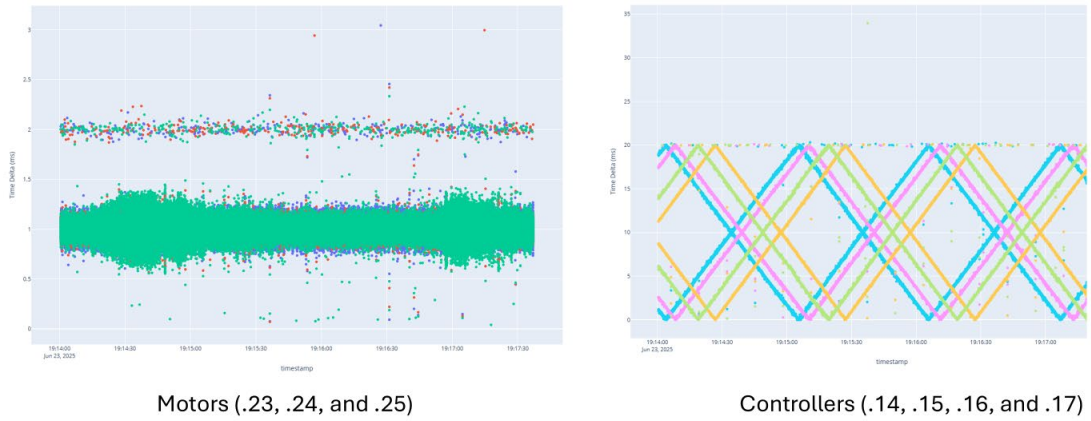

Figure 18. IPAT patterns distinguish classes of devices.

Being able to tell devices on the network apart by just looking at IPAT is very useful. Even though the function of a device is not easy to determine from its network characteristics alone, knowing that it is one of a group of similar devices allows us the expectation that whatever we learn about any member of the group may apply to them all.

Jitter was even more interesting. We proceeded on the thesis that changes in system function would require slightly different communications, some of which may vary in length and timing features and thus would force changes in jitter. There appear to be some correlations between operator-noted events (see Table 2 for an event listing) and jitter variations in packet flows (see Figure 19). However, our working definition of high jitter is arbitrary (99th percentile) and there are many more high-jitter occurrences than operator events. A large part of the problem is that operator events are very sparsely occurring while a score of spikes and troughs in jitter occur in any 30-minute window.

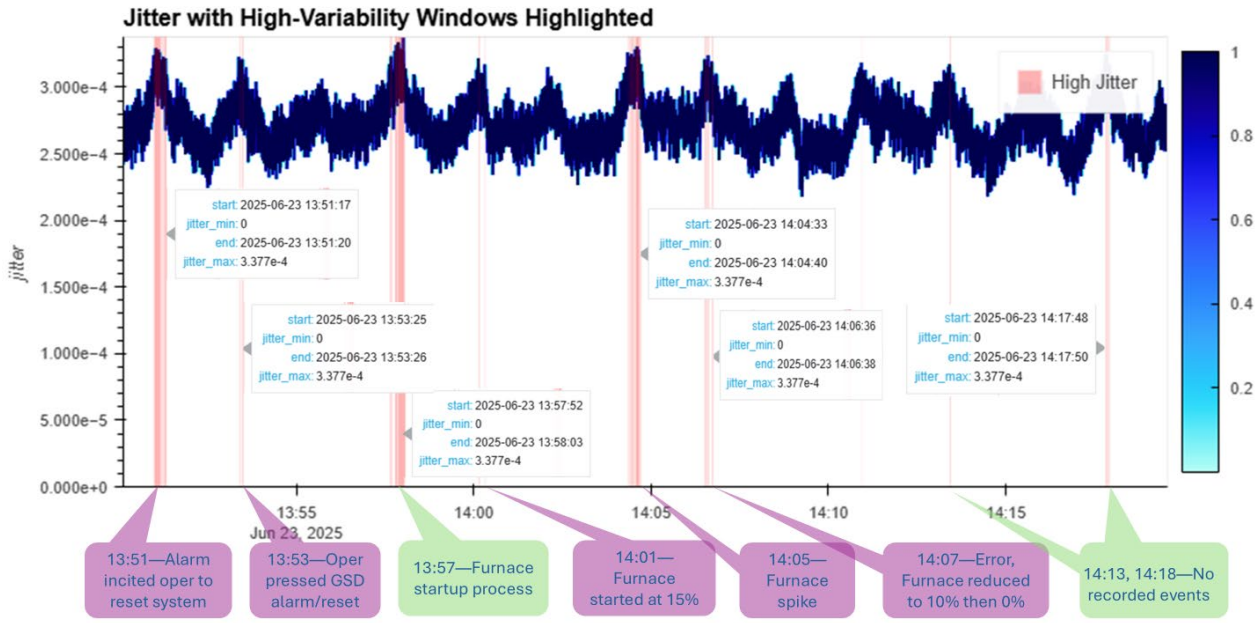

Figure 19. High-Jitter apparent correlation to operator events.

For most connections, jitter was very low, as would be expected. But some components of the Melter seemed to have much higher jitter than the norm. The HEPA system and the furnace were two subsystems where jitter was observed to be greatest. Interestingly, changes in jitter in the HEPA PLC talking to its fan motors coincided somewhat with the furnace start events (Figure 20). This may be an example of an event correlation, but if it is one would expect a similar pattern in the furnace PLC and HMI communications, but it does not appear to do so (Figure 21).

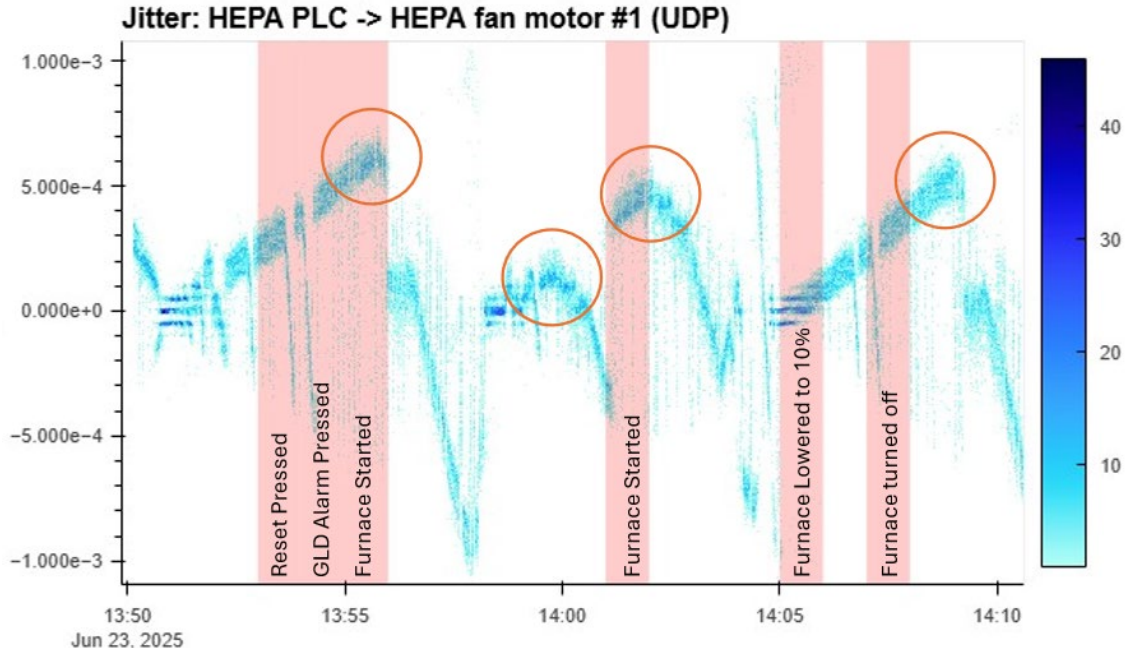

Figure 20. HEPA Jitter falls after furnace restart events.

One possible reason for this mismatch is that the furnace HMI and PLC talk to each other in Transmission Control Protocol (TCP) rather than User Datagram Protocol (UDP). TCP is a connection-oriented protocol and is much more sparsely found throughout the data (less than 0.001%). If the jitter between these components is indicative of system changes, it will be much less apparent than for UDP connections.

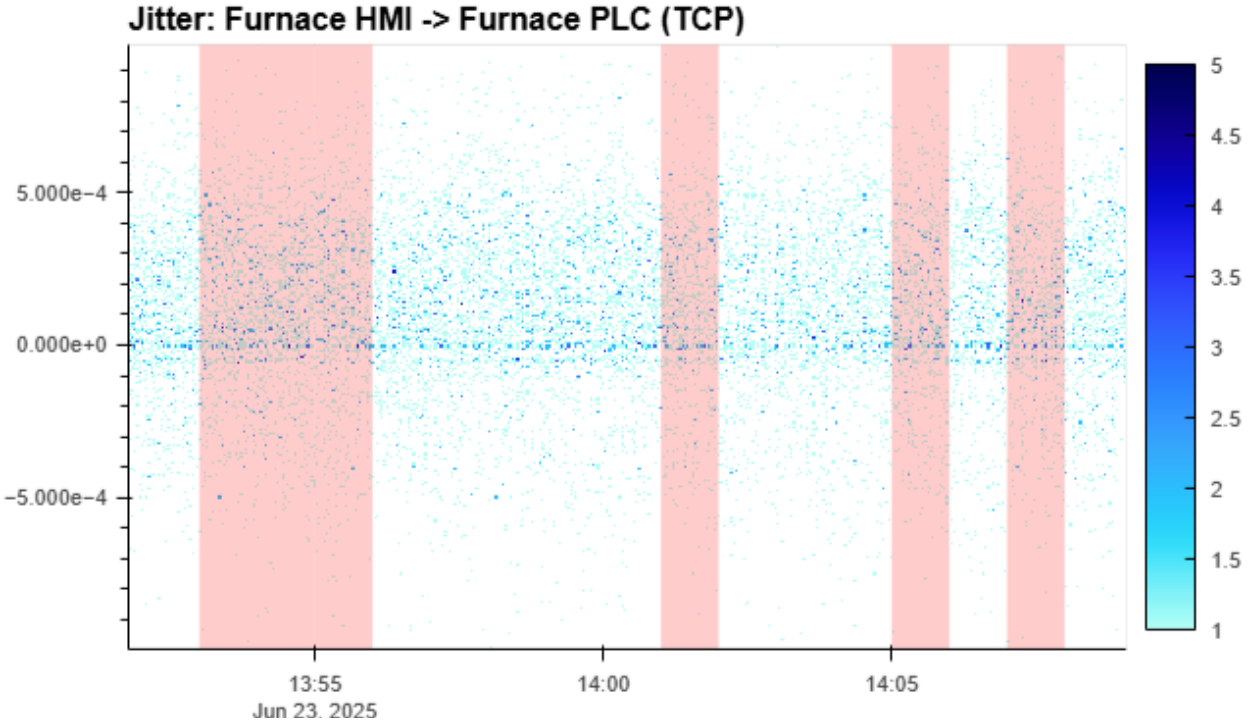

Figure 21: Furnace HMI and PLC communication does not correlate to furnace restart events.

At this point we believe much more is possible and more may be done only if we move past stage 3 (Table 4) of our a priori knowledge and take some semantic meaning into account. This does not force us to move all the way to stage 4 where we have a white-box view of all data, understanding all protocols, instead out next step will be to use large language models on the packet contents to look for predictive correlations at the byte level.

In addition to more conventional cybersecurity analysis approaches discussed earlier, some of the following approaches will be pursued in FY26:

1) State-discovery that includes multiple methods of deriving the state of the system using inferred physical data trends (e.g., [Moor2018]) or command-linkages using association rule mining. Once the states are discovered, then inputs or outputs that are consistent with the current state are detectable anomalies.
2) Converting network traffic profiles into 2-D images that can be used with image-based AI/ML techniques, ([Vann2017], [Moor2019]). In this case, the methodology focuses on translating 1-D data into 2-D images that can leverage the rich world of CNN/DNN-based image classification.
3) OT traffic statistics can be used to detect signal injection attacks as in [Moor2017].

3.2 Out-of-Band Sensing Algorithm Development

This data collection resulted from an unexpected confluence of this project's goals with another demonstration that benefitted this project. Thus, it represents an additional delivery beyond the planned milestones and deliverables. Although the data was collected in FY25, only initial analysis could be performed before the end of the FY. Thus, this section will mainly describe the efforts to use the collected data to develop a cyber mitigation algorithm which will be performed in FY26.

In general, the approach that will be explored in FY26 has been established as part of analyzing or identifying the "hidden" or unknown states of other nuclear related operational processes. There are two main stages – identifying separable "states" such as the ML-based clusters shown in Figure 22 (see work of co-author Vann in [Tu2024]), and the subsequent mapping of those clusters over the timeline of the process of interest as shown in Figure 23. Although this tool was developed as a remote characterization capability, it also lends itself to anomaly detection when addressing cyber-controlled physical processes or cyberphysical systems.

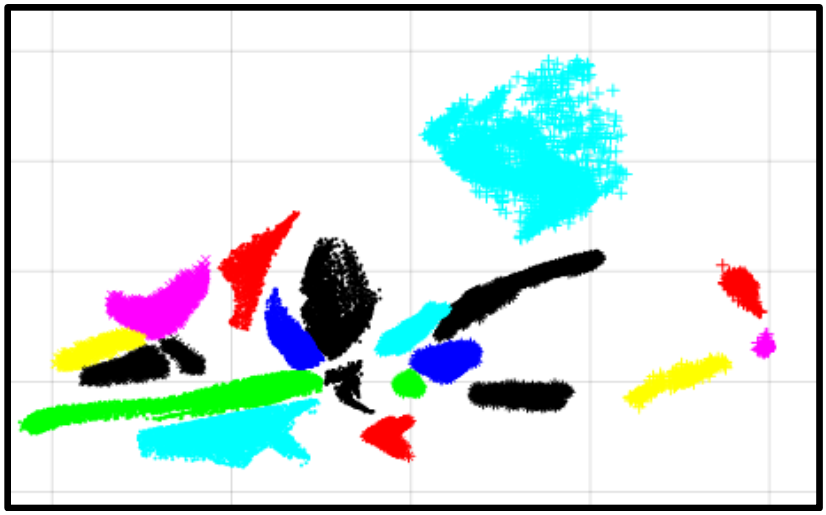

Figure 22- **An example of a Code layer scatter plot using only reconstruction loss for training. NOTE: The cluster colors and labeling were manually assigned by a human analyst.**

As shown in Figure 23, analysis of labeled emitter types over time can be exploiting using associative rules mining algorithms to determine the relationship, whether operational, technical, hardware platforms, or hierarchical, between different sensor modalities or even non-physical information. Additionally, research will be conducted to further separate RF events determined to be the same waveform (i.e., label or class) into additional clusters representing each "serial number". Often referred to as "fingerprinting", each sub-cluster formation would represent a different hardware-based radio transmitting this same waveform.

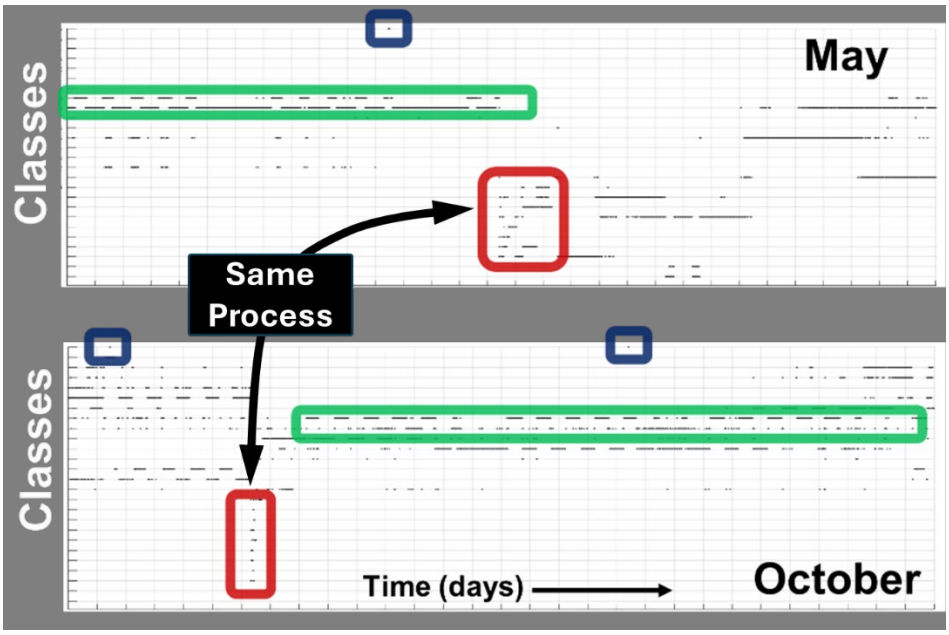

Figure 23 –Analysis of a blind dataset from a test facility is shown as a scatter plot of human labeled (i.e., "classes") cluster points versus time.

3.3 State Based AI – Enabling the Integration of In-Band and Out-of-Band Techniques

Although this work is mainly targeted at future efforts. Some of the groundwork was laid during FY25. State-Based techniques involve identifying or leveraging identifiable states within a process or system. While most state-based analyses (e.g., Hidden Markov Models [HMM] or Partially-Observable Markov Decision Process [POMDP]) start with a set of known states, in [Moor18], it is shown that knowledge of the basic physical nature of a process (e.g., car accelerating versus decelerating) can be determined in a purely data-driven method. Although a supervised learning technique was used in that research, unsupervised approaches can potentially be used in the future. Regardless of whether the process states are pre-defined, derived via ML, or some combination – once the states of a system or process are known, then there is a basis for detecting and mitigating anomalies. For instance, if there is a state called Long-Term Melt that is known to require 60 minutes at a certain temperature, then a) the data (in-band and/or out-of-band) can be used to determine when that state is initiated, b) the state puts bounds on appropriate levels of electrical current, temperature regulation, melter position, etc., and c) the state can be used to determine allowable precedent and antecedent states. Thus, any in-band commands or measurements that are not consistent with the verified state, should lead to immediate investigation by operators and/or messages to cyber protection resources.

Specifically, as in [Moore18] in-band (OT or CAN bus) traffic can be used to learn meaningful state definitions before the system is deployed. Vann's work discussed in 3.2 can be used to learn the behaviors and states independently using the external or out-of-band sensors that are not part of the process control OT traffic. Options exist for both the case where the states of the system are fully or partially known a priori, and for systems that are not well known a priori. The latter situation will require a more thorough develop of Vann's methods and the former will allow the two methods to be used together.

3.3.1 Traffic mapping using Association Rules Mining

While the connections between the command modules (e.g., PLC) and most of the sensors are typically known a priori for government-sponsored nuclear material processes, there are always cases in which the possibility of rogue nodes, nodes that were not properly listed, or malfunctioning nodes force the need to have data driven methods for determining which devices are communicating and when. Note that significant information on this topic has been covered in 2.1 Network Data Description above. This section is meant to supplement those approaches.

As an example exercise, a simple approach was used to label all packets as "unique" that had a unique source-destination combination, but it was assumed that differing packets having the same source and destination could not be determined to be unique unequivocally. Thus, starting with the least knowledge about the meanings of the data fields (designated as Stage 2 in Table 4). The following labels were given to the complete set of source ←→ destination combinations.

Table 7. Connection identifier assignments

Label	Flow
0	Movement PLC to Dolly motor drive
1	Movement PLC to Trolley primary lift motor #2
2	Movement PLC to Trolley secondary lift motor
3	Movement PLC to Trolley primary lift motor #1
4	Movement PLC to OGE fan motor #2
5	Movement PLC to Trolley movement encoder
6	Movement PLC to Trolley movement motor
7	HEPA PLC to HEPA fan motor #2
8	HEPA PLC to HEPA fan motor #1
9	Movement PLC to HEPA PLC
10	Movement PLC to Upper HEPA
11	Movement PLC to Upper Melt Cell
12	Movement PLC to Trolley
13	Movement PLC to OGE fan motor #1
14	Lower HEPA to HEPA PLC
15	Movement PLC to Lower Transfer Module
16	Movement PLC to Furnace door
17	Movement PLC to Lower Melt Cell
18	Movement PLC to Trolley secondary lift encoder
19	Movement PLC to Trolley door
20	Movement PLC to multicast
21	Furnace PLC to Furnace HMI
22	HEPA PLC to HEPA HMI
23	Trolley secondary lift motor to multicast
24	Trolley primary lift motor #2 to multicast
25	HEPA PLC to HEPA fan motor #2
26	HEPA PLC to HEPA fan motor #1
27	Lower HEPA to HEPA PLC
28	HEPA HMI to multicast
29	Movement PLC to OGE fan motor #2
30	Movement PLC to Trolley movement motor
31	Movement PLC to HEPA PLC
32	Movement PLC to Lower Transfer Module
33	Movement PLC to Upper HEPA

34	Movement PLC to Lower Module Controller
35	Movement PLC to Upper MOdule Controller
36	Movement PLC to Trolley
37	Movement PLC to OGE fan motor #1
38	Trolley primary lift motor #1 to multicast
39	Trolley secondary lift motor to multicast
40	PLC Programming/Troubleshooting laptop to multicast
41	Movement PLC to multicast
42	Trolley primary lift motor #2 to multicast
43	Trolley primary lift motor #1 to multicast
44	PLC Programming/Troubleshooting laptop to broadcast
45	PLC Programming/Troubleshooting laptop to multicast

Using actual OT traffic captured as discussed in section 2.1 Network Data Description, Association Rules Mining using the A Priori technique was applied by developing MATLAB scripts. A subset of the results is shown in Table 5 below. Note that the automatically generated text can be used by an LLM or other methods to translate these results for operators and analysts.

Table 8 – Subset of Results from Association Rules Mining

	Antecedent	Consequent	Support	Confidence	Lift
1	0	[0,1,2]	0.1952	0.9986	5.1171
2	[0,1]	[0,1,2]	0.1952	0.9997	5.1226
3	[0,2]	[0,2,3]	0.1953	0.9997	5.1187
4	[0,1,2]	[0,1,2,3]	0.1951	0.9997	5.1242
5	[1,4]	[1,4,6]	0.1161	0.6122	5.2742
6	[1,5]	[1,5,18]	0.1323	0.6768	5.1156
7	[1,2,4]	[1,2,4,6]	0.1150	0.6103	5.3070

Association Rules shown in table format. The apparent closed set of packets labeled 0, 1, 2, and 3 are highlighted. Note the consistent values of support, confidence, and lift.

Based on Rule #1 when the REI PLC sends a packet to the Dolly motor drive there is a 99.86% probability that in addition the REI PLC sends a packet to the Trolley primary lift motor #2 and the REI PLC sends a packet to the Trolley secondary lift motor
Based on Rule #2 when the REI PLC sends a packet to the Dolly motor drive and the REI PLC sends a packet to the Trolley primary lift motor #2 there is a 99.97% probability that in addition the REI PLC sends a packet to the Trolley secondary lift motor
Based on Rule #3 when the REI PLC sends a packet to the Dolly motor drive and the REI PLC sends a packet to the Trolley secondary lift motor there is a 99.97% probability that in addition the REI PLC sends a packet to the Trolley primary lift motor #1
Based on Rule #4 when the REI PLC sends a packet to the Dolly motor drive and the REI PLC sends a packet to the Trolley primary lift motor #2 and the REI PLC sends a packet to the Trolley secondary lift motor there is a 99.97% probability that in addition the REI PLC sends a packet to the Trolley primary lift motor #1

Note that observation of the highlighted rows in Table 5 indicate that there is relatively high support – approximately 20% which means that 1 out of every 5 packets belongs to this group of packets (labeled 0, 1, 2, and 3). Future data driven methods will leverage these types of relationships.

Once the network topology is known (or at least estimated) then methods developed by Vann, et. al., can be used to detect patterns within the other fields of the packets even though they are not explicitly known. It is worth noting that something very similar is done automatically by large language models when they are appropriately trained on a data type: they learn the co-occurrence probabilities of elements within single messages and between successive communications.

3.3.2 Alternative Packet Pattern Mapping Techniques

As an alternative to Association Rule Mining, a similar technique is discussed in 3.1.1 – Packet Bitmaps can be developed by displaying the link labels versus time. As shown in the human-derived pattern identification process below, there are findable patterns in the data that could be used for pattern-based anomaly detection, state-based analyses, and reverse engineering of processes. A hierarchical analysis combined with any a priori knowledge can thus yield rich sources of information.

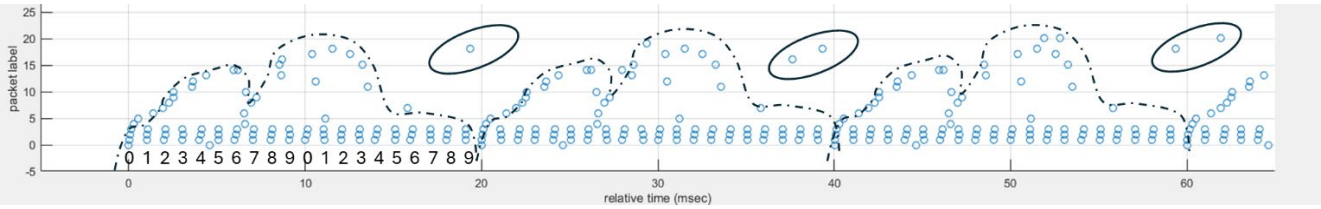

Figure 24 – Example of generation of 2D images from simple packet labeling schemes that lend themselves to multiple levels of AI/ML

3.4 Identifying sensor measurement trends and inferring process states

As shown in [Moor18] and compatible with [Vann17], sensor data values – whether derived from external sensors or internal network traffic – can be used to identify states within an unknown or poorly documented process. By using physics-driven governing equations (such as the product of speed and time yields distance traveled) these sensor value trends can be used to bootstrap knowledge of a system, determine a state (e.g., acceleration or heating) and then look for either inputs or outputs that are inconsistent with the state of the process.

4.0 Conclusions

Our initial review of the Melter shows that it is sufficiently complex to serve as a stand-in for studying the cybersecurity of SMRs. The system contains many similar 3S design considerations, and there are opportunities for testing our hypothesis (that events of interest can be detected in the cyber data from similar systems). Although the Melter is not a closed-loop autonomous system, neither will be most SMRs, and the analytics developed with data from this system should readily translate to the real subjects of our investigation. To extend the study to autonomous SMRs we must assume that all that can be observed about the system can be found in examining the cyber data. This assumption serves us for now but will need to be tested when we transition to real SMRs. In any case, where observations of component failures or variances cannot be reported in the cyber data, we will have to find reliable indicators that are so reported.

5.0 Future Work

This section outlines several areas where we hope to conduct follow-on work in out years:

1. **Full-cycle testing:** This year our collection efforts were often stymied by equipment breakage and continual design changes to the Melter. At this writing, the furnace stirring operation is still not reliable, so we have yet to collect data from a full run-through of the system. The operators graciously exercised the system for us so we could collect data on various functions of the components for this study. In the following year we expect to witness the design of a new Melter and its testing and we will collect data from one or more melt operations. This is a task for FY26.

2. **Off-normal testing:** During the acceptance test, we fully expect off-normal conditions to arise. These will give us a chance to determine whether our approach can detect system instabilities and predict Byzantine failures. However, it is our intention to request that abnormal conditions be purposely induced so that the testing can be more robust and to test our metric. While ideally for our work it would be helpful to introduce cyber attacks, we believe permission to attack an operational system is unlikely to be approved. Operational technology is often brittle and easy to irreversibly damage even by subtle cyber interventions. The risk of ruining components prior to a high-profile system delivery is too great. We hope to consider including realistic cyber attacks in the acceptance testing for future versions of the Melter system where the design is more stable and there may be more time to consider these possibilities. For this reason, in FY26 we will shift focus to data collection from Argonne's METL testbed where cyber attacks may be tested against an operating simulation.

3. **Expanded data storage:** Currently, the data collection PC is only able to store two to four hours of full packet capture data. Because melt operations last for about 18 hours it is not currently possible to collect and store all this data. Therefore, we may need to purchase larger hard drives and possibly a management PC or collection device to facilitate long-term collection. The Melter is expected to be deployed permanently to the host country and to have an operating campaign lasting approximately six months. Ideally, we would like to have a data feed back to the laboratory for online diagnosis, but this would change the security posture of the system, which currently relies on an air gap for safety. The future versions of the Melter are being designed now, and in the coming years we expect to see many of them in operation. We hope to have cybersecurity concerns influence the design of future versions so we can safely enable this analytics connection without introducing hazards. This task may be undertaken in FY26 or later.

4. **Online analytics:** For the initial work, we are doing analysis after the fact. Ideally, this analysis would be done online and in near real time. Online analytics will require compute power collocated with the data sources because it is unlikely that all the internal cyber data would be passed to the outside world for analysis. Being collocated with a reactor, there would be no lack of power, but over the long run, storage will be an issue. Additionally, the analytics themselves need to be checked for efficacy periodically because it is possible that data drift over time will make the models less effective. This task is an out-years project.

5. **Expanded data sources:** In the future we would like to include more data sources including the operator log, video, and other components that are not online currently. Although large language models natively can accept any kind of textual input, video data and tabular data will need special consideration. Expanding the types of data used will also require further consideration and expansion of the kinds of models used and the reasoning applied. As no two plants are identical, it is possible that no single analytics solution will accommodate all of them. We plan to start with the Melter and then port our analytics to another nuclear system, perhaps the METL or a research reactor if available, and determine how portable they are. We are aiming for maximal portability of the models, but all systems will individually need a "burn-in" time to ensure the analytics work properly. Expansion is an out-years project.

6. **State-Based Anomaly Detection Capability Using Out-of-Band Sensing:** As discussed in Section 2.4 Out-of-Band Sensing, during FY26 a capability for detecting anomalous behavior based on using AI/ML techniques and out-of-band sensing will be explored.

6.0 References

[Darknet] DarkNet: Lighting up a secure grid communication network, https://www.ornl.gov/blog/darknet-lighting-secure-grid-communication-network

[Bogg2020] Boggs, N., So, B. and Cui, A., 2020, May. Remote attestation of host-based defense via optical channel. In *Cyber Sensing 2020* (Vol. 11417, pp. 16-26). SPIE.

[NSA/CISA] NSA and CISA report on Control System Defense: Know the Opponent, https://media.defense.gov/2022/Sep/22/2003083007/-1/-1/0/CSA_ICS_Know_the_Opponent_.PDF

[Moor2019] Moore, M.R. and Vann, J.M., 2019, January. Anomaly detection of cyber physical network data using 2D images. In *2019 IEEE International Conference on Consumer Electronics (ICCE)* (pp. 1-5). IEEE.

[Moor2018] Moore, M.R., 2018. Machine learning for cyber-physical system protection of can bus enabled vehicles using driver state discovery (Doctoral dissertation, Tennessee Technological University).

[Moor2017] Moore, Michael R., Robert A. Bridges, Frank L. Combs, Michael S. Starr, and Stacy J. Prowell. "Modeling inter-signal arrival times for accurate detection of can bus signal injection attacks: a data-driven approach to in-vehicle intrusion detection." In *Proceedings of the 12th annual conference on cyber and information security research*, pp. 1-4. 2017.

[Tu2024] Tu, J.H., Eichler West, R.M., Ellwein, E.J. and Vann, J.M., 2024. *Augmented Human Analysis (AHA)* (No. PNNL-35888). Pacific Northwest National Laboratory (PNNL), Richland, WA (United States).

[Vann2017] Vann, J.M., Karnowski, T.P., Kerekes, R., Cooke, C.D. and Anderson, A.L., 2017. A dimensionally aligned signal projection for classification of unintended radiated emissions. *IEEE transactions on electromagnetic compatibility*, *60*(1), pp.122-131.