ORNL/TM-2024/3239
CRADA/ NFE-21-08826

CRADA Final Report: CRADA Number NFE-21-08826 with Sentinel Devices LLC

Forrest Shriver

Date January 9, 2024

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE US DEPARTMENT OF ENERGY

ORNL/TM-2024/3239
CRADA/ NFE-21-08826

CRADA Final Report: CRADA Number NFE-21-08826 with Sentinel Devices LLC

Forrest Shriver

Date Published:
January 9, 2024

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

Approved for Public Release

1. Abstract

Industrial equipment is a critical component of virtually all at-scale manufacturing and infrastructure. Broadly speaking, modern equipment is predominantly controlled using hardened digital controllers – computers designed to be able to operate continuously for sometimes extremely long periods of time, with little to no maintenance. Due to the nature of how these digital controllers have evolved, they are solely optimized to execute a single task, and do not have the capabilities or resources to monitor or analyze their internal state beyond simple execution of their program. As a result, for many "common sense" situations where individual data points can be easily determined to be out-of-normal, the controllers are unable to identify these incorrect operational modes unless a human has explicitly programmed in detection of this degradation. This project seeks to develop an AI/ML system which can identify incorrect or anomalous trends in industrial data streams, of exactly the kind that would be produced and seen by these digital controllers, with a minimal amount of computing resources. The benefits produced by developing this system would ultimately be self-monitoring and self-reporting infrastructure, capable of identifying and alerting humans to issues as soon as they happen, potentially long before they have the chance to impact the industrial process.

2. Statement of Objectives

Modern industrial data analytics is focused on working on large amounts of data collected from across different areas of a facility, drawing upon a centralized data historian for information and assuming that the historian has cleaned and logged the appropriate data streams. This data must be worked on by human experts, with feature engineering and normalization done afterwards. A patent-pending anomaly detection technology has been licensed from the University of Tennessee which enables anomaly detection at the "edge" (digital controller level) of an industrial network, however this technology is limited by the need for a curated database of valid states. Since many industrial processes are unique, it is effectively impossible to pre-generate this database. The goal of this project will be to develop technology to autonomously discover, classify and store trends in cyber-physical data at the controller level in near-real-time. This project will develop a method for performing this "on-the-fly" (OTF) database construction for industrial data streams, where these streams may be subject to adverse conditions (high variability, missing data, high/low frequency of change). The results of this effort will be a fundamental change in how industrial data collection is performed, enabling the deployment of novel technologies at the industrial edge for autonomous, real-time anomaly detection and cybersecurity monitoring.

The research task to be carried out under this CRADA is to develop, test and benchmark an OTF database construction technology on an ORNL testbed. This will involve the examination and use of testbed generated datasets under both normal and transient operations. The technology development will be incrementally modified, integrated, and benchmarked with the anomaly detection technology developed by the Sentinel Devices team, to satisfy timing, accuracy and hardware constraints. Following this initial R&D phase, the selected OTF database construction technology will be integrated and implemented alongside the anomaly detection technology in prototype hardware and will be evaluated in physical/simulated industrial environments. The feedback from this evaluation will be used to iteratively improve and refine the OTF database construction technology as well as existing anomaly detection technology. Following this testing and evaluation phase, a "rapid real-world prototyping" phase will follow where the prototype technology will be evaluated on existing testbeds and equipment in live industrial environments. The feedback from this prototyping will be used to further improve the OTF database construction and anomaly detection technologies, this time in the context of long-running and harsh real-world conditions.

Objective 1: OTF database construction technology R&D

Description

The first milestone will be the completion of an OTF database construction software algorithm. This will correspond to a robust analysis of available testbed datasets and the performance of the selected technology in the context of data storage size, runtime requirements, potential weaknesses and performance when integrated with the existing anomaly detection technology.

Task 1 Summary:

Task 1.1 – Development of OTF database construction technology using existing or developed datasets from ORNL testbed.
Task 1.2 – Integration and evaluation of developed OTF database construction technology with existing anomaly detection technology.

Objective 2: Deployment and evaluation of integrated prototype on existing physical and simulated industrial testbeds

Description

The second milestone will be integration of the selected OTF database construction software algorithm with the developed anomaly detection technology, and the deployment of this technology on existing and simulated industrial testbed, as described in the initial R&D effort from the first milestone. This will correspond to a robust analysis of the performance of the prototype implemented on target hardware and the ability of the integrated technology to perform its functions under different conditions. During this work potential weaknesses in the integrated technology with regards to common industrial constraints, such as "maxed out" digital controllers and flooded communication networks, will also be analyzed to improve the theoretical efficacy of the prototype.

Task 2 Summary:

Task 2.1 – Evaluation of integrated prototype on testbeds under normal and transient conditions.
Task 2.2 – Evaluation of integrated prototype under industrial scenarios.

Objective 3: Integration and deployment of prototypes in a near real-world industrial system and/or testbed

Description

The third milestone will be to validate and verify functional performance of the prototype in a near real-world industrial system. This may include the collection and evaluation of real-world data.

3. Benefits to the Funding DOE Office's Mission

Industrial processes globally represent some of the most intensive and impactful processes carried out by humans, both producing goods and services to support society at large as well as impacting the environment. Making sure that these facilities and processes operate efficiently is key to ensuring not only the safety and security of U.S. critical infrastructure, but also in reducing the environmental and ecological impact that these facilities ultimately have by improving their performance and reducing the occurrence of accidents and other issues. This research sought to significantly advance the intelligence available within common industrial infrastructure, which will be the foundation for reaching carbon neutrality and energy independence goals of the EERE and U.S. Government now and in the future.

4. Technical Discussion of Work Performed by All Parties

The goal of this project was to develop small, domain-specific AI/ML technology for identifying and analyzing anomalies in industrial time series data, with the resulting technology being capable of operating on a minimal compute budget with no connection to any kind of external server or cloud. There are multiple reasons why developing "zero-cloud" technology such as this is desirable: one is that security-sensitive environments such as power plants or defense manufacturing typically require that there be few or no connections between critical control networks and the outside world. Another is that developing technology which is asset-specific, and capable of self-training, means enabling specific per-machine monitoring and diagnostics which is not feasible using current centralized & cloud-focused technologies. Developing AI/ML technology, which is able to operate on minimal compute resources, such that the AI can be readily deployed as a piece of hardware alongside the equipment being monitored, enables the scalable and secure monitoring of industrial equipment without compromising any of the existing security policies of critical infrastructure.

To do this, a multi-pronged research approach was adopted, where real-world hardware testing and validation was performed in parallel to algorithms and software framework development. The result was a strong and performant software runtime which combined multiple different pieces of software, including multiple algorithms, into a single cohesive system that can be used to interactively & rapidly train an AI/ML model using only information collected from a specific machine or piece of equipment. This runtime is able to collect data from a large subset of industrial digital controllers existing today, thereby making it applicable to most industrial facilities and processes that are currently in operation.

Participant was responsible for identifying and integrating all algorithms and software modules which would ultimately be used within the software runtime, and in implementing this runtime such that it was functional on the hardware being targeted for deployment. Participant conducted several surveys of data collection technologies, data storage technologies, and software products, identified commonalities and differences in these products, and tested/developed initial proof-ofconcept implementations of these integrations to prove they could work and to identify any weaknesses in their implementation/usage. Participant then implemented multiple algorithms for anomaly detection and analysis, and also developed a software framework and testing harness which was used to iterate and operationalize these algorithms on the target hardware. Finally, Participant deployed the software runtime on multiple hardware devices which were deployed in the field, and closely monitored and experimented with these "live" iterations to identify issues in connectivity, deployment, and reliability under continuous operation. Below is a screenshot taken from a presentation on the results of the developed system, run on openly-available data taken from a train – the fault identification capabilities of the technology are displayed, where the system is able to pull out five relevant signals that the system thinks the operators should look at first/are most "suspiciously" deviating from normal operation.

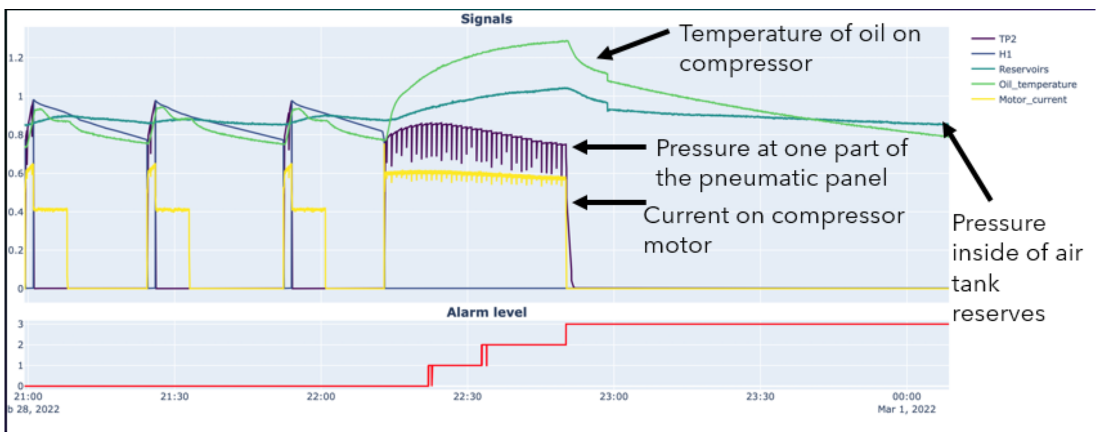

ORNL, as the Contractor, helped facilitate multiple opportunities for testing and validation of the technology. In particular, Contractor provided two testing facilities for deployment and evaluation of the technology: a water loop simulating core components of a high-performance computer cooling system, with many system complexities inherent, and a real physical water loop containing sensors and instrumentation similar to those found in a traditional power plant. Below is a picture from this deployment, where the device (shown in orange) was connected into the control system to collect and monitor signals. Additionally, the Contractor provided significant staff time to facilitate and enable these deployments, as well as de-risk initial exploration of viable deployment strategies and frameworks for enabling communication in an industrial environment. The Contractor also provided multiple sets of data for testing and validation of the technology's underlying algorithms, which are still being actively worked on and evaluated for the potential for future publication work.

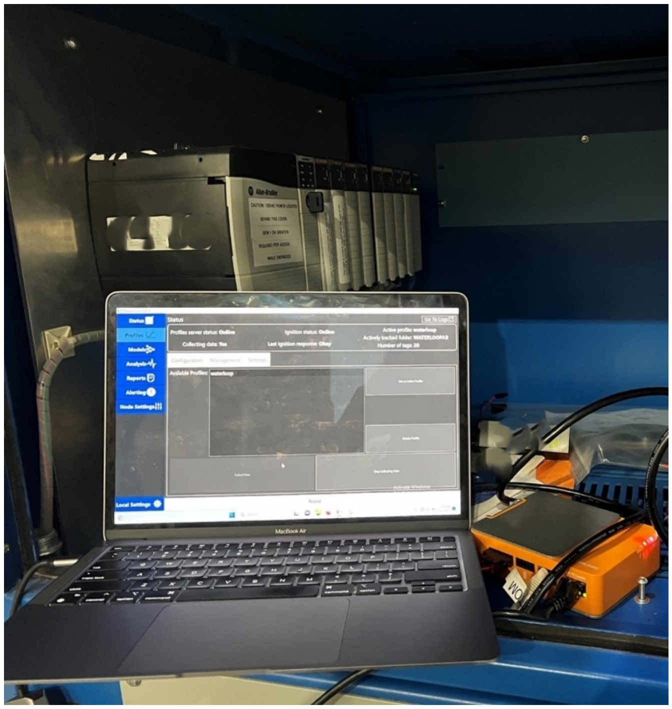

5. Subject Inventions (As defined in the CRADA)

None

6. Commercialization Possibilities

During this time, Participant received several awards and recognition for cutting-edge work in equipment monitoring, including receiving the "Defense Innovation Award 2022" for applications of the technology to defense manufacturing.

7. Plans for Future Collaboration

Contractor's tools and facilities enabled the project to rapidly validate and invalidate hypotheses around what real operational deployments would look like, thus greatly informing the development work and giving strong early feedback on what areas needed the most focus. Participant and Contractor are evaluating several opportunities that would allow for continued development.

8. Conclusion

This project greatly impacted Participant's understanding of operations and maintenance in manufacturing and the energy sector and gave Participant a strong basis for future work applying AI/ML and edge compute to these areas.