ORNL/TM-2023/3093

Unified Software Architecture for Advanced Materials and Manufacturing Technologies Data Management and Processing: FY 2023 Multidimensional Data Correlation Platform

Stephanie Cooper
William Halsey
Vladimir Orlyanchik
Vincent Paquit
Luke Scime
Zackary Snow
Amir Ziabari

September 2023

OAK RIDGE
National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

ORNL/TM-2023/3093

Advanced Materials and Manufacturing Technologies

UNIFIED SOFTWARE ARCHITECTURE FOR ADVANCED MATERIALS AND MANUFACTURING TECHNOLOGIES DATA MANAGEMENT AND PROCESSING: FY 2023 MULTIDIMENSIONAL DATA CORRELATION PLATFORM

Stephanie Cooper
William Halsey
Vladimir Orlyanchik
Vincent Paquit
Luke Scime
Zackary Snow
Amir Ziabari

September 2023

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

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

AI	artificial intelligence
AM	additive manufacturing
AMMT	Advanced Materials and Manufacturing Technologies
API	application programming interface
ASTM	ASTM International, formerly American Society for Testing and Materials
CRADA	Cooperative Research and Development Agreement
DMSCNN	dynamic multilabel segmentation convolutional neural network
DOE	US Department of Energy
DP	digital platform
EDM	electrical discharge machining
IFOV	instantaneous field of view
ITAR	International Traffic in Arms Regulations
JSON	JavaScript object notation
LPBF	laser powder bed fusion
LWIR	long-wave infrared
MDDC	multidimensional data correlation
MDF	Manufacturing Demonstration Facility
MP	megapixel
M:M	many-to-many
NDI	nondestructive inspection
NIR	near-infrared
OD	outer diameter
ORNL	Oak Ridge National Laboratory
TCR	Transformational Challenge Reactor
TI-NIR	temporally integrated near-infrared
UV	ultraviolet

GLOSSARY OF TERMS

The following key terms used throughout this report are defined here for disambiguation. Where possible, the additive manufacturing (AM) terminology used herein complies with ISO/ASTM 52900:2021, *Additive manufacturing - General principles - Fundamentals and vocabulary* [1].

artificial intelligence (AI): An umbrella term referring to any computer algorithm that makes decisions intended to mimic those made by a human.

database: The collection of metadata that comprise the digital threads for all the components manufactured at the Manufacturing Demonstration Facility (MDF) at Oak Ridge National Laboratory (ORNL). Strictly, the heavier process data (e.g., in situ images) are not stored in the database; they are stored within the designated file systems, and the database only moderates and facilitates data upload and retrieval by using application programing interfaces.

deep learning: A class of machine learning algorithms in which features are learned instead of designed.

digital platform (DP): The cyber-physical infrastructure under development at the MDF that is enabling novel design and qualification paradigms for advanced manufacturing.

digital thread: The totality of the design intent information, data, and metadata collected during the fabrication of a component. The digital thread enables the instantiation of a corresponding digital twin.

Damara Tern: A web-based client developed at the MDF to facilitate metadata collection for the different operations with the goal of creating a digital thread for each manufactured component. This application is set to become the primary interface between most users and the DP database.

digital twin: A computer representation and model of a real object, assembly, or system. This model is updated based on data collected from its physical twin; this model contrasts with a computer model that relies solely on aggregated or representative data from many identical components.

ex situ data: Data collected outside of some manufacturing operation. For example, post-build x-ray computed tomography data.

integrated computational materials engineering: A methodology for designing material and component design that links experiments and simulations across multiple length and time scales.

in situ data: Data collected during some manufacturing operation. For example, sensor data collected in the AM process.

machine learning: a subset of AI algorithms that are trained on data to generate models capable of performing complex tasks

operation: Each manufacturing process is decomposed into a sequence of operations. A digital thread is merely a list of all operations that were performed to fabricate the component, along with links to the associated data and metadata.

Peregrine: An ORNL-developed software tool designed to provide a comprehensive suite of data collection, analysis, and visualization capabilities for powder bed AM systems.

Simurgh: An ORNL-developed software tool for performing x-ray computed tomography reconstructions using deep learning.

software tool: Any custom software used to facilitate a digital workflow or specific operation.

trackable: A physical or digital component that can undergo an operation. The term applies to any component that can be catalogued along the digital thread, including AM materials, final or in-process components, and digital models.

ABSTRACT

This report details the various digital manufacturing activities ongoing at Oak Ridge National Laboratory (ORNL) as part of the Advanced Materials and Manufacturing Technologies (AMMT) Program. The AMMT Program is exploring a data-driven approach to demonstrate the use of additive manufacturing for the fabrication of components for nuclear applications, with the goal of providing a greater understanding of manufacturing quality outcomes that would pave the way toward the development of standards for certification and qualification. The objective of this work package is to establish a digital manufacturing discipline common to all participants of the AMMT Program to improve the performance, reliability, and lifetime of nuclear components. As part of this effort, a unified software architecture will be developed for AMMT data management and processing, the digital platform will be deployed across AMMT participants' facilities, and pedigreed data sets will be generated in a common format across multiple labs and facilities. To this end, the multidimensional data correlation work package has focused on three activities during FY 2023. First, the Manufacturing Demonstration Facility Digital Tool was overhauled to better serve the needs of the AMMT Program. Next, multiple laser powder bed fusion systems at the Manufacturing Demonstration Facility were upgraded to a common sensor package for collecting comparable in situ data across machines. Finally, various improvements relevant to the AMMT Program were implemented in the ORNL-developed software tool, Peregrine. **This report marks the completion of FY 2023 milestone M3CR-22OR0403051: Report Describing the Architecture of the Digital Platform to Support AMMT Activities.**

1. INTRODUCTION

Additive manufacturing (AM) continues to promise opportunities for optimized component designs, localized microstructural control, and on-demand manufacturing. However, concern regarding the variability observed in AM material properties and machine-to-machine performance has, to date, prevented widespread adoption of the technology in nuclear and other risk-averse industries. Furthermore, the complex geometries often associated with AM complicate postbuild nondestructive inspection (NDI), presenting additional challenges for qualification.

Despite these challenges, the layer-by-layer nature of AM also provides an unprecedented opportunity for in situ process monitoring that can bolster traditional NDI techniques. For example, in situ images collected on a layer-wise basis can provide detailed information related to part quality in a manner that is robust to material and component size, which are both historical limitations for traditional NDI. Multiscale, physics-based simulations of thermomechanical phenomena and microstructural development may also supplement these data streams, presenting a complete digital data package of the manufacturing process. In this sense, AM and other advanced manufacturing technologies can produce extensive data sets containing valuable information pertinent to component quality at every stage of the manufacturing workflow not afforded by traditional NDI. Through the collection, structuring, and analysis of these disparate data streams, this report argues that such data can be used to understand, optimize, and validate advanced manufacturing processes in a manner congruent with the requirements of qualification and certification. To date, no process- and material-agnostic methodologies exist to rapidly develop standards supporting the qualification of additively manufactured components, and there exists a need for pedigreed data sets for data-driven qualification of AM because existing data sets are dispersed and not recorded in standard formats.

The Advanced Materials and Manufacturing Technologies (AMMT) Program is exploring a data-driven approach to demonstrate the use of AM for the fabrication of components for nuclear applications, with the goal of providing a greater understanding of manufacturing quality outcomes that would pave the way toward the development of standards for certification and qualification. To be successful, this approach requires the creation of large pedigree data sets that would serve as references to validate data analytics solutions for component quality control, failure analysis, and material properties prediction. To build such a database across all activities within the program, the program must establish a digital discipline common to all participants of the AMMT Program that defines the protocol for data collection and processing at every step of the manufacturing process, during subsequent characterization and testing phases, and ultimately throughout a component's life cycle. The collection of such streams for a given part is defined as its *digital thread*.

Leveraging preliminary work conducted as part of the Transformational Challenge Reactor (TCR) Program, the objective of this work package is to establish a digital manufacturing discipline common to all participants of the AMMT Program to improve the performance, reliability, and lifetime of nuclear components. This task will be accomplished by (1) reinforcing the existing Digital Tool architecture at the Manufacturing Demonstration Facility (MDF) to capture additional information pertaining to the manufacturing and testing processes considered by the AMMT Program; (2) providing a unified data architecture deployable at the AMMT members' facilities for data management consistency; (3) generating guidance on how data and metadata generated by this program should be formatted to integrate the main database, with data coming from AMMT members or Nuclear Energy University Program participants; and (4) developing new modules for software tools that extend the current predictive capabilities. The collection and storage of manufacturing process data within this architecture will facilitate data tracking, processing, and exchange between US Department of Energy (DOE) labs; simplify the release of pedigreed datasets; and enable advanced data analytics to support AM component qualification. This new qualification approach, known as the *multidimensional data correlation* (MDDC) framework, will capitalize on the wealth of digital manufacturing data; integrated computational materials engineering tools; artificial intelligence (AI) tools; and accelerated, high-throughput testing and characterization techniques.

To this end, the MDDC work package has focused on three activities during FY 2023. First, the MDF Digital Tool was overhauled to better serve the needs of the AMMT Program. The new framework, called *Damara Tern*, will succeed the previous MDF Digital Tool and related relational database initially deployed under the TCR Program, and its primary function is to structure and organize the metadata collected along the digital thread. It also serves as a library, allowing users to save, explore, and retrieve data associated with each operation and component involved in the manufacturing process. Next, multiple laser powder bed fusion (LPBF) systems at the MDF were upgraded to a common sensor package for collecting in situ data across machines, allowing for better comparisons of process stability and material quality across different builds, alloys, machines, and printing facilities. The Digital Handbooks, a set of guidelines for capturing digital threads for a given manufacturing operation, were also updated to reflect these hardware modifications. Finally, various improvements relevant to the AMMT Program were made to Peregrine, an Oak Ridge National Laboratory (ORNL)-developed software tool for advanced data analytics of powder bed AM processes. These improvements include updates to the real-time data acquisition methodology and the creation of the dynamic multilabel segmentation convolutional neural network (DMSCNN), which is the deep learning algorithm that forms the foundation of Peregrine.

This report is organized as follows. First, an overview of the MDDC approach is presented in Section 2, followed by details of the design of the new digital tool, Damara Tern, in Section 3. Section 4 discusses the various sensing and instrumentation upgrades made, and Sections 5 and 6 detail relevant changes made to the Peregrine software tool. Finally, detailed tables related to the Damara Tern database and updated digital handbook entries for AMMT machines can be found in the appendices.

2. THE MULTIDIMENSIONAL DATA CORRELATION APPROACH

The MDDC framework under development at ORNL seeks to enable the qualification of critical components based on *instance-specific* data. This method contrasts with traditional qualification schemas, which rely on data aggregated from the characterization and historical performance of many identical components combined with a restrictive set of *locked down* manufacturing processes and material feedstocks. This paradigm shift is ultimately necessary for highly regulated industries to fully leverage the latest advanced and digital manufacturing processes, such as metal and ceramic AM. The MDDC framework supports these efforts at multiple levels along a continuum of capabilities. First defined for the TCR Program, Table 1 updates the descriptions of four capability levels [2] as they relate to the AMMT Program. The TCR Program began in 2019 with data-driven qualification capabilities for powder bed AM processes at Level I [3], [4], and the AMMT Program has been developing capabilities at Level III since 2021 [5], [6].

Table 1. Descriptions of the capability levels at which the MDDC framework supports instance-specific qualification of additively manufactured components. The AMMT Program is currently developing capabilities at Level III for powder bed AM systems.

I. **Record keeping**	**II.** **Enhanced understanding**	**III.** **Property prediction**	**IV.** **Accelerated design**
• Data and metadata are recorded primarily for provenance purposes. • Data are painstakingly reviewed manually, layer-by-layer, to identify problems. • In situ sensor data are generally of low quality and minimal volume. • Historical data sets are difficult to access and may be inconsistently formatted.	• Data and metadata are consistently recorded for every build and are used to better understand the process. • AI is used to automatically analyze in situ data and identify anomalies and flaw indications. • In situ sensor data are of high quality and require large storage volumes. • Sensors and algorithms are sufficiently robust to determine if a given build was printed under nominal conditions.	• In situ data are spatially registered with ex situ characterization data at scale. • Process simulations are performed at scale and linked to in situ data. • The correlation between anomalies (indications) and flaws is understood with statistical methods. • AI and physics-based modeling are used to predict local material properties. • Physics-based models are used to predict part performance based on the in situ data (i.e., is a flaw truly a defect).	• Each part's digital thread can be used to simulate its digital twin. • In situ data, process simulations, and local property predictions are leveraged during the design process. • AI is used to automatically iterate both the part design and the manufacturing process steps.

In addition to achieving these technical milestones, the AMMT Program must address the depth, breadth, and scale of the proposed qualification paradigm (see Figure 1). Each of these three dimensions increases the size of the problem space and are discussed in the following paragraphs.

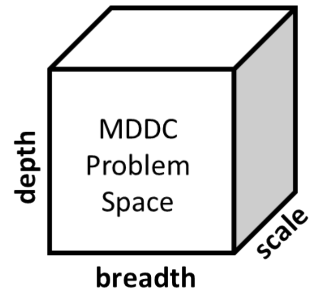

Figure 1. Visualization of the dimensions of the MDDC problem space. *Depth* refers to the technical capabilities of a given digital manufacturing system, and *breadth* refers to the application of technical achievements from each focus area across multiple systems. *Scale* refers to the ability to implement MDDC concepts in production settings with the requisite depth and breadth.

Achieving each subsequent level of data-driven component qualification requires research and development efforts to increase the technical depth across multiple aspects. Key data aspects include (1) AI algorithms [7]; (2) high-resolution, multimodal in situ sensing [8]; (3) process modeling and simulation; (4) nondestructive evaluation techniques [9]; and (5) cyber-physical infrastructure [10]. Note that transitioning between different capability levels requires different amounts of effort across each topical area. For example, although improving in situ sensing capabilities is extremely important to move from Level I to Level II, advances in spatial registration of nondestructive evaluation results are critical to move from Level II to Level III. Notably, creating a robust cyber-physical infrastructure is necessary to bind together all the other digital capabilities. Figure 2 summarizes the relative efforts required in each area along the MDDC capability continuum.

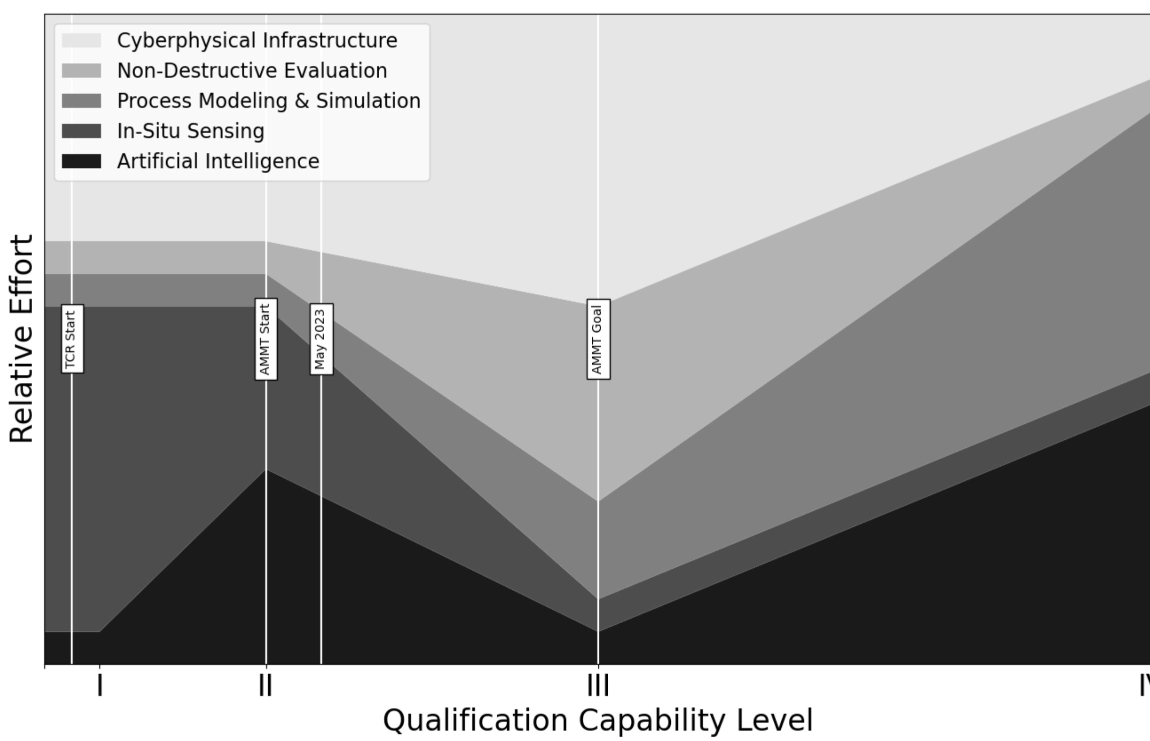

Figure 2. Representation of the order-of-magnitude levels of research effort anticipated in each technical area to move between MDDC capability levels. As an example, the research focus shifts from in situ sensor development at the early levels to AI and modeling at the later levels, and cyber-physical infrastructure remains a focus throughout the entire development process.

The breadth of the problem relates to the application of technical achievements from each focus area across multiple systems. Because this dimension is extremely resource-intensive, ORNL has focused on developing generalizable software tools, such as Peregrine, which are printer-, sensor-, and material-agnostic. This approach enables the rapid transfer of advancements between different printer–sensor–material combinations but requires access to many diverse scenarios for development and testing of each new capability. To that end, Peregrine has been deployed on 19 powder bed printers at ORNL and is now being used across the US Government and multiple industry and academic partners. This platform allows technical capabilities developed at ORNL to be tested under unique conditions and used for a wide variety of use cases. As part of this effort, the DOE laboratories participating in the AMMT Program will share development of not only Peregrine but also the MDF Digital Platform. Figure 3 summarizes the deployment of Peregrine across the United States, as well as recent improvements to Peregrine and the MDF Digital Platform.

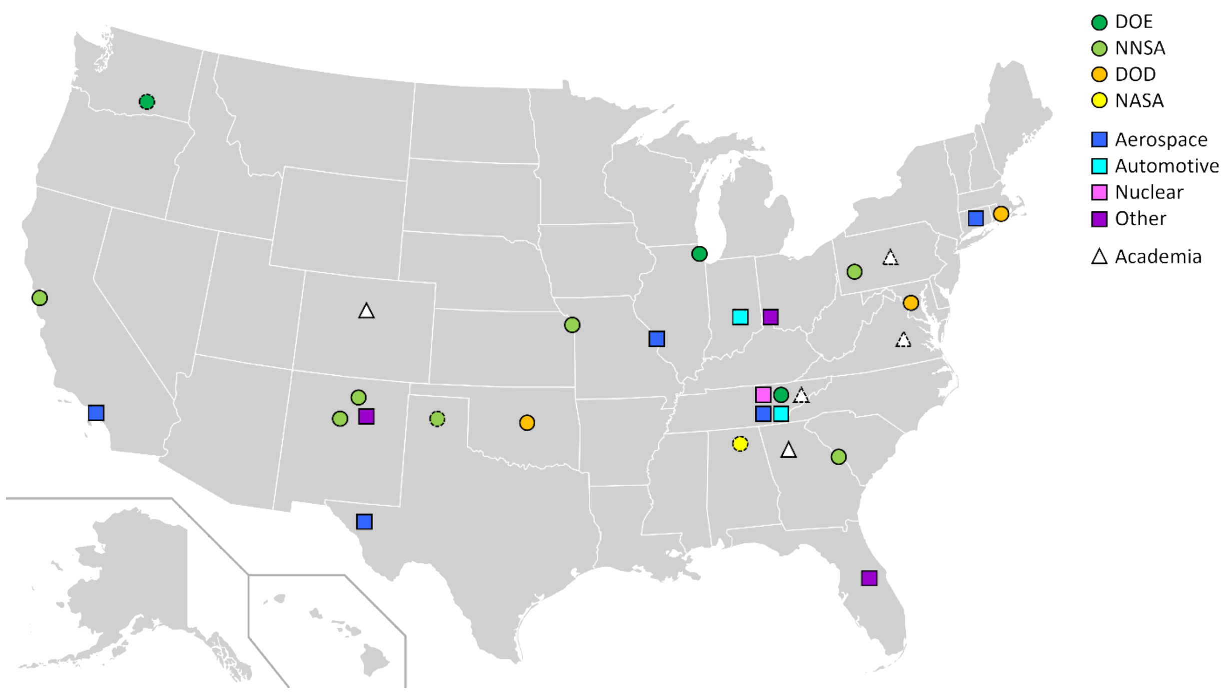

Figure 3. Peregrine licenses across the United States as of May 2023. Government facilities, industry partners, and academic institutions are represented as circles, squares, and triangles, respectively. Licenses that are still in progress but expected to be approved within FY 2023 are indicated with dotted lines.

Finally, implementation of the MDDC qualification paradigm at scale in industry-relevant production environments is considered the third dimension of this problem space. Under AMMT and other programs, ORNL is engaging with industry partners to transfer data-driven qualification technologies and best practices to the private sector. Beyond journal publications and conference presentations, engagement efforts have included (1) a digital handbook serving as a living document recording best practices learned during development of the MDDC, (2) research licenses of the Peregrine and Simurgh software tools, (3) Cooperative Research and Development Agreements and other technical collaborations focused on the role of data in powder bed qualification, (4) the release of open data sets [11], [12], (5) regularly providing the powder bed in situ sensing section of the annual *Wohlers Report* [13], (6) communication with standards organizations such as ASTM International, and (7) ongoing communication with the US Nuclear Regulatory Commission under TCR and now AMMT.

3. CONCEPT AND DESIGN OF THE NEW DIGITAL PLATFORM DATABASE AND INTERFACE FRAMEWORK

Damara Tern is the new framework developed to access the MDF Digital Platform data. Its primary function is to structure and organize the metadata collected along the digital thread and to serve as a library, allowing users to save, explore, and retrieve the data associated with each operation and component (i.e., trackable) involved in the manufacturing process. The framework succeeds the MDF Digital Tool and related relational database initially deployed under the TCR Program for accessing and exploring the data stored within the MDF Digital Platform. Although its predecessor successfully serves this function and enables component tracking during the manufacturing and testing processes, it relies on a printer-centric model that induces specificity and poses significant limitations in terms of extensibility and automation. To overcome these limitations, the Damara Tern framework relies on a generic and flexible operation-centric structure that accommodates a diverse range of machines, manufacturing and testing operations, and trackable components. Additionally, the framework aims to enhance search and view functionalities that reflect the digital thread for each manufacturing process.

3.1 IMPLEMENTATION CONSIDERATIONS

The main objective of the MDDC framework is to provide a unified data management platform that is accessible to all members of the AMMT Program. This framework should enable (1) nationwide tracking of physical components and digital assets across multiple sites and (2) the creation of a database adhering to the findable, accessible, interoperable, and reusable best practices to support the research and development activities of the program. In contrast to the implementation for the TCR Program, the platform described here requires using a data management strategy that preserves the integrity of the digital thread of each component across multiple physical locations. This important requirement resulted in the following considerations.

- **Database location:** Data will not be hosted at a single location, nor will there be replicas of the complete database at each national laboratory. Instead, AMMT participants will host a database of the information they will produce, and those databases will be interconnected using a cross-referencing mechanism described in Section 3.2.

- **Database access:** Initially, users will access the database using credentials provided by the host national lab. As new data exchange and linkage functionalities of the platform are developed, web-tools will be provided that allow direct access to data regardless of geographic location.

- **Information retrieval:** Rarely, users will need access to the entire database to work on scientific problems. In this case, they will retrieve subsets of the data based on scoped queries, for which application programming interface (API) functionalities will be developed to make intelligent and focused queries to the database.

- **Knowledge extraction and data visualization:** The selected Django framework provides access to numerous Python libraries for advanced data visualization and processing. Domain experts will be expected to create their own data processing pipelines to access the database and format the data for their specific applications.

3.2 OPERATION-TRACKABLE DATA MODEL

The data model employed for the development of the framework revolves around the concepts of operations and trackables, which are defined as follows:

- **Operation:** Any action performed via the use of a machine (such as print, cure, heat treatment, blue light scan, microscopy analysis, tensile test, etc.) or human interaction (procurement, annotation, registration, manipulation, etc.). Operations can be performed prior to, during, or after the manufacturing process itself to extend the data context.

- **Trackable:** Any physical or digital component that can be subjected to an operation (such as parts, builds, materials, etc.)

This model relies on the manufacturing process flow and provides the necessary framework for storing and organizing metadata collected along the digital thread. The database is implemented along this model, where each physical or digital component (trackable) undergoes a series of tests or actions (operations). Each operation supports the collection of metadata or substantial in situ data and can (but not necessarily) lead to the creation, alteration, combination, or transformation of the trackables.

By collecting the data at each operation level and preserving the traceability of relationships between the trackables and applied operations, it becomes possible to recreate the digital thread for each trackable. This pathway reflects the entire manufacturing process, accompanied by the comprehensive context of collected data and metadata gathered throughout the entire process.

Figure 4 gives an example of the operation/trackable-based model of the digital thread.

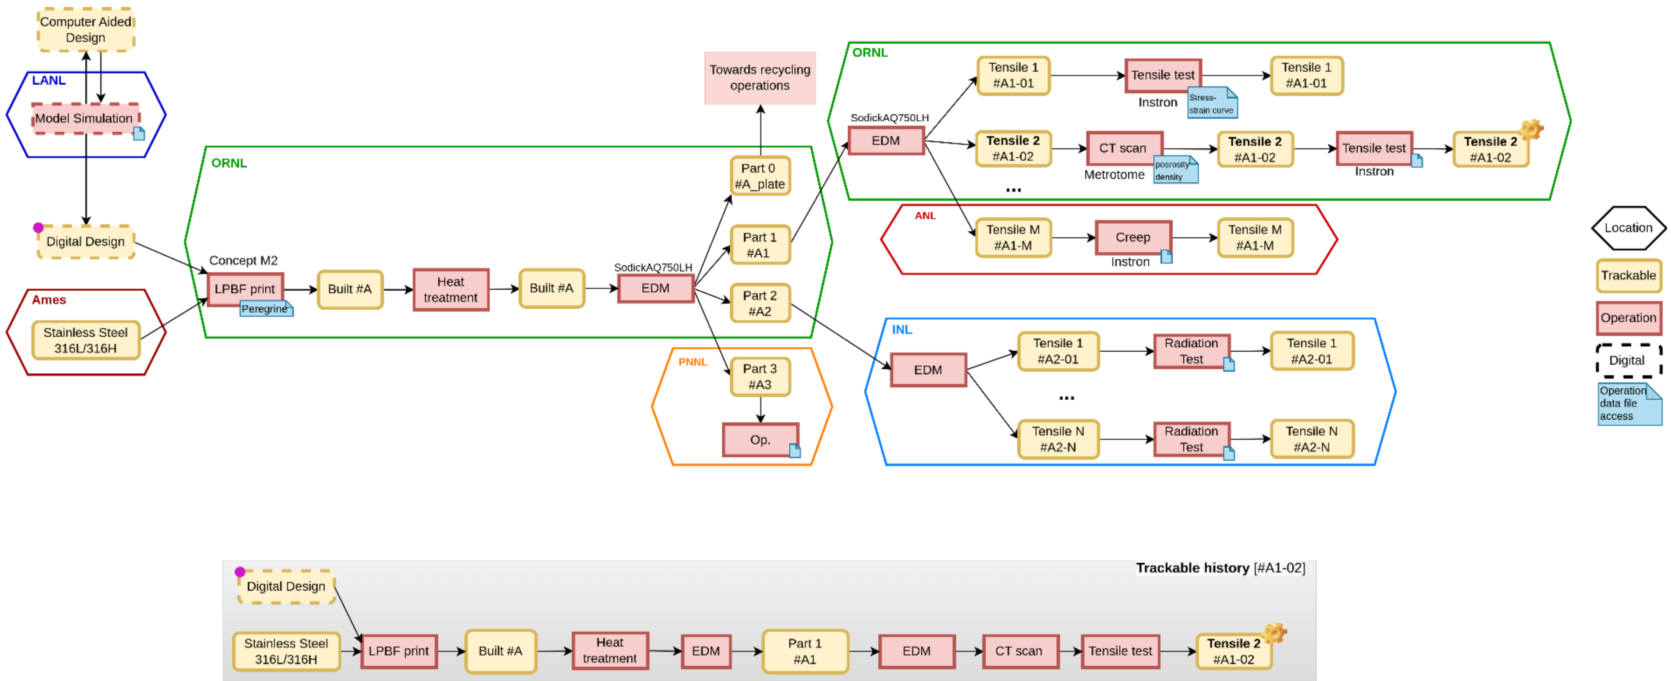

Figure 4. Representation of a multisite digital thread in the context of the operation-trackable data model. In this example, the trackable specimen "Tensile 2" history of operations and parent trackables are highlighted. The trackable entire creation context can be retrieved by retracing this history and gathering the data and metadata collected along each operation. *EDM* stands for electrical discharge machining.

3.3 DAMARA TERN FRAMEWORK: SOFTWARE AND TECHNOLOGY

The Damara Tern framework combines a web interface for exploring, entering, and accessing data within the MDF Digital Platform and the underlying MDF metadata database. This section provides a summary of the software and technology selected for its implementation within the MDF.

3.3.1 Database Implementation

The initial version of the database MDF Digital Platform, launched in FY 2020, relies on a printer-centric implementation. Each build (or print) is recorded in a dedicated table within a machine-specific PostgreSQL schema. This structure allows for the inclusion of all the machine-specific characteristics of the builds as fields within the build table. Additionally, other machine-related information can be stored in custom tables within the machine schema. This design customizes the metadata collection to the equipment and facilitates efficient machine-oriented data exploration. However, a significant drawback of this original model is its lack of extensibility and automation. Specifically, it mandates the creation of a new schema with a hand-crafted set of fields and tables for every new machine addition.

To address this challenge and ensure the platform's longevity and scalability, the development team initiated a comprehensive overhaul of the MDF Digital Platform in the last two quarters of FY 2023, building upon the operation-trackable data model detailed earlier. The new database architecture adheres to a semistructured approach by including only generic fields in the tables and storing the remaining fields within JavaScript object notation (JSON) field structures. As a result, this structure offers an abstraction of the specifics associated with distinct equipment types, materials, and operations, thereby broadening the scope of trackables and operations registered within the framework.

The database is implemented using the PostgreSQL relational database management system. This technology choice was driven by the development team's existing knowledge and experience. Nevertheless, alternative relational database management systems such as MariaDB or MySQL could potentially be substituted for PostgreSQL.

Figure 5 presents an overview of the database structure divided into five logical blocks: Operation, Trackable, Machine, Related Data and People and Affiliation. Each block comprises the main tables responsible for providing information within its designated scope as follows:

- **Operation:** Represents performed operations, including their types and specificities that define the contexts in which they are carried out

- **Trackable:** Encompasses catalogued physical or digital components subject to operations, along with the types and specifications used for their categorization

- **Machine:** Enumerates printers and other pieces of equipment employed for operations, including contextual details such as calibration and dedicated features

- **Related Data:** Involves information pertaining to substantial process data (such as images, videos, documents, etc.) stored within file systems or accessible remotely via URLs. These data are essential for retrieval and access purposes.

- **People and Affiliation:** Encompasses all information regarding users, institutions, and projects required to manage access and affiliations

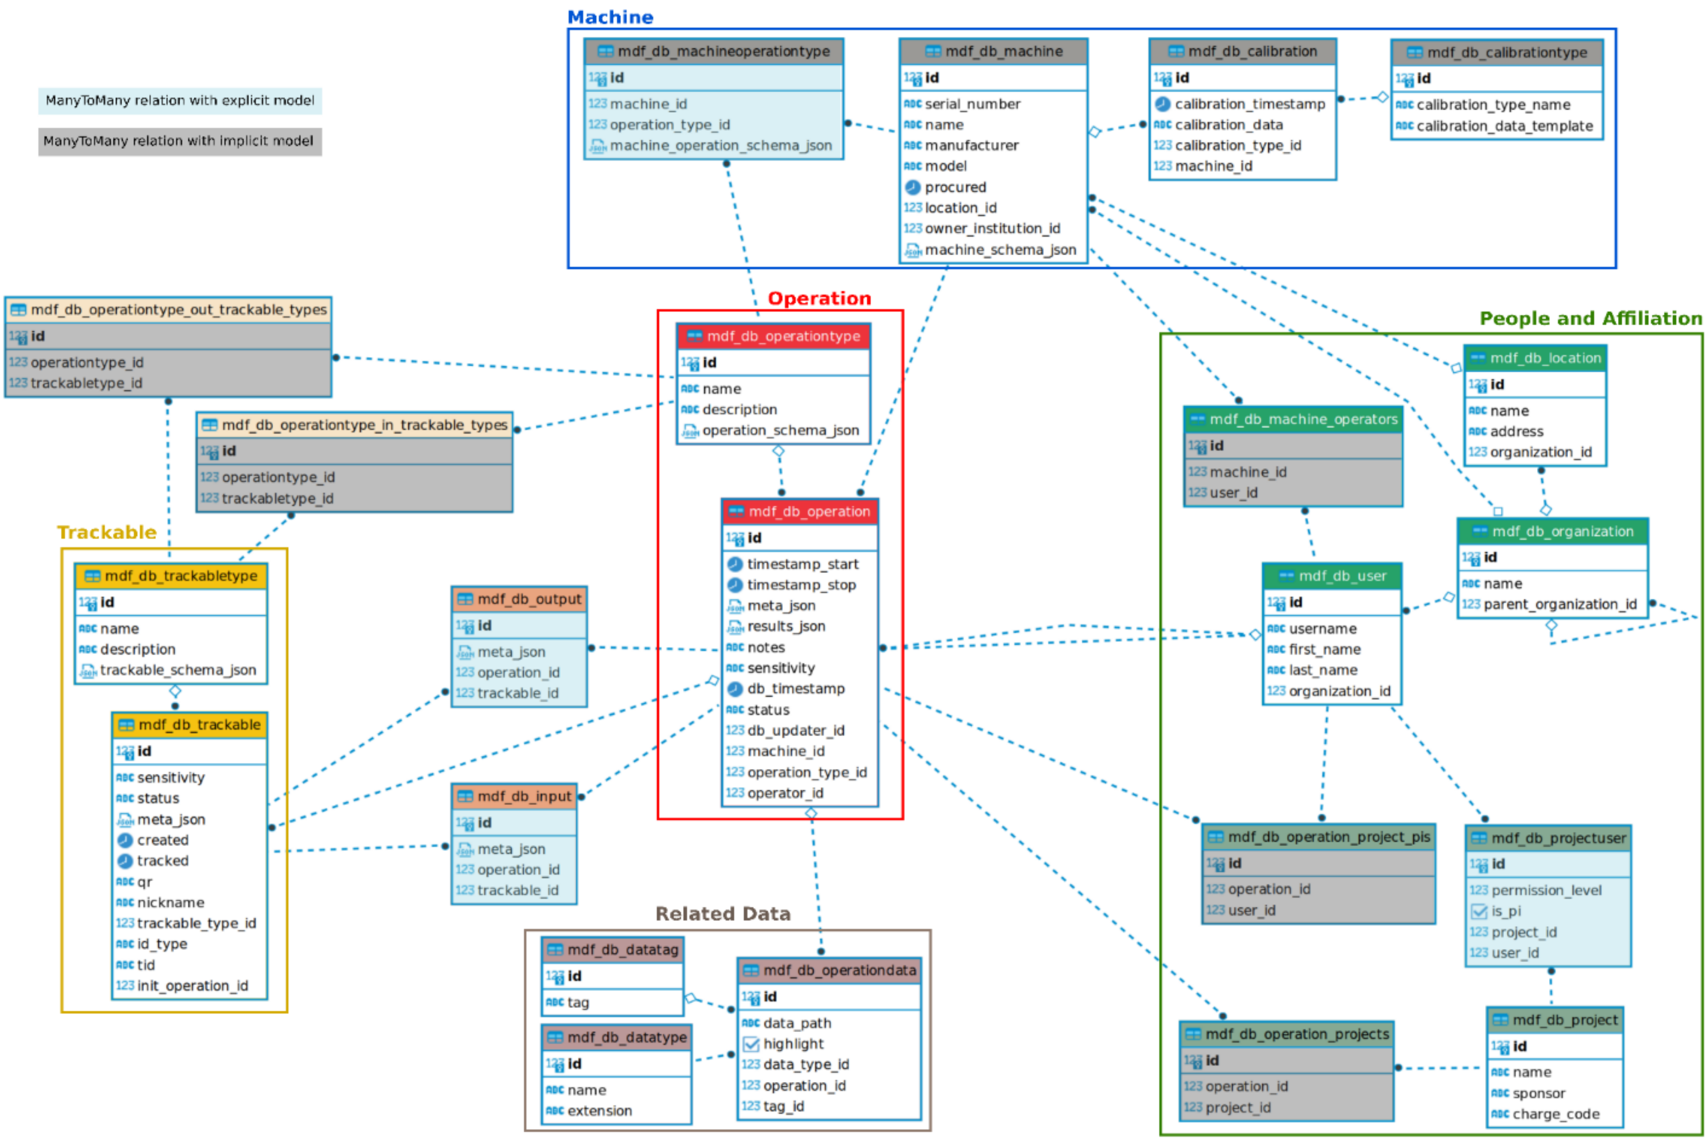

Figure 5. Overview of the database structure organized into five logical blocks: Operation, Trackable, Machine, People and Affiliation, and Related Data.

The following is a list functions for the primary tables corresponding to each logical group. For further information on the specific fields associated with these tables, please refer to Appendix A.

- **Operation:** Stores operation records and related metadata
- **Operation_Type:** Names and describes the supported operation types and provides form templates for operation-specific metadata collection
- **Trackable:** Stores trackable records and related metadata
- **Trackable_Type:** Names and describes trackable types and provides form templates for trackable-specific metadata collection
- **User:** Lists the MDF Digital Platform users
- **Organization:** Lists the institutional organizations for the users
- **Location:** Defines locations and addresses for organization and vendor sites
- **Project:** Lists the projects and sponsors supporting the purchases and manufacturing operations
- **Machine:** Identifies machines and equipment, along with their related metadata
- **Calibration:** Records details of the machines and equipment calibration over time
- **Calibration_Type:** Provisional table used to categorize calibrations and provide form templates for calibration-specific metadata collection
- **Related_Data:** Records the data paths, URLs, or references to the collected data stored remotely or on file system
- **Data_Type:** Defines the data type referenced by the Related_Data entries
- **Data_Tag:** Provides additional contextual tags to categorize the data referenced by the Related_Data entries

In addition to the main tables listed here, the database contains several association tables that store the many-to-many (M:M) relationships required by this data model design. Simple tables (highlighted in blue in Figure 5) operate on an implicit model with three fields for each entry: a primary key ID and a foreign key for each of the tables involved in the relation. Other association tables (Input, Output, Project_User, and Machine_Operation_Type) have been incorporated through explicit declarations to include additional information along with the M:M relationship.

3.3.2 Damara Tern Web Interface

With the redesign of the database, a new user interface was required to align with the underlying structure. The development of the Damara Tern web interface relies on the Django web framework, which adheres to a model view controller architecture. This choice was influenced by its out-of-the-box features, versatility, use of Python as the implementation language, and its proven success in prominent projects (e.g., Instagram, Mozilla, National Geographic).

The Django project is built upon a file structure organized into *apps*. Each app is thoughtfully designed to partition the project by core functionalities. The current project structure, illustrated in Figure 6, consists of four custom Django apps: home, mdf_db, operation, and trackable.

The mdf_data app handles the interfacing and connection with the metadata database described in Section 3.3.1. It provides the Python declarations for the database tables and fields, along with database-level constraints and validation rules. This app also defines the list fields and forms accessible through the admin interface and supports the LDAP authentication backend.

The operation app encompasses all the functionalities and classes required to manage the operation entries of the database model. This app includes various forms for creation; views for listing and displaying; and search filters that allow navigation, viewing, and retrieval of the operations and their related stored data.

The trackable app provides similar functionalities, including creation forms, list and detail views, and filter-based search capabilities, which are tailored for the management of trackable entries.

The home app includes fundamental functionalities tied to the website, such as home and login/logout pages. Additionally, the remaining structural components, labeled as "dt" and "static," house the essential settings, configurations, and static files required to bind the project.

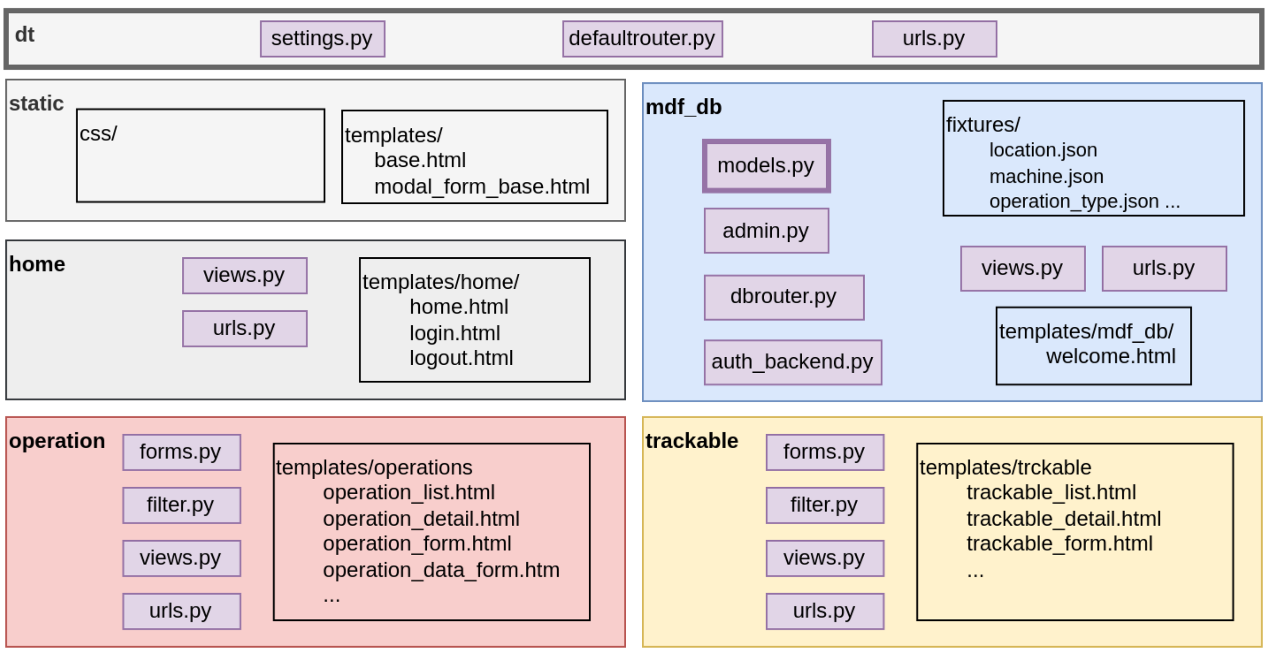

Figure 6. Damara Tern web interface project structure. The files are structured around the concept of apps, organizing the project by core functionalities.

For testing and development purposes, the prototype Damara Tern interface has been deployed on multiple ORNL internal servers accessible via the intranet. The current implementation provides operation and trackable lists and details views, fundamental navigation and search functionalities, and creation forms. Some of these features are illustrated in the interface screenshots presented in Figure 7.

The ongoing development efforts for this interface focus on replicating the functionality of the digital tool interface and implementing backend validation features to ensure the integrity of the database. Other improvement efforts target enhancing the views, filtering, and display to facilitate the user experience when exploring the data through the interface.

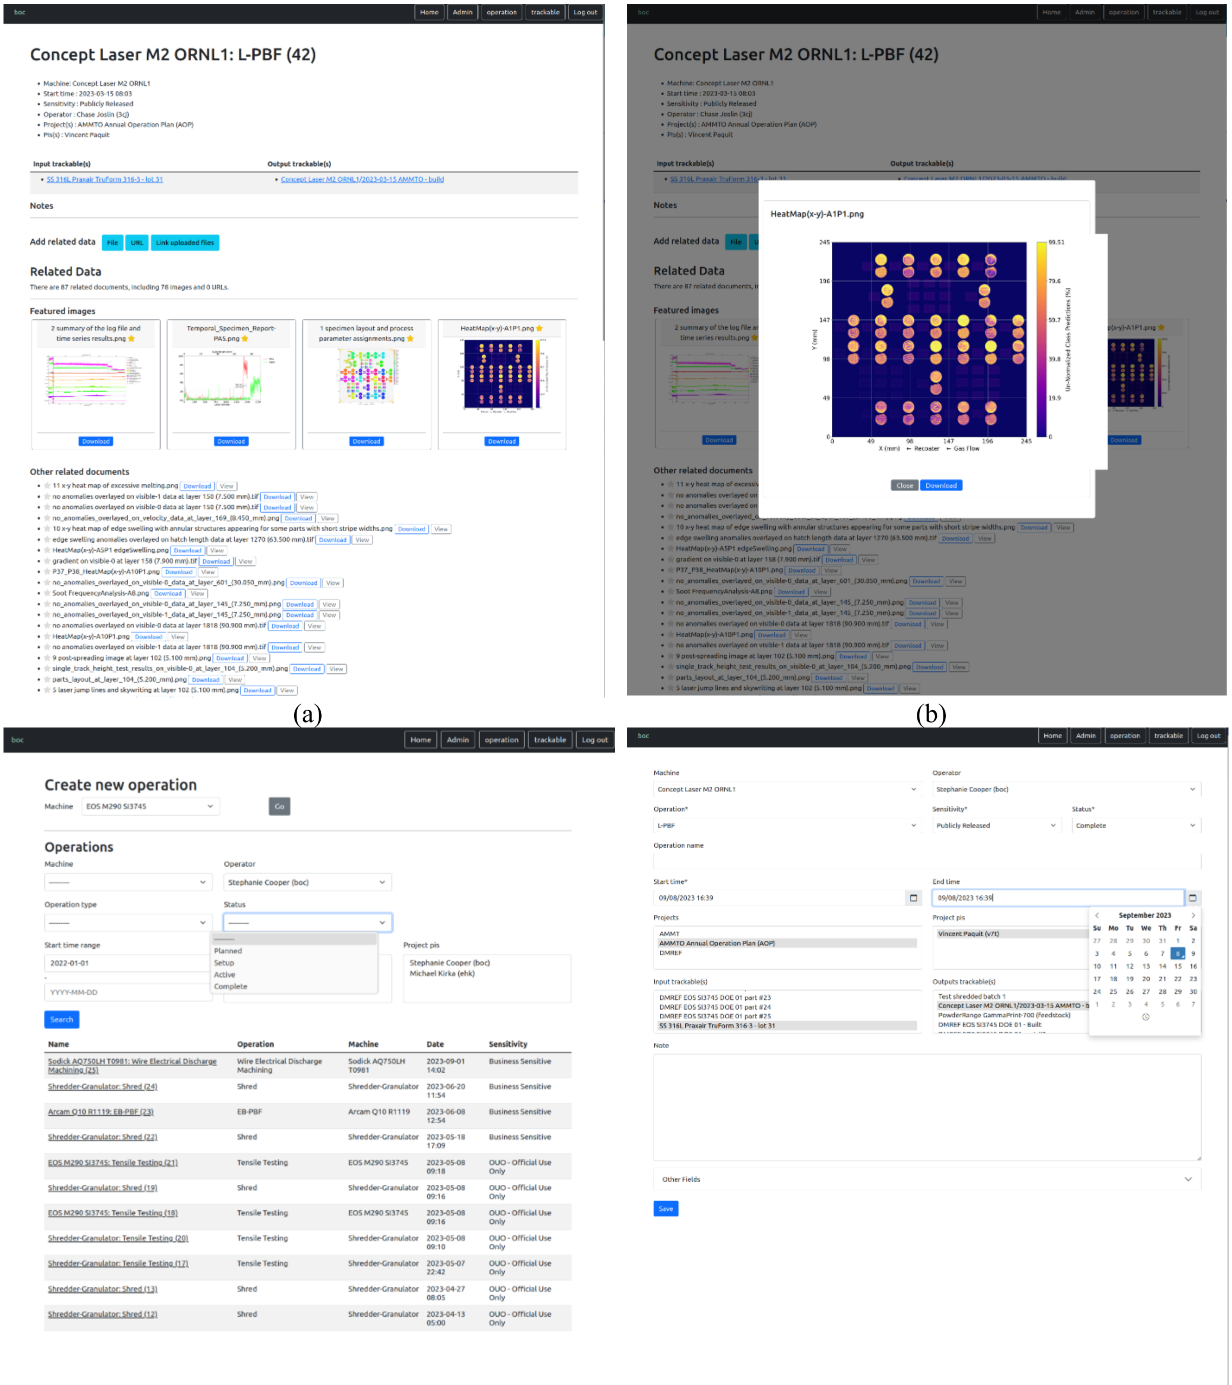

(a) (b)

(c) (d)

Figure 7. Screenshots of the Damara Tern prototype interface illustrating several available functionalities to manage the model's operations. The interface includes operation detail views featuring (a) related file lists, highlights upload/download options, and (b) basic visualization features. (c) Operation lists are accessible through the main menu and include filtering options to streamline searches, along with links toward the detail views. (d) New operations can be created using dedicated forms. Note that some data in these figures are derived from practice examples and do not represent real data sets (e.g., no Office Use Only data have been shown).

4. SENSING AND INSTRUMENTATION UPGRADES OF LASER-BASED POWDER BED FUSION SYSTEMS

To support the activities of other work packages during FY 2023 and beyond, the sensors of multiple LPBF systems were upgraded. Efforts were made to install a common sensor package for in situ monitoring data across these machines, which will also be extended to facilities across the AMMT partner labs. This common sensor package, which includes a combination of high-resolution, visible-light images and temporally integrated near-infrared (TI-NIR) images, will better allow comparisons of process stability and material quality across different builds, materials, LPBF machines, and printing facilities. The data generated by these sensors have been integrated into the Peregrine software tool and will be made available to the other work packages to support ex situ characterization efforts.

4.1 RENISHAW AM250

Three new sensors were installed on the Renishaw AM250: a 20 megapixel (MP) visible light camera (Basler acA5472-17um), a 4.2 MP near-infrared (NIR)-sensitive camera (Pixelink PL-D734MU-NIR-T), and a 0.33 MP long-wave infrared (LWIR) camera (Teledyne FLIR Boson 640). The visible-light camera is outfitted with an ultraviolet (UV)/NIR cutoff filter to protect the camera detector from the laser. Similarly, the NIR-sensitive camera has a narrow band-pass filter centered in the NIR range (808 ± 10 nm), as well as neutral density filters (outer diameter [OD] ≈ 1.8) to prevent image saturation. Each camera (a Peregrine "explicit channel") can produce multiple images per layer (or Peregrine "frames"), which may be analyzed by a trained DMSCNN (Section 6). As shown in Figure 8, all three cameras were placed inside the build chamber given the physical constraints of the Renishaw AM250's viewports. Cables for the cameras run through three cable passthroughs in the top of the build chamber, and cord grips ensure that the build chamber atmosphere is maintained.

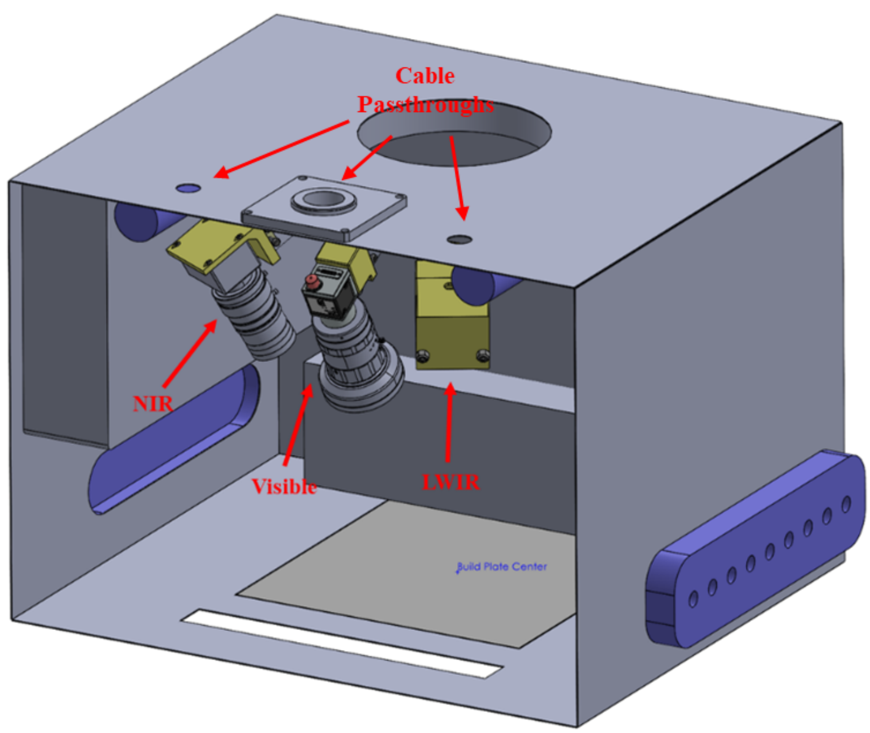

Figure 8. CAD model representation of the Renishaw AM250's build chamber showing the three cameras installed inside of the machine. Available data modalities include visible-light imaging, as well as TI-NIR and LWIR images.

A total of seven frames are captured each print layer. The visible-light camera, which has an instantaneous field of view (IFOV) of approximately 55 µm, captures a post-melt and a post-recoat image for each layer. The NIR camera, which has an approximate IFOV of 125 µm, produces a video buffer that is dynamically analyzed at the end of a layer to produce a total of three images engineered to extract relevant process dynamics for the layer. These images include temporally integrated sum, max, and argmax images. Finally, the LWIR camera, which has an IFOV of approximately 140 µm, currently captures a post-melt and post-recoat image but may be used to explore other temporally integrated imaging modalities. The LWIR has a reduced field of view and therefore cannot image the entire build plate. Representative examples of each of these data modalities are shown in Figure 9.

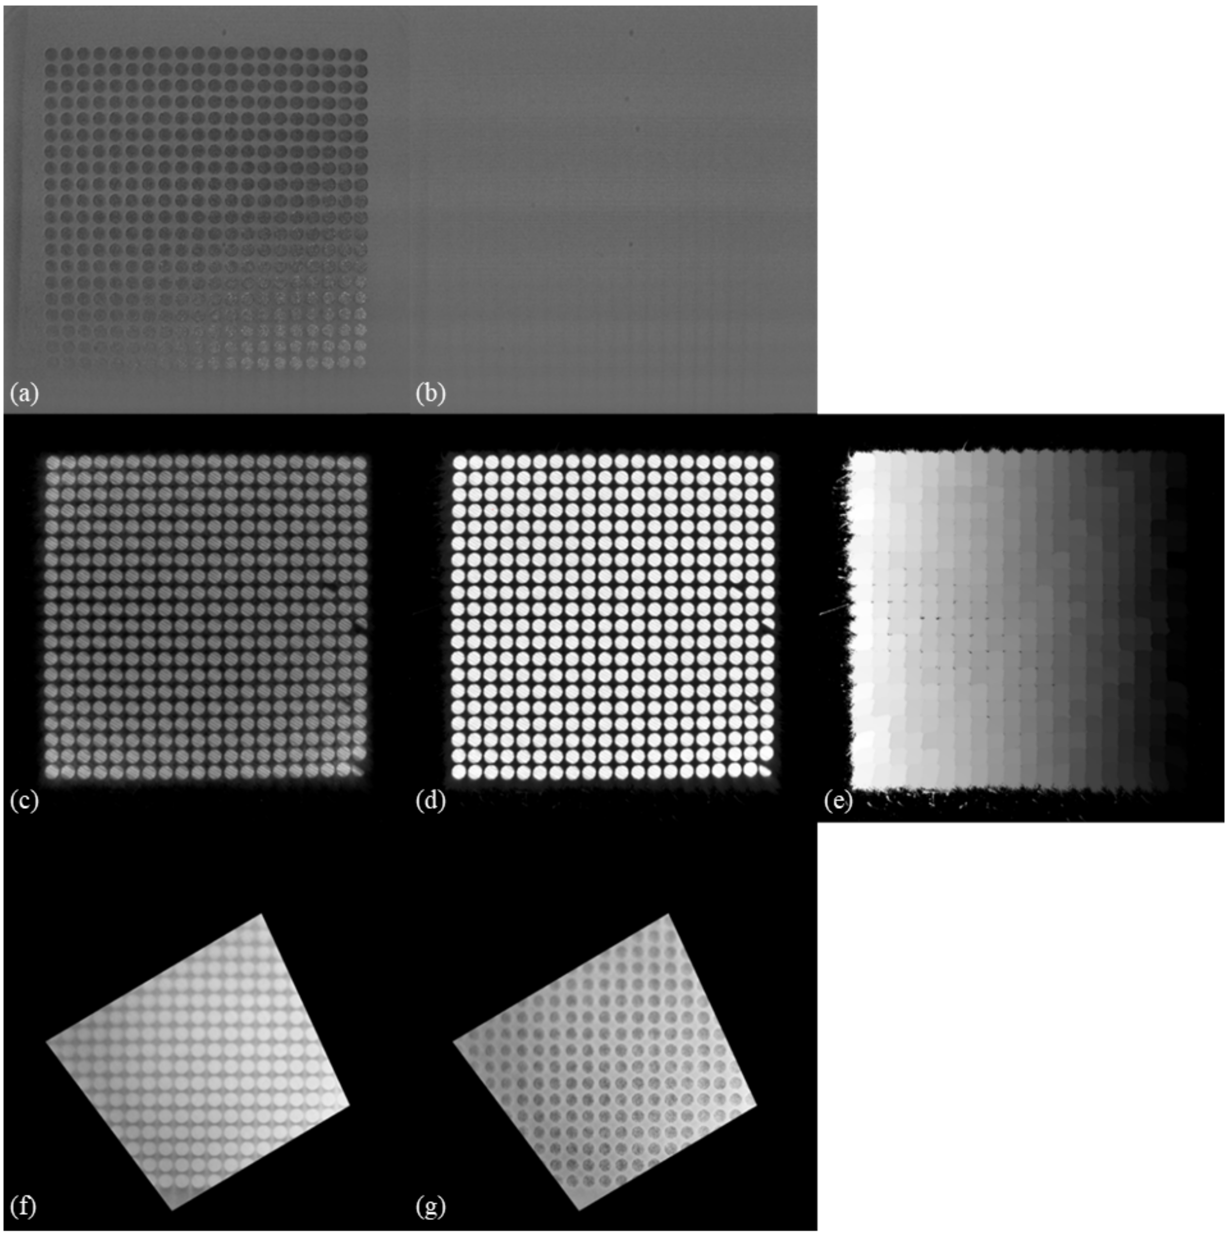

Figure 9. Representative examples of each of the seven data modalities captured on a given layer for the Renishaw AM250. The visible camera collects (a) post-melt and (b) post-recoat snapshots, and the NIR camera produced three temporally integrated images that summarize the process dynamics: (c) sum, (d) max, and (e) argmax. The LWIR camera also collects a (f) post-melt and (g) post-recoat snapshot.

4.2 RENISHAW AM400

Whereas the Renishaw AM250 only received sensor upgrades, the Renishaw AM400 was a brand new connection to the MDF Digital Platform. Two sensors were installed on the Renishaw AM400: a 20 MP visible-light camera (Basler acA5472-17um) and a 4.2 MP NIR-sensitive camera (Pixelink PL-D734MU-NIR-T). The Renishaw AM250 and AM400 share nearly identical build chamber and viewport designs, so a similar approach was taken to instrument the AM400. As shown in Figure 10, both cameras were mounted inside of the build chamber, and the camera cables were passed through the top of the build chamber and secured with cord grips. The visible-light camera is outfitted with a UV/NIR cutoff filter to protect the camera detector from the laser, and the NIR-sensitive camera has neutral density filters (OD ≈ 2.4) and a narrow band-pass filter centered in the NIR range (808 ± 10 nm).

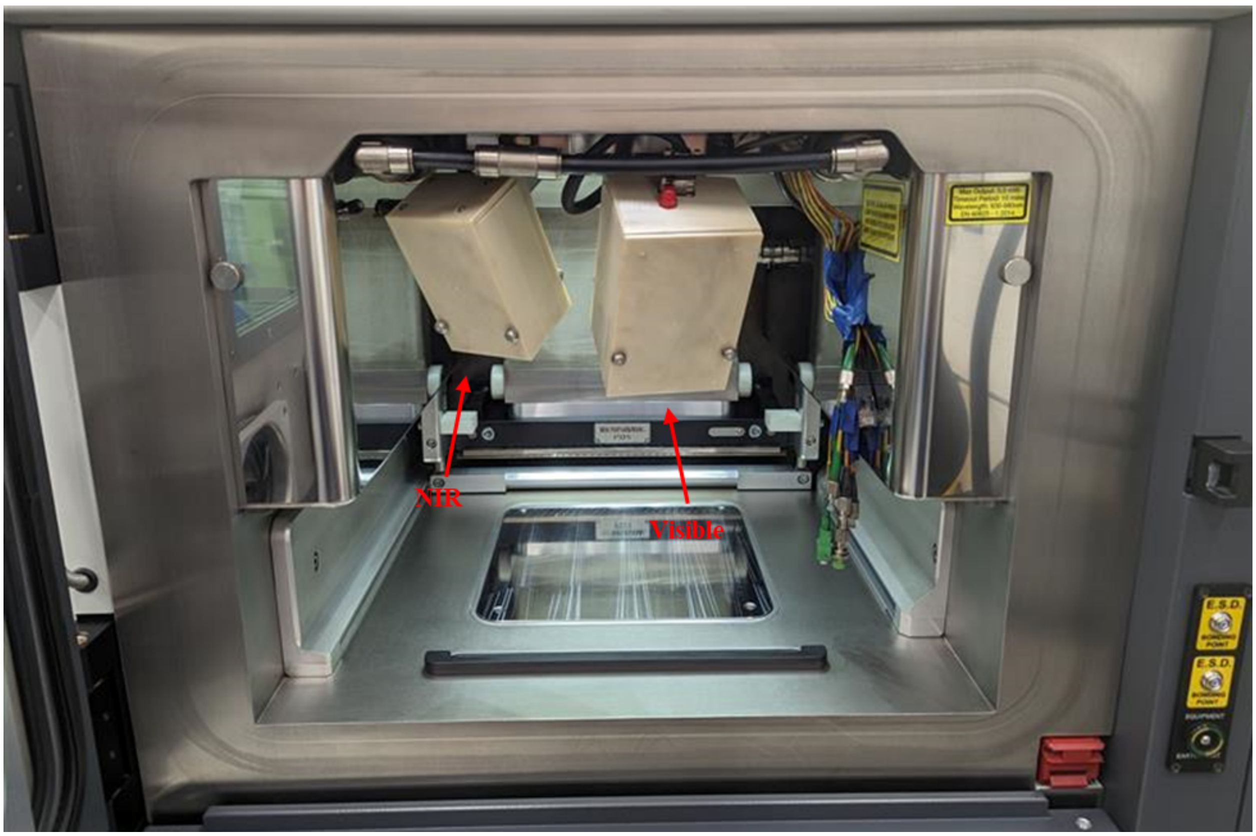

Figure 10. View of the Renishaw AM400 build chamber with the NIR and visible-light cameras installed. Both cameras are placed inside of the build chamber.

For the Renishaw AM400, five Peregrine frames are captured each print layer. The visible-light camera, which has an IFOV of approximately 55 µm, captures a post-melt and a post-recoat image each layer. The NIR camera, which has an approximate IFOV of 125 µm, produces integrated sum, max, and argmax images. Representative examples of each of these data modalities are shown in Figure 11.

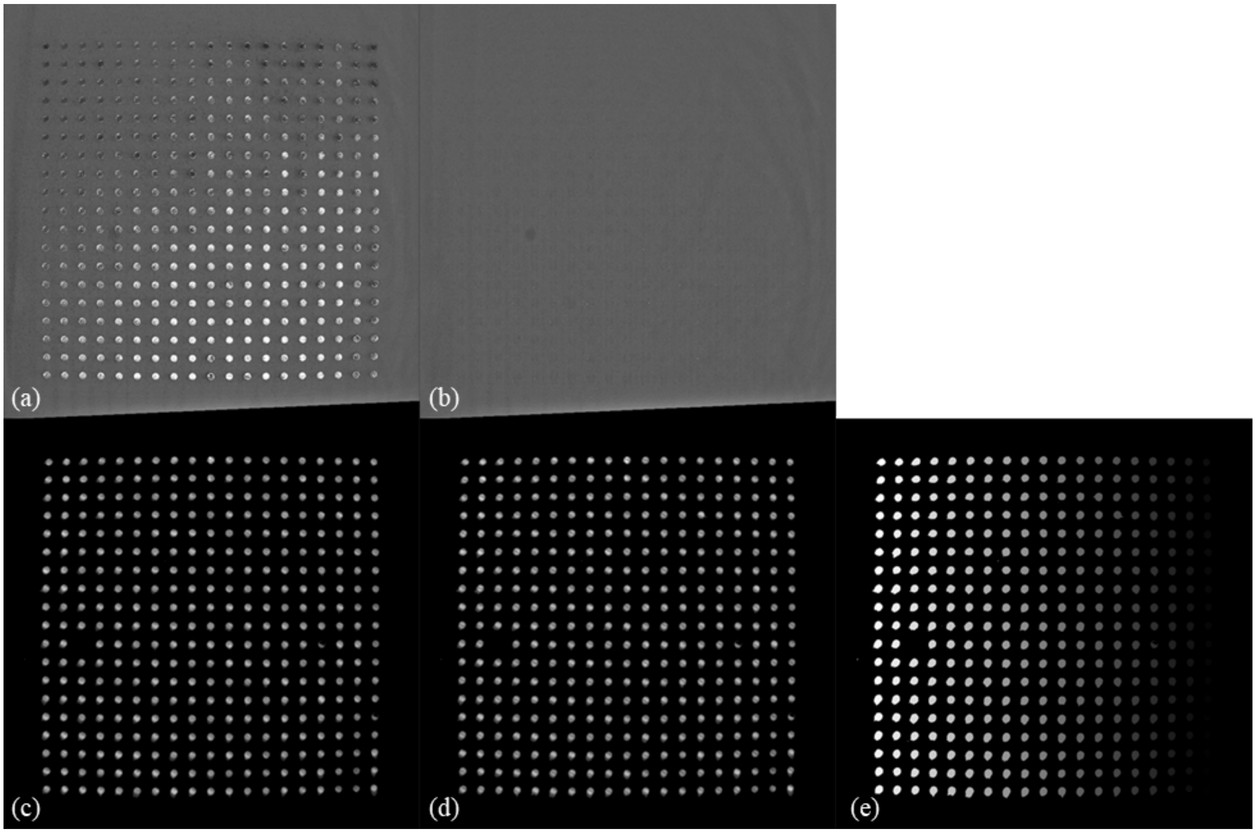

Figure 11. Representative examples of each of the five data modalities captured on a given layer for the Renishaw AM400. The visible camera collects (a) post-melt and (b) post-recoat snapshots, and the NIR camera produced three temporally integrated images that summarize process dynamics: (c) sum, (d) max, and (e) argmax.

4.3 CONCEPT LASER M2

4.3.1 Off-Axis Cameras

A 20 MP visible-light camera (Basler acA5472-17um) and a 4.2 MP NIR-sensitive camera (Pixelink PL-D734MU-NIR-T) have been installed in the Concept Laser M2. The visible-light camera is mounted to the back viewport of the machine, replacing the 5 MP visible-light camera that had been installed in that location previously. The new visible-light camera has an IFOV of approximately 60 µm. Conversely, the NIR camera, a new sensing modality for this machine, is installed inside of the build chamber to maximize the spatial resolution and improve the camera's viewing angle; other mounting locations would have resulted in significant perspective distortions and nonuniform spatial resolutions. The NIR camera has an IFOV of approximately 140 µm.

4.3.2 On-Axis Photodiodes

The Concept Laser M2 used by the AMMT Program at ORNL is equipped with an on-axis, photodiode-based sensor suite originally installed by Concept Laser [14]. During the TCR Program, several limitations with this system were discovered; however, only a subset of these issues were addressed at the time [3]. In FY 2023, the remainder of the known issues were investigated and either corrected or sufficiently characterized such that the on-axis photodiodes can be used as a valid sensor signal in FY 2024. First, the photodiodes on both laser modules were changed to the PDA100A2 model (initially tested in FY 2020) because it has a larger detector area. At the same time, the sensor gains were lowered to

60 dB to prevent saturation of the signal and decrease the sensor's rise time. Next, the signal inversion observed on one of the two laser modules was corrected by independently identifying the background sensor signal and using it as a known baseline intensity. Because it is unknown if this issue (potentially caused by incorrect wiring at the analog-to-digital converter) is common among other Concept Laser M2 printers, the code was written to automatically detect and correct for this issue on a laser-by-laser basis. Figure 12 shows an example of the photodiode data before and after the inversion corrections are applied.

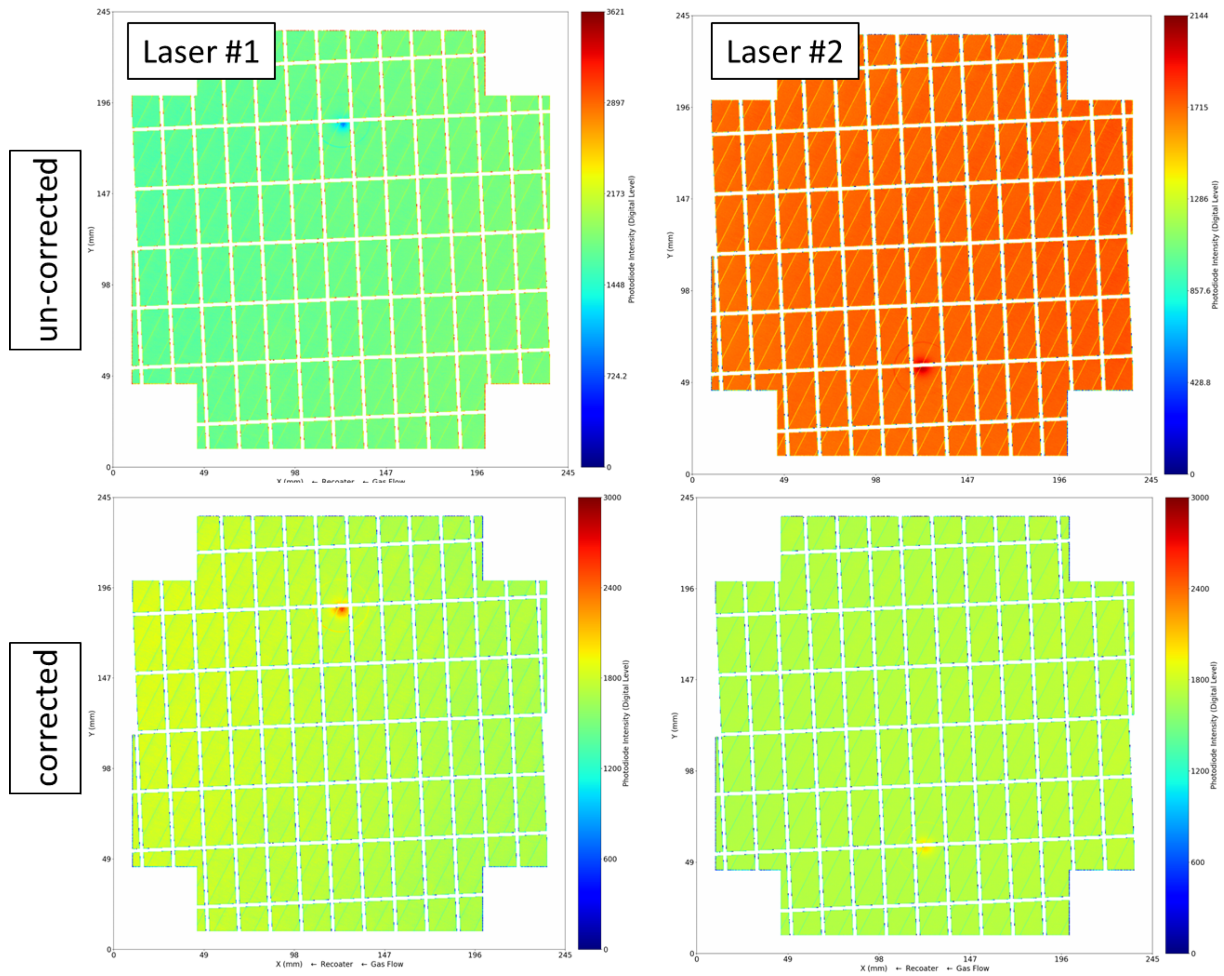

Figure 12. Photodiode data from the first column were collected from laser module #1, and data in the second column were collected from laser module #2. Despite observing nominally identical melting conditions, note that the uncorrected data in the first row report significantly different signal magnitudes. Furthermore, the gradients between the two sensors are inverted, with the signal from laser #1 being highest when the laser is turned off and lowest when the laser is actively melting material. The corrected data in the second row report similar magnitudes and identical behaviors when observing similar laser behaviors.

The final outstanding issue is the presence of ring-shaped artifacts apparent in the reconstructed photodiode sensor data. By exposing the entire print area with both lasers, these artifacts were determined to be confined to regions located directly underneath each laser module and are most likely caused by reflections within the laser optics. At this time, no corrective actions are planned, and the AMMT Program will avoid printing test coupons in these regions for studies that rely upon the photodiode data. Figure 13 shows an example of fully corrected photodiode data.

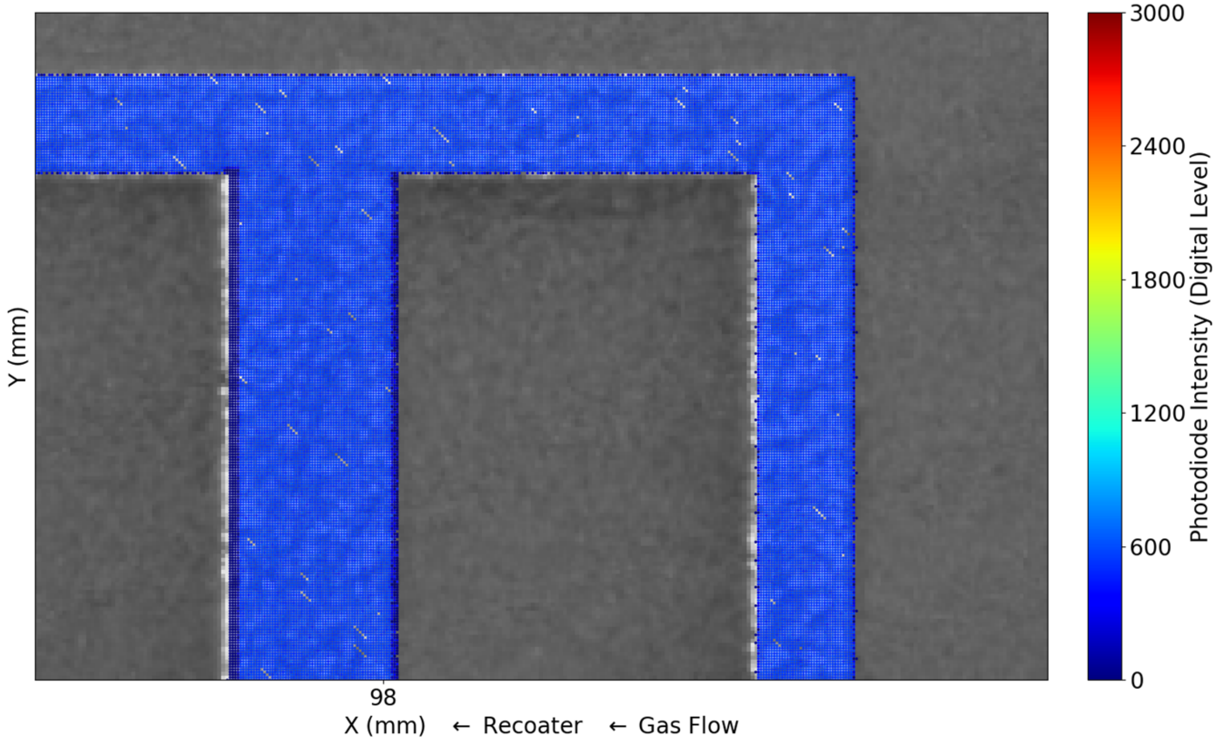

Figure 13. An example of fully corrected on-axis photodiode data collected from a Concept Laser M2 printer and overlaid on top of a visible-light image captured after the layer was melted. The applied corrections include (1) a lossless bit shift to correct for an incorrect upstream datatype conversion, (2) installation of a high-pass filter to block reflected laser light, (3) installation of a photodiode with a larger detector area and a modified sensor gain setting, (4) a temporal shift used to synchronize the sensor signals with the recorded scanner position, (5) algorithmic filtering to remove laser jump vectors, (6) using the background sensor signal to correct for an inverted electrical response, and (7) avoiding regions of the print area susceptible to reflections within the optics.

4.4 SUPPORT FOR SENSING AT OTHER NATIONAL LABS

In addition to instrumenting LPBF machines at ORNL, efforts have been made to install similar sensor packages at the other national labs partnered under the AMMT Program. Specifically, Argonne National Laboratory has a Renishaw AM250, which is being instrumented in the same way as ORNL's Renishaw AM400. Specific hardware requirements and design files have been detailed and exchanged to ensure that the sensing configuration is as similar to ORNL's as possible. Similarly, Los Alamos National Laboratory is planning to purchase the same visible-light camera to install on their EOS M290. Having analogous sensing systems installed both across machines inside of ORNL and between national labs will facilitate comparisons between in situ data collected from different machines at different sites and could possibly be used for round robin studies between labs. Such efforts could help explain machine-to-machine variability and inform data requirements for AM qualification.

5. REAL-TIME DATA ACQUISITION

The Peregrine software tool collects layer-wise imaging data for the LPBF printers used under the AMMT Program, including the Concept Laser M2, the Renishaw AM250, and the Renishaw AM400. The real-time data acquisition code was originally developed in FY19 and had seen only incremental improvements in the intervening years. Therefore, this portion of the Peregrine code base was overhauled in FY2023 to support the newly developed and installed sensor packages described in Sections 4.1, 4.2, and 4.3.1. The primary objectives of the overhaul included (1) improved camera connectivity and cross-camera synchronization and stitching, (2) support for sensor noise and view factor corrections, (3) support for video capture and long-exposure compositing, (4) reduced layer-wise computation time, and (5) improved modularity for ease of future modification.

First, all the camera hardware-specific code was moved into a Python class so that the camera connection, data transfer, and error checking methods could be properly abstracted. Hardware support is currently implemented for Basler, Pixelink, and OpenCV-compatible cameras, and the new software architecture enables the trivial addition of additional camera-specific software developer kits and application programming interfaces (APIs). This flexibility will be particularly important in the future as other AMMT partner labs explore cameras different from those used at ORNL.

A second Python class was created as a wrapper around the camera hardware layer. Each imaging modality is assigned a unique, threaded instance of this class which maintains the camera connection by automatically recovering from intermittent camera connectivity issues, synchronizing the clocks across multiple cameras, and stitching multiple cameras of the same type together to cover a larger field of view. The sensor noise suppression and view factor corrections required for TI-NIR imaging can also be applied from within this class. Figure 14 shows a summary of the operations performed by each threaded instance of this Python class.

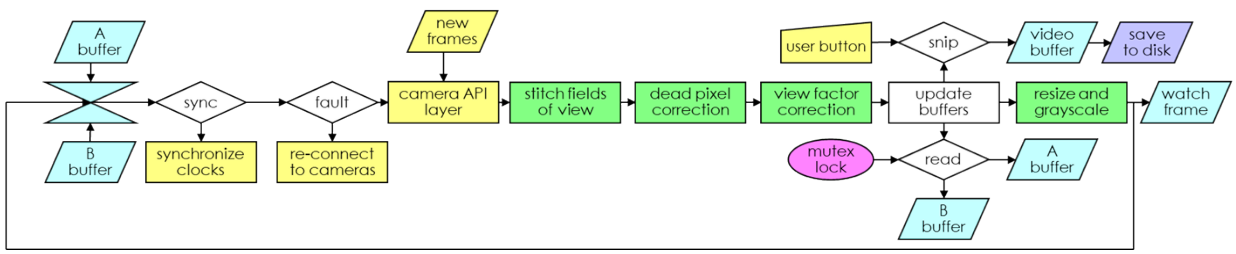

Figure 14. A representation of the data capture thread, where each camera imaging modality is assigned a unique instance of this Python class. Yellow elements indicate hardware or user interfaces, green elements indicate classical computer vision calculations, and blue elements represent data storage buffers. Note that the pink mutex lock is triggered from a parallel thread and determines whether new data should be written to the A-Buffer or the B-Buffer. For simplicity, the video buffer triggered by the user is shown in this diagram; however, in its actual implementation, this operation is spawned as an independent thread.

The most significant change from previous iterations of the Peregrine data acquisition architecture is the introduction of an advanced data buffering system. These buffers allow over 100 GB of video data per camera to be stored in volatile computer memory. Using these buffers, arbitrary calculations can be performed on the entire video stream to create composite images encoding information about the entire print layer. For example, when performing TI-NIR imaging, the pixel intensities can be summed through the entire video stream, as discussed in Sections 4.1, 4.2, and 4.3.1. The time stamps marking the beginning and end of each print layer are determined by a state machine running in a parallel thread, as shown in Figure 15. The buffers themselves are processed in the same thread that processes the final layer-wise image stack using the DMSCNN segmentation model.

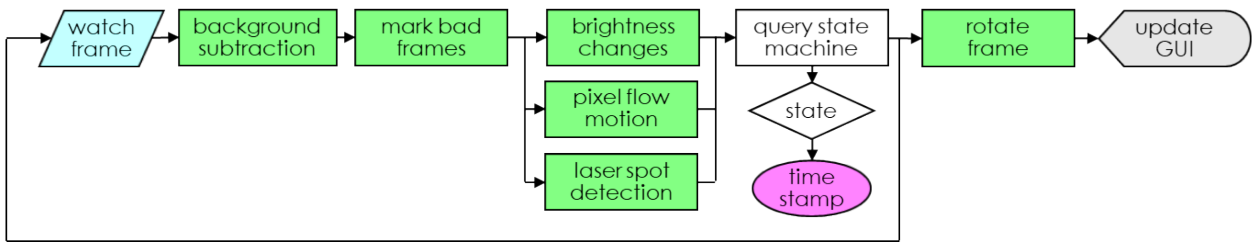

Figure 15. A representation of the state machine thread. The video stream from a single user-selected camera is observed by a configurable state machine to determine the beginning and end of each print layer. These time stamps are shared between threads and are used to extract only the portions of each camera video buffer that are associated with the current print layer.

These major improvements, along with a number of additional modifications, slightly reduced the layer-wise analysis time and substantially increased the flexibility of the code base. Figure 16 shows a screenshot of Peregrine's user interface during a test of the real-time data acquisition. The software changes implemented in FY 2023 have enabled the collection of TI-NIR data across three new LPBF printers at ORNL and will facilitate the deployment of Peregrine at the other AMMT partner labs.

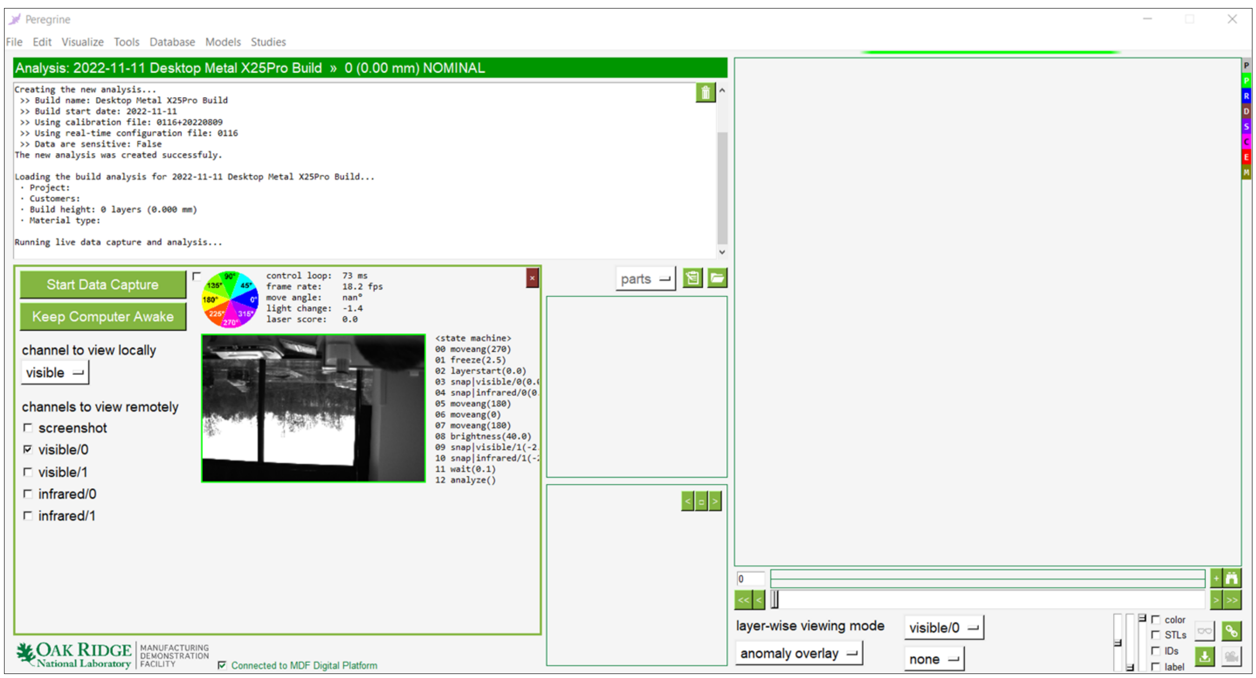

Figure 16. A representative screenshot of the Peregrine user interface during real-time data acquisition.

6. DYNAMIC MULTILABEL SEGMENTATION CONVOLUTIONAL NEURAL NETWORK

In FY 2023, Peregrine's pixel-wise segmentation model was substantially upgraded to support multilabel classification. That is, the new DMSCNN is capable of predicting multiple anomaly classes at each individual pixel in an image stack. Although this effort was primarily supported through a technical collaboration between ORNL and RTX (also known as Raytheon Technologies), as well as the DOE Advanced Materials and Manufacturing Technologies Office Digital Factory annual operating plan, this upgrade has three significant implications for the AMMT Program. First, because the anomaly classes are no longer mutually exclusive, it is possible for the DMSCNN to better represent different laser processing conditions (Figure 17). This capability will be critical when the in situ imaging data are further leveraged for developing optimized process parameter sets for stainless steel 316H in FY 2024. Second, because the DMSCNN effectively operates as a set of jointly-learned binary classifiers, pixels can be ignored during training for specific classes—a crucial requirement for using x-ray computed tomography data as training ground truths for detecting spatter-induced lack-of-fusion porosity with TI-NIR imaging [8]. Finally, the new architecture also allows the DMSCNN to report that a pixel is "unknown" (i.e., it cannot confidently predict its state given its current training). This is a first step toward uncertainty quantification and providing justifications and confidence levels alongside AI-based predictions.

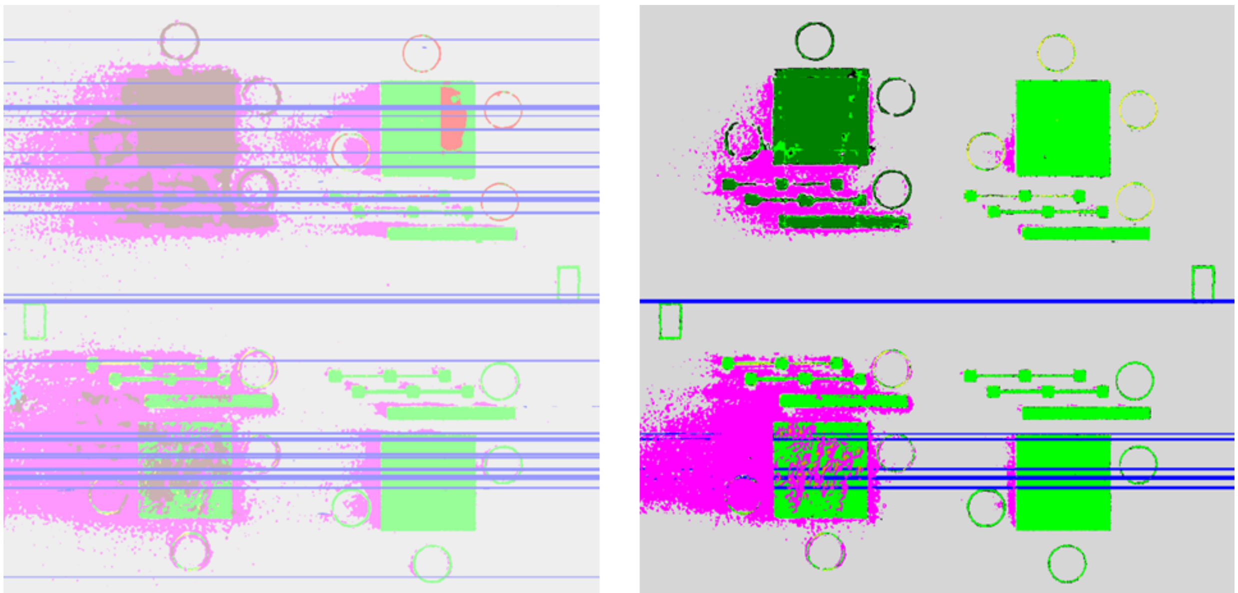

Figure 17. A single layer of anomaly segmentation results from a TCR build printed on a Concept Laser M2 LPBF system. The results on the left were produced by a trained deep stochastic convolutional neural network model, and the results on the right were produced by the new, multilabel DMSCNN model. Note that the DMSCNN can now correctly differentiate between lack-of-fusion process parameters (dark green in the upper left quadrant) and keyholing process parameters (orange in the lower left quadrant) instead of classifying both off-nominal melting conditions as debris (brown in the left-hand image). Also note that the DMSCNN no longer falsely predicts super-elevation (red) for the melting parameters in the upper right quadrant.

CONCLUSIONS

The work package seeks to create a digital manufacturing discipline common to all participants of the AMMT Program to improve the performance of nuclear components. The MDDC work package has focused on three activities during FY 2023. First, the Damara Tern framework was created, combining a web interface for exploring, entering, and accessing data within the MDF Digital Platform and the underlying metadata database. Next, sensor upgrades and standardized sensor packages were implemented across three AM machines at the MDF, ensuring comparability between builds, machines, and sites, and avenues for installing similar sensors at the other AMMT labs were also explored. Finally, improvements were made to Peregrine's real-time data acquisition and deep learning frameworks. Copies of the Peregrine software were also provided to the AMMT partner labs. **This report marks the completion of FY 2023 milestone M3CR-22OR0403051: Report Describing the Architecture of the Digital Platform to Support AMMT Activities.**

REFERENCES

[1] "Additive manufacturing - General principles - Fundamentals and vocabulary," ISO/ASTM, 2021. doi: 10.3403/30448424.

[2] L. Scime *et al.*, "Report on Progress of correlation of in-situ and ex-situ data and the use of artificial intelligence to predict defects," Oak Ridge, TN (United States), Oct. 2020. doi: 10.2172/1684671.

[3] L. Scime *et al.*, "Development of Monitoring Techniques for Laser Powder Bed Additive Manufacturing of Metal Structures (Progress Report)," Oak Ridge, TN (United States), Sep. 2020. doi: 10.2172/1675044.

[4] L. Scime *et al.*, "Development of Monitoring Techniques for Binderjet Additive Manufacturing of Silicon Carbide Structures," Oak Ridge, TN (United States), Oct. 2020. doi: 10.2172/1671401.

[5] L. Scime *et al.*, "Report on Diagnostic and Predictive Capabilities of the TCR Digital Platform," Oak Ridge, TN (United States), 2021.

[6] L. Scime *et al.*, "Diagnostic and Predictive Capabilities of the AMMT Digital Platform," Oak Ridge, TN (United States), 2022.

[7] L. Scime, D. Siddel, S. Baird, and V. Paquit, "Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation," *Addit. Manuf.*, vol. 36, p. 101453, Dec. 2020, doi: 10.1016/j.addma.2020.101453.

[8] Z. Snow, L. Scime, A. Ziabari, B. Fisher, and V. Paquit, "Observation of spatter-induced stochastic lack-of-fusion in laser powder bed fusion using in situ process monitoring," *Addit. Manuf.*, vol. 61, p. 103298, Jan. 2023, doi: 10.1016/j.addma.2022.103298.

[9] A. Ziabari *et al.*, "High Resolution X-Ray CT Reconstruction of Additively Manufactured Metal Parts using Generative Adversarial Network-based Domain Adaptation in AI-CT," *Microsc. Microanal.*, vol. 27, no. S1, pp. 2940–2942, Aug. 2021, doi: 10.1017/S1431927621010254.

[10] L. Scime, A. Singh, and V. Paquit, "A scalable digital platform for the use of digital twins in additive manufacturing," *Manuf. Lett.*, vol. 31, Jun. 2021, doi: 10.1016/j.mfglet.2021.05.007.

[11] L. Scime, V. Paquit, C. Joslin, D. Richardson, D. Goldsby, and L. Lowe, "Layer-wise Imaging Dataset from Powder Bed Additive Manufacturing Processes for Machine Learning Applications (Peregrine v2021-03)." Oak Ridge National Laboratory, Oak Ridge, TN (United States), 2021. doi: 10.13139/ORNLNCCS/1779073.

[12] L. Scime *et al.*, "Layer-wise Imaging Dataset from Powder Bed Additive Manufacturing Processes for Machine Learning Applications (Peregrine v2022-10)." Oak Ridge National Laboratory, Oak Ridge, TN (United States), 2022. doi: 10.13139/ORNLNCCS/1896716.

[13] T. Wohlers, N. Mostow, I. Campbell, O. Diegel, J. Kowen, and I. Fidan, *Wohlers Report 2022: 3D Printing and Additive Manufacturing Global State of the Industry*. Fort Collins, CO, 2022.

[14] ConceptLaser, "Quality Management." Accessed: Dec. 14, 2020. [Online]. Available: https://www.ge.com/additive/sites/default/files/2018-02/1708_QM-pm_EN_update_1__lowres_einzel.pdf

APPENDIX A. DATABASE TABLE DETAILS

This appendix presents the primary tables of the operation-based database that store metadata for the MDF Digital Platform. Each table presented here is explicitly defined through a Django model, enabling Damara Tern to interface with the database. The following tables display the database tables listed in Section 3.3.1 and employ the same color code convention used to identify the logical block organization in the header. They list the fields, data types, and field descriptions.

Operation		
Description:	Table of operation entries	
Note:	Operation entries are meant to be created by machine operators via webform.	
Field	Data type	Description
id	int 8	Unique ID to identify the operation record, primary key
timestamp_start	timestamptz	Time at which the operation started
timestamp_stop	timestamptz	Time at which the operation ended
meta_json	jsonb	JavaScript object notation (JSON) field holding the contextual metadata related to the machine in use and the type of operation performed, as well as any additional contextual data identified
results_json	jsonb	Provisional field to store and structure metadata related to the specific entry
notes	text	Any comments/notes by technicians or contributing users
sensitivity	varchar(1)	Data sensitivity: P, B, O, C, E, or I, indicating Publicly Released, Business Sensitive, Official Use Only, Cooperative Research and Development Agreement (CRADA) Protected, Export Controlled, or International Traffic in Arms Regulations (ITAR), respectively
db_timestamp	timestamptz	Time at which the record was added to the database (automatically entered)
status	varchar(1)	Operation status: P, S, A, or C, representing Planned, Setup, Active, and Complete, respectively
db_updater_id	int 8	ID of the user who created the record; foreign key referencing the Users table
machine_id	int 8	ID of the equipment used for printing; foreign key referencing the Machine table
operation_type_id	int 8	Type of operation performed (digitally or by the equipment); foreign key referencing the Operation Type table
operator_id	int 8	ID of the technician who performed the operation; foreign key referencing the User table

id	int 8	Unique ID to identify the record, primary key
name	varchar(50)	Name of the operation type
description	text	Description of the operation type
operation_schema_json	text	Field schema as JSON string: used by the interface to generate a form entry with structured metadata fields tailored to the specific operation type. The JSON is stored as text to maintain the defined dictionary order.

Trackable		
Description:	Entries for trackables	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
id_type	varchar(1)	Type of identification used for human-led inventory (e.g., serial number, name)
tid	varchar(255)	ID used for human-led inventory
nickname	varchar(50)	Human-readable name given to a trackable
trackable_type_id	int 8	Type of trackable; foreign key referencing the Trackable Type table
sensitivity	varchar(1)	Data sensitivity for the trackable: P, B, O, C, E, or I, indicating Publicly Released, Business Sensitive, Official Use Only, CRADA Protected, Export Controlled, or ITAR, respectively
status	varchar(1)	Trackable status: P, T, E, C, or V, representing Planned, Transition, Exists, Consumed, and Virtual, respectively
meta_json	jsonb	Provisional field to store and structure metadata related to the specific entry
created	timestamptz	Creation date of the trackable referenced by the record
tracked	timestamptz	Time for which the trackable was last tracked as entry or result of an operation. Set to the starting time of the last operation referencing the trackable as input or output.
qr	text	Unique QR code within the database used to identify the trackable
init_operation_id	int 8	First operation referencing the trackable

Trackable type		
Description:	List of trackable types	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
name	varchar(50)	Name of the trackable type
description	text	Description of the trackable type
trackable_schema_json	text	Field schema as JSON String: used by the interface to generate a form entry with structured metadata fields tailored to the specific trackable type and displayed in the trackable creation form. The JSON is stored as text to maintain the defined dictionary order.

Machine		
Description:	Entries for machines and equipment	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
serial_number	varchar(50)	Serial number or string used to identify the machine
name	text	Human-readable name given to the machine
manufacturer	varchar(50)	Manufacturer name
model	varchar(50)	Model name
procured	date	Date the machine was procured or delivered
location_id	int 8	Location site hosting the machine; foreign key referencing the Location table.
owner_institution_id	int 8	Institution that owns or hosts the machine; foreign key referencing the Organization table.

Machine_Operation_Type		
Description:	Association table for the operation types provided for each machine; explicit declaration of many-to-many relation field.	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
machine_id	int 8	Machine record ID; foreign key referencing the Machine table.
operation_type_id	int 8	Operation type record ID; foreign key referencing the Operation_Type table.
meta_json	text	Field schema as JSON string: used by the interface to generate a form entry containing structured metadata fields customized to the specific operation type conducted by a particular machine. It overwrites the schema outlined in the operation_schema_json and machine_schema_json fields to adapt the metadata form structure according to the specific association of the operation type on the designated machine. The JSON is stored as text to maintain the defined dictionary order.

Machine_Calibration		
Description:	Entries for the machines/equipment calibration over time	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
calibration_timestamp	timestamptz	Time at which the new calibration was applied to the machine
calibration_data	jsonb	Measurements and calibration fields stored as JSON
calibration_type_id	int 8	calibration_type record ID; foreign key referencing the Machine table.
machine_id	int 8	Machine record ID; foreign key referencing the Machine table.

Machine_Calibration_Type		
Description:	List of calibration types	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
calibration_type_name	varchar(50)	Name for the calibration type
calibration_data_template	text	Field schema as JSON string: used by the interface to generate a form entry containing structured metadata fields customized to the specific calibration_types. The JSON is stored as text to maintain the dictionary order.

Operation_Data		
Description:	Entries for the data files stored, URLs, and references access for each operation	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
operation_id	int 8	Operation record ID; foreign key referencing the Operation table
data_path	varchar(255)	Measurements and calibration fields stored as JSON
data_type_id	int 8	Data type record ID; foreign key referencing the Data_type table
tag_id	int 8	Data tag record ID; foreign key referencing the Data_Tag table
highlight	bool	Boolean to label the referenced data as a highlight for the operation

Data_Tag		
Description:	List of tags to classify operation data by categories	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
tag	varchar(25)	Tag name

Data_Type		
Description:	List of data types and extension	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
name	varchar(25)	Data type name (e.g., PDF, Image)
extension	varchar(25)	File extension, if relevant

User		
Description:	Users entries	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
username	varchar(25)	User’s username
first_name	varchar(50)	User’s first name
last_name	varchar(50)	User’s last name
organization_id	int 8	Record ID of the user’s main institution; foreign key referencing the Organization table

Organization

Description:	Entries for the user’s institutions	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
name	varchar(50)	Institution’s name
parent_organization	int 8	Parent institution record ID if the institution pertains to a division, department, or other component of an institution. Reference a primary key of the table.

Location		
Description:	Entries for the users	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
tag	varchar(25)	Tag name

Project		
Description:	List of the projects sponsoring the operations	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
name	varchar(255)	Project’s name
sponsor	text	Name of the project’s sponsor
charge_code	varchar(225)	Charge code for the project

Project_User		
Description:	Users/projects association table to identify the user’s involvement in projects and define access permission; explicit declaration of many-to-many relation field	
Field	Data type	Description
id	int 8	Unique ID to identify the record, primary key
project_id	int 8	Project record ID; foreign key referencing the Project table
user_id	int 8	User record ID; foreign key referencing the User table
is_pi	int 4	Boolean defining if the user is one of the principal investigators for the project
permission_level	int 8	Provisional permission level to grant to the user view, edit, or management privileges for the project and related operations/trackables

APPENDIX B. RENISHAW AM250 DIGITAL HANDBOOK ENTRY

Renishaw AM250

The Renishaw AM250 is a single-laser laser powder bed fusion printer with a build volume of 250 × 250 × 365 mm. This additive manufacturing operation acts on powdered metal feedstock and instantiates new parts within the MDF Digital Platform. Oak Ridge National Laboratory (ORNL) typically prints a wide range of metal alloys on this system.

Category	**Value**
Digital Platform Tag(s)	RenishawAM250-008W73
Digital Point of Contact	Zackary Snow
Approximate Data Volume per *Operation* (GB)	250 GB
Approximate Number of *Operations* per Year	25
Associated Software Tools	Peregrine, Magics

Data Collection and Transfer

Peregrine is installed at the edge on a local compute node with a user display (i.e., a desktop computer on a rolling cart). The Peregrine computer is able to communicate with the printer control computer through WinSCP, an open-source SSH File Transfer Protocol, allowing the Peregrine computer to directly access folders on the printer control computer to distribute build files and retrieve log files. The Peregrine computer is also connected via USB to three cameras (a 20 megapixel [MP] Basler acA5472-17um, a 4.2 MP Pixelink PL-D734MU-NIR-T, and a 0.33 MP FLIR Boson 640), which observe the powder bed. Automated analysis of the live video stream from the Basler camera is used to trigger layer-wise image capture for all three cameras. Each layer-wise image is then stored on the Savitar file storage system, and the edge instance of Peregrine locally analyzes the data using the neural networks and collates the data for the entire build. After the build is complete, the operator loads the log file onto Savitar via the edge Peregrine instance. Metadata are entered into either Peregrine or the Digital Tool before, during, and after each build.

Powder Bed Imaging

A 20 MP grayscale camera (Basler acA5472-17um), sensitive to light in the visible spectrum, captures an image of the entire print area immediately after powder fusion and after powder spreading for each layer. The Basler is outfitted with an ultraviolet/near-infrared (NIR) cutoff filter to protect the camera detector from the laser. A 4.2 MP camera (Pixelink PL-D734MU-NIR-T) captures temporally integrated thermal NIR images of the entire build area, resulting in three images: integrated sum, integrated max, and integrated time-of-max. The Pixelink has a narrow bandpass filter centered in the NIR range (808 ± 10 nm), as well as neutral density filters (outer diameter ≈ 1.8) to prevent image saturation. Finally, a 0.33 MP FLIR Boson 640 captures post-melt and post-recoat images with a reduced field of view. All three imaging systems were designed and installed by ORNL. Anomalies observable with this system include damage to the recoating mechanism, improper spreading of the powder, swelling or distortion of the part geometry, damage to the as-printed components, abnormal generation of spatter and other melt pool ejecta, and improper fusion of the powder.

Machine Health Data

At the end of each build, a log file is produced, which reports various machine error states, as well as temporal sensor streams including build chamber gas (argon) flow rates, build chamber oxygen concentrations, powder roller loading conditions, and the temperatures of selected components in the laser optic trains. This system is installed and maintained by Renishaw, and the resulting log files are loaded into Peregrine at the end of each build.

Data Visualization and Analysis

In situ data analyses are currently performed by the Peregrine software tool, as well as embedded versions of the Pterodactyl package.

APPENDIX C. RENISHAW AM400 DIGITAL HANDBOOK ENTRY

Renishaw AM400

The Renishaw AM400 is a single-laser laser powder bed fusion printer with a build volume of 250 × 250 × 365 mm. This additive manufacturing operation acts on powdered metal feedstock and instantiates new parts within the MDF Digital Platform. Oak Ridge National Laboratory (ORNL) typically prints a wide range of metal alloys on this system.

Category	Value
Digital Platform Tag(s)	RenishawAM400-1CVT28
Digital Point of Contact	Zackary Snow
Approximate Data Volume per *Operation* (GB)	250 GB
Approximate Number of *Operations* per Year	75
Associated Software Tools	Peregrine, Magics

Data Collection and Transfer

Peregrine is installed at the edge on a local compute node with a user display (i.e., a desktop computer on a rolling cart). The Peregrine computer is able to communicate with the printer control computer through a file transfer protocol, allowing the Peregrine computer to directly access folders on the printer control computer to distribute build files and retrieve log files. The Peregrine computer is also connected via USB to two cameras (a 20 megapixel [MP] Basler acA5472-17um and a 4.2 MP Pixelink PL-D734MU-NIR-T), which observe the powder bed. Automated analysis of the live video stream from the Basler camera is used to trigger layer-wise image capture for both cameras. Each layer-wise image is then stored on the Savitar file storage system, and the edge instance of Peregrine locally analyzes the data using the neural networks and collates the data for the entire build. After the build is complete, the operator loads the log file onto Savitar via the edge Peregrine instance. Metadata are entered into either Peregrine or the Digital Tool before, during, and after each build.

Powder Bed Imaging

A 20 MP grayscale camera (Basler acA5472-17um), sensitive to light in the visible spectrum, captures an image of the entire print area immediately after powder fusion and after powder spreading for each layer.

The Basler is outfitted with an ultraviolet/near-infrared (NIR) cutoff filter to protect the camera detector from the laser. A 4.2 MP camera (Pixelink PL-D734MU-NIR-T) captures temporally integrated thermal images of the entire build area, resulting in three images: integrated sum, integrated max, and integrated time-of-max. The Pixelink has a narrow band-pass filter centered in the NIR range (808 ± 10 nm), as well as neutral density filters (outer diameter ≈ 2.4) to prevent image saturation. Both imaging systems were designed and installed by ORNL. Anomalies observable with this system include damage to the recoating mechanism, improper spreading of the powder, swelling or distortion of the part geometry, damage to the as-printed components, abnormal generation of spatter and other melt pool ejecta, and improper fusion of the powder.

Machine Health Data

At the end of each build, a log file is produced, which reports various machine error states, as well as temporal sensor streams including build chamber gas (argon) flow rates, build chamber oxygen concentrations, powder roller loading conditions, and the temperatures of selected components in the laser optic trains. This system is installed and maintained by Renishaw, and the resulting log files are loaded into Peregrine at the end of each build.

Data Visualization and Analysis

In situ data analyses are currently performed by the Peregrine software tool, as well as embedded versions of the Pterodactyl package.

APPENDIX D. CONCEPT LASER M2 DIGITAL HANDBOOK ENTRY

Concept Laser M2

The Concept Laser M2 is a dual-laser laser powder bed fusion printer with a build volume of 245 × 245 × 350 mm. This additive manufacturing operation acts on powdered metal feedstock and typically instantiates new parts within the MDF Digital Platform. Oak Ridge National Laboratory (ORNL) typically prints materials such as stainless steel 316L or 316H on this system.

Category	Value
Digital Platform Tag(s)	ConceptLaserM2-ORNL1
Digital Point of Contact	Luke Scime, Zackary Snow
Approximate Data Volume per *Operation* (GB)	250
Approximate Number of *Operations* per Year	50
Associated Software Tools	Peregrine, SWAN, Magics, WRX Control

Data Collection and Transfer

Peregrine is installed at the edge on a local compute node with a user display (i.e., a desktop computer on a rolling cart). Both the Peregrine computer and the printer control computer are connected to the same subnetwork, allowing the Peregrine computer to directly access any folders on the printer control computer. The Peregrine computer is also connected via USB to two cameras (a 20 megapixel [MP] Basler acA5472-17um and a 4.2 MP Pixelink PL-D734MU-NIR-T), which observe the powder bed. Automated analysis of the live video stream from the Basler camera is used to trigger layer-wise image capture for both cameras. Each layer-wise image is then stored on the Savitar file storage system, and the edge instance of Peregrine locally analyzes the data using the neural networks and collates the data for the entire build. After the build is complete, the operator loads the log file and the on-axis melt pool data onto Savitar via the edge Peregrine instance. Metadata are entered into either Peregrine or the Digital Tool before, during, and after each build.

Powder Bed Imaging

A 20 MP grayscale camera (Basler acA5472-17um), sensitive to light in the visible spectrum, captures an image of the entire print area immediately after powder fusion and after powder spreading for each layer.

The Basler is outfitted with a blue light band-pass filter to protect the camera detector from the laser. A 4.2 MP camera (Pixelink PL-D734MU-NIR-T) captures temporally integrated thermal images of the entire build area, resulting in three images: integrated sum, integrated max, and integrated time-of-max. The Pixelink has a narrow band-pass filter centered in the near-infrared range (808 ± 10 nm), as well as neutral density filters (outer diameter ≈ 1.8) to prevent image saturation. Both imaging systems were designed and installed by ORNL. Anomalies observable with this system include damage to the recoating mechanism, improper spreading of the powder, swelling or distortion of the part geometry, damage to the as-printed components, abnormal generation of spatter and other melt pool ejecta, and improper fusion of the powder.

QM Meltpool

This system consists of a photodiode and a high-speed camera, which are coaxially aligned with each laser module such that they continually observe the molten pool throughout the build. This system also indirectly provides information concerning the laser scan path. This system is installed and maintained by Concept Laser and is duplicated for each of the machine's two laser modules. As installed by the OEM, the on-axis photodiode captures light intensity data in the region around the melt pool at a rate of 10–100 kHz and is sensitive in the 350–1,100 nm range. Extensive testing revealed that the overwhelming majority of the sensor signal can be attributed to reflected laser light at 1,060 nm. This signal is not ideal because the thermal emissions from the molten pool and the just-solidified material are more likely to be correlated with subsurface porosity and other salient printing defects than reflected laser light. Therefore, a high-pass optical filter was installed to block out the 1,060 nm laser light.

ORNL has made several additional modifications to this system. In FY 2023, the photodiodes on both laser modules were changed to the PDA100A2 model (initially tested in FY 2020) because it has a larger detector area. At the same time, the sensor gains were lowered to 60 dB to prevent saturation of the signal and decrease the sensor rise time. Next, the signal inversion observed on one of the two laser modules was corrected by independently identifying the background sensor signal and using it as a known baseline intensity. Because it is unknown if this issue (potentially caused by incorrect wiring at the analog-to-digital converter) is common among other Concept Laser M2 printers, the code was written to automatically detect and correct for this issue on a laser-by-laser basis. Other ORNL implemented corrections include (1) a lossless bit shift to correct for an incorrect upstream datatype conversion, (2) a temporal shift used to synchronize the sensor signals with the recorded scanner position, (3) algorithmic filtering to remove laser jump vectors, and (4) avoiding regions of the print area that are susceptible to reflections within the optics.

Machine Health Data

At the end of each build, a log file is produced that reports various machine error states, as well as temporal sensor streams, including build chamber gas (argon) flow rates, build chamber oxygen concentrations, build plate temperature, and the temperatures of selected components in the laser optic trains. This system is installed and maintained by Concept Laser.

Data Visualization and Analysis

All in situ data analyses are currently performed by the Peregrine software tool and the embedded SWAN and Pterodactyl packages. Various visualizations and analytics tools are available for the layer-wise imaging, QM Meltpool, and machine health monitoring data.