ORNL/TM-2020/1724

Development of Monitoring Techniques for Binderjet Additive Manufacturing of Silicon Carbide Structures

Luke Scime
James Haley
William Halsey
Alka Singh
Michael Sprayberry
Amir Ziabari
Vincent Paquit

September 2020

M3TC-20OR0403016

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

Transformational Challenge Reactor

DEVELOPMENT OF MONITORING TECHNIQUES FOR BINDERJET ADDITIVE MANUFACTURING OF SILICON CARBIDE STRUCTURES

Luke Scime
James Haley
William Halsey
Alka Singh
Michael Sprayberry
Amir Ziabari
Vincent Paquit

Date Published: September 2020

M3TC-20OR0403016

CONTENTS

LIST OF FIGURES

ACRONYMS

AI	augmented intelligence
AM	additive manufacturing
API	application programming interface
BH	beam hardening
CAD	computer-aided design
CVI	chemical vapor infiltration
DED	directed energy deposition
DIC	digital image correlation
DL	deep learning
DSCNN	dynamic segmentation convolutional neural network
DSCNN-P	DSCNN-Perceptron
FDK	uncorrected standard analytical
FPGA	field programmable gate array
GPU	graphics processing unit
IoT	Internet of Things
LPBF	laser powder bed fusion
MBIR	model-based iterative reconstruction
MDF	Manufacturing Demonstration Facility
ML	machine learning
MP	megapixel
MWIR	mid-wave infrared
ORNL	Oak Ridge National Laboratory
STL	standard triangle language
TCR	Transformational Challenge Reactor
TIP	Technology Innovation Program
XCT	x-ray computed tomography

ABSTRACT

The Transformational Challenge Reactor (TCR) program is leveraging additive manufacturing (AM) technologies to fabricate nuclear components to be assembled into a fully functional microreactor core. Compared with traditional manufacturing technologies, AM technologies allow (1) observation of the manufacturing process at a much higher resolution in real-time using in situ monitoring technologies to capture the sensor signature that scientifically describes each event occurring over time and space and (2) validation of the manufacturing process quality using domain-informed data analytics techniques as a potential qualification and certification methodology for the final component.

This report provides an update on the program work on binder jetting in situ process monitoring and associated data analytics results, as well as sample placement and tracking for the subsequent chemical vapor infiltration (CVI) process. Examples are provided to illustrate the progress. Elements of the Digital Thread and data management are discussed in the main document, and an extensive supplemental material section is provided detailing the Digital Platform, as well as its implementation and components. In conclusion the path forward for the next fiscal year is discussed.

1. INTRODUCTION

The FY 2019 Transformational Challenge Reactor (TCR) report, *Monitoring for Additive Manufacturing Technologies: Report on Progress, Achievements and Limitations of Monitoring Techniques,* establishes that the implementation and deployment of the Digital Platform requires three elements: (1) a hardware architecture for data exchange and storage, (2) a software platform for advanced data analytics, and (3) an information-rich manufacturing and material database for domain discovery. The previous report also provides an evaluation of the existing sensing modalities installed on selected additive manufacturing (AM) machines and an overview of image processing results. Based on these preliminary results, a path forward is established to improve the Digital Platform through enhancement of the hardware and software components.

This report provides an update of the envisioned implementation for binder jetting technology. A companion report will discuss monitoring of the laser powder bed fusion (LPBF) AM process also used by the TCR program. This document is structured as follows: the first section gives a brief overview of the two binder jet machines, explaining the manufacturing process, the dimensions of the build chamber, and the type of instrumentation provided on the systems. The second section discusses the in situ monitoring results in FY 2020, addressing both hardware and software components and the system developed for tracking samples from the binder jet machines to the chemical vapor infiltration (CVI) process. All results are illustrated with examples of data collection and processing results, followed by the conclusion.

Note that several images included in this document are from parts not produced for the TCR. This choice was to alleviate export control concerns and produce a document that is publicly available.

2. BINDER JET TECHNOLOGIES

2.1 EXONE – INNOVENT AND INNOVENT+

The Innovent system (Figure 1) is a metal printer using binder jetting technology to manufacture parts. It has a build volume of $160 \times 65 \times 65$ mm^3 and is equipped with a liquid deposition head similar to a traditional ink jet printer to deposit the binder used to glue together the metal particles. The completed "green part" is a solid but fragile mixture of binder and metal powder that must be post-processed to

attain a higher final density. The printing process has a build rate of 166 cm^3/h. The manufacturing process is done in five steps: (1) each layer is configured with a specific set of process parameters; (2) a rolling mechanism covers the build plate with a layer of powder; (3) the binder deposition process operates similarly to the ink jet printing process: the print head moves in a pre-defined pattern over the build plate, depositing binder only when needed; (4) a radiative heating element scans the build plate to dry the binder; (5) once the layer completes, the stage supporting the build plate lowers to the desired layer thickness (typically 50 to 200 μm), and the process repeats until completion of the geometry. By default, the Innovent system is only equipped with standard sensing technologies to ensure its correct mechanical operation. There is no standard image base in situ monitoring available to assess the layer of powder or the quality of the binder deposition.

The Innovent+ system is nearly identical to the Innovent system. The primary difference between the two systems are their powder depositions mechanisms. The Innovent uses an eccentricly weighted motor to vibrate the powder hopper while the Innovent+ uses an ultrasonic driver instead.

Figure 1. ExOne – BinderJet Innovent.	Figure 2. ExOne – BinderJet M-Flex.

2.2 EXONE – M-FLEX

The M-Flex system uses the same technology and manufacturing process as the Innovent system (Figure 2). However, the M-Flex has a build volume of $400 \times 250 \times 250$ mm^3 (about 37 times the volume of the Innovent). The printing process is very fast with a build rate of 1600 cm^3/h (10 times the build rate of the Innovent and 13 times the build rate of the X-Line). Like the Innovent, the M-Flex has limited instrumentation.

3. IN SITU PROCESS MONITORING

3.1 EXONE INNOVENT AND INNOVENT+

The ExOne Innovent and Innovent+ systems are Binder Jet printers. One ExOne Innovent and one ExOne Innovent+ are located on Main Campus in Building 4508. Their internal designations are "ExOneInnovent-0137" and "ExOneInnovent-0166," respectively. These machines produce several in situ data streams which are collected and analyzed as part of the TCR program. The following sub-

sections describe each of these data streams, highlighting any changes in hardware, software, or analysis methodologies in FY 2020. The in situ data are analyzed and visualized, together, using the Peregrine software package, which is described in the Digital Platform supplementary material.

3.1.1   Layer-wise Powder Bed Imaging

A 20 mega-pixel (MP) 8-bit grayscale camera, sensitive in the visible spectrum, captures an image of the entire print area immediately after binder deposition and powder spreading for each layer. A top-mounted light strip provides illumination. This system was installed and is maintained, by ORNL. Image capture sequencing is controlled by Peregrine. Examples of a post-binder image and a post-spreading image for the same layer of a build are shown in Figure 3.
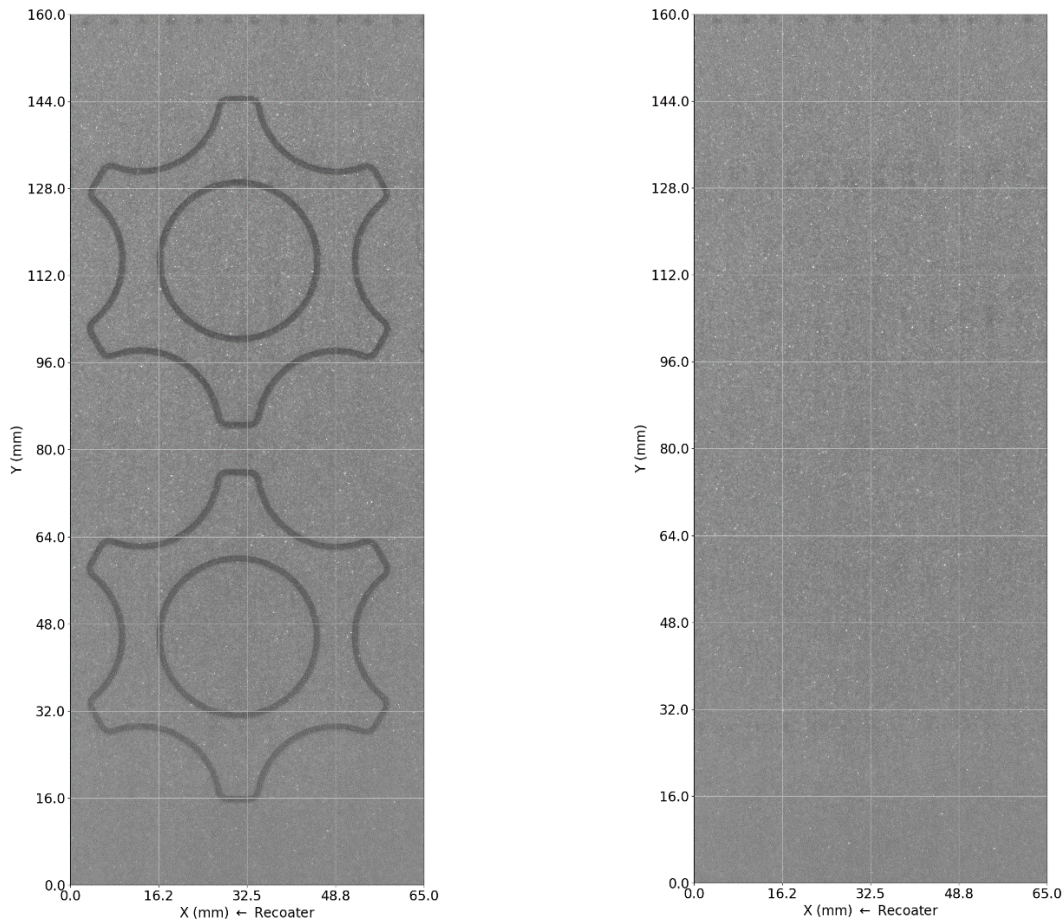
Figure 3. An image taken after binder deposition (left) and an image taken after powder spreading (right).

This imaging system has not changed since FY 2019 and has an effective resolving power of approximately 90 μm. Neural network analysis of these powder bed images, performed by Peregrine, enables the detection of several types of defects and process flaws [1] as well as comparisons between the as-printed geometry and the intended geometry. Detectable anomalies include contamination of the recoating roller, improper powder spreading, and dragging of the bound powder (due to over saturation). Ongoing, early-stage research (not supported by TCR) at the Manufacturing Demonstration Facility

(MDF) is seeking to directly quantify binder saturation levels based on imaging data. If successful, this technique may provide enhanced capabilities for the TCR program in the future.

For comparison to the intended geometry, computer-aided design (CAD) information is first converted into a standard triangle language (STL) format and then "sliced" into layer images by the Arcam Build Assembler software. Development of an alternative ORNL-developed slicing software was started in FY 2020 and will continue in FY 2021. Note that the accuracy of the geometrical comparisons is limited by the resolving power of the camera as well as the registration error between the camera data and the geometry data. Due to the mounting configuration of the camera, data co-registration is less accurate for the ExOne Innovent printers compared with the ConceptLaser M2 printer.

3.1.2   Temporal Log File 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 powder bed surface temperature and various process settings which might be altered by the operator during the build. This system is installed and maintained by ExOne, and new in FY 2020, these data can be viewed using Peregrine, as shown in Figure 4.
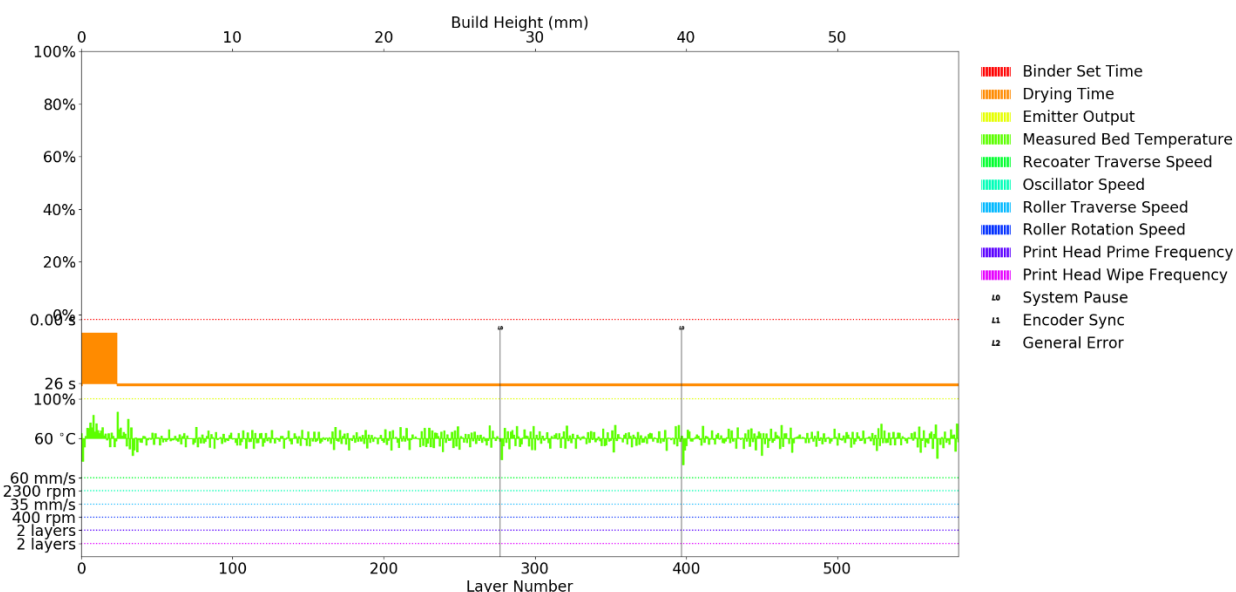
Figure 4. Example log file data displayed as a time series.

3.1.3   Metadata Tracking

Build-specific metadata are tracked using a combination of Peregrine and, new in FY 2020, the Digital Tool web interface. Tracked metadata now include the build name, build start date, build end date, machine-specific processing parameters, data sensitivity levels (e.g., export control restrictions), project and customer information, feedstock material type, feedstock batch information, printer operator information, ambient environmental data, and various operator-input notes. In FY 2021, the Digital Tool's capabilities will continue to be expanded, allowing for increased data search functionality and the ability to view Timelines of feedstock batch utilization and printer calibration and maintenance histories as they relate to each build. A standardized procedure for capturing high-quality post-build images of the printed parts is a goal for FY 2021.

4

3.1.4 Data Transfer

Currently, the log file data are manually recovered from the ExOne Innovent and Innovent+ using an external hard drive, following each build. This hard drive is then hand-delivered to one or more desktop computers located directly adjacent to the printers, and the data are uploaded to the Savitar Data Storage Server. Metadata are tracked using both Peregrine and the Digital Tool web interface. Powder bed imaging data are capture and analyzed in real-time by Peregrine. The raw imaging data and analysis results are upload automatically to the Data Storage Server.

3.1.5 Remote Monitoring

Peregrine allows machine operators to monitor the status of the ExOne Innovent prints remotely through automated email/text alerts as well as a portal (Figure 5) showing the current and previously printed layers with any detected defects highlighted. This system was new at the end of FY 2019 and has been used routinely throughout FY 2020.
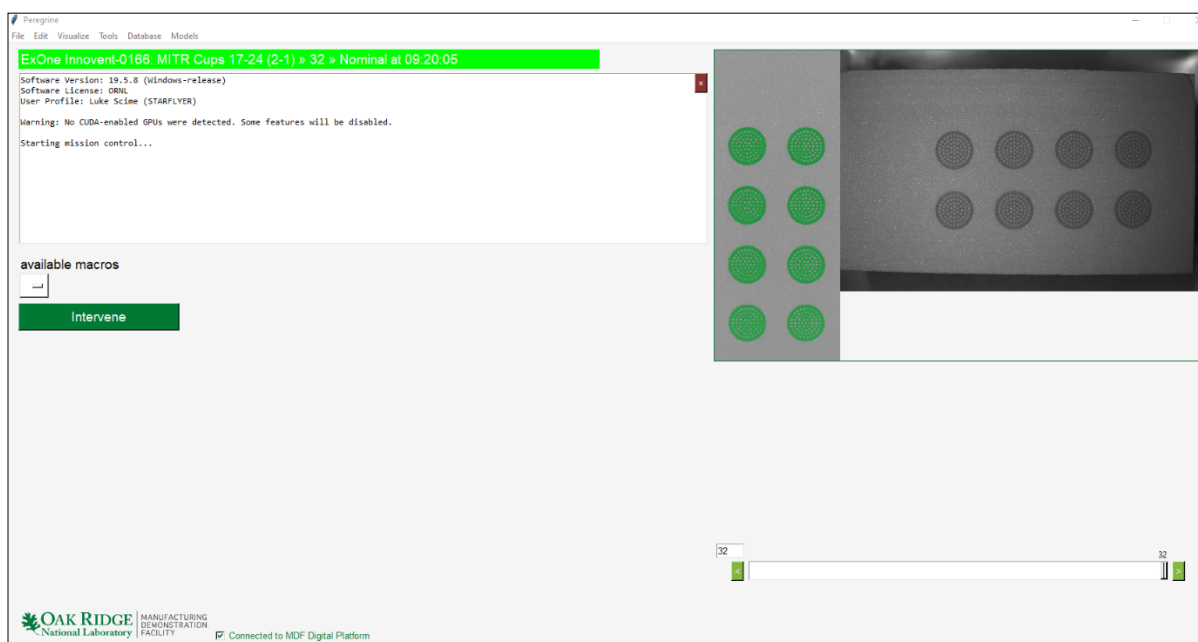
Figure 5. Remote monitoring of an ExOne Innovent build using Peregrine.

3.1.6 Data Visualization and Analysis

The in situ data produced by the ExOne Innovent and Innovent+ printers are analyzed by trained machine operators and data team members using the Peregrine software package. A total of 282 ExOne Innovent builds were analyzed in FY 2020. A broader overview of Peregrine is provided in the Digital Platform supplementary material. To demonstrate the usage of Peregrine in the context of the ExOne Innovent a single case study involving a single build of a prototype reactor fuel element, is discussed below.

Peregrine uses a neural network to analyze and classify each pixel in every layer of powder bed imaging data (Section 3.1.1). One way to visualize these results is by overlaying the classifications on top of the powder bed images, as shown in Figure 6. These classifications can be cumulatively analyzed throughout the entirety of the build and visualized as "heat maps." Figure 7 shows a heat map highlighting

"streaking" caused by contamination of the recoating roller. For this particular machine, this type of streaking is commonly observed in this region of the powder bed. In this build the streaking most likely interfered with one of the fuel elements, designated P3.
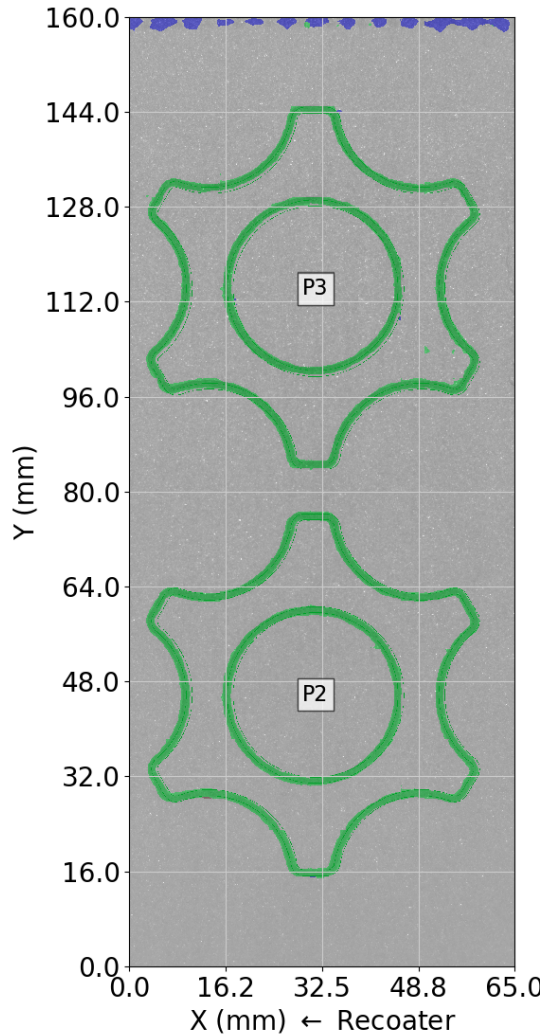
Figure 6. Anomaly detections overlaid on top of a post-binder image of the powder bed.
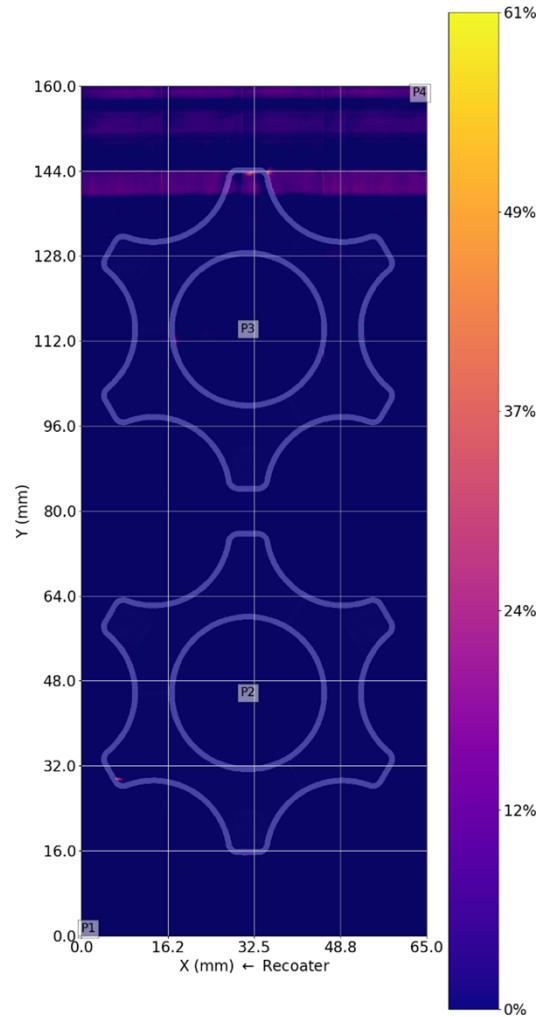
Figure 7. Heat map (projection) of recoater streaking detections throughout the height of the build.

The recoater streaking can also be observed in a 3D reconstruction of the build data (Figure 8). Viewing these data temporally (Figure 9), it can be seen that the build initially started out without any roller contamination, but recoater streaking began to appear at layer 257. The operator attempted to correct this error on two occasions, indicated by the two system pauses; however, the recoater streaking continued to occur. Also note that the powder bed temperature drops by approximately 7°C after each system pause. This could have implications for binder drying rates on subsequent layers.
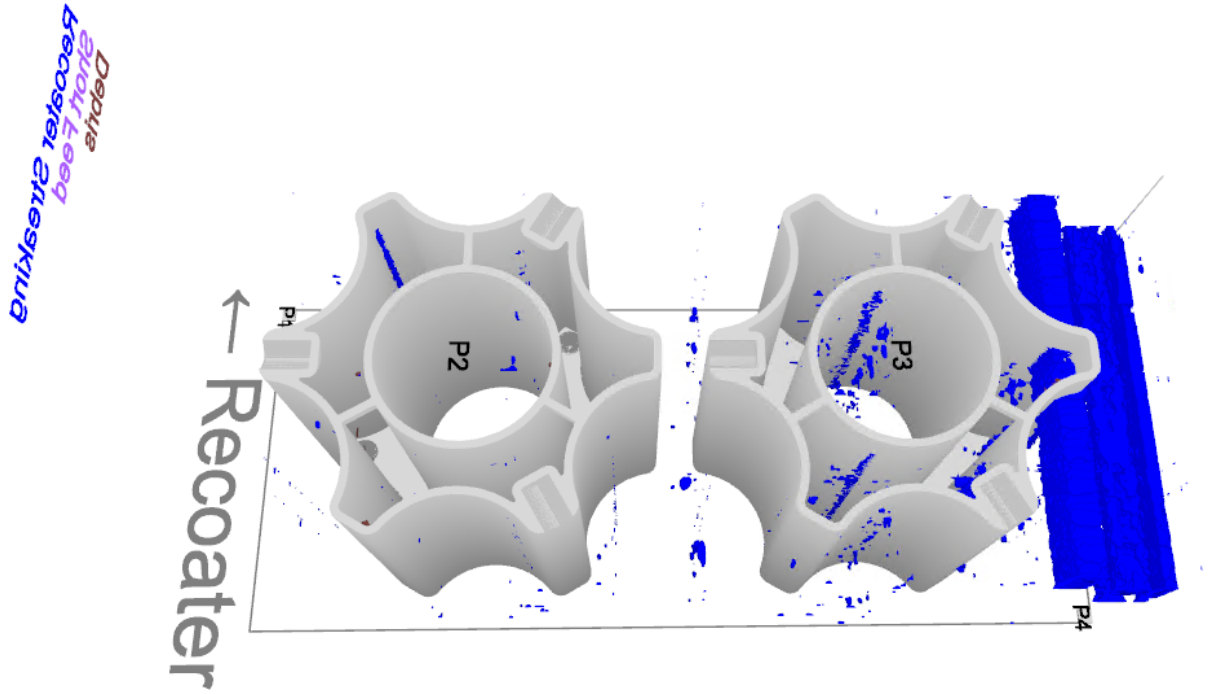
Figure 8. A 3D reconstruction of the entire build.
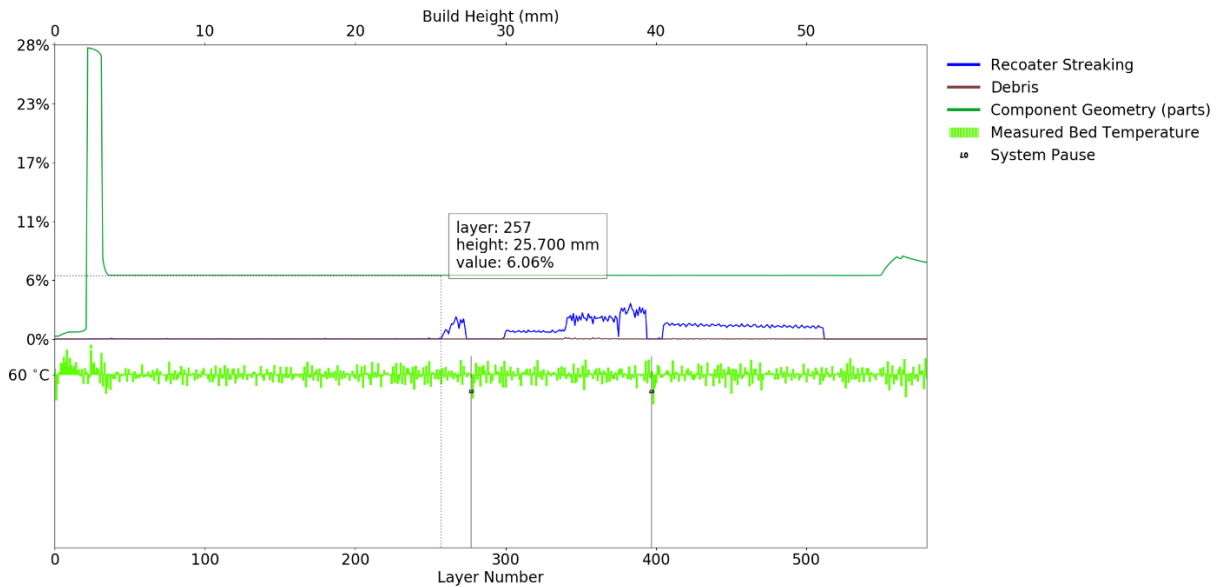
Figure 9. Time series visualization of recoater streaking detections throughout the height of the build.

3.2 EXONE M-FLEX

A single ExOne M-Flex Binder Jet printer is located at the MDF in the low bay. This machine is internally designated "ExOneMFlex-0126" and produces several in situ data streams which are collected and analyzed as part of the TCR program. The following sub-sections describe each of these data streams,

7

highlighting any changes in hardware, software, or analysis methodologies in FY 2020. The in situ data are analyzed and visualized, together, using the Peregrine software package, which is described in the Digital Platform supplementary material.

3.2.1 Layer-wise Powder Bed Imaging

A 10 MP 8-bit grayscale camera, sensitive in the visible spectrum, captures an image of the entire print area immediately after binder deposition and powder spreading for each layer. A top-mounted light strip provides illumination. Additionally, a 0.3 MP 8-bit mid-wave infrared (MWIR) camera captures an image of approximately 80% of the build area immediately after binder deposition and powder spreading for each layer. This system was installed and is maintained by ORNL. Image capture sequencing is controlled by Peregrine.

Due to ORNL environmental health and safety concerns, movement of the M-Flex printer to a new building, COVID-19 restrictions, and TCR funding restrictions, this camera system has not been reinstalled in CY 2020. Reinstallation of the system is a priority for the beginning of FY 2021. During reinstallation, the MWIR camera will be moved to ensure that the entire powder bed is observable, and the system will be modified to protect the optics from the abrasive silicon carbide powder. Examples of a post-binder image and a post-spreading image from each camera for the same layer of a build are shown in Figure 10.

Neural network analysis of these powder bed images, performed by Peregrine, enables the detection of several types of defects and process flaws. Detectable anomalies include contamination of the recoating roller, improper powder spreading, and dragging of the bound powder (due to over saturation). Ongoing, early-stage research (not supported by TCR) at the MDF is seeking to directly quantify binder saturation levels based on imaging data. If successful, this technique may provide enhanced capabilities for the TCR program in the future. Comparison of the as-printed geometry to the intended geometry has not been demonstrated for this printer but is expected to be possible upon reinstallation of the imaging system in FY 2021.
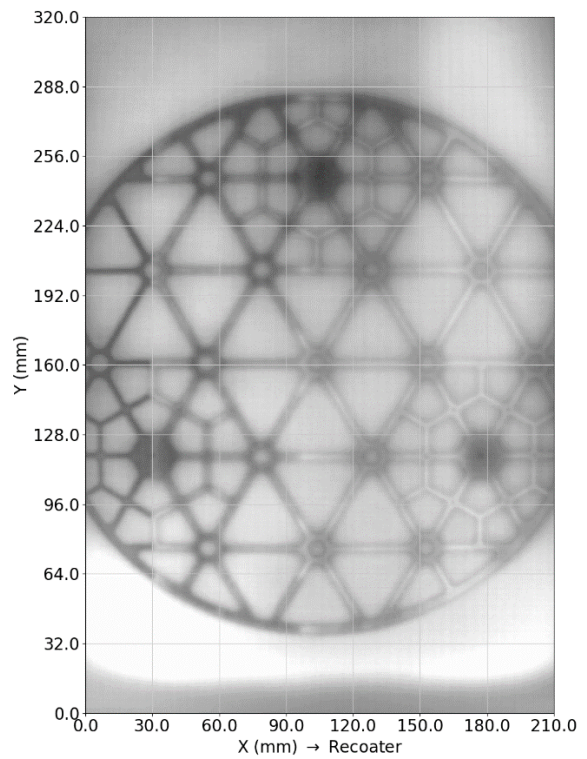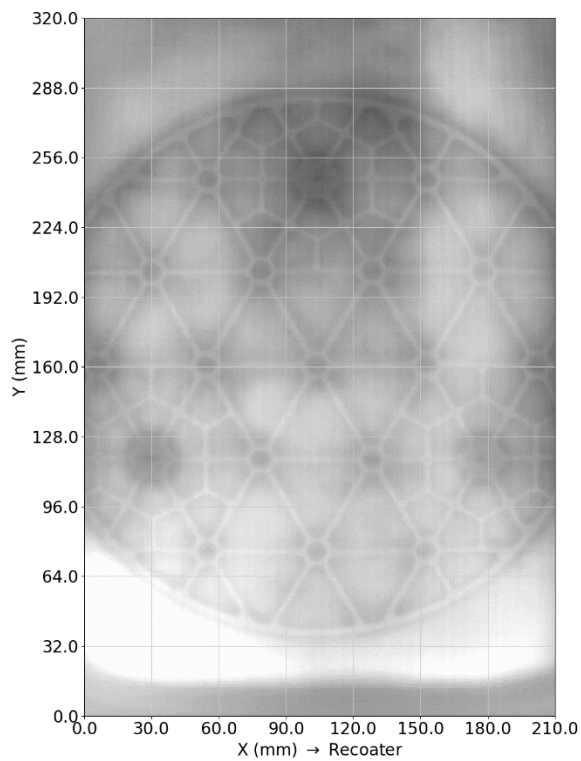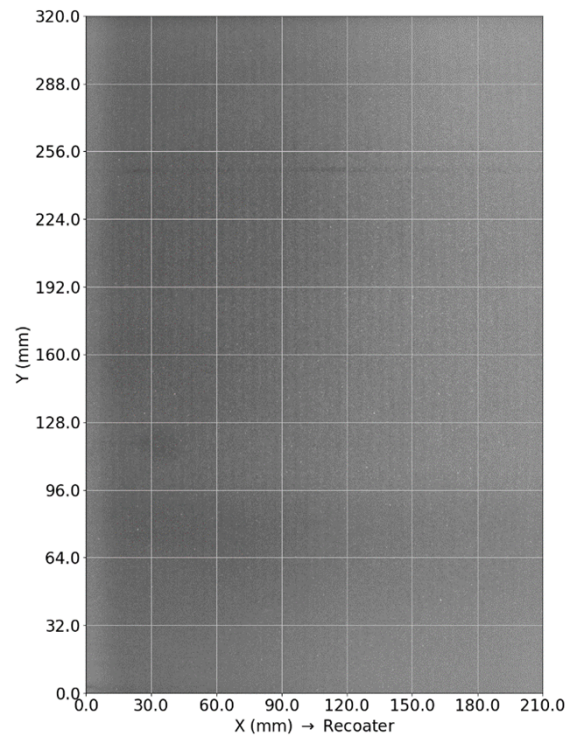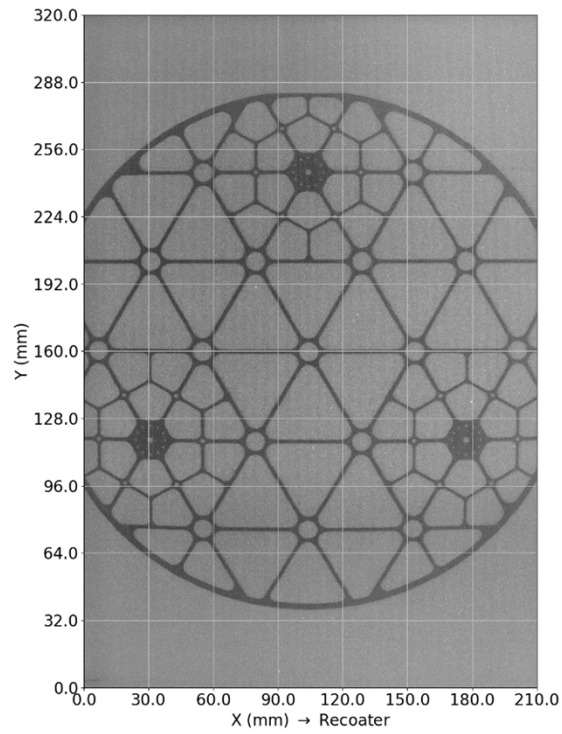
Figure 10. A visible-light image taken after binder deposition (top-left), a visible-light image taken after powder spreading (top-right), a MWIR image taken after binder deposition (bottom-left), and a MWIR image taken after powder spreading (bottom-right).

9

3.2.2   Temporal Log File 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 powder bed surface temperature and various process settings which might be altered by the operator during the build. This system is installed and maintained by ExOne, and new in FY 2020, these data can be viewed using Peregrine, as shown in Figure 11.
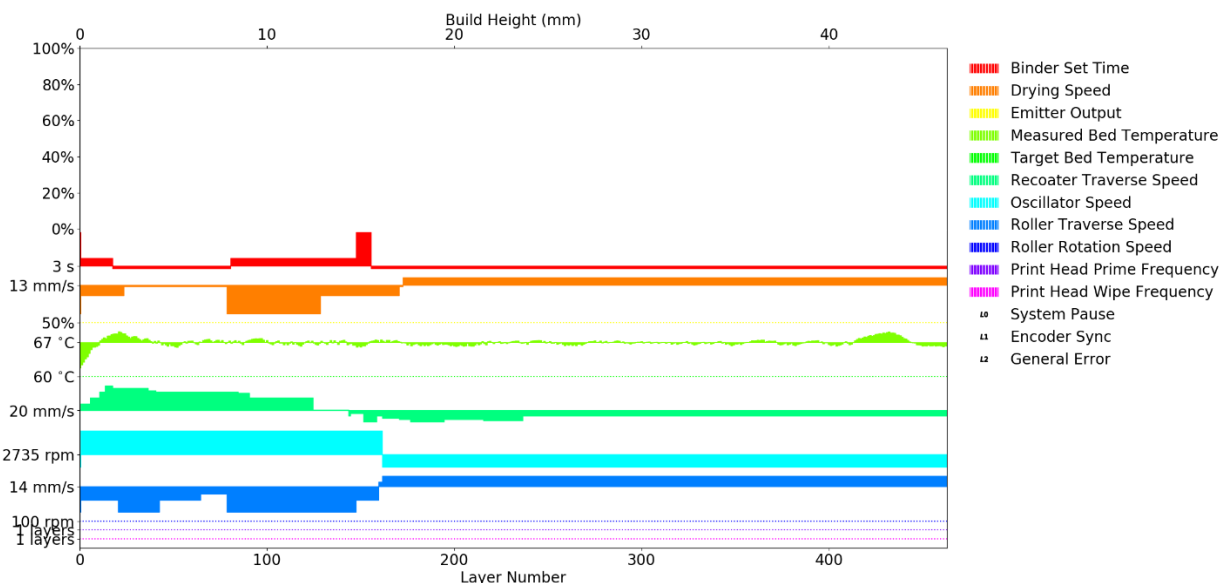
Figure 11. Example log file data displayed as a time series.

3.2.3   Metadata Tracking

Build-specific metadata are tracked using a combination of Peregrine and, new in FY 2020, the Digital Tool web interface. Tracked metadata now include the build name, build start date, build end date, various processing parameters, data sensitivity levels (e.g., export control restrictions), project and customer information, feedstock material type, feedstock batch information, printer operator information, ambient environmental data, and various operator-input notes. In FY 2021, the Digital Tool's capabilities will continue to be expanded, allowing for increased data search functionality and the ability to view Timelines of feedstock batch utilization and printer calibration and maintenance histories as they relate to each build. A standardized procedure for capturing high-quality post-build images of the printed parts is a goal for FY 2021.

3.2.4   Data Transfer

Currently, the log file data are manually recovered from the ExOne M-Flex using an external hard drive, following each build. This hard drive is then hand delivered to a desktop computer located directly adjacent to the printer, and the data are uploaded to the Data Storage Server. Metadata are tracked using both Peregrine and the Digital Tool web interface. Powder bed imaging data are capture and analyzed in real-time by Peregrine. The raw imaging data and analysis results are upload automatically to the Data Storage Server.

3.2.5 Remote Monitoring and Automated Process Control

Peregrine allows machine operators to monitor the status of the ExOne M-Flex prints remotely through automated email/text alerts as well as a portal (Figure 5) showing the current and previously printed layers with any detected defects highlighted. Initial experiments using silicon carbide powder on the ExOne M-Flex suggest that mid-build process parameter modifications will be essential in order to ensure successful prints. Because these prints may be longer than a day in duration, it is not feasible for the technicians to monitor and control these prints in person.

New in FY 2020, Peregrine allows operators to remotely modify the process parameters mid-build for the ExOne M-Flex. Similarly, Peregrine is also capable of autonomously modifying the process parameters during the build. These capabilities were first demonstrated in FY 2019, however, implementation has stalled due to new network security protocols at ORNL. Continuing these remote and autonomous process control efforts is a priority for FY 2021.

3.2.6 Data Visualization and Analysis

The in situ data produced by the ExOne M-Flex are analyzed by trained machine operators and data team members using the Peregrine software package. A total of six ExOne M-Flex builds were analyzed in FY 2020. Because this machine was non-operational for the majority of CY 2020, only limited analysis results are available for FY 2020; however, results and capabilities are expected to be highly similar to those reported for the ExOne Innovent machine. A broader overview of Peregrine is provided in the Digital Platform supplementary material.

3.3 CHEMICAL VAPOR INFILTRATION SYSTEMS

There are currently two operational lab-scale Chemical Vapor Infiltration (CVI) systems at ORNL, internally designated "OakRidgeLabScaleCVI-4515L104" and "OakRidgeLabScaleCVI-4508265." The following discussions regarding data collection, tracking, and analysis relate to these two systems. The first small-scale industrial CVI system will come fully online in FY 2021, with two more industrial CVI systems planned further in the future. The integration of the CVI systems with the Digital Platform is designed to be scalable; however, integration of these larger systems is expected to require additional TCR Digital Team support in subsequent FYs.

3.3.1 Part Location Tracking

Preliminary experiments have suggested that the silicon carbide infiltration rate varies as a function of part location within the CVI furnace. The infiltration rate, and therefore the amount of silicon carbide deposited, directly impacts the density and material properties of the final part. To understand and compensate for this variation, and to maintain tracking of each part throughout the CVI process, it is critical that the locations of the parts within the CVI furnace are recorded in the corresponding Digital Threads. To this end, development of the Software Tool Raven began in FY 2020.

CVI technicians will take multiple pictures of the silicon carbide parts as they are loaded onto the furnace platforms. These images are analyzed by Raven using background subtraction and other classical computer vision algorithms. Raven accesses the Digital Platform to allow the technicians to link each imaged part to the correct tracking ID within the Database. More details about Raven can be found in the supplementary Digital Platform material. Figure 12 shows the prototype imaging setup, and Figure 13 shows example results produced by Raven. These results are uploaded to the Database for integration with Digital Threads via the Digital Tool web interface.

Figure 12. The prototype imaging setup: the camera is suspended below the tripod, and a ring light provides consistent illumination. Several silicon carbide parts are sitting on top of a furnace platform "grid plate" in the center of the assembly. The furnace platform is bounded by a custom calibration target.
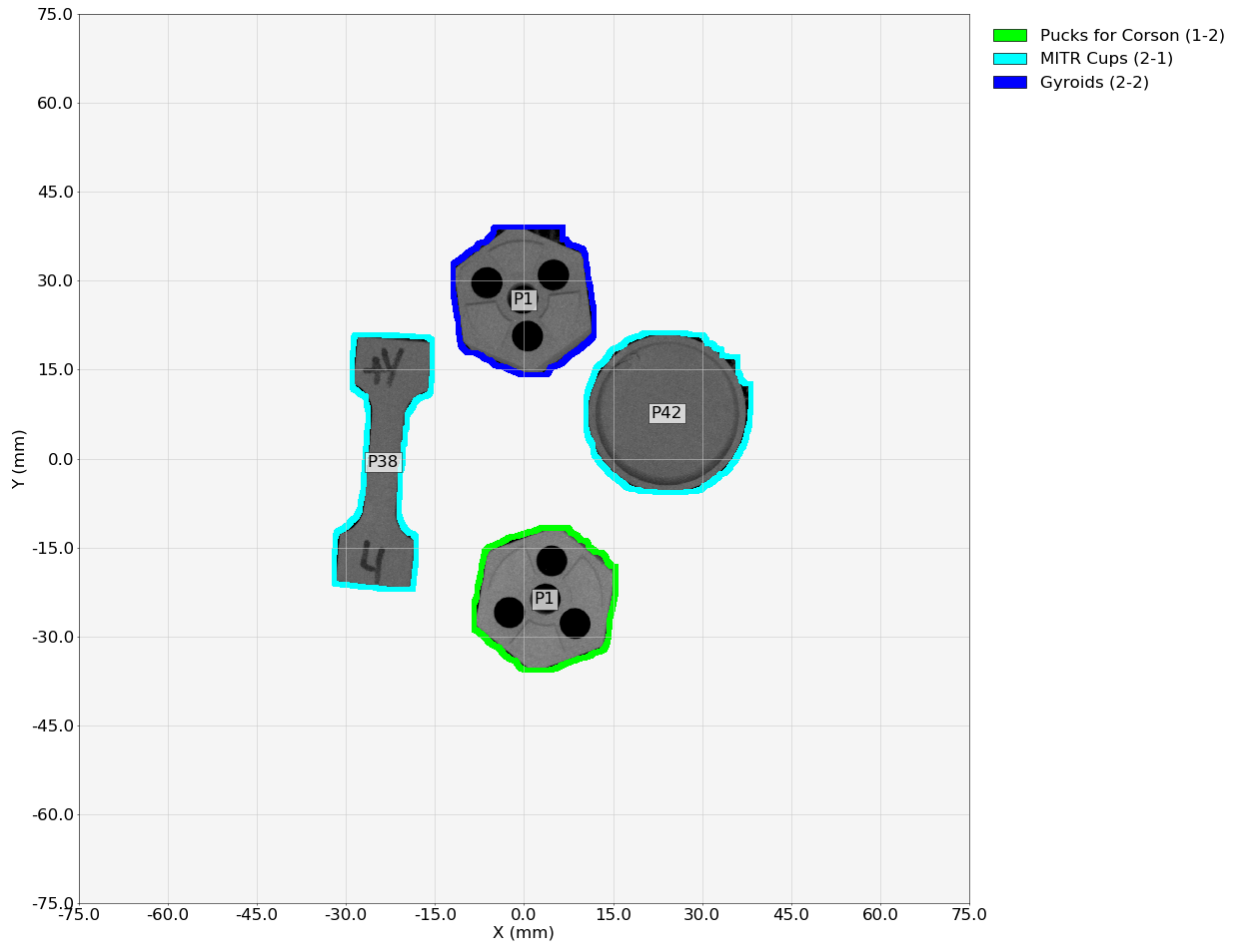
Figure 13. An example Raven output showing parts and part locations extracted from the images. The user has applied the appropriate part tracking IDs from the Database.

In FY 2020, Raven was conceptualized and an initial version was completed and tested on a benchtop. Modifications have also been made to the CVI hardware to facilitate image capture. Assembly and testing of the Raven imaging system was delayed due to COVID-19 restrictions, but the deployment of the Raven systems is a priority for FY 2021.

3.3.2   Temporal Data

A goal of FY 2021 is the collection and tracking of in situ time-series sensing data from several of the TCR CVI systems. These data include chamber pressure, various temperature readings, heating element current draw, and various gas flow rates. These data will be collected at suitable frequencies in the 0.1 Hz to 10 Hz ranges. These data will be uploaded directly to the Digital Platform via a Software Tool with API access. The necessity of developing specific Software Tools for analysis and visualization of these data streams has not been scoped at this point but will be a focus in FY 2021.

3.3.3   Metadata Tracking

To integrate the results from Raven into the Digital Threads, a CVI metadata tracking web form was developed within the Digital Tool. The web form captures and validates the metadata necessary to

13

associate the parts with a particular CVI Operation. The metadata include the CVI machine used, the CVI technician name, the date of the CVI Operation, the duration of the CVI Operation, and any notes from the CVI technician. The results from Raven are also uploaded through the Digital Tool form. Before the backend Database can be updated, a service validates the consistency of the submitted information with the Database and associates the CVI Operation with the parts. Figure 14 shows the metadata tracking web form for the CVI Operation.

Figure 14. CVI metadata entry form, accessed via the Digital Tool.

3.4 PRINTED FUEL ELEMENT DIGITAL WORKFLOW

Manufacture of each silicon carbide fuel element requires a large number of individual steps. Indeed, the TCR team has identified at least nine separate Operations which must be included in each fuel element's Digital Thread. Discounting operations related to feedstock (i.e., silicon carbide powder, binding agents, and TRISO fuel particles) characterization and machine calibrations, the required Operations include the following.

1. Binder Jet printing of the fuel elements on either the ExOne Innovent, the ExOne Innovent+, or the ExOne M-Flex machine
2. Curing of the binding agents in a low-temperature oven
3. Manual depowdering of the green parts
4. An initial mass measurement of each part
5. An initial CVI run of one or more fuel elements
6. Loading the infiltrated fuel elements with TRISO fuel particles and additional silicon carbide filler powder
7. A second mass measurement of each part

8. A second CVI run of one or more fuel elements

9. A final mass measurement of each part

10. Any required metrology or materials characterization Operations

As of FY 2020, the Digital Workflow for Operation #1 is largely complete, as discussed in Sections 3.1 and 3.2 of this report. The Digital Workflows for Operations #3, #4, #7, and #9 are nominally trivial but require substantial extant infrastructure within the Digital Platform. With that infrastructure completed in FY 2020, these workflows will be rapidly tested and deployed in early FY 2021. Initially, the Digital Workflow for Operation #2 will consist only of metadata entry; however, the TCR team is still investigating if in situ temperature measurements via one or more thermocouples is necessary. If necessary, incorporation of these temperature measurements into the Digital Thread is not anticipated to be challenging. The status of data collection for Operations #5 and #8 is described in Section 3.3 of this report. Data collection for Operation #6 has not yet been fully scoped by the TCR team. Overall, testing and deployment of a complete Digital Workflow for fuel element fabrication is a high priority for FY 2021.

4. CONCLUSION

This report proves that image-based defect detection coupled with time-series data analysis in binder jet technology is a very reliable methodology for in situ quality control. Multiple augmented intelligence (AI)–based techniques have been developed to address these problems, and techniques have been successfully deployed on all powder bed systems being used by the TCR program and even other systems that are available at the MDF. In addition, we have developed a computer vision-based solution to reinforce traceability of the items entering the CVI furnace as part of the Digital Platform's Digital Thread. Moving forward, this work will focus on (1) expanding the binder jet database and (2) improving the AI and computer vision algorithms for feature detection.

5. REFERENCES

[1] L. Scime, D. Siddel, S. Baird, 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.* 36 (2020): 101453. https://doi.org/10.1016/j.addma.2020.101453.

APPENDIX A. TCR DIGITAL PLATFORM

A-1. THE DIGITAL PLATFORM

The Digital Platform encompasses the design intent information, networking, and computing infrastructure, in situ sensing hardware, metrology and characterization processes, data storage and management systems, and software and analytics tools necessary to create and effectively leverage a Digital Thread (aka a *digital twin*) for each part and test specimen created for the TCR program. Components of the Digital Platform that have TCR-specific definitions are capitalized throughout this report for clarity. The backbone of the Digital Platform is a searchable Database linking data from component design, in situ sensors, characterization, user entries, machine calibration Timelines, and feedstock utilization Timelines. This information is physically stored on various Data Servers and is accessible via the Digital Tool web interface and associated application programming interfaces (APIs). The Digital Tool and APIs are hosted on the Database Server. Data are processed and visualized on Compute Servers or local user computers (e.g., desktops and laptops) using various ORNL-developed Software Tools; the Software Tools access the Database using standardized APIs. Figure A-1 provides an overview of the Digital Platform.
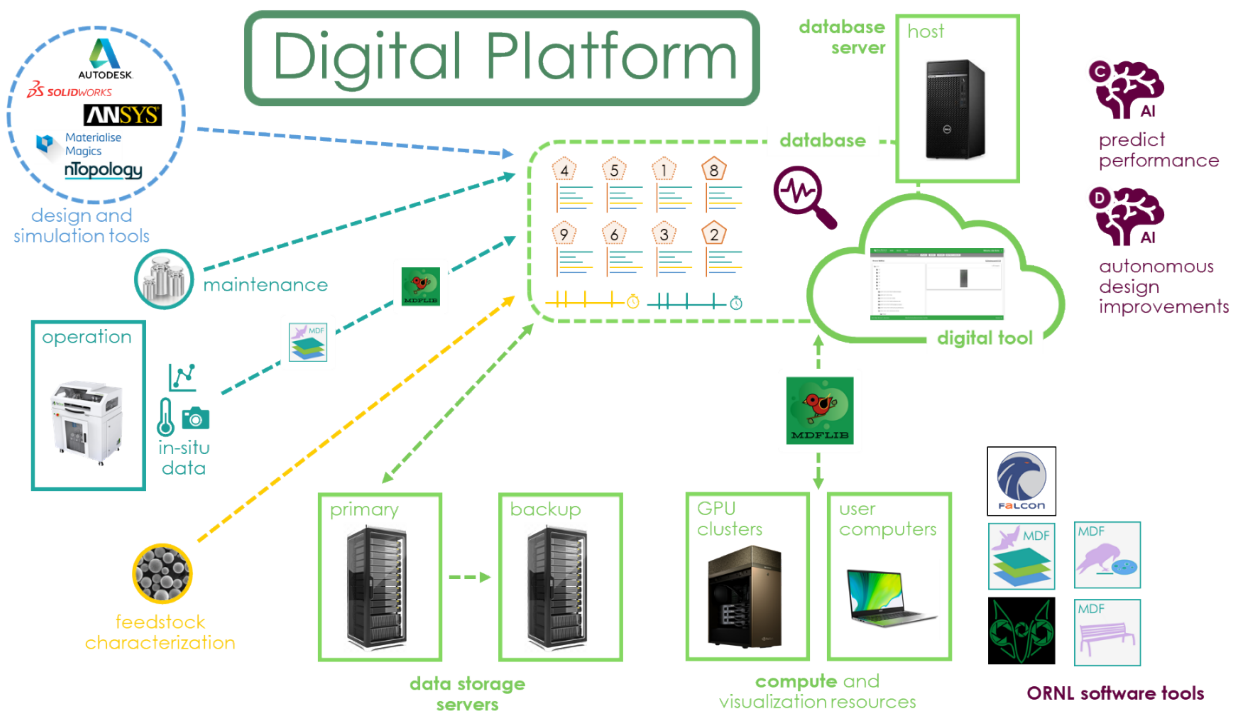
Figure A-1. An overview of the Digital Platform.

Each physical TCR component and test specimen will have a corresponding Digital Thread. Each Digital Thread is composed of building blocks referred to as Operations. Operations may be manufacturing processes, material characterization techniques, maintenance procedures, or purchases. Operations may act on a physical component, material feedstock (e.g., powder or wire), or a machine (e.g., a 3D printer). Example Operations include laser powder bed additive manufacturing (AM), annealing, wire electrical discharge machining, chemical vapor infiltration (CVI), scanning electron microscopy, tensile testing,

purchasing a component from a supplier, and laser calibration. Some Operations may also instantiate new parts or split an existing parent part into multiple child parts. The Digital Thread for a child part includes the thread for its parent.

Functionally, each Digital Thread is a list of Operations that have been performed on a given component and any corollary Operations that compose the machine maintenance and feedstock utilization Timelines. The Database stores these Digital Threads and allows users to access the data (e.g., in situ powder bed images) associated with each Operation. A combination of the Digital Tool web interface and various Software Tools is used to visualize and analyze the Operation data. Data from multiple Operations may be spatially co-registered and visualized using certain Software Tools. Figure A-2 shows a representation of an example Digital Thread.
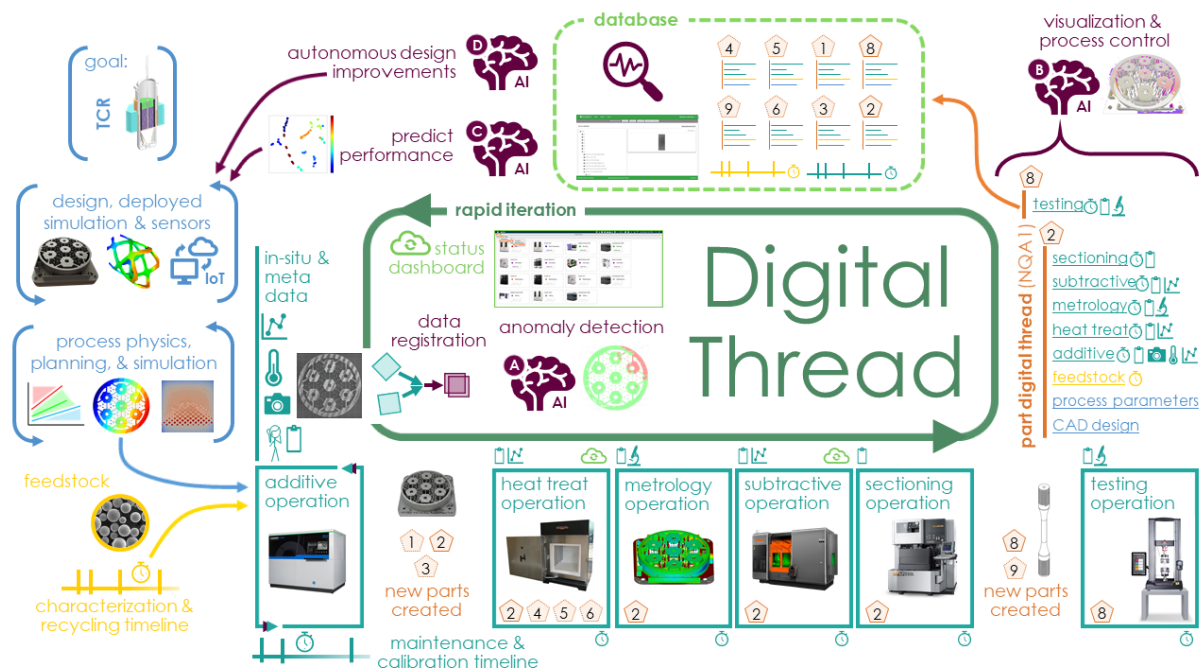
Figure A-2. A representation of an example Digital Thread. The process begins in the upper left-hand corner and follows the arrows counterclockwise.

At a more granular level, information is incorporated into a Digital Thread as part of the Digital Workflow. Many different types of users may interact with a given Digital Workflow, including reactor designers, advanced manufacturing engineers, machine technicians and operators, material scientists, data analysts, program managers, and public communications specialists. The composition of each Digital Thread, and therefore the corresponding Digital Workflow, varies dramatically, depending on the particular component. As an example, Figure A-3 shows a simplified Digital Workflow for a binder jet–printed silicon carbide reactor fuel element.

Figure A-3. A simplified Digital Workflow for the manufacture of a silicon carbide fuel element.

A primary goal of the TCR program is the prediction of reactor-relevant part properties based on in situ data collected during various advanced manufacturing processes. To this end, augmented intelligence (AI) algorithm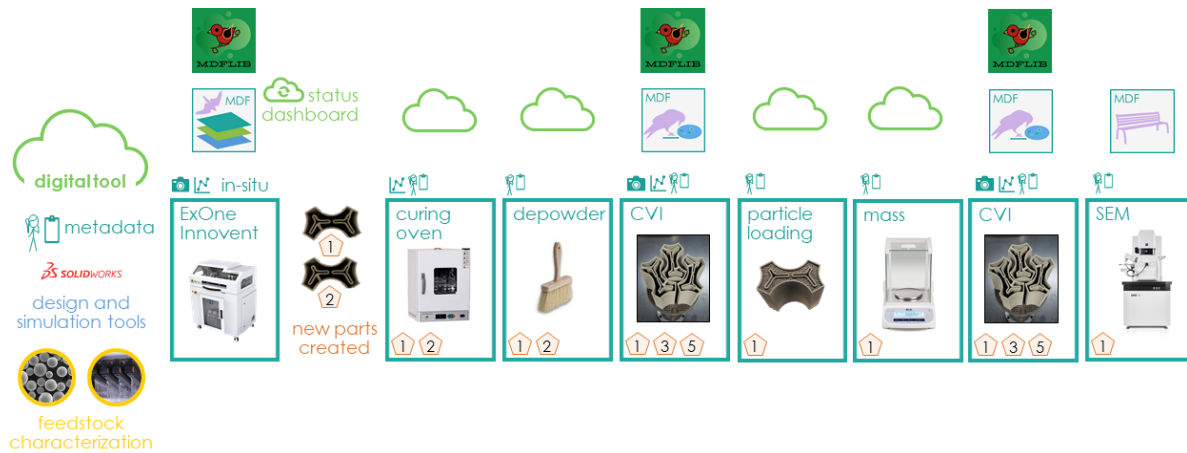s will ingest collections of Digital Threads in order to identify correlations between in situ data and part properties. We deliberately refer to such algorithms as *augmented intelligence* instead of *artificial intelligence* to convey the necessity of keeping expert humans and their physics-based understanding of the world "in the loop" throughout these analyses. Critically, part property prediction will be accomplished through a relay of AI algorithms, with each hand-off leveraging human expertise to inspect the AI performance and to point the subsequent AI in the correct direction. This relay approach is essential to ingest the highly unstructured and extremely high dimensional in situ sensing data and to successfully link it to part property measurements without requiring a prohibitively large number of ex situ characterization experiments. To further clarify the terminology used throughout the TCR reports, *AI* is considered to be a broad umbrella of algorithms which includes heuristics, machine learning (ML), and deep learning (DL). The in situ/ex situ property correlation efforts for FY 2020 are described in a separate TCR report titled *Viability of Data Analytics to Ascertain Component Performance for Additive Manufacturing*. Consistent tracking of each reactor part and test specimen is a critical aspect of this approach. Development of a robust, integrated, and automatic specimen naming scheme was completed in FY 2020 and is illustrated in Figure A-4.
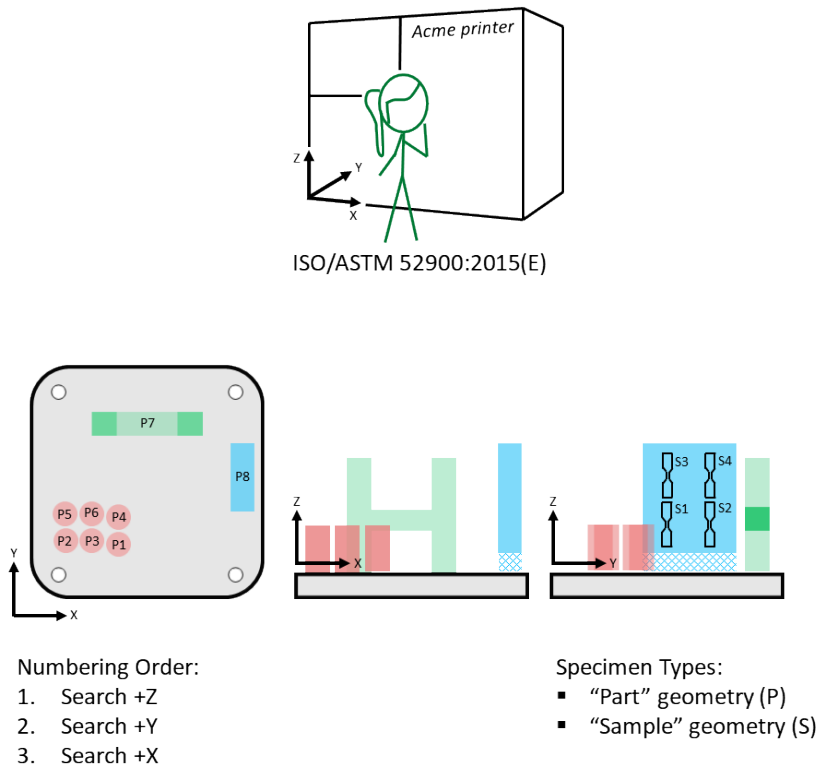
Figure A-4. Specimen numbering scheme following the ASTM 52900:2015(E) standard [1].

Importantly, AI is used for data analysis and visualization throughout the Digital Thread. AI usage for the TCR program can be approximately separated into four different application areas: *descriptive*, *diagnostic*, *predictive*, and *prescriptive*. Descriptive AI is used to analyze raw data, an example being segmentation of powder bed images and identification of process anomalies. Diagnostic AI summarizes large datasets and presents them to a human user or another algorithm. A diagnostic AI might flag certain layers of an AM Operation for closer inspection by a human, or it might decide on an autonomous process intervention to attempt to correct a detected defect. A predictive AI seeks to predict part properties (e.g., fracture toughness) based on in situ data, process parameter information, and part geometry. Demonstration of a predictive AI is an end goal of the TCR program. Finally, a prescriptive AI might autonomously modify a part design in order to improve the predicted performance; this class of AI is beyond the scope of the TCR program, but it is the ultimate goal for the Digital Platform.

In FY 2020, the Digital Platform transitioned from a conceptualization framework and a set of stand-alone tools into a highly flexible, scalable architecture with increasing integration between the Software Tools and the Digital Tool. The physical infrastructure needed to support the Digital Platform includes (1) a optical fiber–based dedicated network to connect all machines to the facility digital backbone, (2) dedicated Wi-Fi networks for Internet of Things (IoT) sensors, (3) redundant storage systems, and (4) multi-GPU (graphics processing unit) compute systems for AI model training and AI at the edge. FY 2020 also saw the standardization of the onboarding process for new Operations to be recognized by the Database. The remaining sections of this supplementary appendix describe the various components of the Digital Platform, including the ORNL-developed Software Tools. Changes and improvements made in FY 2020 are highlighted.

A-2. DATA STORAGE INFRASTRUCTURE AND DATABASE ARCHITECTURE

There is a diverse variety of data produced by a Digital Workflow during different Operations. The data produced vary with respect to the data source, data type, data format, and the scale of the data volume. The standardization of this wide variety of complex data is a challenging problem to solve. To address this complex data organization, the data are structured, organized, and stored on a dedicated Storage Server. This standardization facilitates users and services with improved data search, management, access, and retrieval functionality. The data are organized by the data source (i.e., Operation) and are further organized by the date of the Operation. The dedicated data Storage Server is used as the primary location for storing and accessing the data, whereas a secondary data Storage Server is used for monthly data backups. Some of the benefits that are achieved using a dedicated Storage Server for TCR data include the following.

1. The Data Storage Server acts as a centralized data repository, and the data are available to all the computers within the same network.
2. Multiple users can access the same dataset, eliminating the need to make duplicate copies of data for different users.
3. Data integrity is maintained because the latest datasets are centralized and accessible.
4. The server provides better data management and implementation of security protocols, and it allows data access protocols to ensure that TCR data can only be accessed by authorized users.

A Database Server was installed, and a relational database was developed to store and organize the metadata produced from Operations and associated with the Digital Threads. PostgreSQL was selected as the standard database protocol; PostgreSQL is an open-source, object-relational database system. It is one of the most widely used database systems because of its strong community support, scalability, reliability, and consistent performance. PostgreSQL also provides a robust set of features to support complicated Digital Workflows.

The purpose of the PostgreSQL Database is to facilitate metadata collection, data organization, and storage, along with search and retrieval on the datasets stored on the Data Storage Server. To achieve these requirements, the Database was designed so that relevant data are stored in individual schemas. Some of the schemas present in the current database are as follows.

1. Users: contains information about authorized users of the Database
2. Roles: contains information about which users have access to each Operation and dataset
3. Projects: contains information about different projects at the Manufacturing Demonstration Facility (MDF)
4. Machines: contains information about different manufacturing equipment
5. Materials: contains information about the feedstock materials and their manufacturers
6. Operations: contains information about Operations supported by the Digital Workflows
7. Parts: contains information about all manufactured specimens and the Operations which compose their associated Digital Threads

In addition to the schemas listed above, there are additional AM machine-specific schemas to store information about build datasets and the files in each dataset. These schemas contain metadata that populate the Digital Thread of each manufactured component.

A-3. DIGITAL TOOL (WEB INTERFACE AND APIs)

The Digital Tool is a web-based application developed in Vue.js at ORNL. Vue.js was selected for the development of this platform because of its small size, high performance, and excellent flexibility. Vue.js

also has detailed documentation available, and it supports microservices architectures for developing web applications. Microservices architectures support incremental and modular growth of the application, thus making it a suitable framework for developing scalable infrastructure. Figure A-5 illustrates a microservices architecture.
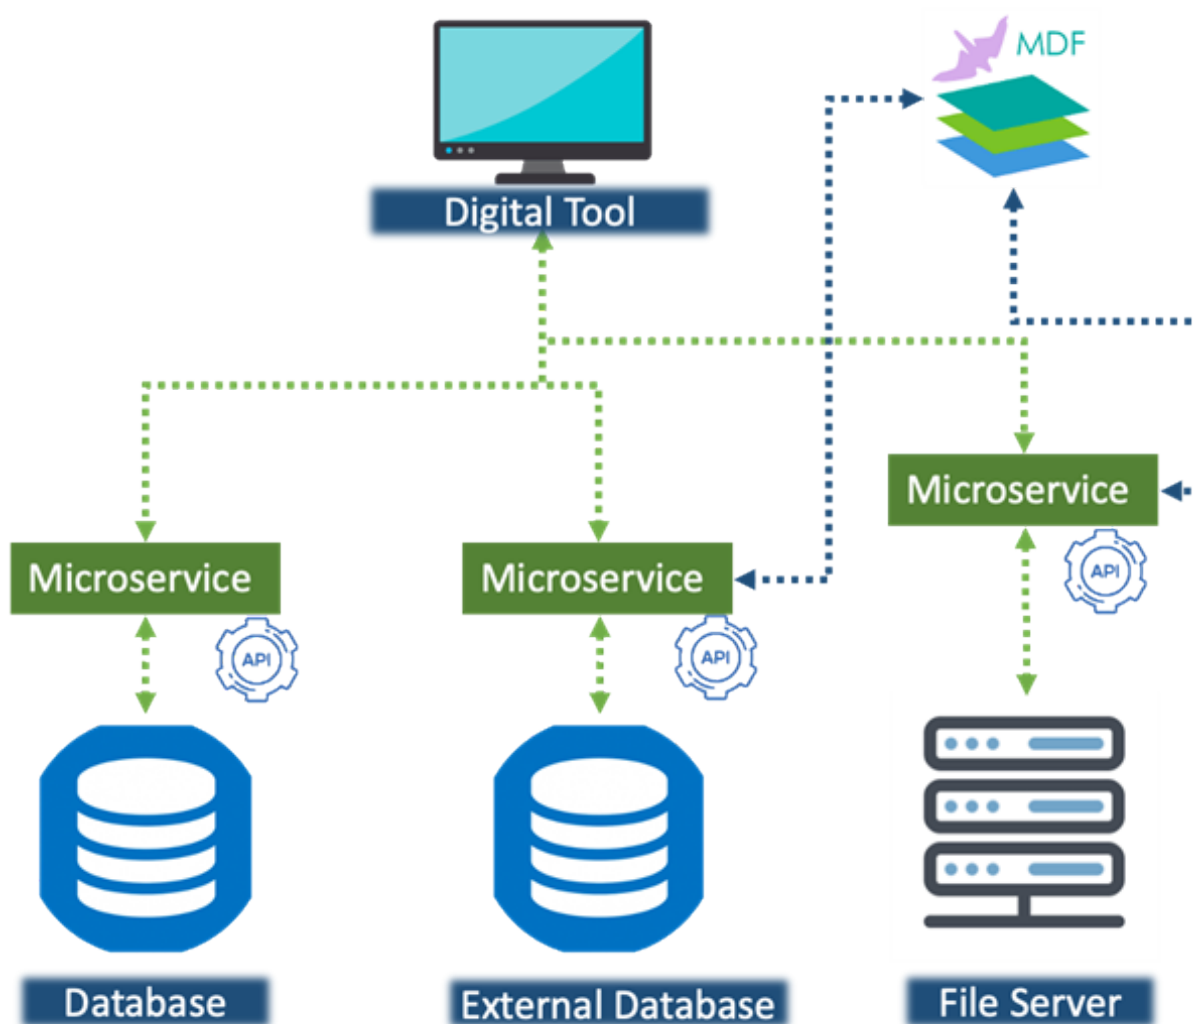

Figure A-5. Representation of the microservices-based architecture utilized by the Digital Tool.

The main components of the Digital Tool microservices architecture are as follows.

Digital Tool: the web application that is developed using JavaScript and html

Microservices: RESTful services developed in Python flask. These services run independently as a service and can communicate and pass information between data sources and the Digital Tool. Additionally, these services can be used by various Software Tools developed by the TCR data team. Currently, there are multiple independent microservices running which vary in functionality from querying databases, updating databases, and performing data transfer actions.

Data sources: microservices act as an interface between the applications and data sources. In this architecture, the data sources are the PostgreSQL Database and central Data Storage Server.

The Digital Tool provides online forms (Figure A-6) to help technicians record metadata before, during, or after an Operation. These metadata tracking forms were developed based on feedback from technicians and the ultimate data users. The forms are validated to ensure that the metadata collected are standardized, complete, and correct.

Figure A-6. An example metadata entry form for an Operation.

The Digital Tool allows users to search (Figure A-7) the existing datasets based on the recorded metadata. This feature enables users to narrow down the desired datasets based on specific metadata, as well as less structured criteria such as the processing parameters used for an additive build.

Figure A-7. An example outcome of a search action.

The Digital Tool also allows users to explore the existing AM build datasets contained on the Data Storage Server. There are multiple data viewers (Figure A-8) to support exploring different data types including image viewers, STL viewers, PDF viewers, and log file viewers. This feature allows multiple users to access data simultaneously without requiring direct access to the stored data, thereby preventing accidental data modification.
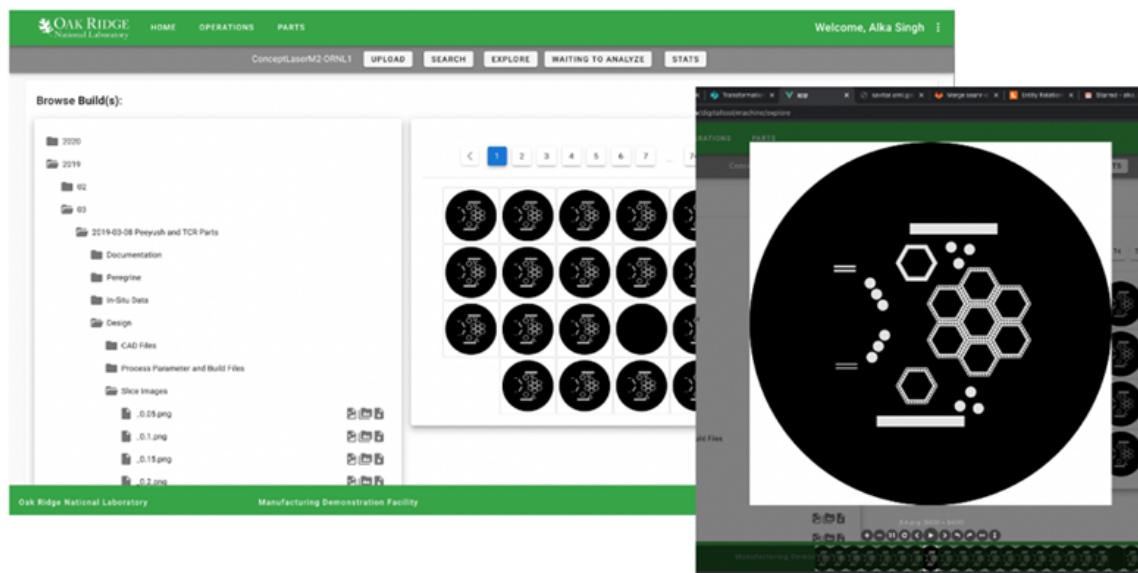

Figure A-8. An example of the image viewer feature.

A-4.   PEREGRINE

Peregrine is a Software Tool developed in Python at ORNL. Peregrine is intended for use with powder bed AM machines (i.e., printers), including laser powder bed fusion (LPBF), electron beam powder bed fusion (EB-PBF), and binder jetting technologies. In its fullest implementation, Peregrine autonomously collects and analyzes layer-wise imaging data, provides remote monitoring and process intervention capabilities, tracks metadata and part information, produces advanced visualizations of both the underlying data and the AI analysis results, and enables identification of correlations between in situ data and process parameters or ex situ measurements.

Peregrine's core capability is pixel-wise anomaly detection based on layer-wise images of the powder bed. This semantic segmentation of the data is achieved using a novel DL algorithm referred to as a *dynamic segmentation convolutional neural network* (DSCNN). The DSCNN architecture was designed specifically to overcome many of the challenges common to powder bed imaging data and is fully described in Scime et al. [2]. Beyond this core capability, Peregrine also empowers users with several advanced analytics and data tracking tools.

1. Each layer in a build can be flagged based on the DSCNN pixel segmentation results using either a set of user-defined heuristics (i.e., rules) or a learned ML model, referred to as the *DSCNN-Perceptron* (DSCNN-P).
2. Peregrine provides advanced data visualization, allowing the user to visualize the raw data, DSCNN, and DSCNN-P results in various formats. These visualizations include image overlays, plots of time series data, frequency analyses, specimen quality scoring, data projections in each orthogonal plane, 3D reconstructions, and time-lapse videos of the build.
3. Peregrine provides automatic specimen numbering and statistics. Up to two sets of computer-aided design (CAD) geometry information can be associated with each build, including the as-built geometries, which are referred to as *parts*, and the geometries of any *samples* removed from the components after the build is complete, such as a tensile bar. The term *specimen* is used by Peregrine to refer to either a *part* or a *sample*.
4. Metadata can be tracked for each build, along with reference images, log files (supported printers only), coaxial time series data (supported printers only), and registered ex situ data. Peregrine also includes 2.5D registration capability, allowing users to overlay ex situ data such as x-ray computed tomography (XCT) slice images on top of in situ data and DSCNN results.
5. Peregrine is capable of direct data collection and real-time analysis using most USB and ethernet cameras. During a *live* analysis, Peregrine can send alert emails to technicians, execute customizable macros to interface with a printer's user interface to effect process interventions, and provide remote viewing of the build status via a companion instance of Peregrine.
6. A limited number of statistics tools are also provided to enable the user to investigate potential correlations between Peregrine's in situ results and process parameters or ex situ measurements.

A high-level overview of Peregrine's capabilities is presented in Figure A-9, and Peregrine's internal data flow is summarized in Figure A-10. The figures illustrate bundling of all of the layer-wise imaging data into a single calibrated image stack, followed by a tiling operation before the data are fed into the DSCNN. The pixel classifications predicted by the DSCNN are then compared with the *template* image, and potentially the raw powder bed imaging data, using pixel heuristics. These comparisons create the final pixel classifications observed by the user. The output of the DSCNN is fed into the DSCNN-P to allow for *labeling* of the entire build layer, or *image stack*. Finally, the *labels* assigned by the DSCNN-P are combined with the global heuristics and presented to the user as layer flags.
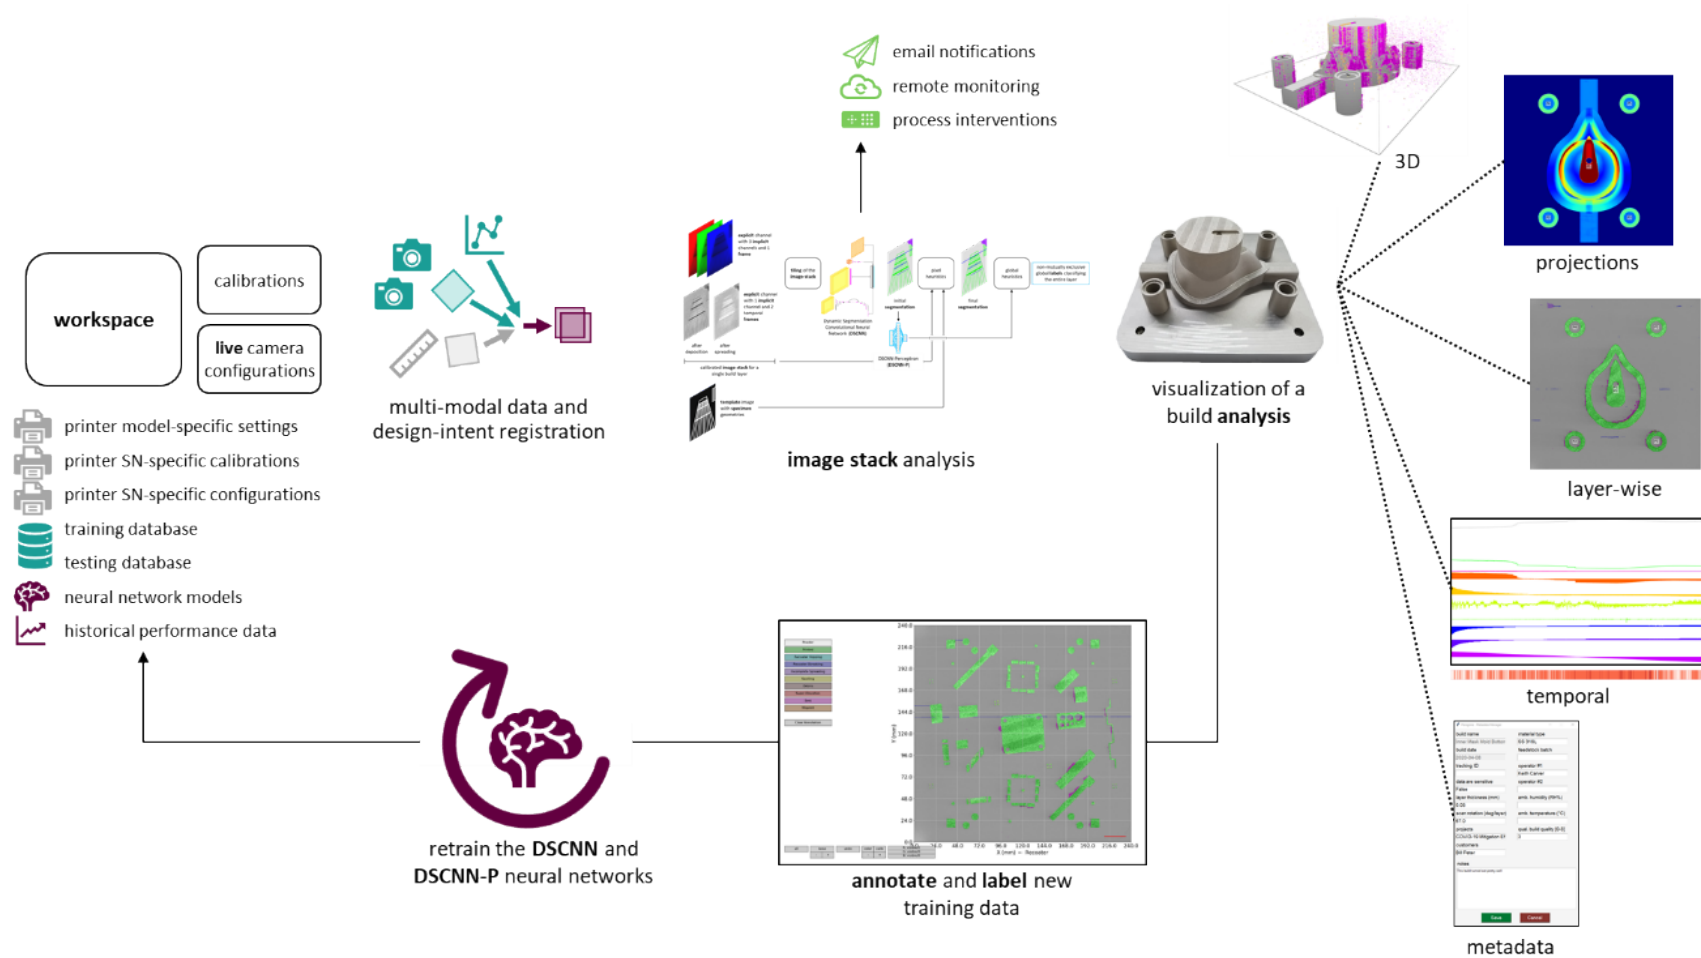
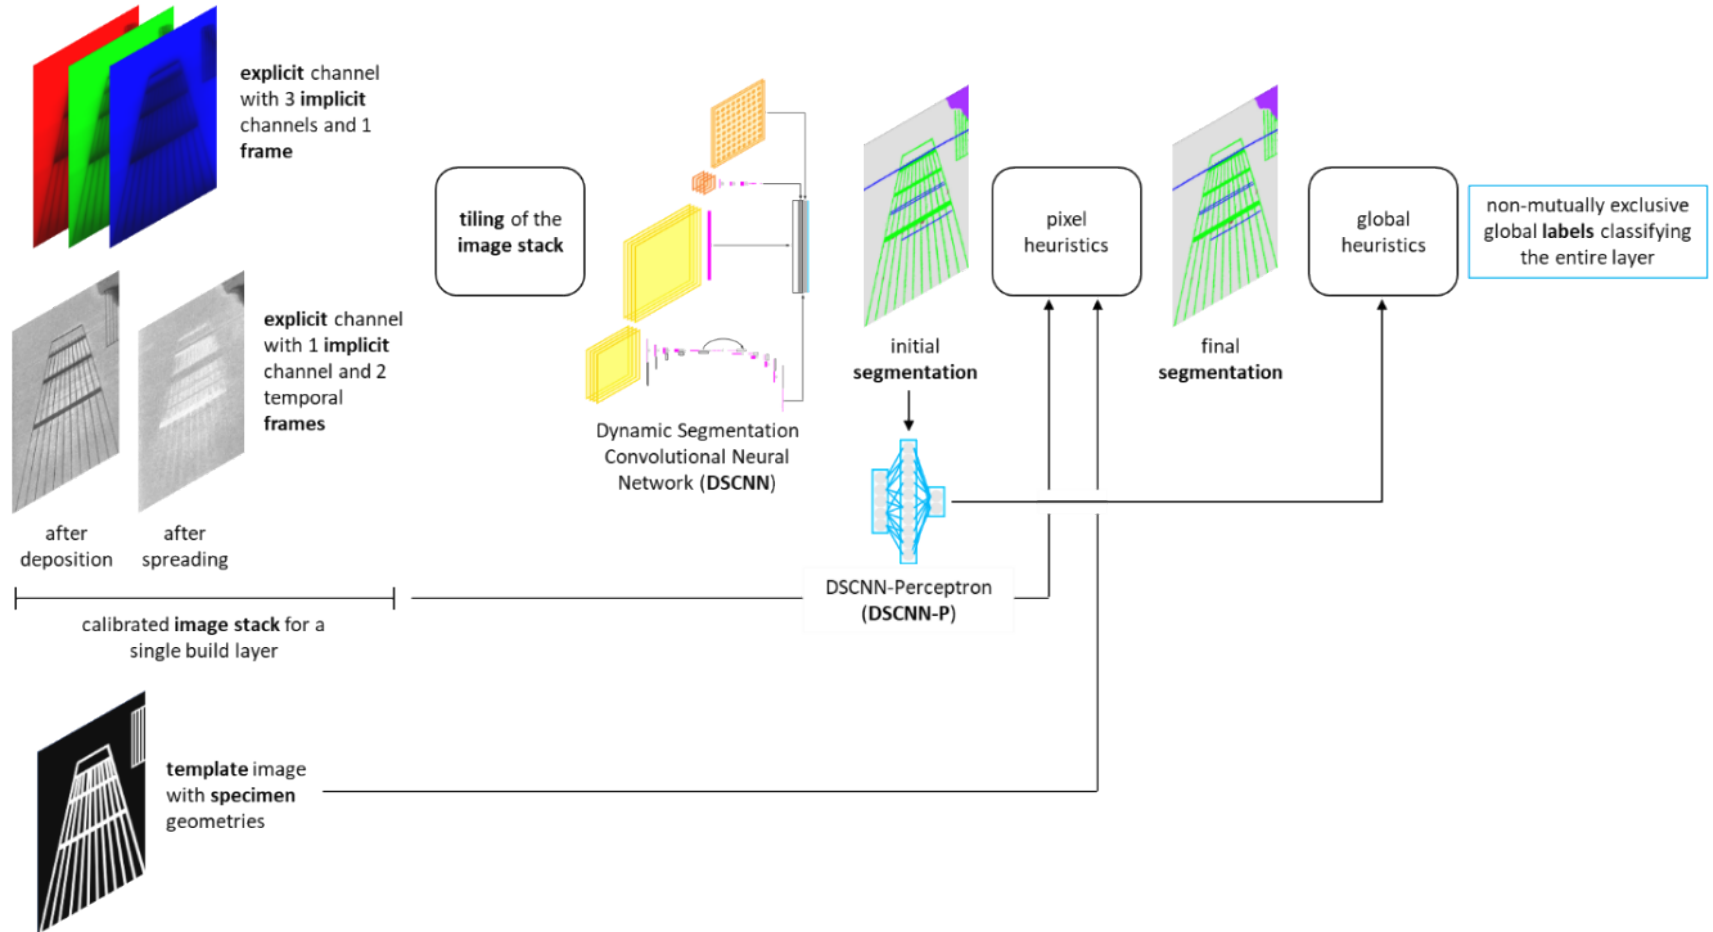**Figure A-9. Overview of Peregrine's capabilities and operation.**

Figure A-10. Peregrine's internal data flow.

Peregrine has been under development since FY 2018. It is currently on major version 20 and is used by approximately 10 members of the TCR team. Users include data scientists, material scientists, and machine technicians. Over 100 pages of documentation and training materials have been written for Peregrine in FY 2020. In FY 2020, ORNL authored a journal paper describing the DSCNN that was published in *Additive Manufacturing*. Peregrine can now be licensed to companies, with over 10 companies already expressing interest. Below is a list of the major Peregrine improvements completed in FY 2020.

1. The software release process was formalized, including documentation, versioning, archiving, and bug reporting.
2. An immutable change log is now created for each build analysis to record all data edits and builds can now be marked as containing sensitive data.
3. Peregrine has been integrated into the Digital Platform, allowing data to be passed to and from the Database. Printer settings, calibrations, trained neural networks, and analyzed builds are hosted on the Data Storage Servers and are transparently accessible to all users at ORNL and those working remotely. Any trained user can directly update training data to improve AI model performance. Versioning improvements allow changes to be tracked as multiple users modify shared files.
4. A part numbering standard was introduced for the TCR program. Peregrine automatically numbers all *specimens* for each build. Automatic and consistent numbering is absolutely critical for tracking and maintaining the hundreds of thousands of Digital Threads to be stored on the Digital Platform.
5. The DSCNN architecture was upgraded to DSCNNv2. This architecture improves accuracy, segmentation of curvilinear structures, and segmentation of smaller structures, and it also reduces inference time and *tiling* artifacts. DSCNNv2 dramatically improves training flexibility and follows more industry standards for neural network design.
6. The entire data loading architecture was rewritten to allow for color images, 16-bit images, a truly arbitrary number of cameras, an arbitrary number of temporal *frames*, and more complex inter-layer lighting differences. These changes make Peregrine much more flexible as the imaging systems on several of the printers are improved in FY 2021.
7. XCT data can now be co-registered to the in situ imaging data and overlaid on top of the DSCNN predictions within Peregrine.
8. Data collected *live* are fully analyzed and accessible in real-time—this is critical for handling the massive TCR data burden.
9. The DSCNN-Perceptron can now be used to flag layers to bring them to the attention of a user or machine operator, and early results using this capability are promising. A training data collection and labeling campaign began in FY 2020 and will continue in FY 2021 to improve the utility of the DSCNN-P predictions.
10. The entire user interface was rebuilt based on over a year of user feedback. The primary goal was to enable faster, more efficient review of the analyzed build data and to enable crowdsourcing for collection of new training data.
11. The ConceptLaser X-Line 2000R machine type is now recognized by Peregrine.
12. Several artifacts can now be removed from the QM Meltpool data.
13. Log files can now be parsed for the ExOne Innovent and ExOne M-Flex machines.

A-5. RAVEN

Raven is a Software Tool developed in Python at ORNL. Raven is intended for use with the CVI process. Specifically, it allows the user to track the locations of the binder jet printed parts prior to insertion into the CVI furnace. CVI technicians will take a series of images of the parts as they are loaded onto the furnace platforms, or "grid plates." These images are taken under controlled lighting conditions and

background subtraction, along with other classical computer vision techniques. The images are used to determine the boundaries of each part. Each part is then assigned an identifier number produced by the Peregrine Software Tool during the instantiating printing process. The part locations and their corresponding identifiers are saved to a file which is uploaded to the Database via the Digital Tool web interface.

Development of Raven began in FY 2020, and it is currently on major version 1. Raven has not yet been deployed to users because COVID-19 restrictions delayed assembly and testing of the physical imaging system. Bringing Raven fully online is a priority for early FY 2021. Approximately 10 pages of documentation and training materials were written for Raven in FY 2020. The intended users of Raven will be the CVI machine technicians.

A-6.   SCOPS

SCOPS is a set of imaging systems and an analysis method developed to analyze the complex thermomechanical history of parts printed with directed energy deposition (DED). The system uses stereo digital image correlation (DIC) to 3D map surface roughness features on parts as they are printed and tracks their motion over time. This allows for direct, in situ, noninvasive deformation and strain measurements over the entire part. This is combined with infrared thermal imaging to measure material expansion and contraction as a function of heat flow through the part. The SCOPS system is summarized graphically in Figure A-11.

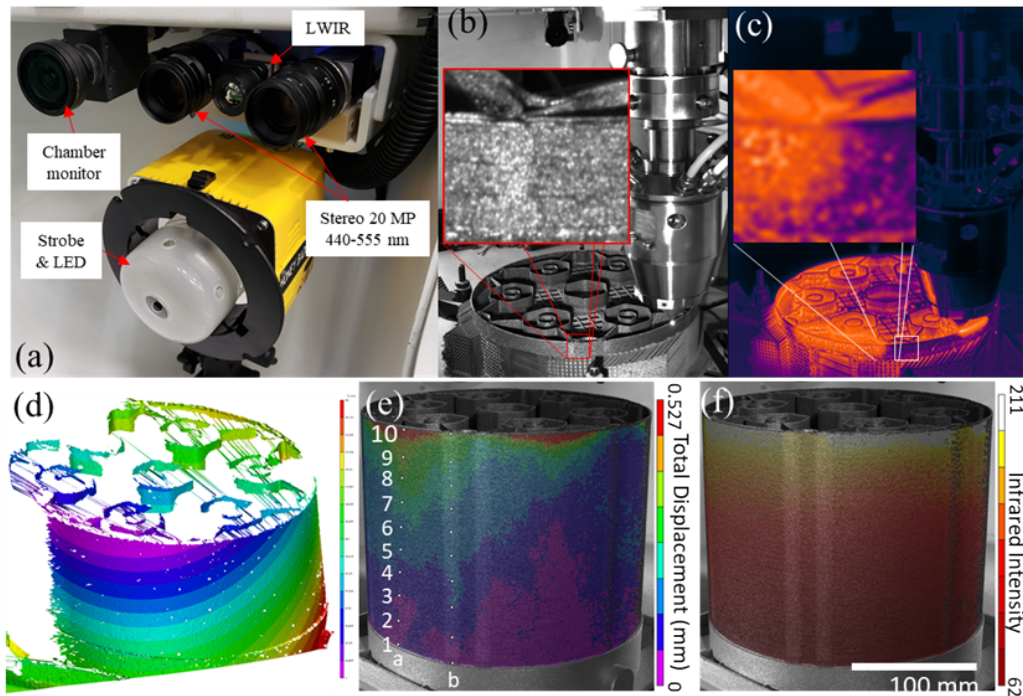**Figure A-11. (a) SCOPS imaging system, two stereo high-resolution visible cameras and an infrared camera, (b) visible trackable surface roughness of representative deposition, (c) infrared imaging showing thermal dissipation, (d) 3D map constructed from surface features, (e) tracked deformation induced from a combination of residual strain and thermal expansion from (f) 3D mapped infrared signal showing hot regions.**

In FY 2020, two DED machines were outfitted with SCOPS: the BeAM modulo 400 and GKN cell 1. System hardware configuration was developed and optimized, along with supporting software for automated video capture and archival. Blackbody calibration was performed, and thermocouple validation of infrared imaging was automated. A validation experiment comparing DIC with neutron diffraction residual stress measurements was performed. Information from SCOPS provides an in situ measure of a previously inaccessible information stream, and it enables direct validation of finite element analysis (FEA) models, enhancing the accuracy of residual stress and microstructure simulation.

A-7. PIGEON

Pigeon was developed at ORNL and is both a Software Tool and a Python API. Pigeon is primarily an image segmentation software tool powered by a DSCNN. For more information regarding the DSCNN, users are referred to Scime et al. [2]. Pigeon's data flow is essentially identical to that represented in Figure A-10 for Peregrine. Some of Pigeon's key features include the following.

1. Pigeon accepts an arbitrary number of input *channels*, working seamlessly with both RGB and single-channel images. Pigeon also supports an arbitrary number of temporal *frames* associated with each *explicit channel*. For example, if two spatially registered color images, each with two temporal *frames*, are provided, then there are two *explicit channels* and 12 *implicit channels*.
2. Pigeon accepts images of arbitrary bit depth (up to 16 bits).
3. Pigeon accepts images of arbitrary resolution while ensuring that the output segmentation is always at the resolution of the input image.
4. The DSCNN is designed to work well with very high-resolution images of 20 MP or greater, even with limited GPU RAM available. It also provides sublinear analysis times, so doubling the image resolution will not double the inference time.
5. A low-resolution leg of the DSCNN always has a receptive field equal in size to the entire input image. This helps the DSCNN to learn to be robust to lighting variations and global scene changes.
6. The medium-resolution leg of the DSCNN has a receptive field which can be dynamically configured by the user. The size of this *tile* should be based on domain knowledge such that it includes the appropriate amount of contextual information.
7. The DSCNN can learn prior probability distributions that are a function of pixel location within the global image.
8. Class-wise imbalances are not a significant issue, because automatic rescaling of the loss function ensures that the DSCNN learns even extremely rare classes.
9. Full transfer learning from any DSCNN to any other DSCNN is supported.
10. The user can create multiple *workspaces* for working on different problems.
11. The DSCNN also outputs a deep feature vector which can be used by a perceptron to classify the entire *image stack*.

Development of Pigeon started in FY 2020. It is currently on major version 1 and is used by several members of the TCR team. Over 30 pages of documentation and training materials have been written for Pigeon in FY 2020. Users include data and material scientists. Pigeon is intended primarily as a development platform for creating different characterization tools. For example, Pigeon can serve as the framework for multiple Software Tools, such as one dedicated to characterizing powder feedstock and another focused on analyzing AM microstructures. Efforts are already under way to use Pigeon for a number of different aspects of the TCR Digital Platform, with a primary focus on automated material characterization.

A-8. SIMURGH

XCT is a process in which a 3D volume of an object of interest that represents attenuation coefficients of its comprising materials is reconstructed from 2D projections of the object acquired at different angles (views). XCT plays a critical role in nondestructive evaluation (NDE) and therefore certification/qualification of metal AM components. The quality of the 3D reconstruction algorithms used for dense metallic parts in complex systems and structures can be compromised by noise and artifacts such as streaks associated with metals, and a confounding effect called beam hardening (BH). Such artifacts and noise complicate the process of detection of salient defects (e.g., pores and cracks) in XCT images.

ORNL has developed an AI-based technique, *Simurgh*, that leverages CAD models of the parts, along with physics-based modeling, that shows significant improvement in XCT resolution by surpassing the state of the art in suppressing BH, metal artifacts, and noise in preliminary tests on both synthetic and real data. Figure A-12 shows a block diagram of the Simurgh technique. Figure A-13 and Figure A-14 show example results on both synthetic and real data sets, respectively. For development and testing of the method, the data currently available at the MDF were used: XCT scans of printed airplane engine turbine blades. In FY 2021, these methods will be evaluated and developed further for application to TCR-relevant geometries.
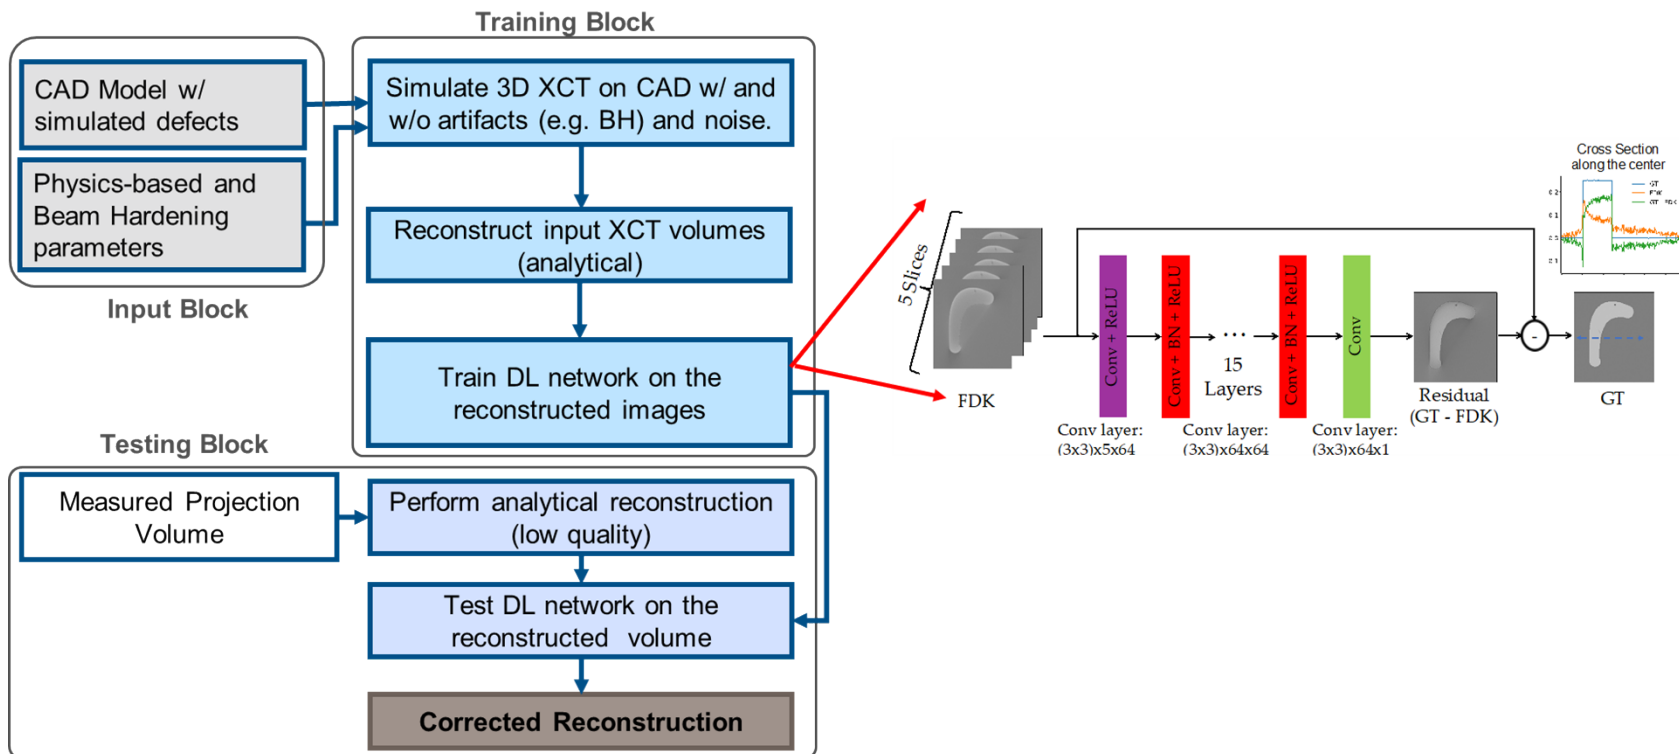
Figure A-12. Block diagram of the Simurgh framework. ORNL trains a convolutional neural network (AI-CT, top right) on the simulated (synthetic) data developed using CAD models and physics-based parameters.
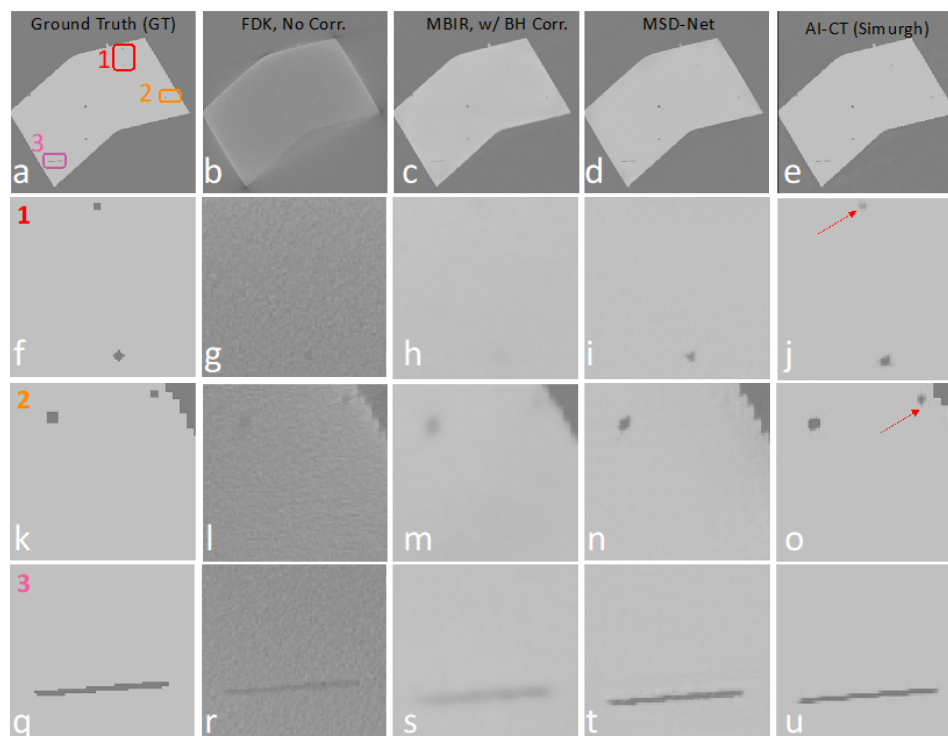
Figure A-13. Example test results on a synthetic data set. Two different neural networks (MSD-Net [3] and AI-CT [4], [5]) to emphasize the modularity of the approach. The method is also compared against uncorrected standard analytical (FDK) approach [6] and state-of-the-art model-based iterative reconstruction (MBIR) [7]. Several areas of interest are highlighted to demonstrate the superiority of the work against the state of the art.
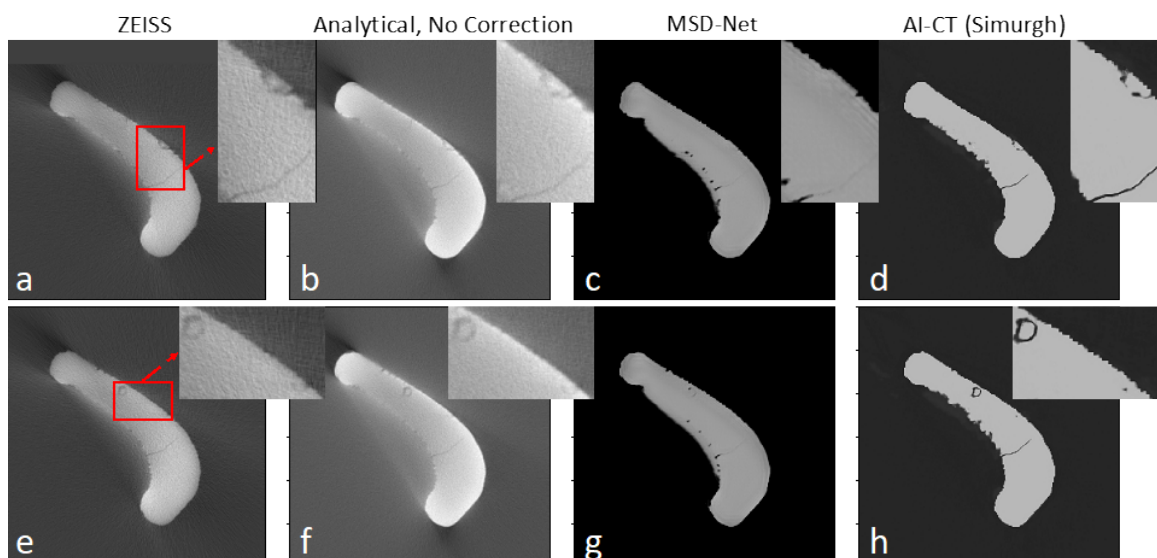

Figure A-14. Example test results on real data set. The proposed method is compared against the standard output from the ZEISS XCT system standard reconstruction, as well as the uncorrected standard analytical (FDK) approach and MSD-Net as a second neural network. Expanded views of areas of interest are highlighted in the inset images.

Functionally, a convolutional neural network (denoted as *AI-CT* in Figure A-12) is trained on the simulated (synthetic) data developed using CAD models and physics-based parameters. Then it is tested on synthetic and real data sets, as shown in Figure A-13 and Figure A-14, respectively. Two different neural networks—MSD-Net [3] and AI-CT [4], [5]—are tested to emphasize the modularity of this approach. The method is also compared against the standard output from the XCT system, as well as the uncorrected standard analytical (FDK) approach [6] and the state-of-the-art model-based iterative reconstruction (MBIR) approach [7].

ORNL has published a peer-reviewed conference paper on this topic, and a provisional patent has been filed by ORNL. A Technology Innovation Program (TIP) proposal was submitted to the Technology Transfer Office, and since its acceptance, the full proposal is in development. The Simurgh method is supported by ZEISS Quality Solutions, and they are considering licensing the technology from ORNL.

A-9.   REFERENCES

[1]   ASTM. "Standard Terminology for Additive Manufacturing." ISO / ASTM52900 - 15, 2015.
[2]   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.* 36 (Dec. 2020):101453.
[3]   D. Pelt, K. Batenburg, and J. Sethian. "Improving Tomographic Reconstruction from Limited Data Using Mixed-Scale Dense Convolutional Neural Networks." *J. Imaging* 4, no. 11 (Oct. 2018):128.
[4]   A. Ziabari, D. H. Ye, S. Srivastava, K. D. Sauer, J. B. Thibault, and C. A. Bouman. "2.5D Deep Learning for CT Image Reconstruction Using A Multi-GPU Implementation." *Conf. Rec. - Asilomar Conf. Signals, Syst. Comput.* 2018-Octob (2019): 2044–2049.
[5]   A. Ziabari et al. "Beam Hardening Artifact Reduction in X-ray CT Reconstruction of 3D Printed Metal Parts Leveraging Deep Learning and CAD Models." In *ASME IMECE 2020*, 2020.
[6]   L. A. Feldkamp, L. C. Davis, and J. W. Kress. "Practical Cone-Beam Algorithm." *J. Opt. Soc. Am.* 1, issue 6 (1984) 612–619.
[7]   P. Jin, C. A. Bouman, and K. D. Sauer. "A Model-Based Image Reconstruction Algorithm with Simultaneous Beam Hardening Correction for X-Ray CT." *IEEE Trans. Comput. Imaging* 1, no. 3 (Sept. 2015): 200–216.