LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-CONF-852123

Uncertainty Quantification of Metal Additive Manufacturing Processing Conditions Through the use of Exascale Computing

R. Carson, M. Rolchigo, J. Coleman, M. Titov, J. Belak, M. Bement

July 25, 2023

SC23: International Conference for High Performance Computing, Networking, Storage and Analysis
Denver, CO, United States
November 12, 2023 through November 17, 2023

Disclaimer

Uncertainty Quantification of Metal Additive Manufacturing Processing Conditions Through the use of Exascale Computing

Robert Carson
Lawrence Livermore National Laboratory
Livermore, CA, USA
carson16@llnl.gov

Matt Rolchigo
Oak Ridge National Laboratory
Oak Ridge, TN, USA
rolchigomr@ornl.gov

John Coleman
Oak Ridge National Laboratory
Oak Ridge, TN, USA
colemanjs@ornl.gov

Mikhail Titov
Brookhaven National Laboratory
Upton, NY, USA
mtitov@bnl.gov

Jim Belak
Lawrence Livermore National Laboratory
Livermore, CA, USA
belak1@llnl.gov

Matt Bement
Oak Ridge National Laboratory
Oak Ridge, TN, USA
bementmt@ornl.gov

ABSTRACT

Metal additive manufacturing (AM) is a disruptive manufacturing technology that opens the design space for parts outside those possible from traditional manufacturing methods. In order to accelerate industry and R&D needs to certify AM parts, the Exascale Additive Manufacturing project (ExaAM) has developed a suite of exascale-ready computational tools to model the process-to-structure-to-properties (PSP) relationship for additively manufactured metal components. One such tool is an uncertainty quantification (UQ) pipeline to quantify the effect that uncertainty in processing conditions has on local mechanical responses. We present an overview of this pipeline and its required simulation and workflow codes. Using the Oak Ridge National Laboratory's (ORNL) exascale computer, Frontier, we utilize this pipeline to cross multiple length and time scales to predict the local mechanical response of a location within a complex AM bridge part, AMB2018-01 produced by the National Institute of Standards and Technology (NIST) as part of their 2018 AM-Bench test series. Our results are then compared to experimental mechanical tests of parts from the NIST build to quantify the error in the ExaAM UQ workflow.

CCS CONCEPTS

• **Applied computing** → **Engineering**; • **Computing methodologies** → **Uncertainty quantification**.

KEYWORDS

additive manufacturing, uncertainty quantification, exascale computing

SC-W 2023, November 12–17, 2023, Denver, CO, USA

ACM ISBN 979-8-4007-0785-8/23/11.
https://doi.org/10.1145/3624062.3624103

ACM Reference Format:
Robert Carson, Matt Rolchigo, John Coleman, Mikhail Titov, Jim Belak, and Matt Bement. 2023. Uncertainty Quantification of Metal Additive Manufacturing Processing Conditions Through the use of Exascale Computing. In *Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W 2023), November 12–17, 2023, Denver, CO, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3624062.3624103

1 INTRODUCTION

Additive manufacturing (AM) is an emerging and growing field that offers the opportunity to manufacture complex parts on the fly. Given the wide number of available manufacturing platforms, AM provides the opportunity to manufacture parts with a wide range of materials such as metals, ceramics, polymers, and etc. In particular, the laser powder bed fusion (LPBF) platform for metals is of interest as it offers the possibility to construct parts with local variation in thermo-mechanical properties by spatially controlling the processing parameters. However, these optimizations are non-trivial given how these properties involve a complex multiscale relationship between material parameters and processing conditions. While experiments can be used to determine insights into these optimizations, a non-negligible amount of experiments and design work can be required to nail down the desired performance and later certify AM parts for use.

The ExaAM project, as part of the Exascale Computing Project (ECP), has developed and extended a number of computational codes and workflows to provide self-consistent high-fidelity simulations of the LPBF platform from the melt pool through the part scale response [11]. A typical ExaAM workflow to model the PSP relationship of LPBF builds can be broken into three stages (and one additional optional stage). Stage 1 is concerned with the process-structure relationship. This involves simulating the melt pool physics across a part build to capture the necessary thermal history required to evolve the underlying microstructure of the part, followed by simulating the as-solidified grain structure. An optional stage, Stage 2, can be performed after the first stage and conducts the simulation of late time microstructure evolution, typically associated with something akin to the heat treatment of a part. The next stage, Stage 3, connects a representative volume

element (RVE) of the microstructure to the local material properties necessary for a part build simulation. The final stage, Stage 4, is a full thermo-mechanical part build simulation to predict residual stresses in the part, and this simulation makes use of self-consistent local material properties from earlier stages of the build. Within the ExaAM project, the core physics components are related to Stages 1 and 3 and those are all simulated using open-source codes developed by this project. Additionally, these codes were also designed with being able to run on exascale systems in mind. The melt pool physics are simulated using AdditiveFOAM [5, 7], an extension of OpenFOAM for AM processes. The microstructure is evolved using ExaCA, a GPU-friendly cellular automata code [9]. Then the local property simulations are performed using ExaConstit, a GPU-friendly crystal plasticity finite element code [4].

ExaAM has leveraged the NIST Additive Manufacturing Benchmark Test Series (AM-Bench) [8] and more specifically the data associated with the Inconel 625 AMB2018-01 complex AM bridge part as a means to validate the codes and workflows [6, 7, 9, 11]. Given uncertainty in AM builds, the ExaAM project developed a UQ pipeline that forwards variations in the processing and microstructural evolution conditions through the part scale. In this paper, we will provide an overview of the UQ pipeline, the codes involved, and the necessary parameters for all simulations. Finally, we will go over the results from an 8000 node run on ORNL's Frontier machine for Stages 1-3 of the pipeline and its validation against experimental macroscopic mechanical results.

2 EXAAM UNCERTAINTY QUANTIFICATION PIPELINE

To accelerate AM part certification and design, the ExaAM project has developed a self-consistent UQ workflow to help determine acceptable machine parameters for a part build. This workflow can be found at [2], and an overview is provided in Figure 1. Through this workflow, ExaAM will demonstrate the capability to propagate uncertainties through the chain of coupled models to understand the potential variability in the overall model prediction. Since these codes are computationally intensive and designed for large computing resources, direct sampling of specific model outputs to obtain statistically significant distributions of the outputs of interest is not possible. Instead, we utilize a grid-based approach using the TASMANIAN library [10] to sample a few responses to the model input space, construct a surrogate model on the output space, and then sample from that surrogate to produce probability distributions of the outputs. Since multiple codes must be connected for this purpose, each containing its own inputs and associated uncertainties, the overall workflow to properly perform the UQ analysis is complex. It should be noted that the UQ grids are only generated for Stage 1 parameters as further extending them to Stage 3 would lead to a computationally intractable problem even on current day exascale computing systems such as ORNL's Frontier machine.

A more in-depth overview of Figure 1 is now described. Stage 0 involves selecting a location(s) within the part build and generating the necessary UQ grid for later stages. Within Stage 1, l melt pool simulations will be performed based on the earlier generated UQ grid. The generated thermal data will be fed into the microstructure evolution simulations, which has its own m size UQ grid, to generate $l \times m$ microstructures. These microstructures are handed off to Stage 3 which might down-select the number of RVE microstructures to use resulting in n RVEs. Stage 3 then runs $o \times p$ simulations per RVE which corresponds to the o number of loading conditions and p number of temperatures being considered. After all $n \times o \times p$ simulations have completed, an optimization procedure occurs on the $n \times o \times p$ simulations' data to generate n local macroscopic material model parameters to potentially be used in Stage 4. Stage 4 might do a further down-selection in-order to run t full part build simulations.

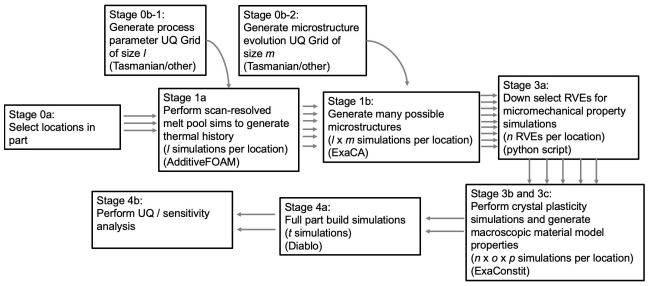

Figure 1: A self-consistent AM UQ workflow to quantify the effects of machine parameters on AM part performance.

3 SIMULATION SETUP

For the ExaAM validation test case on ORNL's Frontier machine, we are only concerned with 1 location in the AMB2018-01 part. The workflow first launches a series of AdditiveFOAM simulations for 5 values of $D4\sigma$, the ISO international standard definition for beam width, based on the produced Tasmanian grid to generate thermal histories for ExaCA. AdditiveFOAM was calibrated to the AMB2018-02 melt pools from experiment, where it was found that a significant discontinuity in melt pool dimensions existed at $D4\sigma$ = 100 μm, which was the value reported from the AMB2018-01 build. To appropriately cover the processing space on both sides of this discontinuity, $D4\sigma$ values spanning 80-120 μm was selected. Once the AdditiveFOAM jobs are complete, the workflow manager launches a series of ExaCA simulations for each location to generate explicit microstructures for ExaConstit. For the ExaCA simulations, a UQ grid was produced for the heterogeneous nucleation density N_0 and mean nucleation undercooling ΔT_N (though we note that the standard deviation of the nucleation undercooling, ΔT_σ, was linked to the mean nucleation undercooling via the expression $\Delta T_\sigma = \Delta T_N / 3$ to more accurately represent nucleation in an undercooled alloy). The workflow produces a total of 25 independent ExaCA simulations (varied N_0 and ΔT_N) per AdditiveFOAM simulation (varied $D4\sigma$). N_0 spanning the order of magnitude 10^{14} - 10^{15} m^{-3} were selected based on the large experimental uncertainty in the parameter along with the model observation that nucleation densities smaller than 10^{14} has little effect on the final microstructure, while nucleation densities larger than 10^{15} resulted in an unrealistic number of small grains. Similarly, a large experimental uncertainty in the mean nucleation undercooling existed as well, and values from 3-30 K were used in the ExaCA simulations (as nucleation at or above the liquidus temperature is unrealistic, as is a mean

nucleation undercooling larger than the dendrite tip velocities observed in the simulations). The 5 permutations of AdditiveFOAM datasets and 25 permutations of ExaCA inputs yielded a total of 125 simulated microstructures which adequately covered the expected processing space of the AMB2018-01 part build. These Stage 1a and 1b runs made use of 125 nodes on Frontier and took in total approximately 4 hours to simulate a 1.0 mm^3 middle portion of the L7 leg of the AMB2018-01 part. Once all the ExaCA simulations are completed, the 125 voxelated microstructures are then passed onto another sub-workflow manager that handles Stage 3.

A 0.5 mm^3 representative volume element (RVE) in the center of the simulated microstructure was taken from each of the 125 ExaCA simulations to be used in the Stage 3 runs. For the Stage 3 runs, a total of 7875 independent ExaConstit simulations would need to be run on Frontier with each simulation requiring 8 nodes of Frontier using 64 MPI ranks with the typical 1 CPU to 1 GPU mapping for each MPI rank. In order to set-up these runs, the *chal_prob_full.py* python script at [2] was set-up to drive the Stage 3 workflow. In this script, the 125 RVEs from Stage 1b are read and the voxel data is coarsened from a 400^3 cubic mesh down to a 200^3 linear hexahedron element mesh and each of the various grains is associated with their unique lattice orientations. The script then generates all necessary directories, input files, and options files for each unique loading condition and temperature for the simulations to be able to construct anisotropic yield surface parameters for the Stage 4 simulations. The material model parameters were parameterized ahead of time by using an AdditiveFOAM-ExaCA produced microstructure that most closely matched the experimentally observed crystallographic texture and matching the LLNL experimental compression test [3] for the room temperature build direction sample out to -30% true strain. The 7875 simulation runs utilized RADICAL-EnTK [1], a job-ensemble manager, to coordinate all the runs on 8000 nodes of Frontier, and these simulations took approximately 2 hours and 15 minutes to complete. While optimization of the macroscopic material model parameter sets and running of the full part build simulations is still ongoing work, the UQ pipeline results can still be verified against the experimental data from AM-Bench.

4 RESULTS

Out of the 7875 simulations run for Stage 3, we are concerned with 250 for the purposes of validating our UQ workflow. These 250 simulations took place at room temperature and monotonically compressed the sample either along the build direction (BD) or transverse direction (TD). We were looking at those specific cases as we had multiple repeat experimental tests corresponding to those conditions as summarized in Table 1.

Load Direction	Yield Stress (MPa)	Stress at -5% strain (MPa)
BD	[-750, -741, -741, -686]	[-951, -967, -951, -916]
TD	[-740, -726, -721, -676]	[-915, -921, -908, -858]

Table 1: Summary of experimental macroscopic stress-strain data used to compare against the simulations.

While entire stress-strain curves do exist from the experiments, we only had two points of interest from each experiment: the stress at which yielding occurs, of importance for the macroscopic model, and the stress at -5% strain, which was the largest strain measured in some of the experiments. Lastly due to the high computational cost of running Stage 3 using high-fidelity meshes, we also examined the effects that varying coarsening levels on the initial microstructure RVE have on the comparison to experiments. A summary of these coarsening levels and computational requirements on Frontier is given in Table 2. It should be noted that on average each simulation takes at most 20 minutes to run with the given resource sets and the workflow will run in under 1 hour once the job has been allocated.

Coarsening Level	Mesh Elements	MPI ranks	Number of Nodes
2	8×10^6	64	1000
4	1×10^6	8	250
5	5.12×10^5	4	125
8	1.25×10^5	1	32
10	6.4×10^4	1	32

Table 2: Summary of different coarsening levels on each microstructure, the number of mesh elements for each simulation, the number of MPI ranks required for a simulation, and finally the total number of nodes required to run on Frontier.

In order to compare the simulations to experiments, we calculate the minimum difference between the simulated yield and stress at -5% strain against the experimental data in Table 1 for each load direction. Afterwards, we can calculate the relative error associated with this minimum point for the yield stress and stress at -5% for the different loading directions leading to two relative errors per loading direction and four per RVE. To further simplify our analysis, we make use of a root mean square relative error (RMSRE) measure:

$$RMSRE = \sqrt{\frac{1}{NPTS}\Sigma_{i=1}^{NPTS} RE_i^2}$$

where RE is the earlier calculated relative error and $NPTS$ is the number of RE per RVE. We utilize kernel density estimation (KDE) plots to summarize the results in Figure 2. Qualitatively, the KDE plots suggest that all the coarsening values produce similar distributions in the data. However, we make use of the first through fourth moments of distribution to terse out differences as summarized in Table 3. The first two moments are similar between all coarsening levels, but we start to see differences in the skewness of the distributions. Coarsening levels 2 and 4 produce more desirable distributions in that the errors are more concentrated to the smaller values, but the higher levels of coarsening shift that balance the other way. The kurtosis reinforces the observable in Figure 2 of the longer tail distributions on the right portion of the curves. Overall, these results validate the ExaAM UQ pipeline as it shows we can accurately predict the mechanical response of a part by modelling from the melt pool on upwards. While the coarsening results might not carry over to all the different loading directions, they do highlight a potential path forward in extending a UQ grid to Stage 3. Additionally, future work will use first-order second-moment methods for local sensitivity analysis of the UQ pipeline.

Outside the observations made on the total RMSRE distributions, we can also qualitatively see how the build and transverse directions contribute to the total RMRSE values. The KDE plots of the RSMSE associated with only the build or transverse direction are provided

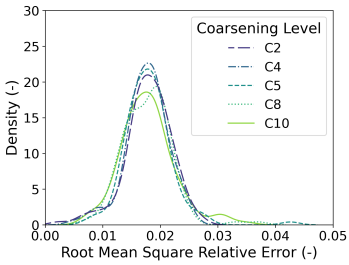

Figure 2: KDE plots of RMSRE values of simulations versus experimental data for both build and transverse directions across all coarsening levels.

Coarsening Level	Mean	Variance	Skewness	Kurtosis
2	0.01787	1.65e-5	-0.83	1.74
4	0.01762	1.38e-5	-0.64	1.10
5	0.01794	1.73e-5	1.67	8.97
8	0.01775	1.76e-5	0.85	3.52
10	0.01781	2.17e-5	0.74	1.56

Table 3: Summary of first through fourth moments of distribution of RMSRE for the various coarsening levels.

in Figures 3a and 3b , respectively. The build direction KDE plots suggest that the errors are on average higher, but the spread in data is smaller. The transverse direction we qualitatively have smaller error in predicting the response, but the spread in data is greater.

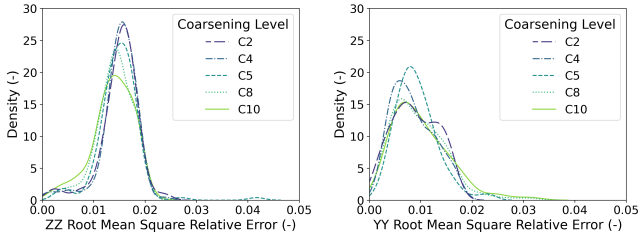

(a) KDE plots of RMSE associated with only BD values

(b) KDE plots of RMSRE associated with only TD values

Figure 3: KDE plots of RMSRE values associated with either the BD or TD across all coarsening levels.

5 CONCLUSION

The ExaAM project has developed a number of high performance codes to model the metal additive manufacturing process. As part of this project, we developed a UQ pipeline to help drive the determination of acceptable machine parameters for AM part builds. We leveraged this pipeline to model one location within the AMB2018-01 part and validated our workflow against experimental macroscopic compression tests. It was found that we were able to get below 5% RMSRE compared to the experiments across all simulated microstructures. Additionally, we observed that varying the coarsening levels of the microstructure produced similar mean RMSRE values. However, the coarser microstructures led to distributions more skewed towards larger errors. Finally, the ExaAM project will extend this workflow to full part builds to determine the uncertainty in residual stress related to varying build parameters.

ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. DOE Office of Science and the NNSA. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This manuscript has been in part authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (https://www.energy.gov/doe-public-access-plan). LLNL-CONF-852123

REFERENCES

[1] Vivekanandan Balasubramanian, Antons Treikalis, Ole Weidner, and Shantenu Jha. 2016. Ensemble Toolkit: Scalable and Flexible Execution of Ensembles of Tasks. arXiv:arXiv:1602.00678

[2] Robert Carson, John Coleman, and Matt Rolchigo. 2023. *ExaAM UQ Workflow.* Retrieved July 20, 2023 from https://github.com/ExascaleAM/Workflows

[3] Robert Carson and Lyle Levine. 2022. Macroscale Compression at Different Temperatures and Orientations (CHAL-AMB2022-04-MaCTO). https://doi.org/10.18434/MDS2-2681

[4] Robert A. Carson, Steven R. Wopschall, and Jamie A. Bramwell. 2019. ExaConstit. [Computer Software] https://doi.org/10.11578/dc.20191024.2. https://doi.org/10.11578/dc.20191024.2

[5] John Coleman, Kellis Kincaid, Gerald L. Knapp, Benjamin Stump, and Alexander J. Plotkowski. 2023. AdditiveFOAM. [Computer Software] https://doi.org/10.5281/zenodo.8034098. https://doi.org/10.5281/zenodo.8034098

[6] Rishi K. Ganeriwala, Neil E. Hodge, and Jerome M. Solberg. 2021. Towards improved speed and accuracy of laser powder bed fusion simulations via multiscale spatial representations. *Computational Materials Science* 187 (Feb 2021), 110112. https://doi.org/10.1016/j.commatsci.2020.110112

[7] G.L. Knapp, J. Coleman, M. Rolchigo, M. Stoyanov, and A. Plotkowski. 2023. Calibrating uncertain parameters in melt pool simulations of additive manufacturing. *Computational Materials Science* 218 (Feb 2023), 111904. https://doi.org/10.1016/j.commatsci.2022.111904

[8] Lyle Levine, Brandon Lane, Jarred Heigel, Kalman Migler, Mark Stoudt, Thien Phan, Richard Ricker, Maria Strantza, Michael Hill, Fan Zhang, Jonathan Seppala, Edward Garboczi, Erich Bain, Daniel Cole, Andrew Allen, Jason Fox, and Carelyn Campbell. 2020. Outcomes and Conclusions from the 2018 AM-Bench Measurements, Challenge Problems, Modeling Submissions, and Conference. *Integrating Materials and Manufacturing Innovation* 9, 1 (Feb 2020), 1–15. https://doi.org/10.1007/s40192-019-00164-1

[9] Matt Rolchigo, Samuel Temple Reeve, Benjamin Stump, Gerald L. Knapp, John Coleman, Alex Plotkowski, and James Belak. 2022. ExaCA: A performance portable exascale cellular automata application for alloy solidification modeling. *Computational Materials Science* 214 (nov 2022), 111692. https://doi.org/10.1016/j.commatsci.2022.111692

[10] Miroslav Stoyanov, Damien Lebrun-Grandie, John Burkardt, and Drayton Munster. 2013. Tasmanian. https://doi.org/10.11578/dc.20171025.on.1087

[11] John A Turner, James Belak, Nathan Barton, Matthew Bement, Neil Carlson, Robert Carson, Stephen DeWitt, Jean-Luc Fattebert, Neil Hodge, Zechariah Jibben, Wayne King, Lyle Levine, Christopher Newman, Alex Plotkowski, Balasubramaniam Radhakrishnan, Samuel Temple Reeve, Matthew Rolchigo, Adrian Sabau, Stuart Slattery, and Benjamin Stump. 2022. ExaAM: Metal additive manufacturing simulation at the fidelity of the microstructure. *The International Journal of High Performance Computing Applications* 36, 1 (Jan 2022), 13–39. https://doi.org/10.1177/10943420211042558