SANDIA REPORT

SAND2020-9870
Unlimited Release
Printed September 2020

Multiscale Approach to Fast ModSim for Laser Processing of Metals for Future Nuclear Deterrence Environments

Daniel Moser, Kyle Johnson, Theron Rodgers, Mario Martinez

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energys National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2020-9870
Unlimited Release
Printed September 2020

Multiscale Approach to Fast ModSim for Laser Processing of Metals for Future Nuclear Deterrence Environments

Daniel Moser
Fluid and Reactive Processes Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-9999
dmoser@sandia.gov

Kyle Johnson
Materials and Failure Modeling Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-9999
kyljohn@sandia.gov

Theron Rodgers
Computational Materials and Data Science Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-9999
trodger@sandia.gov

Mario Martinez
Fluid and Reactive Processes Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-9999
mjmarti@sandia.gov

Abstract

Predicting performance of parts produced using laser-metal processing remains an outstanding challenge. While many computational models exist, they are generally too computationally expensive to simulate the build of an engineering-scale part. This work develops a reduced order thermal model of a laser-metal system using analytical Green's function solutions to the linear heat equation, representing a step towards achieving a full part performance prediction in an "overnight" time frame. The developed model is able to calculate a thermal history for an example problem 72 times faster than a traditional FEM method. The model parameters are calibrated using a non-linear solution and microstructures and residual stresses calculated and compared to a non-linear case. The calibrated model shows promising agreement with a non-linear solution.

Contents

Figures

Tables

Nomenclature

c_p Specific Heat

CALPHAD CALculation of PHAse Diagrams

CPU Central Processing Unit

FEM Finite Element Method

GPU Graphics Processing Unit

k Thermal Conductivity

LENS Laser Engineered Net Shaping

MCMC Monte Carlo Markov Chain

P Laser Power

ρ Density

σ Standard Deviation

T Temperature

1 Introduction

In recent years, there have been many works modeling all aspects of laser metal manufacturing processes. Models have ranged from detailed, high fidelity models of melt pool physics such as Trembacki et al. [20] for LENS, and Khairallah et al. [11], Moser and Martinez [13], and Beghini et al. [2] for laser powder bed additive manufacturing. Such models incorporate detailed physics at the scale of the laser spot and have had some success in explaining and predicting small-scale physical phenomena, melt track morphologies, and even residual stress. However, their high computational cost makes the modeling of even a millimeter laser scan challenging, much less predicting the performance of an entire part, and very little work has been performed on leveraging these models to inform faster, larger scale models.

Other groups have purposed simpler models that neglect the complexities of the melt pool and model heat transfer using conduction only. Often, these models are coupled to solid mechanics or microstructure models in order to predict mechanical properties, e.g. Johnson et al. [10] who used a conduction model to predict microstructure and residual stress in the LENS process, Smith el al. [17] who coupled a conduction model to a CALPHAD microstructure model, and Hodge et al. [9] who implemented a thermo-mechanical model for additive manufacturing. These models are substantially faster than melt pool scale models, but still require very large amounts of compute time due primarily to the fact that the time stepping must be fine enough to resolve the motion of the laser. This means that a single simulation for a small part build takes days, making multiple simulations to perform design optimization or uncertainty quantification intractable, and simulating the build of an engineering scale part on the order of several cm^3 remains out of reach.

Currently, the only methods that are able to model mechanical performance for engineering-scale parts require the use of inherent strain methods, in which the heat transfer problem is bypassed and the solid mechanics problem solved directly by imposing estimated plastic strains, or heat source agglomeration, in which the fine time-scale motion of the laser heat source is lumped together in some way to allow larger time steps. Ganeriwala et al. [8] used an agglomerated heat source mechanism in a thermo-mechanical model of laser powder bed additive manufacturing, whereas Bugatti et al. [4] applied the inherent strain method to the same process. Both methods can be computationally inexpensive enough to perform multiple modeling iterations for engineering-scale parts, but accuracy becomes a challenge. Calibration against experimental data is needed for both methods. However, ranges of model validity are not well defined when models calibrated on a particular material, set of processing conditions, and part geometry are applied to make predictions under new conditions, making it difficult to trust their predictive capability.

This work uses an analytic, Green's function based, heat transfer model to calculate thermal histories of laser metal manufacturing processes. The use of Greens functions in laser process modeling is not entirely new (e.g. Schwalbach et al [16] and Farwell et al [7] used Greens functions to calculate thermal histories for single layers in laser powder bed additive manufacturing). However, this work develops a novel framework for applying Greens functions that can be used to quickly calculate the thermal history of an entire build that can be used on both CPU and GPU architectures.

Unlike FEM, a Greens function solution does not need to be sequentially time stepped forward from an initial condition, but rather can be constructed on the fly at any point in time and space by numerically evaluating an integral. This property is used to overcome the performance challenges posed by conventional FEM models with resolved laser source terms. As the fine time scale heat transfer only occurs in the locality of the laser in time and space, an adaptive space-time grid is employed to allow the solution to be selectively sampled in these areas with high spatial and temporal fidelity, and sampled only coarsely elsewhere. Thus, a thermal history is constructed using as few evaluations as possible. Solution evaluations are also embarrassingly parallel, unlike in mesh and time-step based methods, making the problem amenable to highly parallel GPU computing architectures. Using the Kokkos performance portability software package [6], this work is able to achieve a 72 times wall clock speedup as compared to a conventional FEM code on an example laser processing simulation using GPUs.

The Greens function solution requires making simplifying assumptions to the problem: linearity and a semi-infinite domain. This work uses the Dakota toolkit [1] to calibrate the model parameters (material properties and source term) against a non-linear solution using objective functions based on maximum temperature and temperature norms, and estimate the uncertainties in the solution due to the assumptions. Results from the calibrated model are input into existing models for predicting microstructure and residual stress and compared to results calculated from a non-linear, FEM temperature solution. Although differences are found between the two predictions, there are promising similarities that show the applicability of the developed method for rapidly predicting detailed temperature histories that can be used as a basis for engineering-scale part performance predictions.

2 Methodology

2.1 Green's Function Solution

An analytical Green's function solution is available for partial-differential equations of the form given in Equation 1 for an infinite domain.

$$\frac{\partial u}{\partial t} - \alpha \Delta u = S \tag{1}$$

The heat equation can be cast in this form by assuming constant material properties (ρ, c_p, and k) and a temperature-independent source, and letting $\alpha = \frac{k}{\rho c_p}$. The workpiece to be simulated is assumed to be a semi-infinite medium with the zero-flux boundary condition on the surface on which the laser is applied (z-direction). This is accomplished by using a source term that is symmetric across the z-plane. The Green's function for this equation is given in Equation 2.

$$G = \Theta(t) \left(\frac{1}{4\pi\alpha t} \right)^{3/2} e^{\frac{-x^2-y^2-z^2}{4\alpha t}} \tag{2}$$

$\Theta(t)$ is the Heaviside step function.

For this work, we use a general ellipsoidal Gaussian to model the laser source, as given in Equation 3.

$$S = \frac{P}{\pi^{3/2}\sigma_x\sigma_y\sigma_z} e^{-\frac{(x-x_l)^2}{\sigma_x^2} - \frac{(y-y_l)^2}{\sigma_y^2} - \frac{(z-z_l)^2}{\sigma_z^2}} \tag{3}$$

x_l, y_l, and z_l represent the coordinates of the laser center, which can vary with time. This functional form of the source term is a useful choice as it can be integrated against the Green's function in Equation 2 analytically in space, leaving an integral in time only to be performed numerically. The full temperature solution is given in Equation 4.

$$T = T_i + \frac{2P}{\pi^{3/2}\rho c_p} \int_0^t \frac{e^{-\frac{(x-x_l(s))^2}{\sigma_x^2+4\alpha(t-s)} - \frac{(y-y_l(s))^2}{\sigma_y^2+4\alpha(t-s)} - \frac{(z-z_l(s))^2}{\sigma_z^2+4\alpha(t-s)}}}{\sqrt{\sigma_x^2+4\alpha(t-s)}\sqrt{\sigma_y^2+4\alpha(t-s)}\sqrt{\sigma_z^2+4\alpha(t-s)}} ds \tag{4}$$

T_i is the specified initial temperature of the domain. It is also possible to develop analytical solutions for arbitrary-order derivatives of temperature with respect to spatial dimension and time by simply differentiating Equation 4 with respect to the variable of interest. Thus, the Green's function method allows the calculation of analytical temperature gradients and heating/cooling rates along with temperature.

Equation 4 is amenable to an adaptive quadrature routine, as the integrand varies rapidly with time when the laser center is close to the evaluation point in both space and time, but is close to constant otherwise. Commonly used adaptive quadrature algorithms use interval sub-division and recursion to calculate a solution. These algorithms perform poorly on the GPU as the sub-division is a sequential operation and the memory requirements for recursion mean that relatively few threads can be used.

Instead, an adaptive Clenshaw-Curtis algorithm is applied. Clenshaw-Curtis quadrature is selected as it can be infinitely nested - the evaluation points of an order 2n quadrature rule contain all the evaluation points of an order n quadrature rule, meaning that function evaluations can be reused. Additionally, all evaluation points are known ahead of time, so evaluations of the integrand can be performed in parallel. By parallelizing over both the physical points at which the temperature solution is to be evaluated and the evaluation points of the integrand, many more threads are able to be used, improving GPU performance.

Adaptivity is achieved by comparing the results of the integral performed using an order n and order 2n rule. When the results differ by more than a prescribed tolerance, the integral is computed using an order 4n rule, which is compared to the order 2n result. Refinement continues until either the tolerance is achieved or a maximum order is reached. For cases where the maximum order is reached, the interval is sub-divided and the process repeated. Thus, performance can be tuned for either CPU or GPU by specifying a large maximum order to expose more parallelism for the GPU, or a small one to reduce the number of integrand evaluations on the CPU.

2.2 Adaptive Space-Time Grid

The computational costs, both in time and disk storage, of representing the temperature history of an engineering-scale part using uniform grids and time step size at the resolution required to resolve the action of the laser are prohibitive. Therefore, an adaptive, 4D space-time grid is developed to reduce computational and storage requirements. The grid begins as a series of uniform hyper-rectangular cells, with each vertex representing a point in space and time at which the solution will be calculated and stored. The calculation of the solution at each vertex is done in parallel as there is no spatial or temporal dependence between solution values with the Green's function method.

Once the solution is computed at all vertices, each cell is split into 16 child cells by bisecting the parent cell evenly along each of the 4 dimensions. The solution is then computed at all newly created vertices using the Green's function. An approximation of the solution is also computed by using linear interpolation to estimate the value of the solution at each of the new vertices using only the solution values at the parent vertices. If the approximate solution value differs from the Green's function value by more than a prescribed tolerance, each of the 16 child cells are split. The process repeats until no further refinement is required or a specified maximum refinement level is reached.

The developed grid allows local adaptivity in all 4 dimensions, including time, allowing it to reduce the number of required vertices beyond what is achievable even with adaptive time stepping

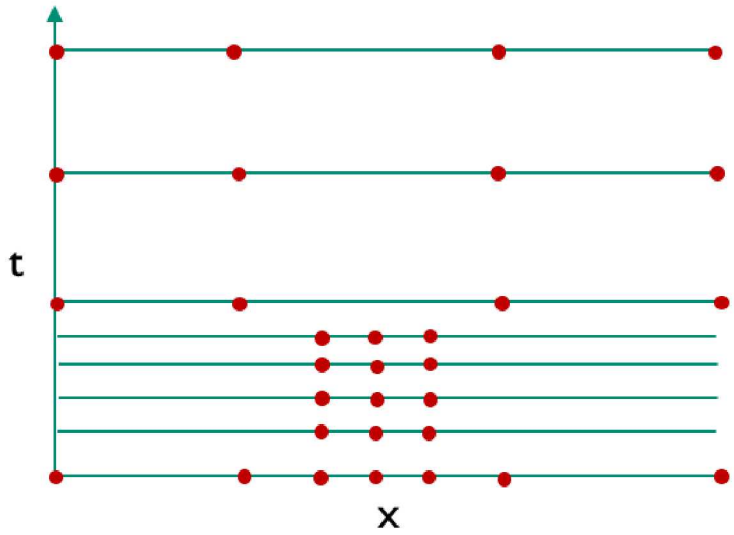

Figure 1. 2D (x-t) grid showing local time refinement

and adaptive spatial grids. With these schemes, the adaptive time step size must be selected for the entire domain, meaning that some points will be unnecessarily represented at a finer time scale than is required since the entire domain must be stored at each time step. With the 4D grid, the solution can be sampled at different time increments at different places in the domain. This is shown in Figure 1 for a 2D (x-t) grid. The red dots represent vertices. Note how not all time steps need to include the extent of the x-domain.

2.3 Calibration and Uncertainty Quantification

In order to approximate the solution to a non-linear problem with a linear solution, the input parameters to the Green's function model are calibrated using the nl2sol method of the Dakota toolkit [1]. This employs a generalized Levenberg-Marquardt algorithm to find a set of model parameters that minimize a set of objective functions. Two objective functions are used in the calibration: the maximum temperature and the L1 norm of the difference in temperature solutions outside the melt region. These are chosen as the maximum temperature is important for predicting the formation of porosity, and the temperature in the solid region governs residual stress and microstructure development.

The Green's function model is calibrated against a non-linear, FEM temperature solution for a laser raster on a flat plate computed using the Aria module of the SIERRA code suite [19]. The non-linear model includes temperature-dependent material properties, latent heat of fusion, and convective/radiative boundary conditions on the surface on which the laser is applied. A single time slice is taken from both the Green's function and Aria models to compute the objective

functions. A time is chosen after the laser has completed several rasters, as this is found to give the best calibration results. Five input parameters are chosen for calibration: k, P, ρc_p, σ_z, and σ_x/σ_y, which are constrained to be the same value.

For purposes of performing uncertainty quantification, a Bayesian calibration of the input parameters is also performed using a MCMC method [1]. Due to the large number of samples required by the MCMC technique, an inexpensive surrogate of the Green's function model is constructed using a Gaussian Process [1], which is sampled in lieu of the full Green's function model. The Bayesian calibration provides an estimate of the probability distribution of the input parameters given the calibration data. Uniform priors are used for all input parameters. The resulting input parameter probability distributions are propagated through the Green's function model to the output predictions using a Latin Hypercube sampling technique [1] for an estimate of the uncertainty in the temperature prediction due to the assumptions required in using the Green's function solution.

2.4 Microstructure and Residual Stress

Temperature histories computed by both the Green's function and non-linear Aria models are used as inputs to existing, Sandia-developed, models for microstructure evolution and solid mechanics in order to predict as-built microstructures and residual stress profiles. The microstructure model uses the SPPARKS Kinetic Monte Carlo software package [15]. An under-cooling based approach is used to simulate the solidification of microstructural grains and a kinetic Monte Carlo method used to simulate grain growth.

The solid mechanics model uses the SolidMechanics module of the SIERRA code suite [18]. A temperature-dependent Bammann-Chiesa-Johnson isotropic elasto-viscoplastic material model, as described in [10], is used to simulate the evolution of the residual stress and distortion.

Table 1. Timing Data

Solver	Platform	Nodes	Time
Aria	Eclipse	10	6 hrs
Green	Eclipse	5	2.9 hrs
Green	Eclipse	10	33 mins
Green	Eclipse	20	22 mins
Green	Vortex	4	4.3 mins

3 Results and Discussion

3.1 Timing Data

Wall clock timings are compared between the Green's function model and the Aria model on two platforms. The first of which is Sandia's Eclipse machine, which has two 18 core 2.1 GHz Intel Broadwell processors per node. The second is Sandia's Vortex machine, which has two 22 core IBM Power9 processors and four NVIDIA Tesla V100 5120 core GPUs per node.

The results of the timing study are show in Table 1. The Green's function model is over 10 times faster than the Aria model on the same platform and number of nodes. Increasing the number of nodes on Eclipse beyond 10 results in only a moderate decrease in wall time due to inadequate parallel load balancing. Some processes get stuck on "hard" portions of the grid requiring lots of refinement. Future work will implement dynamic load balancing, which should improve this issue. As expected, the Green's function model performs particularly well on the GPU machines, achieving a 72 times speedup as compared to the Aria model. This is due to the massively parallel nature of the Green's function solution in which evaluating the solution at any point in space and time is independent of all other evaluations. This highlights the need for the development of GPU-friendly solution algorithms in order to take full advantage of these emerging systems.

3.2 Temperature Comparison

The Green's function model is calibrated against the Aria non-linear model for a range of laser parameter inputs: laser power from 18-28W and speed from 0.05-0.13 m/s. A total of nine calibration points are chosen within this range: $\{P = 18, 23, 28\ W\}$ x $\{v = 0.05, 0.09, 0.13\ m/s\}$. Using the two chosen objective functions, maximum temperature and L1 norm of temperature in the solid material, this gives a total of 18 objective functions for the optimization. A narrow parameter range is chosen for this test case, although the eventual goal is to calibrate over a range of process inputs representative of what is achievable in a given machine.

The non-linear material properties used in the Aria model as well as the laser parameters are summarized in Table 2. The simulated build is a 2mm x 2mm square on a flat plate. The time

Table 2. Material Properties

Property	Value
k	20 W/(m K)
ρ	3715 kg/m^3
c_p	320.3+0.379T J/(kg K)
Solidus Temperature	1648 K
Liquidus Temperature	1673 K
Latent Heat of Fusion	270000 J/kg
σ_x/σ_y	100 μm
σ_z	200 μm
Hatch Spacing	70 μm

Table 3. Calibrated Parameters

Parameter	Mean	Standard Deviation
k	12 W/(m K)	4.5 W/(m K)
ρc_p	45000 J/K	16000 J/K
P	40 W	23 W
σ_x/σ_y	198 μm	68 μm
σ_z	181 μm	67 μm

plane selected for calibration is t=0.1s, a time at which at least one down and back laser scan is completed for all speeds. A Gaussian Process is constructed for each of the objective functions using a total of 1000 samples from the Green's function model. These surrogates are then sampled in the MCMC Bayesian calibration procedure using 100000 samples. The estimated distributions of the calibrated parameters are given in Table 3. The means represent the most likely values of the parameters and correspond to the values found using the Levenberg-Marquardt optimization algorithm. The standard deviations represent the uncertainty in the parameter values due to the inability of the non-linear result data to fully inform them.

The calibrated Green's function model is tested against the Aria model for a parameter set that is within the calibration range, but not one of the calibration points used: P = 20W, v = 0.1 m/s. First, the average values of the calibrated parameters are used and the difference in temperature fields between the Green's function and non-linear model are computed for a time t=0.1s. The results are shown in Figure 2. As is expected, the largest errors occur in the melt region trailing the laser, as the large temperature gradients lead to a highly variable specific heat as well as latent heat effects that the linear model is not able to capture. The linear model results improve moving away from the laser where temperature gradients are less pronounced.

In order to assess the estimates of parameter uncertainty, a Latin Hypercube study is performed using 32 samples drawn from the calibrated input parameter distributions. Temperature histories

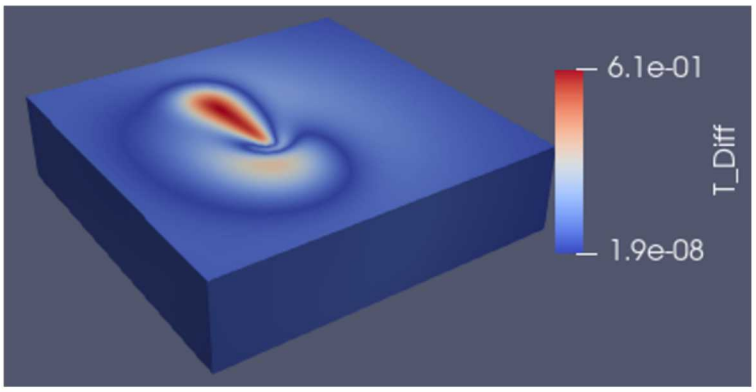

Figure 2. Normalized temperature difference between non-linear and Green's function models

are calculated for each parameter set and means and standard deviations of the temperature at each point are calculated by averaging across the sample results. Figure 3 shows the field of temperature differences between the two models now normalized by the computed temperature standard deviation at each point. As can be seen, the uncertainty estimates do account for the difference in temperature prediction between the two models, as all results lie within two standard deviations. However, the computed standard deviations are quite large, indicating a large uncertainty in the Green's function model prediction. Further work is needed on improving the calibration procedure to reduce this uncertainty.

3.3 Microstructure Comparison

The temperature history computed using the average values of the calibrated parameters is used as input to microstructure model and a predicted microstructure is calculated along with a microstructure computed using the non-linear Aria model results. The two predicted microstructures are shown side-by-side in Figure 4. Results look qualitatively similar and show a very fine grain structure typical of a small melt pool with limited re-melting between scan lines. For a more quantitative comparison, grain size distributions are shown for the two cases in Figure 5. The Green's function model shows a slight shift towards larger grain sizes as compared to the Aria model. This is likely due to differences in thermal gradients close to the melt pool where latent heat and variable specific heat effects are greatest. However, given the large differences in temperature solution observed in these areas of the domain, the difference in microstructure is fairly small.

3.4 Residual Stress Comparison

The temperature histories from the two models are used as inputs to the solid mechanics model and residual stress profiles predicted. Residual stress values for the two cases, after the part has

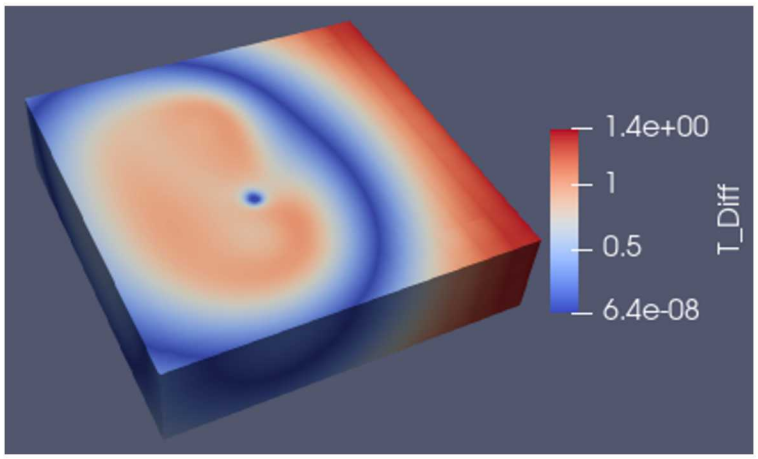

Figure 3. Temperature difference between non-linear and Green's function models normalized by standard deviation

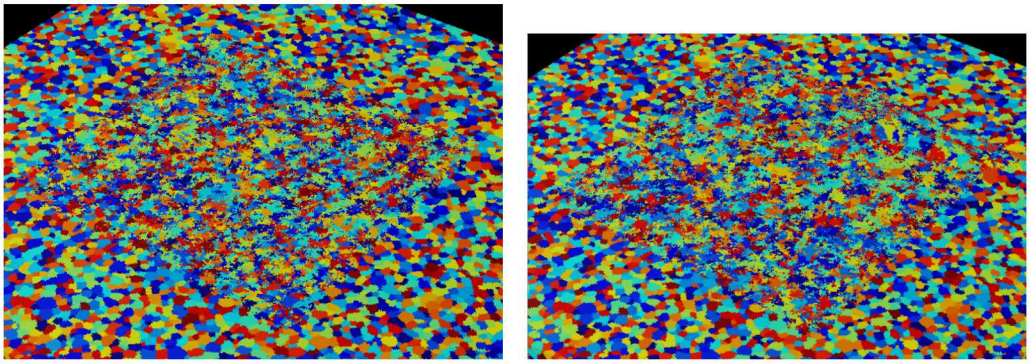

(a) Non-linear **(b)** Green's Function

Figure 4. Comparison between predicted microstructures from Aria and Green's function temperature histories

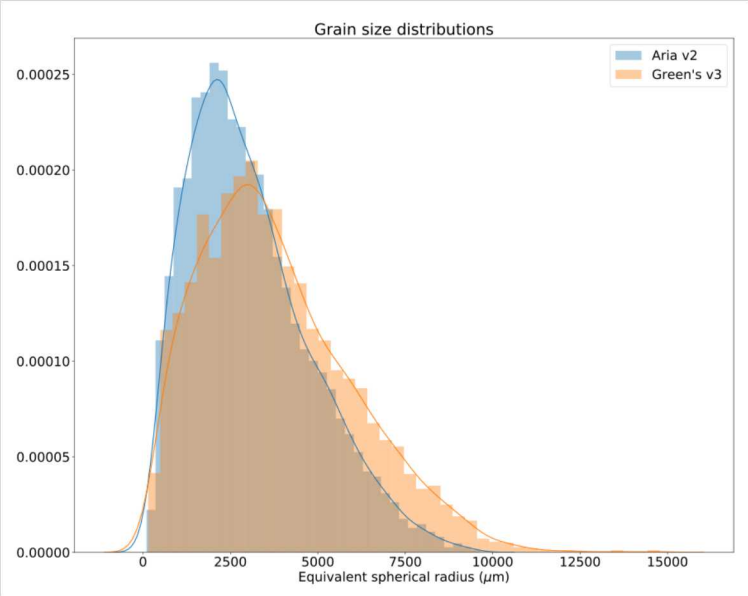

Figure 5. Comparison of grain size distributions between Aria and Green's function models

completely cooled, are shown in Figures 6 and 7. Stress profiles are shown for the XX and YY directions along a line running across the domain perpendicular to the laser scan direction. Profiles are shown at three different depths in the part: the bottom, the top and halfway between bottom and top.

The agreement between the two results is relatively good in the interior of the part. The largest differences are seen at the edges of the part, particularly with the YY stress which at some locations switches from tension to compression. This is likely due to the difference in boundary conditions between the Aria and Green's function models. The Green's function model assumes a semi-infinite domain, whereas this is not possible to impose in the Aria FEM model. Thus, zero-flux conditions are imposed at the domain edges. For a large enough domain, the two boundary conditions will give the same result, but the size of the Aria model domain is practically limited by the computational cost of the model. Given that parts are generally built on top of a base plate much larger than the part size, the semi-infinite assumption is more representative of the physical system than the zero-flux condition. Thus, for this case, the large difference between the stress at the part edges is likely due to a deficiency in the non-linear Aria model, not the Green's function assumptions required for the Green's function model. Initial work has been performed tying the boundaries of the Aria model to a much larger surrounding coarse mesh through contact in order to approximate a semi-infinite boundary condition and initial results show better agreement at the part edges.

Overall, the residual stress results, similar to the microstructure results, show better agreement between the two cases than the initial temperature fields.

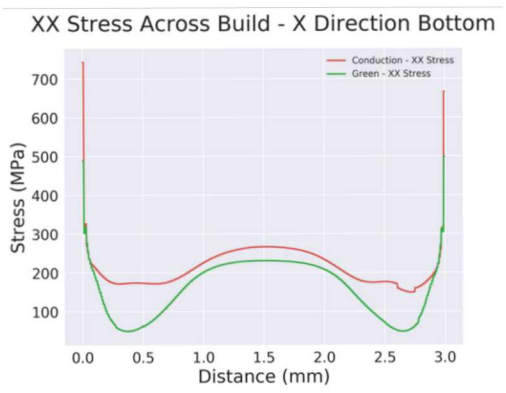

(a) Bottom

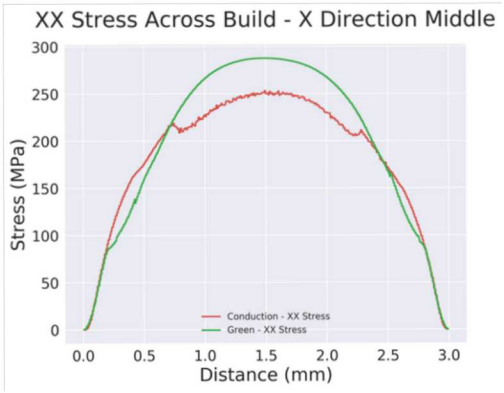

(b) Middle

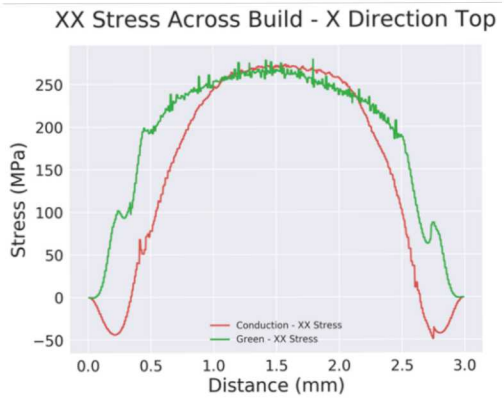

(c) Top

Figure 6. Comparison between predicted residual stress profiles from Aria and Green's function temperature histories - XX Stress

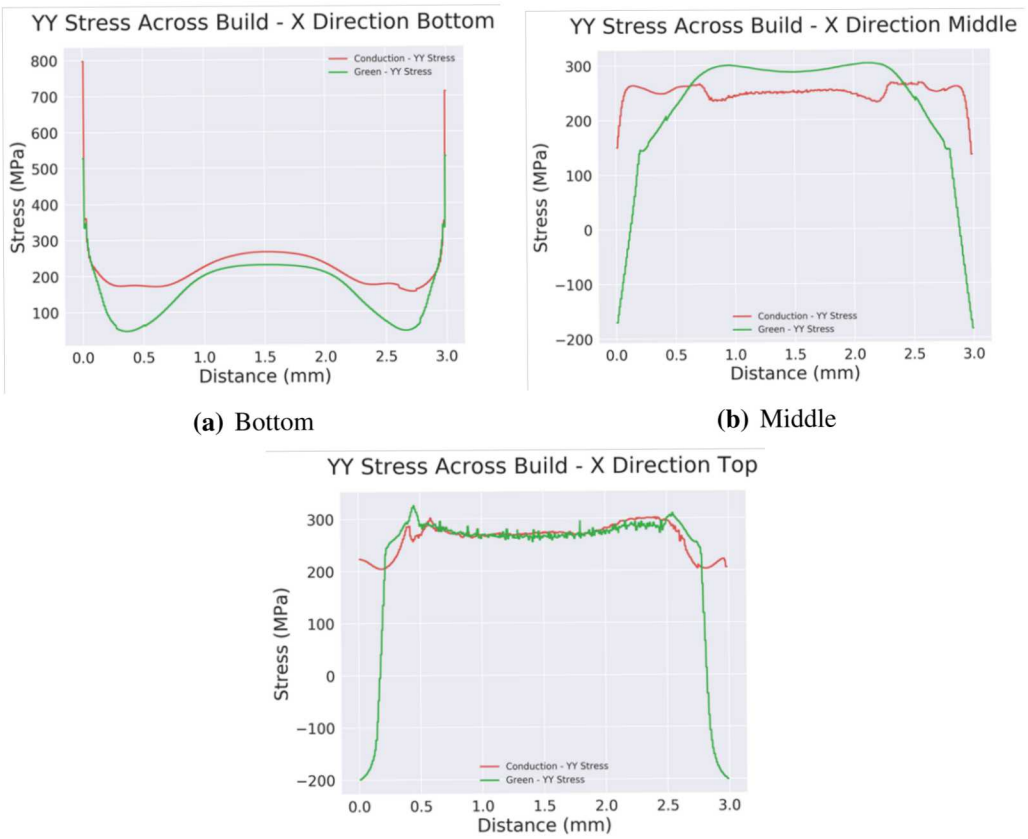

(a) Bottom

(b) Middle

(c) Top

Figure 7. Comparison between predicted residual stress profiles from Aria and Green's function temperature histories - YY Stress

4 Anticipated Outcomes and Impacts

This work represents a first step towards a modeling framework capable of predicting final mechanical performance of laser processed metal parts within a well-defined range of input process parameters with estimated uncertainties in hours-days as opposed to days-weeks. As performance is ultimately governed by the thermal history a part undergoes during processing, a rapid method for predicting thermal history is key to a rapid performance prediction. Having demonstrated the capability to calculate thermal histories 72 times faster than traditional FEM, using a methodology which performs well on new, highly parallel, GPU architectures, this work is well suited to serve as the thermal component of a rapid performance prediction capability.

Further work is needed in order to complete development of this capability. First, the thermal model developed here must be enhanced to better handle non-linear effects. This could be accomplished through sequential Green's functions solutions to iteratively correct for non-linear effects, or potentially using the developed Green's function solution along with the adaptive space-time grid to solve a Green's function pre-conditioned finite elements in time system representing the full non-linear problem. Additionally, the calibration procedure needs to be improved to reduce uncertainties in the calibrated input parameters and expanded to include experimental data as well as higher-fidelity model results.

A capability to model defects also needs to be developed by correlating formation of defects, either predicted by high-fidelity, flow-resolved models, or measured experimentally, with the Green's function temperature solution. This would allow approximate pore densities to be quickly predicted based on the results of the thermal model.

The performance of the microstructure and solid mechanics models also need to be improved to bring their run times in line with the thermal model. Analytical expressions for predicting bulk microstructural features are available in the literature and could be easily incorporated with the thermal model predictions for a coarse estimate of part microstructure. Additionally, newly developed capabilities in the SPPARKS code [12] could enable rapid, fully resolved microstructure predictions by integrating with the Green's function model. For solid mechanics, existing inherent strain models [4] could be enhanced with thermal history predictions from the Green's function model to rapidly predict part residual stress profiles while alleviating some of the known inaccuracies arising from the fact that a full temperature history of the part is not generally available. Additionally, rapid time integration methods such as wavelet relaxation methods [14], large time incremental [3], wavelet transformation multi-time scaling [5], and FFT methods can be investigated for developing a fully time-resolved solid mechanics model.

Finally, the uncertainty quantification work should be expanded to include the propagation of uncertainties from the thermal model through the solid mechanics, microstructure, and defect models so that estimates of uncertainty in the part performance prediction can be made. Many of these efforts will be pursued as part of an ongoing ASC PE&M Advanced Manufacturing project.

Once an "overnight" time frame capability for predicting part performance from laser-metal processing systems is developed, the resulting predictions could then be used to inform design and

qualification decisions for high-consequence applications such as ND. Further investment from ND could enable the use of these novel manufacturing technologies for improving stockpile responsiveness by shortening the design-build-test cycle.

5 Conclusion

The work develops a fully GPU-capable Green's function based solver for laser-metal processing systems along with an adaptive space-time grid for representing temperature solutions of highly-localized phenomena such as laser heating. The solver is found to be 72 times faster than a conventional FEM method for computing the temperature history of an example problem. Bayesian calibration techniques for calibrating the inputs to the Green's function model from non-linear model data are explored. While the estimated uncertainties due to the Green's function assumptions do account for the differences seen in temperature predictions between the non-linear and Green's function methods, further work is needed to reduce the uncertainty in the Green's function parameters.

Using the non-linear and Green's function results as inputs to microstructure and solid mechanics models show that, while the temperature fields have large differences between the two cases, the agreement between the microstructure and residual stress predictions is better than expected given the large temperature differences. This indicates that the downstream models have something of a smoothing effect on the errors incurred by making the assumptions necessary for the use of the Green's function. This is a promising result which suggests that this work could be built upon to develop a rapid part performance capability for laser-metal processing systems. However, additional work is needed improving the calibration procedure, developing a defect model, and improving the performance of the solid mechanics and microstructure models before this can be achieved.

References

[1] Brian M. Adams, Mohamed S. Ebeida, Michael S. Eldred, Gianluca Geraci, John D. Jakeman, Kathryn A. Maupin, Jason A. Monschke, J. Adam Stephens, Laura P. Swiler, Dena M. Vigil, Timothy M. Wildey, William J. Bohnhoff, Keith R. Dalbey, John P. Eddy, Joseph R. Frye, Russell W. Hooper, Kenneth T. Hu, Patricia D. Hough, Mohammad Khalil, Elliott M. Ridgway, Justin G. Winokur, and Ahmad Rushdi. Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.12 User's Manual. Technical report SAND2020-5001, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, May 2020.

[2] Lauren Beghini, Michael Stender, Daniel Moser, Kurtis Ford, and Michael Veilleux. A Coupled Fluid-Mechanical Approach for Additive Manufacturing Process Modeling. In *20th International Conference on Fluid Flow Problems*, 2019.

[3] P.-A. Boucard, P. Ladevze, M. Poss, and P. Rouge. A nonincremental approach for large displacement problems. *Computers and Structures*, 64(1):499 – 508, 1997. Computational Structures Technology.

[4] Matteo Bugatti and Quirico Semeraro. Limitations of the inherent strain method in simulating powder bed fusion processes. *Additive Manufacturing*, 23:329–346, oct 2018.

[5] Pritam Chakraborty and Somnath Ghosh. Accelerating cyclic plasticity simulations using an adaptive wavelet transformation based multitime scaling method. *International Journal for Numerical Methods in Engineering*, 93(13):1425–1454, 2013.

[6] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing*, 74(12):3202 – 3216, 2014. Domain-Specific Languages and High-Level Frameworks for High-Performance Computing.

[7] Nathan Farwell, Mikhail Vorontsov, Dan Moser, and Mario Martinez. Reduced Complexity Heat Transfer Models for Analysis of Powder Bed Multi-beam Laser Fusion. In *20th International Conference on Fluid Flow Problems*, 2019.

[8] R.K. Ganeriwala, M. Strantza, W.E. King, B. Clausen, T.Q. Phan, L.E. Levine, D.W. Brown, and N.E. Hodge. Evaluation of a thermomechanical model for prediction of residual stress during laser powder bed fusion of Ti-6Al-4V. *Additive Manufacturing*, 27:489–502, may 2019.

[9] N E Hodge, R M Ferencz, and J M Solberg. Implementation of a thermomechanical model for the simulation of selective laser melting. *Computational Mechanics*, 54(1):33–51, 2014.

[10] Kyle L Johnson, Theron M Rodgers, Olivia D Underwood, Jonathan D Madison, Kurtis R Ford, Shaun R Whetten, Daryl J Dagel, and Joseph E Bishop. Simulation and experimental

comparison of the thermo-mechanical history and 3D microstructure evolution of 304L stainless steel tubes manufactured using LENS. *Computational Mechanics*, 61(5):559–574, may 2018.

[11] Saad A. Khairallah, Andrew T. Anderson, Alexander Rubenchik, and Wayne E. King. Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones. *Acta Materialia*, 108:36–45, 2016.

[12] Jay Lofstead and John Mitchell. Lazy, Minimal, Eventually Consistent I0 with Stitch. Technical report SAND2018-12554C, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, 2018.

[13] Daniel Moser and Mario Martinez. Particle Scale Melt Modeling of Selective Laser Melting with Uncertainty Quantification. In *Solid Freeform Fabrication Symposium*, Austin, TX, 2018.

[14] Marco Pasetto, Haim Waisman, and J. S. Chen. A waveform relaxation newmark method for structural dynamics problems. *Computational Mechanics*, 63(6):1223 – 1242, 2014.

[15] S. Plimpton, C. Battaile, M. Chandross, L. Holm, A. Thompson, V. Tikare, G. Wagner, E. Webb, X. Zhou, Garcia Cardona, and A. Slepoy. Crossing the Mesoscale No-Man's Land via Parallel Kinetic Monte Carlo. Technical report SAND2009-6226, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, 2009.

[16] Edwin J. Schwalbach, Sean P. Donegan, Michael G. Chapman, Kevin J. Chaput, and Michael A. Groeber. A discrete source model of powder bed fusion additive manufacturing thermal history. *Additive Manufacturing*, 25:485–498, jan 2019.

[17] Jacob Smith, Wei Xiong, Jian Cao, and Wing Kam Liu. Thermodynamically consistent microstructure prediction of additively manufactured materials. *Computational Mechanics*, 57(3):359–370, 2016.

[18] SIERRA Solid Mechanics Team. Sierra/SolidMechanics User's Guide. Technical report SAND2020-5362, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, 2020.

[19] SIERRA Thermal/Fluid Development Team. SIERRA Multimechanics Module: Aria User Manual. Technical report SAND2020-4000, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, 2020.

[20] Bradley Trembacki, David R Noble, Shaun R Whetten, and Mario J Martinez. High-fidelity Mesoscale Thermal/Fluid Modeling of the LENS Additive Manufacturing Process. *The Minerals, Metals & Materials Society (TMS) Meeting*, 2018.

DISTRIBUTION:

1	MS 0899	Technical Library, 1977 (electronic copy)
1	MS 0359	L. Martin, LDRD Office, 1971 (electronic copy)


Sandia National Laboratories