LA-UR-16-26849
Approved for public release; distribution is unlimited.

Title: Developing Materials Processing to Performance Modeling Capabilities and the Need for Exascale Computing Architectures (and Beyond)

Author(s): Schraad, Mark William
Luscher, Darby Jon

Intended for: Report

Issued: 2016-09-08

Developing Materials Processing to Performance Modeling Capabilities and the Need for Exascale Computing Architectures (and Beyond)

Mark W. Schraad (ADX), D. J. Luscher (T-3)

Advanced Simulation and Computing, Physics and Engineering Models

Los Alamos National Laboratory

09/06/16

Additive Manufacturing techniques are presenting the Department of Energy and the NNSA Laboratories with new opportunities to consider novel component production and repair processes, and to manufacture materials with tailored response and optimized performance characteristics. Additive Manufacturing technologies already are being applied to primary NNSA mission areas, including Nuclear Weapons. These mission areas are adapting to these new manufacturing methods, because of potential advantages, such as smaller manufacturing footprints, reduced needs for specialized tooling, an ability to embed sensing, novel part repair options, an ability to accommodate complex geometries, and lighter weight materials.

To realize the full potential of Additive Manufacturing as a game-changing technology for the NNSA's national security missions; however, significant progress must be made in several key technical areas. In addition to advances in engineering design, process optimization and automation, and accelerated feedstock design and manufacture, significant progress must be made in modeling and simulation.

First and foremost, a more mature understanding of the process-structure-property-performance relationships must be developed. Because Additive Manufacturing processes change the nature of a material's structure below the engineering scale, new models are required to predict materials response across the spectrum of relevant length scales, from the atomistic to the continuum. New diagnostics will be required to characterize materials response across these scales. And not just models, but advanced algorithms, next-generation codes, and advanced computer architectures will be required to complement the associated modeling activities. Based on preliminary work in each of these areas, a strong argument for the need for Exascale computing architectures can be made, if a legitimate predictive capability is to be developed.

Other Manufacturing Processes and Aging

Of course, Additive Manufacturing is only one process that affects a material's structure, its properties, and its performance. Any manufacturing process will leave unique fingerprints on a material's structure across the spectrum of length scales from the atomistic to the continuum. Interestingly, material aging will leave a different set of fingerprints, but on the very same components of the material's structure. Therefore, it may be important to note that being predictive with regard to either manufacturing or aging requires the same ability to deal with a material's structure across length scales.

The NNSA Laboratories have had to think about material structure effects in the past, when we evolved our weapons manufacturing from wrought processing to casting. We are currently considering how Additive Manufacturing, and more general Advanced Manufacturing, processes will again require us to think about the influence of material structure on component and systems-level response. And discussions about potential use of microwave casting techniques for manufacturing will only confound these issues in the future.

Fundamentally, however, each of these processes can be used to produce a part for use in an engineering system. For instance, a particular system component (such as an automotive engine piston, an aircraft component, or a component of a nuclear weapon) can be produced using wrought, traditional casting, Additive Manufacturing, or microwave casting technologies. In each of these cases, the component will possess the density, and likely the elastic properties one would expect for the given material used (e.g., stainless steel, aluminum, etc.). However, the details of grain morphology, residual stress state, defect structure, etc. likely will all be different. For applications involving only non-cyclic, small deformations, these details may not matter. However, for deformations through many cycles, such as a piston or an aircraft engine will see, these details may become important, even critical. And when we change the materials and operating conditions to those relevant to nuclear weapons safety and performance, where the pressures, temperatures, strains, and strain states are much more extreme, it becomes imperative to know these details, and their potential effects on component and system-level performance.

Why Material Structure Is Important

If we are to develop a legitimate predictive capability for modeling materials performance; that is, a capability in which we have high confidence in our predictions, then we must have models representative of all the relevant physical processes, across the spectrum of length scales, from the atomistic to the continuum.

The atomic length scale is the fundamental scale that determines the single-crystal properties of metals and metal alloys. It is at this scale where we determine the equation-of-state behavior of these materials; that is, how a material will respond to varying states of pressure and temperature, and when a material will transform from one crystallographic structure to another based on the state.

In addition to a material's atomic structure, the defect structure plays a critical role in determining a material's yield strength, damage response, and ultimate failure under loading. Atomic-level vacancies, impurities, and interstitials, along with initial porosity and micro-crack distributions will be prevalent in even the most carefully processed materials. While the dislocation structure and actual dynamics of dislocation motion dictate the plastic flow response of a material. The defects, along with grain boundary

interfaces, provide impediments to this dislocation motion, thereby actually increasing the material's yield strength. However, these same defects will serve as potential nucleation sites for porosity, thereby decreasing the materials ultimate spall strength. Ultimately, mesoscale models will become critical for understanding which of several potential mechanisms will be triggered in order to accommodate the overall deformation.

Because different manufacturing process will produce different defect structures, dislocation structures, and grain morphologies, these processes will necessarily affect the material's plastic yield strength, damage, and failure response to dynamic loading conditions. Because material aging affects the same material structure details, manufacturing and aging considerations are ultimately coupled.

Predictive Capability Requires More Than Just Material Modeling

Though we have many state-of-the-art models at many of these length scales already, and though we likely will make significant progress in developing more physically based, more accurate, and more numerically robust models in the future, models at each of the individual scales are just some of the necessary "ingredients" of a legitimate predictive capability.

Of course, none of these models will be useful without the accurate experimental data required to guide model development, and to test, calibrate, and validate the models. Current diagnostic capabilities are lacking at the lowest length scales, and nuclear weapons applications often make it difficult to capture data in the regimes of pressure, temperature, and strain-rate space where the model development needs are greatest.

The physical processes of atom vibration, dislocation motion, single-crystal deformation, pore growth and coalescence each require different mathematical representations and different techniques for their solution. It is not possible to simply incorporate all of these physical mechanisms in a "grand unified" model of material response. Nor will a truly predictive capability be developed by simply passing the results of one solution "up the chain" to the next length scale for use. Rather, the individual models must be coupled algorithmically, and thus, appropriate scale-bridging algorithms, or similar techniques, must be developed.

And, of course, appropriate simulation tools must be available that are accepting of this coupled algorithmic machinery. That is, the coupled algorithms and models must be consistent with the native hydrodynamic solution techniques and able to be coupled with the other physics in any given systems-level numerical model.

Such coupled algorithmic modeling treatments will require significantly more computational resources than single-physics, single-scale models, and so the coupled algorithms and simulation tools must be designed for the next generation of computer architectures, and those computer architectures must be large enough to accommodate meaningful component-level or system-level simulations.

An Apparent Need for More Computing Resources

Some researchers are beginning to investigate the development and use of coupled algorithms and advanced models within the ASC IC code base. The results show great promise for our ability to develop more sophisticated models, associated scale-bridging algorithms, and to implement this machinery in our existing codes. However, our ability to run simulations of meaningful size is greatly limited by today's computational resources.

For example, we have conducted coupled multi-physics simulations of shocked polycrystalline metals using LANL's FLAG hydro code. Figure 1 below shows the results of two different simulations. Both represent a 1 mm by 1 mm two-dimensional stochastic volume element. The nominal grain size is 200 μm. A 5 μm computational cell size achieves a resolved numerical solution, which results in a simulation with approximately 40,000 computational cells.

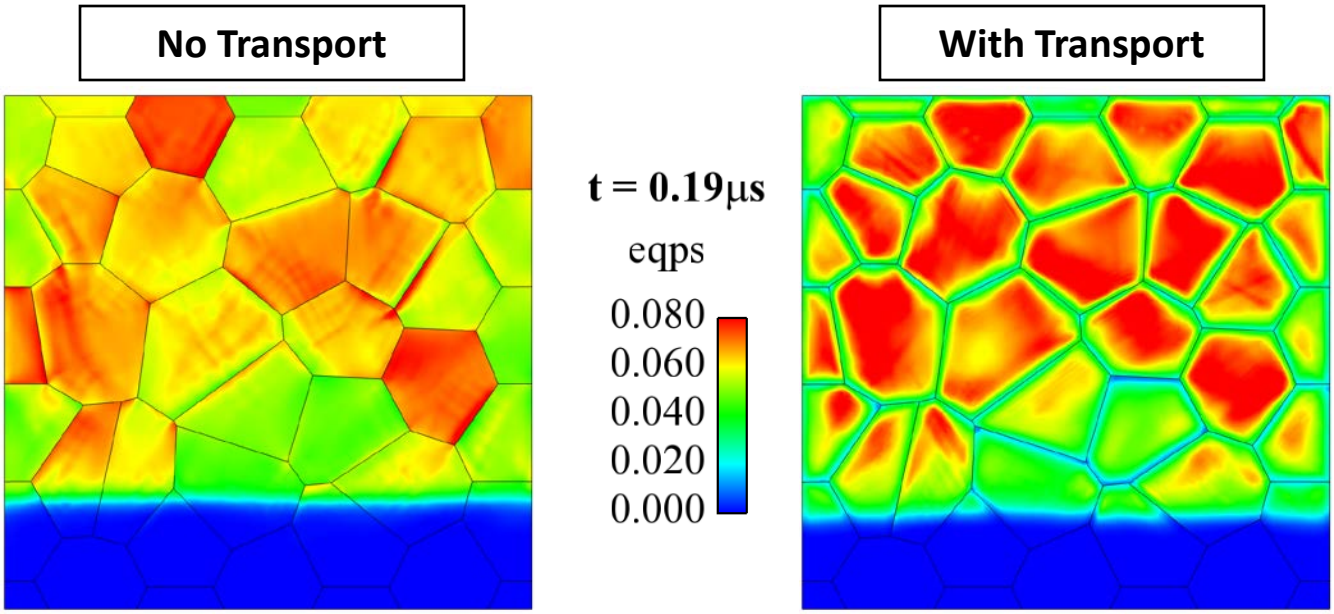

Figure 1. Two-dimensional flyer plate simulations using standard local crystal plasticity and crystal plasticity coupled with the transport of dislocations.

A flyer plate simulation is run with a flyer velocity of 270 m/s, which leads to a shock pressure of approximately 5 GPa. Simulations to 0.55 μs after impact are conducted for standard local crystal plasticity (no transport of dislocations) and for crystal plasticity coupled to dislocation dynamics (with transport of dislocations). Figure 1 shows the difference in equivalent plastic strain at a time of 0.19 μs after impact.

Note the difference in plastic strain level in the interior of the grains, and most notably, the reduced levels of plastic strain at the grain boundaries. The significance of the difference between the results on the left and right of Figure 1 is that by directly coupling the physics pertaining to plasticity within individual crystals

with the accompanying motion of dislocations, the case on the right is able to capture the impeded magnitude of plastic deformation within the vicinity of grain boundaries. This non-local material response is crucial in order to eventually introduce models of specific grain boundary behavior needed to properly model the statistics of damage nucleation at grain boundaries. On the other hand, the standard crystal plasticity model used for the results on the left, knows nothing of the spatial interaction of dislocations with grain boundaries, and consequently, grain size and grain boundary characteristics will not affect the predicted macroscale response.

The standard simulations (no transport) require 5 hours on 128 processors of Moonlight to run to 0.55 μs. The fully coupled simulations require 72 hours on 128 processors of Moonlight. So, just coupling the dislocation dynamics to the polycrystal scale for a small (1 mm by 1 mm) two-dimensional simulation increases computing requirements by more than an order of magnitude. As a point of reference, Moonlight has approximately 5000 processing cores.

Extending this new capability to three-dimensional simulations poses additional computational challenges. Figure 2 below depicts a 1 mm cube stochastic volume element discretized into a computational grid of 20,000 arbitrary polyhedral computational cells. The simulation results generated in FLAG illustrate that this computational grid is far too coarse to be reliable for accurately resolving this behavior. The grid comprised of 500,000 computational cells, shown in Figure 3, better resolves the required accuracy; however, linear scaling suggests that a fully coupled calculation will require either 3 days using over one quarter of Moonlight (1280 processors), or a smaller allocation of 128 processors for nearly one month.

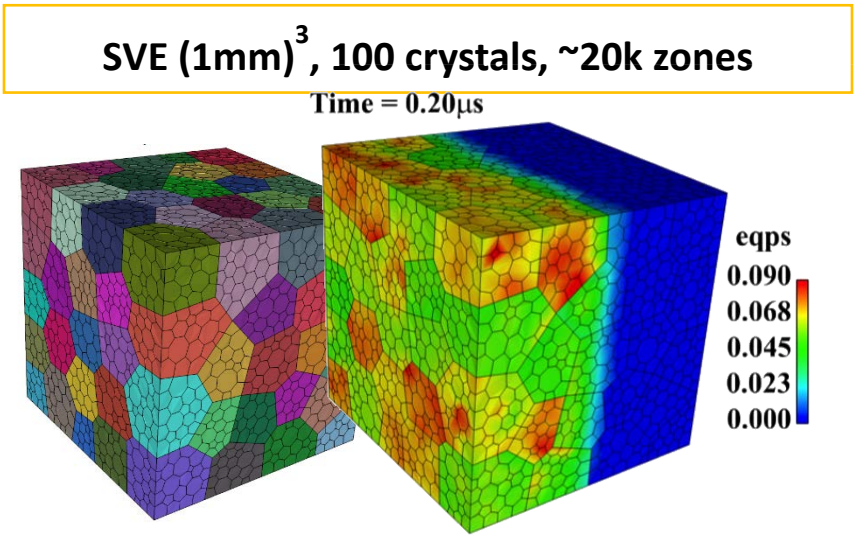

Figure 2. Three-dimensional simulations of flyer-plate impact, analogous to those shown in Figure 1.

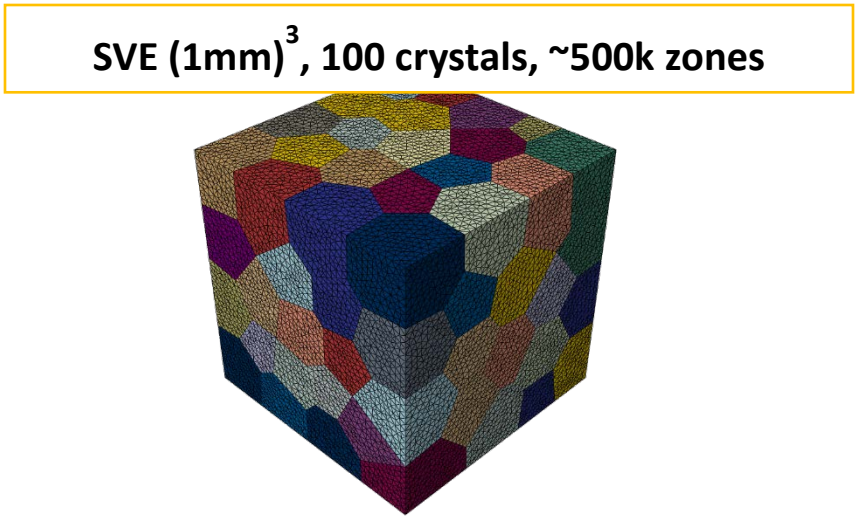

Figure 3. A 500,000 cell computational grid required for accurate resolution of the physics.

Clearly, a credible multiscale modeling strategy to bridge between mesoscale calculations such as these, and macroscale simulations of some component or system, demands reducing the computational time needed for the mesoscale simulations, and there are several paths to achieve such an end. First, the code associated with these models and their coupling can be optimized (acceleration). Second, alternative coupling strategies can be identified to improve numerical stability, and thus attain larger time-step sizes. Third, options can be created for running the fully coupled schemes only in areas where it matters; such as near grain boundaries (adaptive coupling), and these options can be improved for multi-scale applications via adaptive resolution (scale-bridging algorithms). However, even with all of these advances, the target applications may still require Exascale computing, and beyond, for meaningful simulations. Therefore, the final path involves running on many, many more cores using hardware advancements associated with Exascale computing initiatives.

Of course, single-component, single-physics simulations, such as the full-up simulation of a flyer plate experiment, performed in three dimensions, will require spatial dimensions on the order of centimeters, or 10s of centimeters, and thus the simulations shown above in Figure 2 would scale by a factor of 10^3 to 10^6, just to provide the ability to simulate single components. Full systems would require orders of magnitude beyond that. Additionally, a legitimate predictive capability may require coupling across more than one or two length scales. Additive Manufacturing actually introduces additional structures (e.g., the weld bead scalloping and layering) that affect material performance, so it may be necessary to include consideration of physics at many length scales (atomic, dislocation, defect structure, grain boundary, single crystal, polycrystal, scallop, layer, continuum) into any given simulation. If the examples above are any indication, a reasonable estimate of computational cost could be an order of magnitude per length scale per spatial dimension. And finally, the ultimate goal is to couple materials physics simulations with other relevant physics to address NNSA mission needs. And other physical processes may face the same reality; namely, that legitimate predictive capability requires the ability to model lower-length-scale

physical processes and bridge across many length scales. So ultimately, full-system, three-dimensional, coupled physics simulations will require Exascale computing and beyond.

Conclusion

A legitimate predictive capability for assessing material, component, and systems-level performance, based upon the manufacturing processes used, can be achieved. However, such a capability will require continued contributions from the physics modeling, algorithm development, code development, computer platforms, and experimental communities. The requirements of models, algorithms, codes, computer platforms, and experimental data also point to the need for optimized program integration, if we are to meet the grand challenge of material performance prediction based on processing and aging. And while Exascale computing is a necessary ingredient, based upon the examples and arguments provided above, a final comment would be that computing power alone, and indeed, none of these components alone, will achieve the desired objective. But rather a concerted, coordinated effort by all of the components of the ASC Program, along with collaborations with our Science Program colleagues will achieve the desired end state.