

NISTIR 8110

The New Frontiers in Computational Modeling of Material Structures

William Regli
Jarek Rossignac
Vadim Shapiro
Vijay Srinivasan

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.IR.8110>

NIST
**National Institute of
Standards and Technology**
U.S. Department of Commerce

NISTIR 8110

The New Frontiers in Computational Modeling of Material Structures

William Regli
*Defense Sciences Office,
Defense Advanced Research Projects Agency*

Jarek Rossignac
*College of Computing,
Georgia Institute of Technology*

Vadim Shapiro
University of Wisconsin at Madison

Vijay Srinivasan
*Engineering Laboratory,
National Institute of Standards and Technology*

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.IR.8110>

February 2016



U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

The New Frontiers in Computational Modeling of Material Structures

William Regli¹, Jarek Rossignac², Vadim Shapiro³, Vijay Srinivasan⁴

¹Defense Sciences Office, Defense Advanced Research Projects Agency

²College of Computing, Georgia Institute of Technology

³University of Wisconsin at Madison

⁴National Institute of Standards and Technology

Abstract

We are witnessing the emergence of a new paradigm in the modeling of *material structures*. It stems from the digitization of manufacturing and is fueled by advances in additive manufacturing and material science. This paper strives to provide a critical examination of this new paradigm in a historical and technological context and to show that it requires non-trivial extensions and generalizations of the classical theoretical foundation and algorithmic solutions originally developed for solid modeling. Specifically, it requires new models and data-intensive representations for materials, physical behavior, and manufacturing processes across multiple scales. In particular, we argue that most computational tasks that support traditional and emerging manufacturing may be formulated systematically and addressed in terms of relations (*conversions*, *synthesis*, *change propagation updates*, *verification*, and other *harmonization* activities) among four views (manifestations) of an engineered artifact: *Functional*, which captures the design constraints and tolerances on shape, properties, and behavior; *Designed*, which represents a tolerated design that satisfies these constraints; *Planned*, which defines a manufacturing process plan; *Simulated*, which models the expected outcome of the process plan; and a *Real* sample set of physical artifacts produced by executing the process plan on a particular manufacturing technology. Based on this formulation, we outline important directions for a research agenda aimed at enabling, driving, and amplifying further advances in digital design and manufacturing.

1 Introduction

Sutherland’s Sketchpad [46] is often cited as an example of the early CAD systems. Support of 3D rendering, and other early developments in CAD have been driven largely by design and manufacturing applications: Bezier’s curves and surfaces were developed for modeling stamping dies in automotive manufacturing, wireframes were used for modeling aerospace parts, and solid models were originally intended for representing NC-machined parts and mechanical assemblies [49].

With the advent of solid modeling, fundamentally new mathematical theories and representations emerged. They made it possible to capture the complete geometry and topology of manufactured artifacts, as well as some physical properties. They opened the door to a transition from purely visual depiction to computational models of physical artifacts. Over the last four decades, solid modeling has developed into a mature discipline that is based on rigorous foundations [34, 17, 42], supported by a vibrant research community [2], and is at the core of virtually all computer-aided design and manufacturing activities—these, in turn, support over \$10 trillion in global engineering, manufacturing and commerce annually [30]. Solid modeling and




Figure 1: An example of a 3D lattice structure, designed and fabricated by Hughes Research Labs (HRL), representing the new frontiers for materials and design. 99.99% of this metallic truss structure is air. The remaining 0.01% is made of very thin (nanometer, micron and millimeter scale) features.

computer graphics were tightly intertwined till the early 1990s, at which time computer animated movies and video gaming hardware suddenly consumed the attention of the mainstream graphics community. Modeling in support of commercial computer-aided design focused on improved robustness and performance, more powerful design tools, more useful tolerancing, faster and more accurate analysis, and supporting product data and product lifecycle management (PDM/PLM).

Presently, we are witnessing the emergence of the need for new modeling paradigms that go beyond solid modeling. It is largely fueled by the Third Industrial Revolution, which is based on the digitization of manufacturing [9]. Advances in materials science and additive manufacturing make it possible to manufacture artifacts with a complex material structure that for example makes them extremely light and strong (Figure 1). Supporting such models require that we go beyond representing discrete homogeneous parts, sheet metal, and assemblies (as used in the aerospace and automotive industry) or surface and polygonal mesh models (as used in animation and gaming). Although several researchers have proposed theoretical foundations and practical implementations of non-manifold structures (see examples of proposals and surveys in [40, 43, 19, 38, 39]) that extend the representational capabilities of solid modeling, these early attempts do not suffice, by themselves, to address the novel challenges discussed here.

These challenges require the capabilities of modeling embedded microstructures, internal geometry architectures, multi-scale behaviors, and composite multi-material objects, because such artifacts are now physically realizable and widely used. Further, it is now commonly possible to vary internal material properties throughout the artifact, either by using graded microstructures (e.g., lattices) or via fabrication processes that can alter the crystalline structures of metals as they are deposited using 3D printing. The palettes of physical realizations that are now possible are unmatched by the relatively primitive design and modeling capabilities intended to support the mass production systems of the last century.

This paper attempts to describe key aspects of the next frontier for modeling, with a particular focus

on the opportunities (and challenges) emerging from additive manufacturing and the revolution in materials science. Additive manufacturing technologies promise to radically alter production and manufacturing. Elimination of industrial waste, part-count reduction, new forms of multi-functional products and vastly lower material and energy costs are benefits that—at least in principle—will flow from the rapid adoption of additive processes. We examine the set of representational challenges that must be solved in order to support advances in production processes and materials. We also hope that the issues identified in this paper will guide the agenda for research and technology development in CAD, modeling, graphics, and visualization for the coming decade and beyond. In doing so, we follow the spirit of some of the early pioneers in solid modeling [50, 36, 37, 35] and provide a brief context of current technology needs with respect to existing work in geometric and solid modeling.

2 Historical Context

2.1 Geometry-based representations in design and manufacturing

The ability to represent and communicate information about the design and manufacturing of artifacts is at the heart of the modern manufacturing enterprise. Detailed geometric drawings specifying construction of buildings were already in use in ancient Greece [16]. Without the ability to describe and communicate the shape of interchangeable components, manufacturing was largely confined to low-volume and inaccessible artisan activity [6]. The need to describe and communicate the shape of interchangeable mechanical components in assemblies, tooling, and fixtures, as well as the methods of their manufacturing, has led to the wide adaption of standard engineering drawing practices in support of mass production [21].

In spite of enhanced computerization and automation, traditional manufacturing processes, such as CNC machining, casting or forging, have remained largely unchanged for over 50 years. Computer-Aided Design (CAD) emerged as a means of automating mechanical drafting, and Computer-Aided Manufacturing (CAM) relied on the CAD models to directly drive machine tools, to automate tool paths and to facilitate the direct exchange of digital models.

During the 1970s and 1980s, in a major technological paradigm shift, solid modeling emerged in the attempt to create an informationally complete model of a *manufactured shape* that could be used throughout the manufacturing enterprise and support engineering activities throughout the product life cycle. The pioneers of solid modeling also recognized that the notion of informational completeness is not absolute, but is relative to assumed or postulated mathematical models. The latter, in turn, are based on target class of physical artifacts and processes. The instantiation of these techniques in data structures, algorithm and interfaces was a triumph of software industry during this time. Many practical issues related to the underlying representations (constructive solid geometry, non-uniform rational b-splines, winged edge and half edge data structures) and mathematical limits of digital computing (floating point accuracy, error stack up, robustness) were, for most practical purposes, overcome and the resulting companies constitute a \$10B/year industry [52].

The currently accepted mathematical notion of a rigid, internally homogeneous solid was deemed adequate for supporting most (but not all) engineering activities in traditional manufacturing of mass-produced mechanical assemblies and is the basis for all modern commercial CAD systems. Early geometric and solid modeling system were aimed to support NC machining, sheet metal forming, design and planning of mechanisms and assemblies, tolerance analysis, as well as simulation via finite element analysis. Later geometric and solid modeling tools evolved to represent geometry (shapes and operations) associated with




Figure 2: Illustration of an additive manufacturing process

most unit¹ manufacturing processes [29] and became the backbone of the modern PLM (Product Life-cycle Management) systems.

While PLM systems provide “geometry-based” representations of mechanical systems, they aspire to represent a complete virtual product model, including materials, physics (simulated or experimental), and intended behavior (usually in a form of performance specification and testing procedures). Of course, this virtual product model is never truly complete, but is sufficient to effectively support the paper-based and human-centric processes that have been used in most organizations. Many of these processes have evolved over decades and represent best practices, as well as technological limitations, of the modern manufacturing systems. For example, most traditional manufacturing processes assume that material properties of each component are homogeneous (at least within a single part or component of an assembly) and may be represented by a few material constants.

One of the most profound consequences of this evolution is the separation of design and manufacturing activities (and hence modeling practices) that was demanded by efficient and decentralized manufacturing. Thus, modern GD&T (Geometric Dimensioning and Tolerancing) standards discourage explicit representation of manufacturing information in a design model [1], giving rise to distinct notions (and computer representations) of the designed and the manufactured shapes that are manifested by two popular paradigms in solid modeling: design using (manufacturing) features and parametric constrained-based design. Such separate manifestations of the same artifact are often locked in separate software systems, algorithms and digital representations (i.e., the FEA model exists distinctly from the CAD model which is in turn distinct from the tool path model), creating major interoperability problems [18].

2.2 New manufacturing paradigms demand new representations

Recent advances in material science and additive manufacturing completely invalidate many of the assumptions about manufactured shapes that were appropriate for traditional manufacturing. Coincidentally, the same advances serve as a catalyst for the digitization of manufacturing that has already been under way, finally pushing the geometric modeling and CAD systems technology beyond their capabilities. It is clear that the new and emerging manufacturing processes do not fit the class of traditional unit processes for discrete parts, and their manufacturing capacity already exceeds our ability to conceive, design, represent, and simulate artifacts that are manufacturable. In the next section, we attempt to identify the key bottlenecks and

¹Informally, unit processes are individual steps in a manufacturing process that are transform the raw material into a finished product [29].

opportunities in creating the next generation of modeling and design systems that can support and propel the manufacturing revolution currently under way.

To that end, we classify discrete and batch manufacturing processes by how they shape raw materials into finished goods:

- **Formative shaping:** The desired shape is achieved by application of temperature and pressure to a body of raw material. Examples include forging, casting, injection molding, bending, sintering, and compaction.
- **Subtractive shaping:** The desired shape is achieved by selective removal of material. Examples include turning, milling, drilling, eroding, and electro-discharge machining.
- **Additive shaping:** The desired shape is achieved by successive addition of material. This broad definition includes the manufacturing of composite structures, in which layers of fiber-impregnated plastic sheets are stacked on top of each other, usually starting with a tooling surface, and then subjected to pressure and temperature to achieve the final shape. Lately, the phrase “additive manufacturing” (AM) has come to be used to refer to a class of manufacturing processes that join materials, usually deposited as consecutive horizontal slices. Such AM processes (shown in Figure 2) include binder jetting, direct energy deposition, material extrusion, material jetting, powder bed fusion, sheet lamination, and vat photo-polymerization [3, 14].
- **Assembly:** Artifacts of any reasonable complexity are assembled out of many simpler components that are connected to realize the overall functionality of the full design. This processes can be viewed as the reverse of the top down, hierarchical decomposition of a design into subsystems and sub-functions: an automobile is composed of drivetrain, engine, frame, etc., which in turn are composed of other parts and mechanisms. These processes occur over length scales of 3-to-6 orders of magnitude (e.g., millimeters to meters in the case of a vehicle) [47].

The traditional geometric and solid modeling technologies have served us reasonably well to support the formative and subtractive manufacturing processes, but have limitations when applied to additive shaping processes. This is partly due to the fact that the material modeling issues are intricately linked to shape modeling issues in the additive shaping processes. In many cases, this is not by accident. In fact, the main purpose of composites structure manufacturing is to deliberately introduce inhomogeneity and anisotropy of material throughout the solid shape so that the resulting object will have the desired property and behavior in service. In other cases, such as additive manufacturing of some metallic parts that require homogeneous and isotropic material distribution, the peculiarities of the process introduces undesirable characteristics such as porosity and delamination, which require detailed material and shape modeling.

For the sake of examining the underlying shape and material modeling issues, it is useful to separate Composites structure Manufacturing (CM) from Digital Printing (DP) technologies. Let us examine how the current computer-aided modeling systems deal with these two manufacturing technologies. The design and analysis of composite structures is a team effort that involves a close cooperation of design, analysis, and manufacturing engineers well versed in CM. A CAD system that creates a geometric model of a composite structure starts with a base tooling surface and a ply table that is, in essence, drawn from a library of fiber-impregnated polymer sheets. The ply table includes geometric information, such as the relative positions of the sheets and the relative fiber orientations, which can be defined using recently-issued engineering drawing standards. In an iterative design process, the composite structure so created is analyzed to see whether it meets the expected behavior, at which point the iteration stops. As a result, the design model of a composite

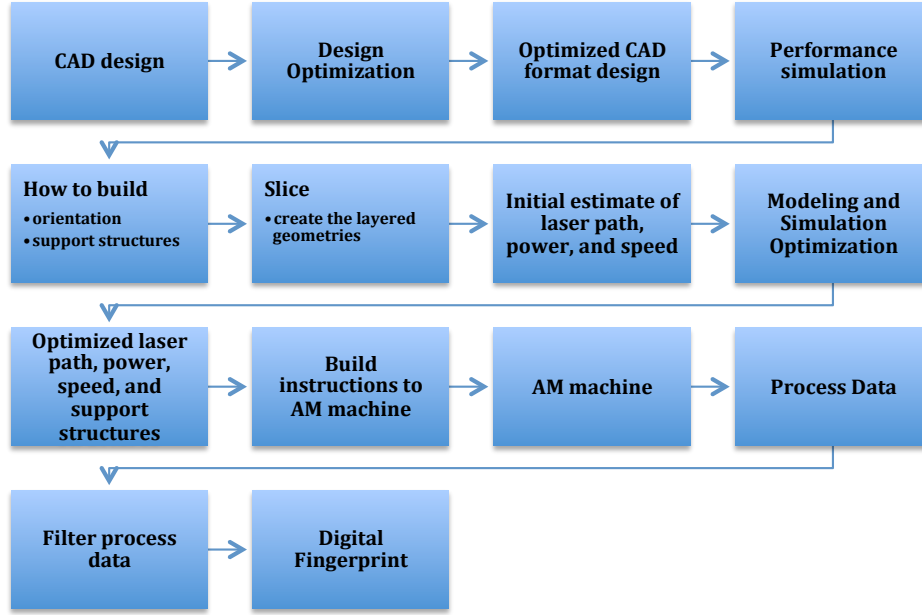


Figure 3: A portion of the digital thread of file formats and representations on the pathway from design model to manufactured part (from W. King of LLNL).

structure also carries part of the manufacturing recipe, i.e., the ply tables that contain the ingredients and the layering sequence needed to bake the laminated structure in an autoclave.

In contrast, the design and analysis of parts for DP take a different approach, that is still dictated by the adopted manufacturing technology. After creating a geometric model of the shape to be manufactured, support structures (to hold the part being built in place) and fins (to dissipate the heat generated in the DP process) are added to the part shape and analyzed. Then a significant effort is spent on computing the planar slices require by the AM processes. Such slices can be generated directly from a 3D model of the part. But, the prevailing approach is to create a triangular tessellation of the boundary of the augmented (e.g., with fins and support structures) 3D model of the part (which may be transferred using a popular stereolithography, or STL, de facto industry standard) and slice it to create a planar polygonal shape. Available since the late 1980s this crude, but simple, surface model format has been a benefit as well as a bane to the advancement of AM. Multiple AM technologies can interoperate with any 3D modeling systems as long as an STL file can be used as the interface. The crudeness of the triangular tessellation is tolerable as long as the discretizations used in the AM process are significantly cruder than the errors introduced by approximating an original geometry by the polygonal model that can be captured in an STL format. The situation is changing as DP processes are getting more precise, and the demand for manufacturing near-finished parts using DP is increasing. Recent efforts, such as AMF standard [23], are addressing the model precision issues, along with ability to capture material variability throughout the part.

These two examples demonstrate the salient features of the new additive manufacturing processes. The notion of a manufacturable shape is now intricately coupled with issues related to materials, physics, geometric representation and computing. The extent to which this coupling is understood and managed determines efficacy of new designs, as well as automation and efficiency of manufacturing.

Perhaps the most important transformational change of additive manufacturing is that it explicitly iden-

tifies manufacturing as a computational process that allows manufacturing artifacts of unprecedented complexity, but does not tolerate ambiguities in assumed models and representations. Figure 3 illustrates the complex data and computational pathway “from art to part” that is currently used for these new processes.

In what follows, we examine the key differences and challenges posed by new manufacturing technologies and propose possible paths towards addressing them through an alternative new paradigm that we call *material structures modeling*. The main purpose of material structures modeling is to ensure that the mathematical models and resulting computational representations are capable of fully supporting and taking advantage of these manufacturing processes and associated engineering activities.

3 The Rise of Material Structures

Three characteristics of emerging additive manufacturing processes that distinguish them from the traditional unit processes are: (1) variable, process-dependent material microstructure, (2) ability to manufacture multi-phase materials (or multi-materials), and (3) ability to produce a broad diversity of material architectures, including those with embedded physical and mechanical functions. Below we discuss how each of these characteristics contributes to emerging concept of a new manufactured shape that we term in this paper *material structure*. Intuitively, we might expect that a material structure is simply a union of homogeneous solids,² but such naive conceptualization does not account for inherent complexity of material structures.

3.1 Material microstructure and process physics

In a metallurgic context a *grain*, or crystallite, is basic unit of crystalline structure that is created during material formation. For example, in metals the grain structure is influenced by factors such as heat, pressure and other processing parameters. The resulting grain sizes, orientations and structures determine material performance. The term “microstructure” refers to the configuration of materials from the atomic scale to the grain structure, as illustrated in Figure 4 from [25]. From a traditional manufacturing viewpoint, such as machining metals, a microstructure is considered to be largely uniform across a workpiece. For example, an ingot or stock workpiece of A2 stainless steel would be assumed to have consistent micro-structure based on its forming processes and its performance and behavior would be predictable from these well-studied metallurgic properties. In contrast, in additive manufacturing, the resulting material microstructure depends explicitly on the energy levels, speed of the energy source, size of the stock powder and other process parameters. Different parameters can alter the subgrain structures and the resulting crystalline lattices they form—thus altering macroscopic behavior of the structures made with the material (see example in Figure 5).

Hence, rather than the uniform nature of the stock material in traditional manufacturing, we now have a final artifact with possible different (i.e., space varying) micro-structures even when they may be created from the same base material. We may wish to control the additive process to create as much uniformity as possible (to mirror a traditional process) or we might want to leverage these variations to create new possibilities for artifact behavior. In the latter case, our understanding of the relationships between process parameters and micro-structure (and hence artifact performance) is limited or, in some cases, even non-existent. For example, as lower layers are partially reheated by subsequent layers of deposition their micro-structure changes, creating perhaps unexpected properties.

Without extensive data and/or deep scientific understanding, predicting the outcome of a manufacturing process and developing process plans to consistently produce desired configurations of microstructure

²Note that traditional notion of “solid” in geometric modeling is a topological property[34] and does preclude the solid object being a fluid, gas, or plastic.

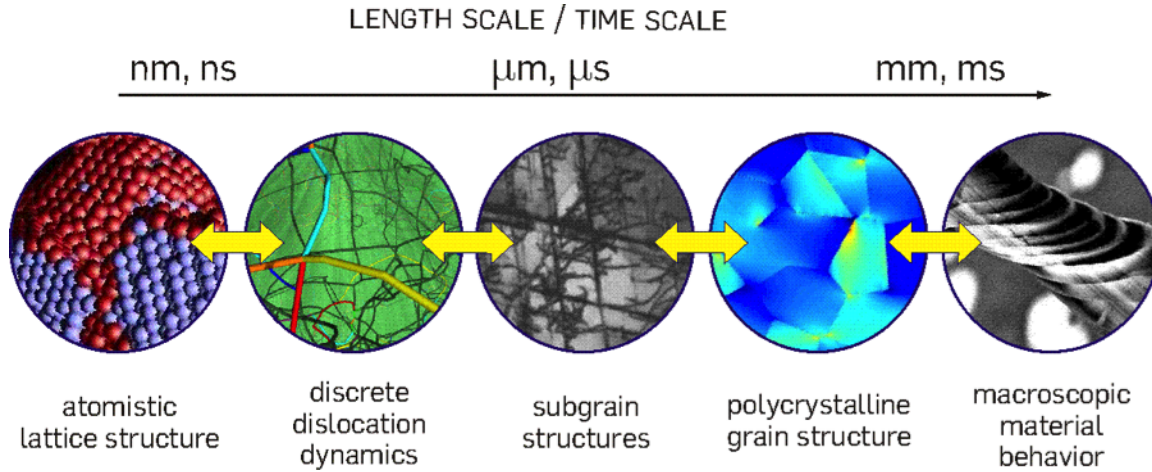


Figure 4: The parameters used in additive manufacturing processes can affect the final material configuration across multiple scales [25].

becomes difficult. Further, present understanding of new additive processes on the stock material is also a matter that is open. For example, in the case of metal laser sintering, our understanding of stock powders, their provenance, structure and how they influence the final product’s properties is nascent, but it is well known that energy level of a laser source drastically changes the resulting microstructure of the final material [20] (shown in Figure 6). Hence, one cannot simply specify “Ti-6Al-4V” and assume the properties of final artifact lasered or electronbeamed out of powder will be the same as that of the same artifact machined out of stock material. Nor can we necessarily assume all powders of the same material will have the same performance. Powders vary in terms of the average size and distribution of the particles, their surface properties (e.g., oxidation) and other issues.

A significant open question for AM is how to make things that we cannot even conceive of fabricating at present. Hence, one might imagine wanting to leverage the ability of these processes to create artifacts with variable internal properties, customized microstructure, or unique configurations for internal stresses in order to achieve novel functional goals.

3.2 Multi-material manufacturing

Many new additive manufacturing processes freely mix and fuse multiple materials. The most common example of such mixing is a great variety of composite materials and structures where fibers and particles of all shapes and sizes may be embedded into a matrix and fused together in order to achieve superior mechanical properties (weight, structural, thermal, etc.). Many of the new 3D printing processes are also capable of printing multi-material structures. For example, processes that use a material wire feed (and energy from an electron beam or laser source) have been shown capable of producing multi-material objects. In these processes, multiple wire feeds, each from a different material source, can be blended or transitioned by the energy source and deposited. This enables production of objects that can be made of multiple materials, or blends of materials.

These and other multi-material configurations can have functional and behavioral properties that span multiple energetic performance domains and temporal regimes: structural, optical, conductivity, heat transfer, etc. Representing such multi-material structures with existing modeling software is challenging, partly

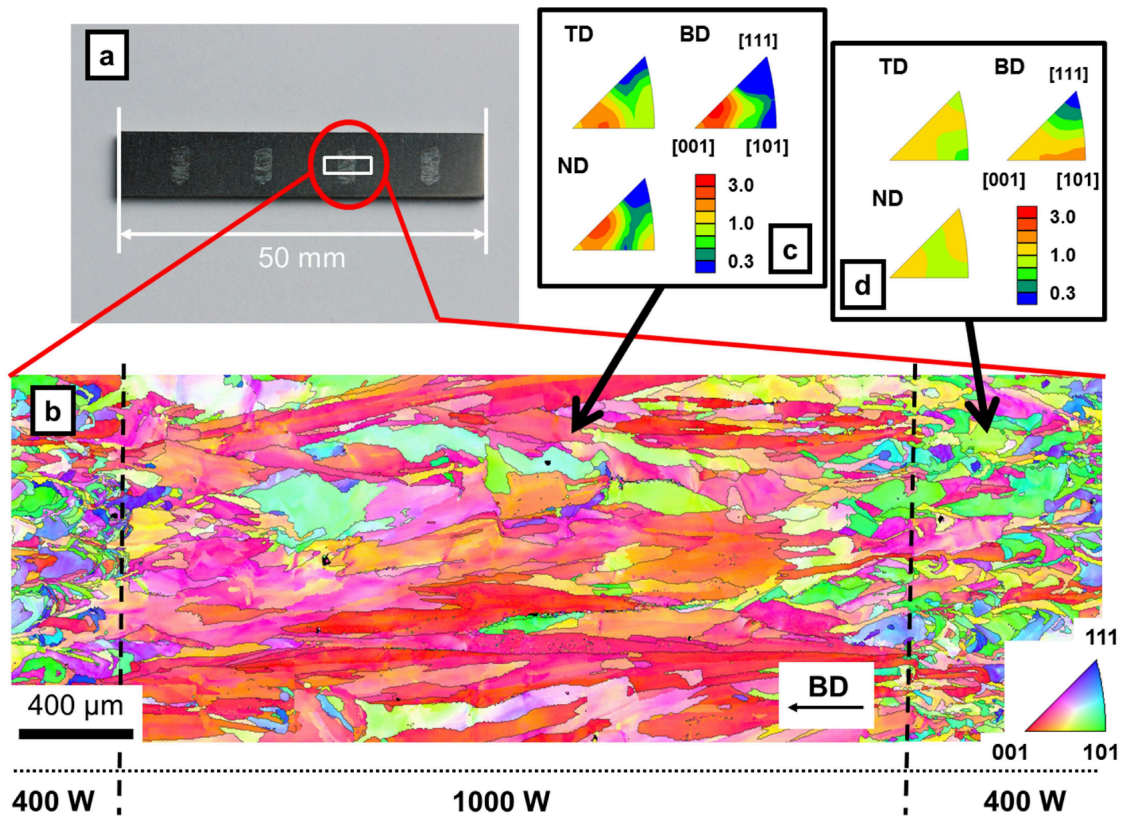


Figure 5: An example of microstructure control from [31], showing variations across scales resulting from different energy parameters.




Figure 6: A specific example from [20] showing how the energy level of a laser source (right-to-left: 115W, 165W, 200W, 275W, 345W and 410W) drastically changes the resulting crystalline structure of 316 stainless steel powders.




Figure 7: An illustration of the microstructure interface between different materials in a multi-material object from [51]. In this example the grain structure alters and blends at the materials interface layer depending on the specific parameters of each material as well as the parameters of the energy source. In this case it is a cross-sectional image of a 7YSZ/Sr₂Zr₂O₇ bi-layer sample containing a Pt layer in the Sm₂Zr₂O₇ part of a coating coating.

due to assumptions of homogeneity, partly due to the order of magnitude increase in complexity, and partly because the precise model depends on the physical process which can be quite complex, especially at the transition layers between materials. For example, predicting shapes and structures at the interface of materials may be challenging, due to the dynamic nature of the melt regions created by the energy source. The example in Figure 7 shows such interfaces for multi-material objects manufactured by additive processes.

3.3 Material Architecture

Material architecture refers to the systematic process of achieving new effective material properties through geometric sub-structuring [10, 24, 32]. Additive processes, especially those with directed laser or electron beams as energy sources, are capable of depositing materials based on different architectures that are tailored to be lightweight and mechanically efficient, thus producing material structures that are optimized for superior performance. In additive manufacturing, these architectures are typically deployed at the mm, sub-mm or even μ level, enabling the additive processes to produce artifacts with vastly different mass and performance parameters from those of parts created from uniform blocks of metal. Examples of such architectures include bio-mimetic architectures (i.e., similar to bone, honeycomb, etc.) as well as those based on more traditional engineering principles (e.g., lattices, grids, trusses, etc.). More generally, note that this systematic substructuring of materials is a well known modeling principle, observed in nature and practiced in engineering, that has been applied at all scales and is the basis for the notion of hierarchical materials [28, 11]. Common examples of such structures are shown in Figure 8 and include wood (scales range from micrometers to centimeters), bones, composite materials, and the Eiffel tower (4 scales) [11].

Many bio-mimetic structures may be naturally stochastic or generative in nature. This means that such structures may be described implicitly by their intended properties, behaviors, or the production processes—and not explicitly by the precise and specific configuration of the final part geometry [33]. While creating an explicit geometric representation of such structures is conceivable and may be necessary for ultimate manufacturing, developing representation schemes and algorithms that can handle the complexity of such models and provide the desired fidelity poses non-trivial research challenges.

Symmetric periodic material architectures, such as those used regular lattices, are more straightforward to represent procedurally or algorithmically. The challenge for these structures lies in the relationship of the




Figure 8: Examples of material architecture: (a) Trabecular micro-architecture (3DS); (b) lattice micro-architecture. (c) trusses-of-trusses at nano-scale; (d) Eiffel tower. Material structures are evaluated in terms of effective (or homogenized) material properties—this is what we usually mean when we say “material properties”—an equivalence class of geometric homogeneous structures with similar averaged material properties.

model to the manufactured product. With present technology, the variations in the manufacturing process at the level of resolution of the model produce (possibly) significant deviations from the idealized geometry. For example, an idealized truss or lattice structure may be specified as a configuration of cylindrical entities. The intersections between the cylindrical boundaries of these entities are represented precisely in CAD systems by smooth curves. In a discretized model, they would be approximated by chains of axis-aligned edges. In the fabricated realization, the members will not be perfect cylinders, but very bumpy versions of an approximating voxel model. The surface roughness of the members will be a function of the grain size in the stock powder and the focal width (and parameters) of the directed energy source, or perhaps a droplet size in a liquid deposition process.

4 Views in Design and Manufacturing

4.1 Four views of an artifact

To effectively support the emerging technologies for manufacturing and material engineering, models of material structures must integrate geometry, effective material properties, process physics, and multi-materials, supporting a variety of material architectures at different levels of scales and fidelity. What exactly constitute a proper mathematical model of a material structure is less clear, because it appears to require very different combinations of geometry, material, physics, and process information at distinct stages of the product life cycle. We present the emerging representational challenges using four different views (Figure 9) of a *Real* artifact: *Functional*, *Designed*, *Planned*, and *Simulated*.

In the spirit of classical solid modeling, each view can be considered a pair (abstraction, representation), where the mathematical abstraction idealizes a class of corresponding physical objects, and the computer representation implements this abstraction on a computer when it is possible. We will not discuss specific abstractions and representations in this article. Rather, we will focus on the essential attributes of the views and use them to discuss the various design, planning, and verification activities within the product life cycle.

Functional. The *Functional* view (F) captures the designer’s intent, or functional specification of artifact’s behavior, usually at a high level. Typically, it describes a set of desired characteristics of the overall shape or of specific interfaces, such as precise surface features that must be present to support contacts in an assembly or a rough shape or topology. It may define a set of geometric or mechanical measures, such as thickness of




Figure 9: Four views used to discuss the design and manufacturing of a real artifact: the arrows represent relations between views and the real physical objects.

a beam, stiffness, or weight, and use them to specify desired (i.e., nominal) values, constraints, or objective functions that the designer is interested in optimizing. The *Functional* view may also capture a description of the desired responsive behavior, by which we mean a representation of the shape deformations and of the changes in mechanical properties that are parameterized by a higher-dimensional model that defines, for example, a choice of acceptable external forces, torques, or pressures and the behavior of the material when these are applied. As a model of material structure, the *Functional* view is incomplete, meaning that it does not completely define a unique artifact, nor provide guidelines for generating one.

Designed. The *Designed* view (D) captures the idealized details of a particular realization of the *Functional* view. These details include a mathematically precise representation of: (1) the overall nominal shape (bounding surface separating the exterior from the interior), (2) the geometric structure that segments the interior of the artifact, and (3) the choice of materials and possibly of their blends and gradations for each chunk (i.e., strut, grain, pore...) of that segmentation. This view represents the material structure at multiple scales and may be deemed complete, to the extent that it contains information that is necessary to determine whether the model represented in the *Designed* view behaves according to the specifications in the *Func-*

tional model. For example, the *Designed* model supports detailed local probing queries, such as: “What is the material mix in the local vicinity of a given point P at a particular scale?”, more abstract queries, such as “What is the thickness of this gorge?”, and integral queries, such as “What is the porosity or what are the aggregate mechanical characteristics of a given cell (portion of the interior)?” Assuming that the *Designed* representation describes an embodiment of the *Functional* model, it must also include allowable variations in form and material properties of the represented material structure, similar to the Geometric Dimensioning and Tolerancing (GD&T) specifications for traditional homogeneous solids [1].

The *Functional* and *Designed* views are deliberately defined in a manner that makes them independent of a particular manufacturing technology and of its current limitations, allowing to choose most efficient and economical ways of manufacturing an artifact. On the other hand, they allow designers and material scientists to design models that may not be manufacturable with existing or even anticipated technology. This option is important both to allow scientists to capture precise models of material structures found in nature and to allow designers to create visionary designs that will challenge material scientists and the additive manufacturing industry. However, practical applications, where the immediate goal is to produce physical artifacts, must translate such idealized representations into manufacturable process plans and take into account the current limits and side effects of the chosen manufacturing technology. Hence, we define below two additional views to represent the manufacturing aspects of material structures.

Planned. The *Planned* view (P) captures the manufacturing process plan, instructions and data that will be used to manufacture the *Designed* model. For example, in the case of 3D printing, these instructions and data may define the *Planned* material structure implicitly by specifying manufacturing steps, choices of materials to be used, G-code³, parameter settings (speed, feed, energy levels, etc.), or more explicitly by providing boundary models (for example, STL files) or slices through the model, or possible voxelized representation as a two- or three-dimensional array of material indices, each identifying which material should be deposited at the corresponding voxel. In practice, many manufacturing processes require a combination of implicit and explicit representations. For instance, in composite manufacturing, the *Planned* view may include models of the tooling surfaces as well as specific instructions for how individual material layers are deposited on that surfaces. Achievable manufacturing precision, uncertainties, and process variations apply to any such *Planned* view. In this sense, the *Planned* view represents a class of objects that are manufacturable by the specific manufacturing process plan.

Simulated. The *Simulated* view (S) attempts to predict a model of the artifact “as manufactured” that is obtained by simulating the process plan. The *Simulated* model is always different from the nominal *Designed* model from which the *Planned* model was derived. The differences come from several sources. All shape forming manufacturing processes have intrinsically limited geometric accuracy, typically ranging from millimeters to microns. For example, CNC machining can produce surface finishes with a guarantee accuracy of about one thousandth of a millimeter, high end 3D printers may achieve resolution of 10 microns, while consumer-oriented inexpensive 3D printers can only guarantee about one tenth of a millimeter. But in case of additive manufacturing, we already identified additional sources of imprecision related to process physics and material properties. Specifically, process phase changes and multiple material make prediction of local geometry material properties difficult. The spatial trajectory of material deposition and/or solidification has direct bearing on directional properties of manufactured materials as well as global shape

³G-code is a numerical control programming language that is commonly used to specify motion of a machine tool, such as cutting tool or a 3D printer head.




Figure 10: An example of a multi-material architectural structure from Lawrence Livermore National Laboratories. Today’s idealized **design** representations useful for analysis cannot account for the deviations in the **real** structure which would be needed to define acceptable levels of process variation.

properties, e.g., warping as a direct consequence of stresses caused by different rates of cooling and solidification. The example from Lawrence Livermore National Laboratories in Figure 10 shows the variation for a micro-stereolithography process in the manufacture of a multi-material truss system.

The design views (*Functional* and *Designed*) and manufacturing views (*Planned* and *Simulated*) represent idealizations of the physical artifact at different stages of the design-manufacturing cycle. These idealizations are connected to the real physical world that is abstracted by the *Real* view in Figure 9.

Real. The *Real* (R) is a collection of physical parts that represent a statistically valid sampling of what is produced by executing the process plan P on a particular manufacturing system. Typically, these parts are not congruent because of the variability of the manufacturing process. In this sense, R is not based on a computer representation of a model of the artifact or of its properties. The *Real* view is always distinct from the *Simulated* view because the latter cannot account for all possible physical phenomena and process parameters.

4.2 Engineering activities as relations between views

Most of the engineering activities in design and manufacturing of material structures may be formally framed and understood in terms of binary relations between the four views and a set R of physical parts introduced above. These relations and activities are depicted as arrows in Figure 9; hence, each arrow represents multiple relations. Below we briefly discuss such relations in two broad categories: *conversions* and *harmonization activities*. We note that the same four views are also relevant in traditional design and manufacturing of rigid homogeneous components, where the same mathematical (solid) model is commonly assumed in *all* views. This assumption allows us to cast the relations between the views in terms of geometric *features* [41]. This approach does not apply in modeling of material structures, for which every view requires a different mathematical model.

Conversions are relations in Figure 9 that represent automated or manual activities that generate a view of an artifact given a view in another category. The horizontal arrow, $F \rightarrow D$, from F to D represents the design synthesis activity, which starts with an initial design of the overall shape, of the internal structure,

and of the distribution and variation of the material used throughout the structure. This design activity is informed by best practices in the application domain, by the capabilities and limitations of the design system, and sometimes by the limitations of the targeted manufacturing technology. Hence, $F \rightarrow D$ represents a conversion process that takes an *F* view and, with the help of a designer and of the available design and automation tools, produces a *D* view.

The vertical arrow, $D \rightarrow P$, represents manufacturing process planning, which is a conversion process that may be fully automatic or may involve human interventions in the form of high-level choices or detailed editing or optimizations. Given a complete specification of the shape and material properties of a material structure, the mapping produces a manufacturing process plan, be it a sequence of NC machined features, G-code, or the slices and scan lines that specify the complete paths and process parameters for 3D printing.

The horizontal, right-to-left arrow, $P \rightarrow S$, represents manufacturing simulation, which is an automatic conversion process that simulates the execution of the process plan and constructs a complete representation of the resulting shape, structure, and material distribution. Given a manufacturing process plan in the *Planned* view, it is evaluated to produce a *Simulated* “as manufactured” representation of shape and material properties. This task may be achieved either by simulation or extensive experimental data that characterize the map $P \rightarrow S$. The latter is sometimes referred to as “process certification”. Certification of processes for use in high-reliability applications, such as aerospace, is currently an area of intense research in the materials science and manufacturing community.

Finally, the thicker purple arrow, $P \Rightarrow R$, represents the actual (physical) result of manufacturing a set of sample artifacts. Together, the four conversions $F \rightarrow D$, $D \rightarrow P$, $P \rightarrow S$, and $P \Rightarrow R$, correspond to the traditional design-manufacturing workflow, where functional requirements *F* are used to design an artifact *D*; *D* is used as a target for manufacturing process plan *P*; and *P* is used to produce a physical part *R*, usually after simulating *P* as a verification step.

The conversions $F \rightarrow D$ and $D \rightarrow P$ involve synthesis and/or planning procedures which formally correspond to set-valued one-to-many maps that range over all feasible designs or manufacturing plans.⁴ Compositions of conversions (when such are defined) often describe alternative workflows. For example, the composite conversion $F \rightarrow P = (D \rightarrow P) \circ (F \rightarrow D) = (F \rightarrow D) \rightarrow P$ is a formalization of (manufacturing) feature-based design, where the *Functional* specification of an artifact is translated directly into manufacturing operations that are abstracted as manufacturing features in the *Planned* representation. Similarly, $F \rightarrow R$ is a composition of three conversions $((F \rightarrow D) \rightarrow P) \Rightarrow R$ that is suggestive of a more futuristic but plausible scenario where the *Functional* representation may be streamed directly to a computationally derived manufacturing process to produce the *Real* shape and material properties.

Harmonization of views The open-loop workflow based on composition of conversions is conceptually simple and appealing; however, it is also unrealistic for at least two reasons. Frequent changes in functional requirements, materials, manufacturing capabilities, and computational tools increase the likelihood of inconsistencies between different views at any given time. Furthermore, all conversions involve heuristics, uncertainty, and approximations. It is virtually impossible to guarantee their correctness, and the results of any such conversion must be validated by comparing them against other views. As the result, more realistic workflows require multiple iterations of comparisons, conversions, modification, and optimization of the views to make them consistent with each other. Collectively, we refer to such processes as *harmonization*. A digital thread, such as that illustrated in Figure 3, is essentially a harmonization process that is “flattened” to show primitive binary relations between involved views. Once again, common binary relations correspond to the arrows in Figure 9, as briefly summarized below.

⁴Strictly speaking, such one-to-many relations are not maps.

The horizontal, red arrow, $F \leftrightarrow D$, represents the validation of D against the constraints of F . It is typically used in an iterative design editing process that adjusts D to ensure that it conforms to the specifications of F . In a typical design cycle, this editing process may involve human decisions and also automated shape/structure/material optimization processes. The red arrow, $D \leftrightarrow S$, represents a harmonization of the *Designed* and *Simulated* views. If the results of this comparison do not meet the prescribed shape and material tolerances, they may indicate how to edit D and then recompute P and S automatically or to tweak P manually. The red arrow, $S \leftrightarrow F$, represents harmonization of the *Functional* and *Simulated* views which takes place in order to validate the design and process plan. Such a validation typically requires analysis and simulation tools. The arrow $P \leftrightarrow D$ is an abstraction for manufacturability analysis, which includes a procedure to determine if a given *Planned* process plan is capable of producing the *Designed* material structure. The latter could be realized as a composite relation by process simulation ($P \rightarrow S$) followed by a comparison $S \leftrightarrow D$. Finally, the arrow $F \leftrightarrow P$ signifies reuse of previously designed and manufactured artifacts P to meet new or modified functional specification F .

Ultimately, all engineering activities must be related to and validated by the physically manufactured *Real* artifacts R . A relation $R \Rightarrow X$ compares a set of measures defined in view X to a statistical representation of the same measures of artifacts in R . The measures may be limited to statistics on specific sets of numbers that represent measurements of the shape geometric, topological, and integral characteristics, locations, sizes, or form factors of features, surface finish and material characteristics, as well as local or global mechanical properties. They provides a critical reality check that allows validation and comparison of other views used in design and manufacturing. For example, $R \Rightarrow F$ is an abstraction of the manufactured product validation to ascertain that *Real* material structures perform as specified by the *Functional* view; $R \Rightarrow D$ corresponds to product inspection and verification, where the *Real* parts are inspected for conformance to shape and material information specified by the *Designed* view; $R \Rightarrow P$ could represent an instance of reverse engineering task, where the *Real* representation of the manufactured artifact is analyzed to determine the *Planned* model of how it was manufactured; it may also be implemented as a composite relation $(R \Rightarrow D) \rightarrow P$; and the arrow $R \Rightarrow S$ corresponds to a comparison between the *Real* as-manufactured components and the predicted *Simulated* models. The latter may also be used by vendors of manufacturing technologies to validate their manufacturing simulator.

4.3 Systematic formulation of computational challenges

The usefulness of the diagram in Figure 9 should now be apparent: it draws direct parallels between the traditional and emerging computational problems in design and manufacturing, while at the same time, it allows us to pinpoint precise challenges and opportunities in developing solutions for modeling of material structures.

For instance, it could be argued that the Oyster system [26] developed in early 80s at IBM Research was a successful example of implementation of a $P \rightarrow S$ conversion where a *Simulated* model of the manufactured VLSI assembly was derived algorithmically from a *Planned* model of the VLSI manufacturing process plan. The resulting *Simulated* representation of an assembly of polyhedral solids was sufficient to support the physical simulation and analysis for manufactured product validation, and hence played the role of a *Designed* view, effectively making the the a $S \leftrightarrow D$ verification unnecessary.

But executing a similar scenario in the context of modern additive manufacturing is difficult or impossible, partly because the $P \rightarrow S$ conversion is not fully understood, and partly because the $S \leftrightarrow D$ comparison or the $S \leftrightarrow F$ validation must deal with representations of geometry and materials that are orders of magnitude more complex than their models used in traditional manufacturing. This explains why a more pragmatic (but still challenging) approach adapted, for now, by the industry is to certify the manufacturing process itself [7]

and to report its variability in surface accuracy and smoothness and in the local mechanical characteristics of the physical part produced. This, in essence, characterizes the properties of the $P \rightarrow S$ conversion itself using a benchmark of *Planned* models. Note that even if this approach is successful, one still have to rely on the $P \Rightarrow R$ relation and $R \Rightarrow F$ validation to verify the integrity of the design of a particular artifact.

More generally, the paths in the four-view diagram describe possible workflows in a manufacturing enterprise and may be used to systematically analyze and solve new problems. For example, in most traditional CAD scenarios, an expert designer has the charge of designing the *Designed* view, based on his knowledge of *Functional* goals and of manufacturing constraints. This requires effective (often application specific) design tools. The design process will typically involve trial-and-error, where the *Designed* model is tested and tweaked until its *Functional* behavior and mechanical characteristics meet the requirements, or until some objective function has reached an acceptable value, while satisfying the manufacturing constraints. In contrast, the casual or novice users from the community of ‘makers’, who are perhaps interested in designing quickly a rough shape that performs a specific function (say, flexible and yet strong structure), are typically not willing, or may not have the expertise, to design the *Designed* or *Planned* model. Instead, they would like to specify the *Functional* model only and have it automatically compiled into the *Planned* representation via the composition of automated $F \rightarrow D$ and $D \rightarrow P$ conversion processes.

To support the latter scenario, we need to: (1) design intuitive and effective representation for the *Functional* model, including direct-manipulation tools for expressing the overall shapes, the constraints, the prescribed behaviors, and the tolerances on these; (2) develop digital representations of the *Designed* view that can be processed by synthesis algorithms, and (3) invent efficient synthesis algorithms that either produce a *Planned* view or that produce first a *Designed* view and then derive a *Planned* view from it.

A particularly popular approach to the challenge of automating the $F \rightarrow P$ conversion is to assume that both the *Designed* and the *Planned* models are represented by voxelized models, with a material index per voxel. Then the synthesis problem reduces to computing the values of the material indices for all the voxels, so that they meet the constraints and minimize the objective cost functions defined in the *Functional* model. Unfortunately, because of the exponential number of possible assignments, we cannot expect to obtain a good solution by trying random assignments of material indices and by selecting the one that minimizes the cost function while satisfying all the prescribed constraints. For example, a binary material assignment (metal or air) for a tiny $5 \times 5 \times 5$ model has 2^{125} possibilities, so it is impossible to try even a minute fractions of these. Yet, for high-resolution representations, we want to model resolutions reaching $1000 \times 1000 \times 1000$ or higher. Topology optimization [4] and various heuristics search strategies are being developed to guide the search of a good solution in this discrete, but huge (in fact infinite in practice) search space.

5 A research agenda

The main goal of this paper is to articulate the technical challenges and research opportunities brought about by the digitization of manufacturing and advances in material science. To that end, we proposed a new framework and paradigm that uses relations and mappings between the four distinct views of solid material structures for systematically identifying and formulating computational modeling problems that span a large portion of the product life cycle. It would be impossible to discuss all such problems in details here. However, it is only fitting to conclude the paper by brief discussion of some outstanding research problems that the authors consider particularly pressing.

5.1 New models and abstractions

Perhaps first and foremost, we must understand and formalize the models and computer representations for the four views of material structures. While we can find examples of such models and representations in every type of manufacturing, none of them seem to generalize to support the full spectrum of digital thread in additive manufacturing.

From the four, the *Functional* view is least understood. Many abstractions of function and behavior have been proposed, for example [13, 48, 45], but the formal semantics of such models remains unclear as is the nature of their mapping to *Designed* representations. SysML [12] is increasingly used to specify functional behavior of *Functional* complex engineering systems, but currently does not provide support for describing complex geometric information or physical behaviors. Furthermore, a designer must be able to specify, edit, and visualize the models represented in these views in high-level languages that may include analog gestures for specifying and editing shapes, patterns, behaviors and may benefit from multi-modal input environments, combining tracked gestures and voice commands and scanned input data. The *Functional* model must also define tolerances on shape, structure, and behavior that are sufficiently generous for a particular printing technologies, so that we will be able to check algorithmically whether a *Designed* or *Planned* model is guaranteed to yield a physical artifacts that meets the requirement specified in the *Functional* model. But how does one define a tolerance on a behavior that is not representable by a standard scalar or tensor measure of mechanical characterization?

A *Designed* view of a material structure may appear to be nothing more than a complex instance of a bonded multi-material assembly of solids and can be represented using classical modeling techniques [22]. The essential new element in material structures is the shape-material interaction that occurs across multiple scales. In the simplest case of a two-scale structure, at the finer scale, the *Designed* model is a pattern of material features (grains, fibers, layers, pores, etc.) that are arranged (semi-)periodically or stochastically over space or a scaffolding. Such patterns may be graded, exhibiting progressive variations in size, location, type and material properties of individual features, contributing to the rich variety of representable structures in nature and engineering. The physical behavior of such a material structure is determined by its effective (homogenized) material properties that are evaluated at a coarser scale. For example, structural properties of a composite laminate panel are determined by a careful arrangement of individual layers in the structure, and bulk material properties of cellular materials (foams, bones, wood) may be predicted from known arrangement of individual cells in the structure [15, 24]. Thus, it should be clear that the two levels of abstraction in such a material structure interact through associated models of physical behavior, which are also the key to establishing the conversion $F \rightarrow D$ and the validation $F \leftrightarrow D$ between the *Functional* and *Designed* representations. In situations where a very specific internal structure must be engineered or cannot be generated by the synthesis algorithms discussed above, the expert designer has the charge of designing the *Designed* model. Hence, we need to provide effective tools to support this delicate design task, which perhaps can be better described as “architecting the artifact.” The process will typically involve trial and error, where a design is simulated and tweaked until its behavior and mechanical characteristics meet the requirements and until the objective function has reached an acceptable value. Lastly, to dispel the illusion of similarity to classical solid modeling, we observe that most material structures involve more than two scales, leading to notion of hierarchical material structures [32].

Significant research progress is also required in understanding of *Planned* view. In stark contrast to traditional unit manufacturing processes, it appears that the parameter space spanned by the additive manufacturing process is an essential part of the design of the material as well as the gross shape of the artifact. This material information is a form of temporal data describing the speed and feed rate of each of the material spools, the intensity of the energy source (laser or electron beam) as well as any motion parameters of

the deposition pallet, part orientation and location on the pallet, or laser or materials feeds. This procedural specification of an artifact goes far beyond just saving the equivalent of the CNC code for a machined part. In this case, *the process parameter space actually forms the part specification for its internal structure*. Hence, producing another part of similar performance would require re-playing this process. But the information required to fully capture an artifact might need to include this complete set of data, recorded at some as yet not specified minimum required level of resolution, and could lead to explosion in size of process data. How much of this data needs to be kept, how to best compress or filter it, and what the appropriate archive formats should be are open questions that affect matters such as lifecycle maintenance, certification and qualification.

The *Simulated* representation inherits all multi-scale challenges of the *Designed* representation, with added complexity of geometric and material uncertainties caused by physical processes underlying manufacturing operations in *Planned* model. Clearly, every part will be different. While this is certainly also the case with traditional manufacturing processes (such as machining metal), what is meant by uniformity in the manufacturing process is not as yet scientifically clear. When machining out of a block of metal, the metallurgic properties of the base material are used to extrapolate performance properties across the final part. In the case of an additive process with complex internal microstructure, the manner of such extrapolation is less clear. To account for variation (possibly significant) across the microstructure, one could build the appropriate statistical models. However, these models may need to be parameterized based on data from the actual artifact. For example, it is possible to use in-process monitoring (image analysis) to assess model porosity. Presumably, a test can be conceived where, if the porosity deviates too significantly from a statistical model that predicts positive performance, the part can be scrapped or additional tests ordered. From a design and modeling perspective, how detailed of a model of the design specification does one need to have? From a manufacturing process, inspection or certification perspective, what needs to be captured about the artifact resulting from the production process?

5.2 Efficient conversions and harmonization

With formal models of the four views in place, representations and computational properties of the view relations become a key issue. While the computational properties of the mapping $F \rightarrow D$ is perhaps the most poorly understood issue, definite characterization of *Simulated* shape and material properties produced by specific manufacturing processes in *Planned* is required to support rational manufacturing process planning of material structures from their *Designed* representations. Otherwise, process planning for additive manufacturing is likely to remain a costly trial-and-error process. One possible approach is to create libraries or databases that enable engineers to correlate manufacturing properties with microstructure properties and resulting effective material performance. Another approach may be to evolve towards standards, akin to the codes on an ALU (Arithmetic Logic Unit) or the graphic primitives on the GPU (Graphics Processing Unit), against which people design structure primitives and specify their properties with tolerances. Lack of such data regarding the performance properties of microstructures makes development of automated topological optimization and synthesis codes difficult; and, even with such datasets, the configuration space is combinatorially vast. At present, modeling these microstructures and associated phenomena is typically the domain of high-performance computing (HPC) and supercomputer codes to achieve physics-based simulation at the micro-scale and below. Such representations are typically dynamic in nature, with a large reliance on initial conditions and experimental parameters to define the material.

The above discussion suggests that relatively complete *Designed* and *Simulated* representations of material structures may require dramatic increase in required computational resources. The data and information required to fully capture an artifact might need to include this complete set of geometry, material, and pro-

cess data. This also suggests that both *Designed* and *Simulated* models should be represented implicitly and evaluated on demand. For example, an implicitly defined *Designed* representation may include only partial information about a material structure (perhaps for some carefully chosen locations, features, and patterns), whereas the rest of the structure can be generated on demand in order to satisfy the constraints indicated by the corresponding *Functional* model. Similarly, *Simulated* representation of the artifact may be generated on demand, perhaps in form of machine instructions, evaluating operations from *Planned* model. Thus, efficient generation and streaming of *Designed* and *Simulated* representations may become a critical issue. The explosion in data becomes particularly dramatic in situations when the process data becomes an integral part of *Simulated* representation; this could conceptually thought of as a discrete, time-sampled vector of the process parameters. If recorded in discrete form, noting that these processes have beam speed rates measured in 1000s of meters per second, this could require hundreds or thousands of vectors per second. Consider an example where we assume the vector of process parameters takes up 1000 bytes and process is sampled at 100 times per second. For a one hour build, this results in 360MB of data about the manufacturing process; and for more complex artifacts one could expect multiple terabytes of data for defining a single artifact. “How much of this needs to be kept?”, “How to best compress or filter it?”, and “What the appropriate archive formats should be?” are largely open questions.

5.3 Standardization and interoperability

Based on our experience with traditional solid modeling, issues related to standardization and interoperability of models, representations, and systems are likely to emerge as more critical bottlenecks than the modeling technology itself. The importance of interoperability in modeling of material structures was also highlighted in panel discussions at recent manufacturing conferences [44]. Broadly, interoperability refers to ability to convert, harmonize, and combine distinct computer representations of material structures between two different views.

Theoretically, standardization and interoperability of different representations of material structures will require a broad agreement on a common formal reference model of what such structures are mathematically. Such an agreement may be difficult or impossible to achieve, given the richness and complexity of the four representations discussed above. Geometric models and representations are relatively well understood in the context of solid modeling [42, 19], and recent proposals for material models include combinatorial models extending geometric complexes to cochains complexes [8], point set models extending manifolds to fiber bundles [27], as well as more pragmatic approaches extending implicit models to material property fields [5], and replacing commonly used STL format by a more general AMF format [23] that supports curved geometry and some material information. However, the issue of what constitutes a suitable formal model of materials, physical behavior, processes, scale and uncertainly remains largely open.

The problems of standardization and interoperability have been extensively studied in design and manufacturing of homogeneous components where all views are based on the common mathematical model of a rigid solid. In this case, the same mathematical model is used in all views, and the interoperability problems become a somewhat simpler (but still challenging) problems of (solid) representation conversions. In practical terms, there are two approaches to standardization and interoperability [18]. A file-centric approach depends on the ability to translate and convert various representations to a common standards (such as STEP) and/or to each other. This approach is widely accepted by traditional CAD industry, but has proved costly, error prone and ineffective in solving interchangeability, interoperability, and integration challenges in digitally-driven manufacturing enterprise. An alternative approach advocated in [18] is to standardize on queries: computable mathematical functions that encapsulate and abstract details of individual representations and allow systematic development of interoperable solutions. For example, the query-based approach

allowed the development of a structural simulation system that is fully interoperable with any solid modeling representation that supports point membership and distance queries.

The query-based approach to interoperability is a natural extension of the concept of interchangeability of mechanical parts in assembly [21]. A key ingredient of both approaches is a notion of equivalence of parts (and queries) with respect to a common mathematical reference. Thus, all equivalent mechanical parts belong to the same classes of toleranced congruent parts as determined by a finite set of measurements; similarly, equivalence of solid models may be defined in terms of tolerances on their geometry and topology as determined by a finite set of computational queries. These observations suggest a promising approach to dealing with more general problems of conversions, harmonization, and integration of material structure views: common mathematical constructs (properties, functions, measures) may be used to define a notion of equivalence between distinct views and representations. This equivalence can then be used to support conversion, validations, or harmonization between the views using view-specific computable queries that implement the specified construct. This approach allows us to deal with the interoperability challenge in an incremental, extensible, and object-oriented fashion.

We conclude by noting that the principles of interchangeable parts in mechanical assemblies has evolved for more than a century, transforming artisan crafts into the modern system of mass production. Similarly, we expect that the research agenda proposed in this paper aims to transform additive manufacturing from its current “artisan state” to its rightful leading position fueling the digital revolution in manufacturing.

6 Acknowledgments

Vadim Shapiro’s research was supported by NSF grant CMMI-1361862 and National Institute of Standards and Technology. Responsibility for errors and omission lies solely with the authors. William Regli’s work supported by NSF grants CMMI-1361862 and OCI-0940841. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF, DARPA or the other supporting government and corporate organizations.

References

- [1] ASME. Y.14.5-2009. *Dimensioning and Tolerancing*, 2009.
- [2] Solid Modeling Association. www.solidmodeling.org.
- [3] ASTM. *ASTM Standard F2792-12a Standard Terminology for Additive Manufacturing Technologies*. ASTM International, West Conshohocken, Pennsylvania, 2012.
- [4] Martin Philip Bendsoe and Ole Sigmund. *Topology optimization: theory, methods and applications*. Springer, 2003.
- [5] Arpan Biswas, Vadim Shapiro, and Igor Tsukanov. Heterogeneous material modeling with distance fields. *Computer Aided Geometric Design*, 21(3):215–242, 2004.
- [6] Peter Jeffrey Booker. *A history of engineering drawing*. 1963.
- [7] David L Bourell, David W Rosen, and Ming C Leu. The roadmap for additive manufacturing and its impact. *3D Printing and Additive Manufacturing*, 1(1):6–9, 2014.

- [8] Antonio DiCarlo, Franco Milicchio, Alberto Paoluzzi, and Vadim Shapiro. Chain-based representations for solid and physical modeling. *Automation Science and Engineering, IEEE Transactions on*, 6(3):454–467, 2009.
- [9] Economist. The third industrial revolution. *April 21 issue on the Third Industrial Revolution*, 2012.
- [10] Peter Fratzl, Dunlop, John WC, and Richard Weinkamer. *Materials Design Inspired by Nature: Function Through Inner Architecture*. Royal Society of Chemistry, 2013.
- [11] Peter Fratzl and Richard Weinkamer. Natures hierarchical materials. *Progress in Materials Science*, 52(8):1263–1334, 2007.
- [12] Sanford Friedenthal, Alan Moore, and Rick Steiner. *A practical guide to SysML: the systems modeling language*. Elsevier, 2011.
- [13] John S Gero and Udo Kannengiesser. The situated function–behaviour–structure framework. *Design studies*, 25(4):373–391, 2004.
- [14] Ian Gibson, David W Rosen, and Brent Stucker. *Additive manufacturing technologies*. Springer, 2010.
- [15] Lorna J Gibson and Michael F Ashby. *Cellular solids: structure and properties*. Cambridge University Press, 1999.
- [16] Lothar Haselberger. The construction plans for the temple of apollo at didyma. *Scientific American*, 253(6):114–122, 1985.
- [17] C. M. Hoffmann. *Geometric and Solid Modeling*. Morgan Kaufmann, San Mateo, Cal., 1989.
- [18] Christoph Hoffmann, Vadim Shapiro, and Vijay Srinivasan. Geometric interoperability via queries. *Computer-Aided Design*, 46:148–159, 2014.
- [19] Christoph M Hoffmann and Jarek R Rossignac. A road map to solid modeling. *Visualization and Computer Graphics, IEEE Transactions on*, 2(1):3–10, 1996.
- [20] William Hofmeister, Melissa Wert, John Smugeresky, Joel A Philliber, Michelle Griffith, and Mark Ensz. Investigation of solidification in the laser engineered net shaping (lenstm) process. *JOM-e, online component of JOM*, 51(7):1–6, 1999.
- [21] David Hounshell. *From the American system to mass production, 1800-1932: The development of manufacturing technology in the United States*. JHU Press, 1985.
- [22] Keith A. Hunten, Allison Barnard Feeney, and Vijay Srinivasan. Recent advances in sharing standardized STEP composite structure design and manufacturing information. *Computer-Aided Design*, 45(10):1215 – 1221, 2013.
- [23] ISO/ASTM. *Standard Specification for Additive Manufacturing File Format (AMF) Version 1.1*. ASTM International, West Conshohocken, Pennsylvania, 2013 E.
- [24] David RH Jones and Michael F Ashby. *Engineering Materials 2: An Introduction to Microstructures, Processing and Design*. Butterworth-Heinemann, 2005.

- [25] California Institute of Technology Kochmann Research Group. Using L^AT_EX for your thesis, 2011. http://www.kochmann.caltech.edu/research_multiscale.html.
- [26] George M Koppelman and Michael A Wesley. Oyster: A study of integrated circuits as three-dimensional structures. *IBM Journal of Research and Development*, 27(2):149–163, 1983.
- [27] Vinod Kumar, D Burns, Debasish Dutta, and C Hoffmann. A framework for object modeling. *Computer-Aided Design*, 31(9):541–556, 1999.
- [28] Roderic Lakes. Materials with structural hierarchy. *Nature*, 361(6412):511–515, 1993.
- [29] Manufacturing Studies Board. *Unit Manufacturing Processes: Issues and Opportunities in Research*. National Academies Press, 1995.
- [30] James Manyika, Jeff Sinclair, Richard Dobbs, Gernot Strube, Louis Rassey, Jan Mischke, Jaana Remes, Charles Roxburgh, Katy George, David OHalloran, and Sreenivas Ramaswamy. Manufacturing the future: The next era of global growth and innovation. Technical report, McKinsey Global Institute, November 2012.
- [31] Thomas Niendorf, Stefan Leuders, Andre Riemer, Florian Brenne, Thomas Trster, Hans Albert Richard, and Dieter Schwarze. Functionally graded alloys obtained by additive manufacturing. *Advanced Engineering Materials*, pages n/a–n/a, 2014.
- [32] Gregory B. Olson. Designing a new material world. *Science*, 288:993–998, May 12 2000.
- [33] William C. Regli, Joseph B. Kopena, and Michael Grauer. On the long-term retention of geometry-centric digital engineering artifacts. *Computer-Aided Design*, 43(7):820 – 837, 2011. Special Issue from the 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling.
- [34] Aristides AG Requicha. Representations for rigid solids: Theory, methods and systems. *ACM Computing Surveys*, 12(4):437–464, December 1980.
- [35] Aristides AG Requicha and Jarek R Rossignac. Solid modeling and beyond. *IEEE Computer Graphics and Applications*, 12(5):31–44, 1992.
- [36] Aristides AG Requicha and Herbert B Voelcker. Solid modeling: a historical summary and contemporary assessment. *Computer Graphics and Applications, IEEE*, 2(2):9–24, 1982.
- [37] Aristides AG Requicha and Herbert B Voelcker. Solid modeling: Current status and research directions. *Computer Graphics and Applications, IEEE*, 3(7):25–37, 1983.
- [38] Jarek Rossignac. Structured topological complexes: A feature-based api for non-manifold topologies. In *Proceedings of the fourth ACM symposium on Solid modeling and applications*, pages 1–9. ACM, 1997.
- [39] Jarek R Rossignac. Through the cracks of the solid modeling milestone. In *From Object Modelling to Advanced Visual Communication*, pages 1–75. Springer, 1994.
- [40] Jarek R Rossignac and Michael A O’Connor. *SGC: A dimension-independent model for pointsets with internal structures and incomplete boundaries*. IBM TJ Watson Research Center, 1989.

- [41] Jami J Shah and Martti Mäntylä. *Parametric and feature-based CAD/CAM: concepts, techniques, and applications*. John Wiley & Sons, 1995.
- [42] V. Shapiro. Solid modeling. In G. Farin, J. Hoschek, and M.-S. Kim, editors, *Handbook of Computer Aided Geometric Design*, pages 473–518. Elsevier Science Publishers, 2002.
- [43] Vadim Shapiro. Representations of semi-algebraic sets in finite algebras generated by space decompositions. 1991.
- [44] Vadim Shapiro and Vijay Srinivasan. Report from a 2013 asme panel on geometric interoperability for advanced manufacturing. *Computer-Aided Design*, (47):A1–A2, 2014.
- [45] Rachuri Sudarsan, Steven J Fenves, Ram D Sriram, and Fujun Wang. A product information modeling framework for product lifecycle management. *Computer-aided design*, 37(13):1399–1411, 2005.
- [46] Ivan E Sutherland. Sketch pad a man-machine graphical communication system. In *Proceedings of the SHARE design automation workshop*, pages 6–329. ACM, 1964.
- [47] D. A. Swyt. Challenges to nist in dimensional metrology: the impact of tightening tolerances in the u.s. discrete-part manufacturing industry. Technical Report NIST-IR-4757, National Institute of Standards and Technology, Gaithersburg, MD, 20899-1020, January 1992.
- [48] Yasushi Umeda, Masaki Ishii, Masaharu Yoshioka, Yoshiki Shimomura, and Tetsuo Tomiyama. Supporting conceptual design based on the function-behavior-state modeler. *Artificial Intelligence for Engineering, Design, Analysis and Manufacturing*, 10(04):275–288, 1996.
- [49] Herbert B Voelcker. Modeling in the design process. *Design and Analysis of Integrated Manufacturing Systems*, pages 167–199, 1988.
- [50] Herbert B Voelcker and Aristides AG Requicha. Geometric modeling of mechanical parts and processes. *Computer*, 10(12):48–57, 1977.
- [51] Haydn Wadley. High temperature coatings: Thermal barrier coatings systems. Technical report, The University of Virginia, 2014. <http://www.virginia.edu/ms/research/wadley/high-temp.html>.
- [52] Terry Whollers. The whollers report. Technical report, The Wollers Group, 2015.