

Efficient Microstructural Design for Additive Manufacturing

Krishnan Suresh

suresh@engr.wisc.edu

Department of Mechanical Engineering
UW-Madison, Madison, Wisconsin 53706, USA

ABSTRACT*

A salient feature of additive manufacturing is that the cost of fabrication, to a large extent, is independent of geometric complexity. This opens new opportunities for custom-designing parts both at a macro and micro-level. An elegant and powerful method of designing custom-parts is through topology optimization.

While the theory of topology optimization is well understood, current methods can be extraordinarily expensive. The focus of this paper is on efficient *microstructural* topology optimization for 3d-printing. In particular, the computational bottle-necks in microstructural topology optimization are identified. Then, a framework that not only eliminates these bottle-necks, but incorporates other significant improvements, is developed. The framework is demonstrated through numerical experiments involving microstructures with millions of degrees of freedom, using multi-core CPUs and NVidia GPU.

1. INTRODUCTION

The most significant benefit of additive manufacturing is that geometric-complexity is 'free', i.e., to a large extent, it costs no more time or money to fabricate a very complicated part than it takes to fabricate a simple one. This opens new opportunities in product design; specifically, one can exploit modern topology optimization methods [1], [2] to translate problem-specifications into optimal, but perhaps complex, topologies, and then directly fabricate these through 3d-printing (see Figure 1). The optimality of the design in Figure 1 ensures that the least amount of material is used, translating into reduced fabrication time.




Figure 1: From problem specification to optimal part.

While Figure 1 exemplifies topology optimization at a macro-scale, one can go further in optimizing the topology at a micro-scale. Microstructural design [3]–[7] can be viewed as a special case of topology optimization [1], [2] where the objective is to

find the distribution of one or more materials that would result in a desired microscopic behavior. For example, Figure 2 illustrates the concept of optimizing both the macro-level geometry, and localized micro-level structure to meet a desired objective. Current 3d-printing processes can typically achieve a 100 micron resolution (or less), and thus the optimal multi-scale structure illustrated can potentially be fabricated.




Figure 2: Optimal microstructural design.

Through microstructural optimization, one can create surprisingly unique behavior at the macro-level; a classic example is that of negative Poisson-ratio bolt-design [8].

Microstructural design can be generalized to the distribution of multiple materials [9]. For example, consider the Lawrence Livermore National Lab (<https://manufacturing.llnl.gov/>) concept in Figure 3 of a limited thermal expansion body-armor; this concept is specifically designed to exploit the capabilities of 3d-printing.




Figure 3: LLNL's concept for a tailored micro-structure (<https://manufacturing.llnl.gov/>)

Unfortunately, an efficient infrastructure for computing such optimal microstructures is lacking. Current methods of microstructural optimization can be extraordinary expensive; for example, to quote a recent paper on microstructural design [7] “[for 3 million degrees of freedom] ... an optimization run with 1000 iterations takes approximately 10 h on 120 CPU cores”. This can be detrimental to the integration of microstructural design into mainstream engineering and 3d-printing-driven applications.

The main focus of this paper is efficient microstructural design to support current and future needs of 3d-printing. Since microstructure design fundamentally rests on the concept of homogenization, current homogenization methods are reviewed in Section 2, followed by a review of current methods for microstructural optimization. The reasons for the high computational cost are also identified in this Section. Then, in Section 3, a simple but efficient topological sensitivity [10] based level-set method is proposed. In Section 4, numerical results are presented, followed by conclusions and open issues in Section 5.

2. LITERATURE REVIEW

2.1 Homogenization

In microstructural design, a common hypothesis is that the microstructure is locally periodic. For example, Figure 5a illustrates a locally periodic microstructure with identical unit cells repeated spatially. Two unit cells, a.k.a. *Representative Volume Elements (RVEs)*, are illustrated in Figure 5b.




Figure 4: (a) Periodic microstructure, (b) two unit cells.

Given a unit cell, the ‘forward’ problem is to find its homogenized elasticity tensor \mathbb{C} [11]. The theory of homogenization rests on five fundamental principles: (1) strain-averaging, (2) stress-averaging, (3) displacement-assumptions consistent with strain-averaging, (4) micro-structural equilibrium, and (5) Hill-Mandel principle of energy consistency. These principles are discussed in detail, for example, in [12], [13].

The theory of homogenization directly leads to computational strategies. For example, a popular strategy for computing \mathbb{C} (associated with a unit-cell) is to impose six-independent periodic boundary conditions, and solve the resulting six finite element problems [14]. The tensor \mathbb{C} is then computed by averaging the resulting stresses. An alternate strategy is discussed later in the paper.

2.2 Microstructural Design

Microstructural design, on the other hand, is an ‘inverse’ problem where the challenge lies in finding an optimal distribution of voids that would result in a desired tensor \mathbb{C} .

The pioneering works of several authors [8], [15]–[17] demonstrated that it is possible to design microstructures with desired properties, such as materials with negative Poisson ratio. The first systematic approach to designing microstructures was proposed by Sigmund [3], where microstructural topology was

captured through discrete structural elements, and then optimized.

The modern approach towards microstructural optimization is to combine continuum finite element analysis, and a topology optimization technique such as Solid Isotropic Material with Penalization (SIMP) or level-set.

In SIMP-based microstructural design, pseudo-density parameters are assigned to finite-elements, and these parameters are directly linked to material properties, and optimized [7], [18]. The primary advantages of SIMP are that the theoretical derivation is simple, and the implementation is straightforward. Indeed, most of the recent advances on microstructural design are based on SIMP [5], [19].

However, in SIMP-based topology optimization [11], the sensitivity is less accurate since it is computed by differentiating the discretized global stiffness matrix with respect to the density parameters. The more serious challenge is that the finite element stiffness matrices become ill-conditioned in the presence of low pseudo-density values [20], [21]. As noted by several authors [7], [22], [21], since direct solvers are impractical for large scale problems, one must rely on iterative solvers, that exhibit poor convergence for such ill-conditioned problems.

We therefore pursue an alternate optimization strategy based on the concept of topological sensitivity. Topological sensitivity captures the first order impact of topological changes within a domain on various quantities of interest. This concept has its roots in the influential paper by Eschenauer [23], and has later been extended and explored by numerous authors [10], [24]–[27] including generalization to arbitrary features [28]–[30].

Topological sensitivity can be computed analytically as opposed to a semi-discrete formulation used in SIMP. In particular, for micro-structural design, the authors of [13] developed closed-form expressions for topological sensitivity of the tensor \mathbb{C} . This was exploited in [31] for 2D microstructural design, and in [4] for 3D microstructural design.

In the above two implementations, the evolving voids within the microstructure were represented using an artificial material with low Young’s modulus. This once again leads to high-condition number of the stiffness matrices ... a barrier to fast iterative analysis.


Here, we exploit the concept of topological sensitivity to avoid the pseudo-density challenges. In addition, the proposed method is designed to trace certain *pareto-optimal* curves (see next Section); this has numerous implications: (1) the underlying stiffness matrix is always well-conditioned, leading to fast iterative solutions, (2) the voids can be excluded from the finite element model without resulting in singularity, leading to shrinking matrix size, and (3) multi-fold symmetry can be imposed and exploited during the optimization process, leading to faster computation.

3. PROPOSED METHOD

3.1 Background: The PareTO Method

The proposed method for microstructural design builds upon the PareTO method described in [21], [32]. PareTO rests on the concept of *topological sensitivity* [10], [25], [29], [33]. To illustrate this concept, consider the 2D structural problem in Figure 6a where the objective is to find the optimal topology of a desired volume fraction with minimal compliance.

Consider now inserting a small *hypothetical* hole, i.e., modifying the topology, as illustrated in Figure 6b. Topological sensitivity is the expected change in any quantity of interest due to this topological change.



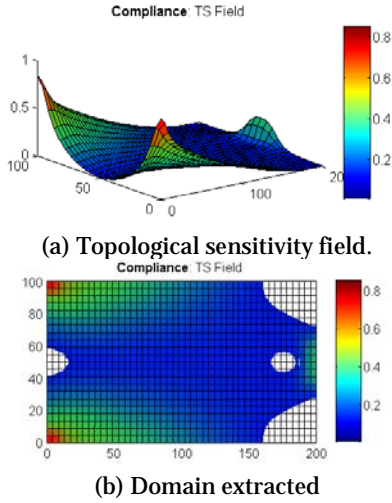
(a) Structural problem. (b) Topological change.

Figure 5: The main concept of topological level-set.

If the quantity of interest is compliance, one can show that topological sensitivity (TS) field in 2-D is given by [34]:

$$\mathcal{T}(p) \equiv \lim_{\varepsilon \rightarrow 0} \frac{q(p; \varepsilon) - q}{\pi \varepsilon^2} = \frac{4}{1 + \nu} \sigma : \varepsilon - \frac{1 - 3\nu}{1 - \nu^2} \text{tr}(\sigma) \text{tr}(\varepsilon) \quad (3.1)$$

The above field can be interpreted as a level-set as illustrated in Figure 7a. By 'slicing' the level-set at an appropriate height τ , one can extract topologies of desired volume fraction, as illustrated in Figure 7b. Similar topological level-sets can be computed in 3-D, and for various other quantities of interest [28].



(a) Topological sensitivity field.

(b) Domain extracted

Figure 6: The main concept of topological level-set.

The unique feature of PareTO is that it traces the pareto-optimal curve governing the desired objective (example: compliance) and the volume fraction, as illustrated in Figure 8; each of the steps is described below.

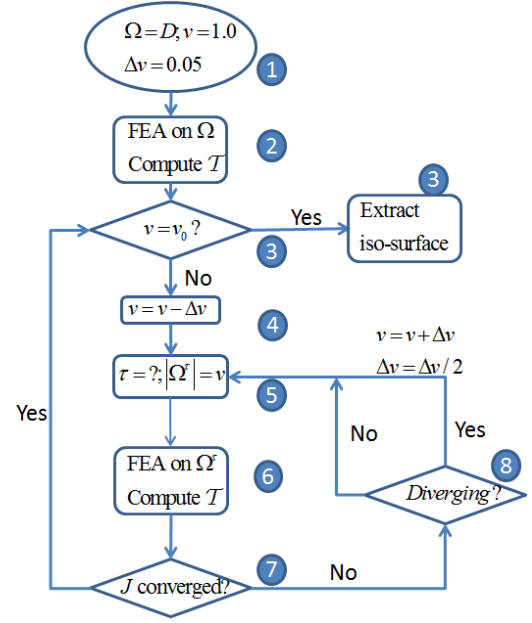


Figure 7: PareTO topology optimization algorithm.

1. Start with $\Omega = D$, i.e., start with a volume fraction of 1.0. Observe that this is different from SIMP, where a constant density ρ is assigned to the design-space D such that the 'apparent' volume fraction v_0 . The cutting-plane parameter τ is initialized to zero.
2. Next, a finite element analysis is carried out on D and the topological sensitivity field, corresponding to the quantity of interest, is computed.
3. If the desired volume fraction has been reached, the iso-surface with the current cutting-plane value τ is extracted. For iso-surface extraction, the classic marching-cubes algorithm described in [35] is used, where the topological sensitivity values at the corner nodes of the hex mesh are used to extract the iso-surface.
4. (Else) The current volume fraction is decremented by Δv ; Δv is initialized to 0.05, and this is controlled in an adaptive fashion (see step 8 below).
5. Given the current topological sensitivity field and the target volume fraction, we seek the parameter τ such that the volume of the extracted from the iso-surface is equal to the target volume fraction. This is a simple binary-search algorithm.
6. Once the desired value of τ has been computed, a finite element analysis is carried out on the extracted topology. Elements that lie outside are not included in the FEA (see discussion below) and, the topological sensitivity field is recomputed.
7. If the τ value has converged (to within user defined accuracy) return to Step-3. If the parameter has not yet converged, return to Step 5, after performing the check below to ensure that the optimization process is not diverging.
8. If a very large step size Δv is specified, the above process may diverge. If this is detected (by diverging values of compliance), the value of Δv is reduced by a factor of 2, prior to returning to Step 5.

Figure 9 illustrates the pareto-optimal curve traced and a set of optimal topologies computed. The optimization process starts at a volume-fraction of 1 (at the bottom right), and the pareto-optimal curve is traced in small decrements of volume fractions. Since the intermediate topologies lie on the pareto-curve, the structure is always connected during the optimization process, and the finite element analysis does not result in a singularity.

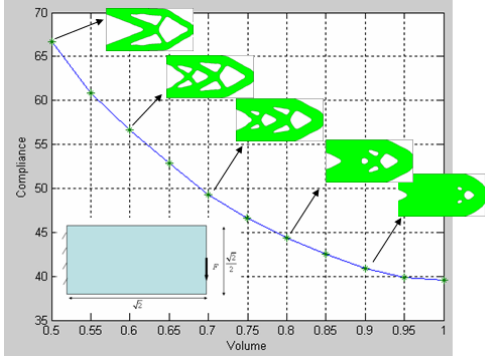



Figure 8: The pareto-optimal curve and topologies.

During the optimization, the ‘out-elements’ are not included in the finite element analysis. However, these are allowed (and may) reenter the optimization process as follows. Consider Figure 10a where the ‘in’ and ‘out’ elements within the current topology are identified; once the topological sensitivities over the ‘in’ elements are computed, the sensitivities at all nodes (see Figure 10b) are computed by averaging the element sensitivities. Finally, prior to the next iteration, an extrapolation procedure is used to predict the topological sensitivities over the ‘out’ elements that are adjacent to the boundary (see Figure 10c). This provides a simple mechanism for the ‘out’ elements to enter the optimization process.



(a) In-elements. (b) Nodal field. (c) Field over ‘out’ elements

Figure 9: Estimating topological sensitivity over ‘out’ elements

3.2 Microstructural Design

In microstructural design, these pareto-curves correspond precisely to the *Hashin-Shtrikman curves* [4], [11], [36], where the x-axis corresponds to the volume fraction, and y-axis corresponds to one of the microstructural objectives, such as bulk-modulus or shear-modulus. For example, it is well known that the upper limit on the bulk-modulus as a function of the volume fraction φ is given by [36]

$$K^*(\varphi) = \frac{K\varphi}{1 + \frac{3K(1-\varphi)}{5G}} \quad (3.2)$$

while the upper limit on the shear modulus is given by:

$$G^*(\varphi) = \frac{G\varphi}{1 + \frac{2(6K+12G)(1-\varphi)}{3(9K+8G)}} \quad (3.3)$$

where K & G are the material bulk and shear modulus respectively. Figure 11, for example, illustrates the Hashin-Shtrikman (pareto-optimal) curve for bulk-modulus.

Given the similarity in concept between Figure 9 and Figure 11, the objective of this paper is to extend the PareTO method for directly tracing the Hashin-Shtrikman curves.

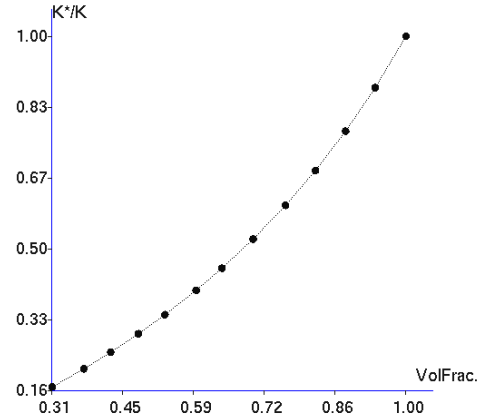


Figure 10: The Hashin-Shtrikman curve for bulk modulus.

3.3 Microstructural Discretization

For the purpose of finite element analysis, the microstructure is discretized into uniform-sized hexahedral elements, i.e., *voxels*, equipped with tri-linear shape functions. As the optimization progresses, voxels are either deleted from, or inserted into the mesh, via the topological level-set, as described later. Figure 12 illustrates two instances of the voxel-mesh.




Figure 11: Initial and an intermediate microstructure voxel mesh. As stated earlier, the underlying material is assumed to be linear isotropic, where the material tensor is given by:

$$\mathbb{D}^{-1} = \frac{1}{E} \begin{bmatrix} 1 & -\nu & -\nu & 0 & 0 & 0 \\ -\nu & 1 & -\nu & 0 & 0 & 0 \\ -\nu & -\nu & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2(1+\nu) & 0 & 0 \\ 0 & 0 & 0 & 0 & 2(1+\nu) & 0 \\ 0 & 0 & 0 & 0 & 0 & 2(1+\nu) \end{bmatrix} \quad (3.4)$$

The symbol \mathbb{D} is used for the underlying material tensor to distinguish it from the homogenized tensor \mathbb{C} of the microstructure. The latter depends both on \mathbb{D} and on the distribution of the voxels/voids.

3.4 Extracting the Homogenized Tensor

Given a microstructure as in Figure 12b, the first task lies in computing the homogenized 4th order tensor \mathbb{C} . While a generic 4th order tensor is defined by 81 independent values, the isotropic material \mathbb{C} satisfies multiple symmetries [37]:

$$\mathbb{C}_{ijkl} = \mathbb{C}_{ijlk} = \mathbb{C}_{jilk} = \mathbb{C}_{kltj} \quad (3.5)$$

Thus, 21 independent values are sufficient to represent \mathbb{C} . Further, if the microstructure exhibits three planes of symmetry,

i.e., if it is orthotropic, the number of independent constants reduces to nine, resulting in:

$$\mathbb{C} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ & C_{22} & C_{24} & 0 & 0 & 0 \\ & & C_{33} & 0 & 0 & 0 \\ & sym & & C_{44} & 0 & 0 \\ & & & & C_{55} & 0 \\ & & & & & C_{66} \end{bmatrix} \quad (3.6)$$

A popular strategy for computing \mathbb{C} is to impose six-independent periodic boundary conditions on the microstructure, and extracting the tensor through stress averaging [14]. While this strategy exhibits fast convergence with respect to mesh size, generating periodic constraints during the optimization process poses several challenges [38]. For example, the resulting Lagrangian system is not amenable to a fast iterative solution.

Therefore an alternate *constant-strain strategy* [14] is adopted. This method is fairly easy to implement, and it requires far fewer conjugate gradient iterations to converge. Here, the following six independent strains are considered:

$$\varepsilon^1 = \begin{Bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{Bmatrix}; \varepsilon^2 = \begin{Bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{Bmatrix}; \varepsilon^3 = \begin{Bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{Bmatrix}; \varepsilon^4 = \begin{Bmatrix} 0 \\ 0 \\ 0 \\ 0.5 \\ 0 \\ 0 \end{Bmatrix}; \varepsilon^5 = \begin{Bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0.5 \\ 0 \end{Bmatrix}; \varepsilon^6 = \begin{Bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0.5 \end{Bmatrix} \quad (3.7)$$

For each of the above six strains, the corresponding displacements are given by:

$$\begin{Bmatrix} u^i \\ v^j \\ w^j \end{Bmatrix} = \begin{bmatrix} \hat{x} & 0 & 0 & \hat{y} & \hat{z} & 0 \\ 0 & \hat{y} & 0 & \hat{x} & 0 & \hat{z} \\ 0 & 0 & \hat{z} & 0 & \hat{x} & \hat{y} \end{bmatrix} \varepsilon^i; i=1\dots6 \quad (3.8)$$

These six sets of displacements are imposed on the *boundary* of the microstructure, leading to six independent finite element problems, with the corresponding stresses:

$$\sigma^i(x, y, z) = \{\sigma_1^i \ \sigma_2^i \ \sigma_3^i \ \sigma_4^i \ \sigma_5^i \ \sigma_6^i\}^T \quad (3.9)$$

Once these stresses are extracted, one can show that the components of the \mathbb{C} matrix are given by [4]:

$$\mathbb{C}(i, j) = \frac{1}{|\Omega|} \int_{\Omega} \sigma_j^i(x, y, z) d\Omega \quad (3.10)$$

Many of these components are identically zero, consistent with Equation (3.6).

3.5 Matrix-Free Finite Element Analysis

We now focus on the efficient solution of the above set of six finite element problems. Observe that, in the voxel mesh, since each voxel possesses eight nodes, and each node is associated with three degrees of freedom, the element stiffness matrix for each voxel is given by a 24x24 matrix [39]:

$$K_e = \int_{\Omega} B^T \mathbb{D} B d\Omega \quad (3.11)$$

Since all elements in the voxel mesh are identical, it is sufficient if the above matrix is computed once at the beginning of the optimization process.

In classic finite element analysis, the element stiffness matrices are typically assembled into a global stiffness matrix K , whose size and structure will depend on the distribution of the voxels. Once the K matrix is assembled, the Dirichlet problems described result in a constrained linear algebra problem of the form:

$$\begin{aligned} \text{Solve: } & Ku^i = 0 \\ \text{s.t: } & Cu^i = b^i \end{aligned} \quad (3.12)$$

In the present work, the explicit construction of these matrices is avoided through an 'assembly-free' (a.k.a. 'matrix-free') approach [40] where neither K nor C are assembled/stored.

Specifically, to solve Equation (3.12), the *conjugate-residual* method [41] is employed. The conjugate-residual method is a generalization of the popular conjugate-gradient method in that it is designed to efficiently handle linear constraints. The critical steps in this method are the matrix-vector multiplication Ku and the constraint-vector multiplication Cu ; please see [41] for details. These are implemented here in a matrix-free form as

$$Ku = \left(\sum_e K_e \right) u = \sum_e K_e u_e \quad (3.13)$$

and

$$Cu = \left(\sum_e C_e \right) u = \sum_e C_e u_e \quad (3.14)$$

The advantages of a matrix-free analysis are:

1. Memory requirements are obviously reduced, and therefore fine resolution microstructures can be analyzed and optimized.
2. Since modern computer architectures are memory-bandwidth limited, memory reduction indirectly translates into increased computational speed [42].
3. As the voxels are deleted during the optimization, only the relevant elements need to be considered in Equation (3.13) and Equation (3.14).
4. Matrix-free multiplication is well suited for parallelization on multi-core architectures [43].

Once the six independent constrained linear algebra problems are solved, the homogenized tensor \mathbb{C} is computed via Equation (3.10).

3.6 Topological Sensitivity Computation

We now focus on the optimization/design process. Two specific design objectives considered in this paper are the average bulk-modulus and the average shear-modulus of the microstructure. Recall that, in linear elasticity, the standard bulk-modulus and shear-modulus are defined as:

$$\begin{aligned} K &= \frac{E}{3(1-2\nu)} \\ G &= \frac{E}{2(1+\nu)} \end{aligned} \quad (3.15)$$

Further, observe that these can be extracted from the tensor \mathbb{D}^{-1} in Equation (3.4) via:

$$K = \left(s^T \mathbb{D}^{-1} s \right)^{-1} \quad (3.16)$$

$$G = \left(\frac{t^T \mathbb{D}^{-1} t}{3} \right)^{-1} \quad (3.17)$$

where:

$$\begin{aligned} s &= \{1 \ 1 \ 1 \ 0 \ 0 \ 0\}^T \\ t &= \{0 \ 0 \ 0 \ 1 \ 1 \ 1\}^T \end{aligned} \quad (3.18)$$

Motivated by Equations (3.16) and (3.17), given a homogenized tensor \mathbb{C}^{-1} associated with a microstructure, the *average bulk modulus* and *average shear modulus* are defined as:

$$\bar{K} = (s^T \mathbb{C}^{-1} s)^{-1} \quad (3.19)$$

$$\bar{G} = \left(\frac{t^T \mathbb{C}^{-1} t}{3} \right)^{-1} \quad (3.20)$$

Since the primary objective in this paper is to maximize \bar{K} or \bar{G} , the topological sensitivity of \bar{K} and \bar{G} must be computed. Towards this end, one must first compute the topological sensitivity of the homogenized \mathbb{C} matrix. It was shown in [13] that the topological sensitivity field of the homogenized \mathbb{C} is given by:

$$\mathbb{T}_{\mathbb{C}}(i, j)(x, y, z) = (\sigma^i)^T \mathbb{H} \sigma^j; i, j = 1 \dots 6 \quad (3.21)$$

where:

$$\mathbb{H} = \begin{bmatrix} \alpha_1 + \alpha_2 & \alpha_2 & \alpha_2 & 0 & 0 & 0 \\ \alpha_2 & \alpha_1 + \alpha_2 & \alpha_2 & 0 & 0 & 0 \\ \alpha_2 & \alpha_2 & \alpha_1 + \alpha_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \alpha_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \alpha_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \alpha_1 \end{bmatrix} \quad (3.22)$$

$$\begin{aligned} \alpha_1 &= \frac{-15(1-\nu)(1+\nu)}{E(7-5\nu)} \\ \alpha_2 &= \frac{3(1-\nu)(1+5\nu)}{2E(7-5\nu)} \end{aligned} \quad (3.23)$$

where the stress fields σ^i & σ^j were defined earlier in Equation (3.9). Observe that $\mathbb{T}_{\mathbb{C}}$ is a 4th order tensor field with the usual symmetries of an orthotropic tensor. Further since:

$$\mathbb{C} \mathbb{C}^{-1} = I \quad (3.24)$$

the topological sensitivity of \mathbb{C}^{-1} is given by:

$$\mathbb{T}_{\mathbb{C}^{-1}} = -\mathbb{C}^{-1} \mathbb{T}_{\mathbb{C}} \mathbb{C}^{-1} \quad (3.25)$$

Exploiting Equation (3.25) and Equation (3.19), the topological sensitivity of the homogenized bulk modulus is given by:

$$\mathcal{T}_{\bar{K}}(x, y, z) = -(\bar{K})^2 s^T \mathbb{T}_{\mathbb{C}^{-1}} s = (\bar{K})^2 s^T \mathbb{C}^{-1} \mathbb{T}_{\mathbb{C}} \mathbb{C}^{-1} s \quad (3.26)$$

while the topological sensitivity of the homogenized shear modulus is given by

$$\mathcal{T}_{\bar{G}}(x, y, z) = -\frac{(\bar{G})^2 t^T \mathbb{T}_{\mathbb{C}^{-1}} t}{3} = \frac{(\bar{G})^2 t^T \mathbb{C}^{-1} \mathbb{T}_{\mathbb{C}} \mathbb{C}^{-1} t}{3} \quad (3.27)$$

3.7 Proposed Method

The above topological sensitivity theory can now be used to carry out a topology optimization as follows:

1. The optimization starts at a volume fraction of 1.0; the microstructure is discretized into N elements, where N ranges from 10,000 to 300,000 in the numerical

experiments discussed in the next Section. A single element stiffness matrix is computed per Equation (3.11). The 'current volume fraction' v is set to 1.0, and 'desired decrement' and the increment Δv , is set to 0.01.

2. The finite element problems are solved over the current mesh, and the stresses are extracted at the center of each element. The tensor \mathbb{C} is computed via Equation (3.10).
3. The topological sensitivity field (either $\mathcal{T}_{\bar{K}}$ or $\mathcal{T}_{\bar{G}}$) is computed at the center of each voxel, and locally smoothed with neighboring elements.
4. Treating the topological sensitivity field as a level-set, a new topology with a volume fraction of $(v - \Delta v)$ is extracted. The quantity of interest (say, bulk modulus) is computed over the new topology. If this quantity has converged for the volume fraction $(v - \Delta v)$, then the optimization moves to the next step, else it returns to step 2.
5. The current volume fraction is set to $(v - \Delta v)$, and the optimization returns to step 2 (until the final volume fraction is reached).

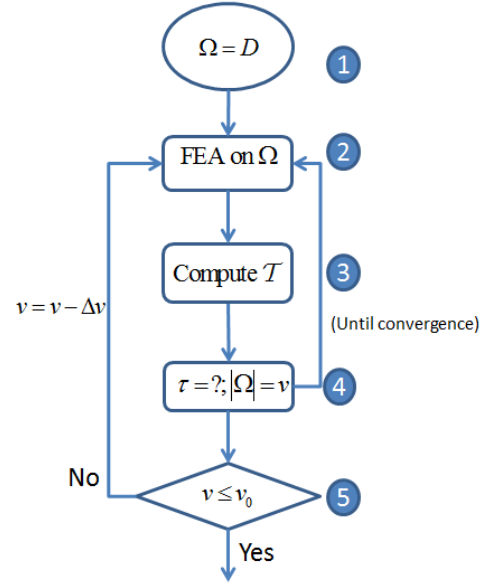


Figure 12: An overview of the algorithm.

3.7 Exploiting Symmetry

One of the significant advantages of tracing the pareto-optimal curve is that since each intermediate microstructure is pareto-optimal, it exhibits a three-fold symmetry (i.e., there is no preferred direction due to material isotropy and symmetry of problem specification). In other words, as the bulk or shear modulus is optimized, the components of \mathbb{C} in Equation (3.6) satisfy the following additional properties:

$$\begin{aligned} C_{11} &= C_{22} = C_{33} \\ C_{12} &= C_{13} = C_{24} \\ C_{44} &= C_{55} = C_{66} \end{aligned} \quad (3.28)$$

Thus, it is sufficient to solve two finite elements problems (instead of the six), namely, a stretching and a shear problem. This further reduces the computation without sacrificing on accuracy.

3.8 Parallelization on CPU and GPU

Much of the computation time is spent in executing the assembly-free sparse matrix-vector multiplication (SpMV) in Equation (3.13). In the CPU, parallelization of this module was attained through OpenMP commands (www.openmp.org). Other supporting modules such as 'vector dot-product', etc., were also accelerated through OpenMP commands.

The GPU implementation relies heavily on the CUDA language [44]. At the start of the optimization process, the following data was transferred to the GPU: (1) mesh connectivity, (2) element stiffness matrix and (3) material properties. Additional memory was allocated for the solution vector and other auxiliary data.

Each node of the mesh is assigned to a scalar processor of the GPU. When a block of threads is launched, observe that a single memory fetch of the element stiffness matrix is sufficient. The three degrees of freedom associated with each node are fetched in parallel, and assembly-free SpMV is implemented. Many of the low-level operations such as vector-product were implemented using CUBLAS library. Once the linear system is solved, the solution is transferred to the CPU for post-processing and sensitivity computation. Despite this overhead, the GPU implementation offers a speed-up for fine-resolution optimization.

4. NUMERICAL EXAMPLES

In this Section, we present results from numerical experiments based on the above algorithm. The material properties are $E = 1$ and $\nu = 0.28$, and the initial microstructure is a cube of 1 mm on each side (unless otherwise noted). All experiments were conducted on a Windows 7 64-bit machine with the following hardware:

- AMD 8-core FX-8350 CPU running at 4.0GHz with 16 GB of memory; parallelization of CPU code was implemented through OpenMP commands.
- The GPU is an NVidia GTX Titan (2688 cores) with 5.6 GB, and CUDA compute capability of 3.5.

4.1 Bulk-Modulus Maximization

In the first experiment, the cube is discretized into 30,000 elements (about 100,000 degrees of freedom), and the bulk modulus is optimized. Figure 14 illustrates the optimal microstructure of maximum bulk-modulus at various volume fractions. The desired volume fraction of 0.2 is reached in 45 seconds, using 152 finite element operations.




Figure 13: Optimal microstructures of maximum bulk modulus at various volume fractions.

Figure 15 illustrates the computed and theoretical relative bulk-moduli as a function of the volume fraction. As one can observe, the proposed method traces the theoretical curve closely.

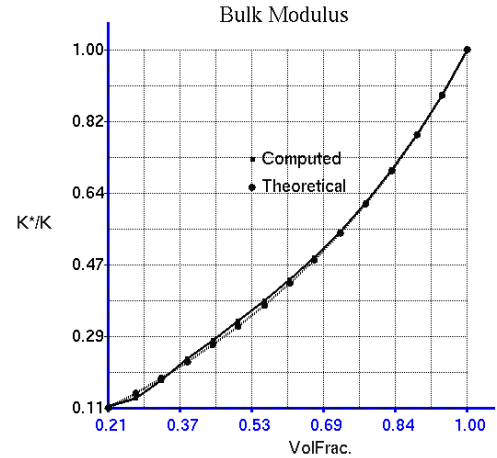


Figure 14: Computed and theoretical Hashin-Shtrikman pareto-optimal curves for bulk-modulus maximization.

Given the unit cell of volume fraction of 0.2 in Figure 14, these can be stacked to create locally-periodic microstructures as illustrated in Figure 16.




Figure 15: Stacking of the unit-cell in Figure 14d.

4.2 Shear-Modulus Maximization

The above set of experiments is now repeated for shear-modulus maximization. Figure 17 illustrates the optimal microstructure of maximum shear modulus at various volume fractions. The desired volume fraction of 0.3 was reached in 40 seconds, using about 100 finite element operations.




Figure 16: Optimal microstructures of maximum shear modulus at various volume fractions..

Figure 18 illustrates the pareto-optimal curves for shear-moduli versus volume fraction.

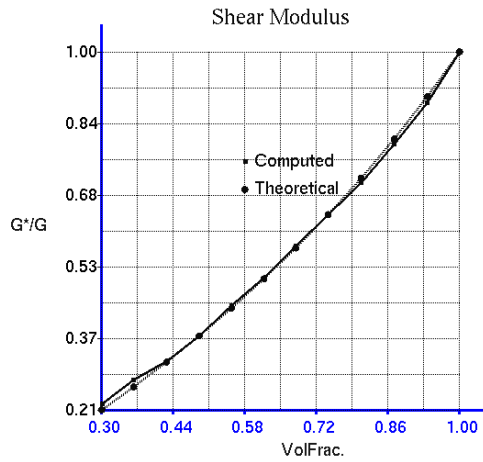


Figure 17: Computed and theoretical Hashin-Shtrikman pareto-optimal curves for shear-modulus maximization.

Figure 19 illustrates the stacked microstructure for the unit-cell in Figure 19d.




Figure 18: Stacking of the unit-cell in Figure 17d.

4.3 Computational Cost

One of the salient features of the proposed method is that microstructures of fine resolution can be optimized efficiently. Figure 20 illustrates the total time taken to optimize the bulk modulus up to a volume fraction of 0.2 as a function of finite element degrees of freedom (DOF). As one can observe, the time taken by the proposed method (on a desktop with 8 cores) for a 3 million degree system is around 4 hours.

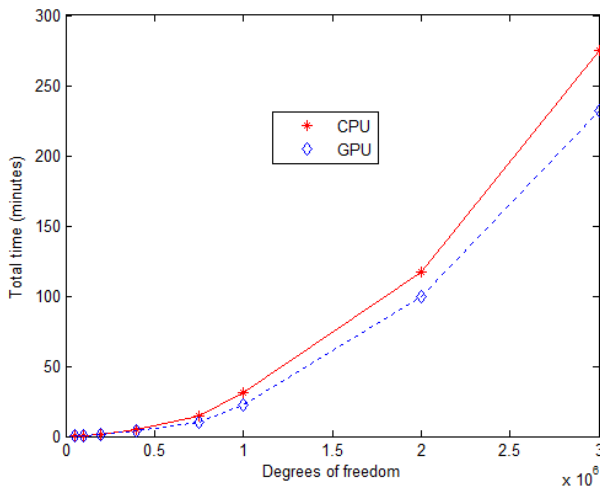


Figure 19: Time taken to optimize the bulk-modulus to a volume fraction of 0.2, in steps of 0.025, as a function of DOF.

5. CONCLUSIONS

The main contribution of the paper is an efficient method for microstructural design. As illustrated, the proposed method

closely traces the Hashin-Shtrikman pareto-optimal curves. Consequently, the underlying stiffness matrices are always well-conditioned, and multi-fold symmetry can be imposed and exploited during the optimization process. The software will be made available through the author's research website www.erl.wisc.edu. This paper serves as a foundation for future work on: (1) negative Poisson ratio materials, (2) multi-level topology optimization, and (3) multi-material microstructural design.

Acknowledgements

The author would like to thank the support of National Science Foundation through grants CMMI-1232508 and CMMI-1161474.

REFERENCES

- [1] H. A. Eschenauer and N. Olhoff, "Topology optimization of continuum structures: A review," *Applied Mechanics Review*, vol. 54, no. 4, pp. 331–389, 2001.
- [2] G. I. N. Rozvany, "A critical review of established methods of structural topology optimization," *Structural and Multidisciplinary Optimization*, vol. 37, no. 3, pp. 217–237, 2009.
- [3] O. Sigmund, "Tailoring materials with prescribed elastic properties," *Mechanics of Materials*, vol. 20, pp. 351–368, 1995.
- [4] I. Ozdemir, "Topological derivative based optimization of 3D porous elastic microstructures," *Computational Materials Science*, vol. 81, pp. 319–325, 2014.
- [5] J. Guest and J. H. Prevost, "Optimizing multifunctional materials: Design of microstructures for maximized stiffness and fluid permeability," *International Journal of Solids and Structures*, vol. 43, pp. 7028–7047, 2006.
- [6] V. J. Challis, A. J. Roberts, and A. H. Wilkins, "Design of three dimensional isotropic microstructures for maximized stiffness and conductivity," *International Journal of Solids and Structures*, vol. 45, pp. 4130–4146, 2008.
- [7] E. Andreassen, B. S. Lazarov, and O. Sigmund, "Design of manufacturable 3D extremal elastic microstructure," *Mechanics of Materials*, vol. 69, pp. 1–10, 2014.
- [8] J. B. Choi and R. S. Lakes, "Design of a fastener based on negative Poisson's ratio foam," *Cellular Polymers*, vol. 10, pp. 205–212, 1991.
- [9] N. de Kruijf, S. Zhou, Q. Li, and Y.-W. Mai, "Topological design of structures and composite materials with multiobjectives," *International Journal of Solids and Structures*, vol. 44, pp. 7092–7109, 2007.
- [10] J. Sokolowski and A. Zochowski, "On Topological Derivative in Shape Optimization," *SIAM journal on control and optimization*, vol. 37, no. 4, pp. 1251–1272, 1999.
- [11] J. Cadman, S. Zhou, Y. Chen, and Q. Li, "On design of multi-functional microstructural materials," *Journal of Materials Science*, vol. 48, pp. 1–16, 2012.
- [12] J. C. Michel, H. Moulinec, and P. Suquet, "Effective properties of composite materials with periodic microstructure: a computational approach," *Comput. Methods Appl. Mech. Eng.*, vol. 172, pp. 109–143, 1999.
- [13] S. M. Giusti, A. A. Novotny, E. A. de Souza Neto, and R. A. Feijoo, "Sensitivity of the macroscopic elasticity tensor to topological microstructural changes," *Journal of the Mechanics and Physics of Solids*, vol. 57, pp. 555–570, 2009.
- [14] K. Terada, M. Hori, T. Kyoya, and N. Kikuchi, "Simulation of the multi-scale convergence in computational homogenization approaches," *International Journal of Solids and Structures*, vol. 37, no. 16, pp. 2285–2311.

- [15] G. W. Milton, "Composite materials with Poisson's ratios close to -1," *J. Mech. Phys. Solids*, vol. 40, no. 5, pp. 1105–1137, 1992.
- [16] K. Evans, "Tailoring the negative Poisson's ratio," *Chem. & Indus.*, vol. 15, 1990.
- [17] R. Lakes, "Foam Structures with negative Poisson's ratio," *Science*, vol. 235, pp. 1038–1040, 1987.
- [18] O. Sigmund, "Topology optimization: a tool for the tailoring of structures and materials," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 358, pp. 211–288, 2000.
- [19] A. Diaz and O. Sigmund, "A topology optimization method for design of negative permeability metamaterials," *Structural and Multidisciplinary Optimization*, vol. 41, pp. 163–177, 2010.
- [20] S. Wang, E. D. Sturler, and G. Paulino, "Large-scale topology optimization using preconditioned Krylov subspace methods with recycling," *International Journal for Numerical Methods in Engineering*, vol. 69, no. 12, pp. 2441–2468, 2007.
- [21] K. Suresh, "Efficient Generation of Large-Scale Pareto-Optimal Topologies," *Structural and Multidisciplinary Optimization*, vol. 47, no. 1, pp. 49–61, 2013.
- [22] P. G. Coelho, J. B. Caddoso, P. R. Fernandes, and H. C. Rodrigues, "Parallel computing techniques applied to the simultaneous design of structure and material," *Advances in Engineering Software*, vol. 42, pp. 219–227, 2011.
- [23] H. A. Eschenauer, V. V. Kobelev, and A. Schumacher, "Bubble method for topology and shape optimization of structures," *Structural Optimization*, vol. 8, pp. 42–51, 1994.
- [24] A. A. Novotny, R. A. Feijóo, C. Padra, and E. Taroco, "Topological Derivative for Linear Elastic Plate Bending Problems," *Control and Cybernetics*, vol. 34, no. 1, pp. 339–361, 2005.
- [25] A. A. Novotny, "Topological-Shape Sensitivity Method: Theory and Applications," *Solid Mechanics and its Applications*, vol. 137, pp. 469–478, 2006.
- [26] A. A. Novotny, R. A. Feijóo, and E. Taroco, "Topological Sensitivity Analysis for Three-dimensional Linear Elasticity Problem," *Computer Methods in Applied Mechanics and Engineering*, vol. 196, no. 41–44, pp. 4354–4364, 2007.
- [27] J. Cea, S. Garreau, P. Guillaume, and M. Masmoudi, "The shape and topological optimization connection," *Computer Methods in Applied Mechanics and Engineering*, vol. 188, no. 4, pp. 713–726, 2000.
- [28] I. Turevsky, S. H. Gopalakrishnan, and K. Suresh, "An Efficient Numerical Method for Computing the Topological Sensitivity of Arbitrary Shaped Features in Plate Bending," *International Journal of Numerical Methods in Engineering*, vol. 79, pp. 1683–1702, 2009.
- [29] I. Turevsky and K. Suresh, "Generalization of Topological Sensitivity and its Application to Defeaturing," in *ASME IDETC Conference*, Las Vegas, 2007.
- [30] S. H. Gopalakrishnan and K. Suresh, "Feature Sensitivity: A Generalization of Topological Sensitivity," *Finite Elements in Analysis and Design*, vol. 44, no. 11, pp. 696–704, 2008.
- [31] S. Amstutz, S. M. Giusti, A. A. Novotny, and E. A. de Souza Neto, "Topological derivative for multi-scale linear elasticity models applied to the synthesis of microstructures," *International Journal for Numerical Methods in Engineering*, vol. 84, pp. 733–756, 2010.
- [32] K. Suresh, "A 199-line Matlab code for Pareto-optimal tracing in topology optimization," *Structural and Multidisciplinary Optimization*, vol. 42, no. 5, pp. 665–679, 2010.
- [33] A. A. Novotny, R. A. Feijóo, and E. Taroco, "Topological Sensitivity Analysis for Three-dimensional Linear Elasticity Problem," *Computer Methods in Applied Mechanics and Engineering*, vol. 196, no. 41–44, pp. 4354–4364, 2007.
- [34] R. A. Feijóo, A. A. Novotny, E. Taroco, and C. Padra, "The topological-shape sensitivity method in two-dimensional linear elasticity topology design," in *Applications of Computational Mechanics in Structures and Fluids*, CIMNE, 2005.
- [35] W. E. Lorensen and H. E. Cline, "Marching Cubes: a high resolution 3D surface reconstruction algorithm.," *Computer Graphics (Proc. of SIGGRAPH)*, vol. 21, no. 4, pp. 163–169, 1987.
- [36] Z. Hashin and S. Shtrikman, "A variational approach to the theory of the elastic behaviour of multiphase materials," *Journal of the Mechanics and Physics of Solids (UK)*, vol. 11, pp. 127–140, 1963.
- [37] Y. C. Fung, *Classical and Computational Solid Mechanics*. River Ridge, NJ: World Scientific, 2001.
- [38] V.-D. Nguyen, E. Bechet, C. Geuzaine, and L. Noels, "Imposing periodic boundary condition on arbitrary meshes by polynomial interpolation," *Computational Materials Science*, vol. 00, pp. 1–28, 2011.
- [39] O. C. Zienkiewicz and R. L. Taylor, *The Finite Element Method for Solid and Structural Mechanics*. Elsevier, 2005.
- [40] C. E. Augarde, A. Ramage, and J. Staudacher, "An element-based displacement preconditioner for linear elasticity problems," *Computers and Structures*, vol. 84, no. 31–32, pp. 2306–2315, 2006.
- [41] D. G. Luenberger, "The conjugate residual method for constrained minimization problems," *SIAM Journal of Numerical Analysis*, vol. 7, no. 3, 1970.
- [42] D. Goddeke, R. Strzodka, and S. Turek, "Performance and accuracy of hardware-oriented native-emulated- and mixed-precision solvers in FEM simulations," *International Journal of Parallel, Emergent and Distributed Systems*, vol. 22, no. 4, pp. 221–256, 2007.
- [43] K. Suresh and P. Yadav, "Large-Scale Modal Analysis on Multi-Core Architectures," in *Proceedings of the ASME 2012 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*, Chicago, IL, 2012.
- [44] NVIDIA Corporation, *NVIDIA CUDA: Compute Unified Device Architecture, Programming Guide*. Santa Clara., 2008.