

Scaling Thermoelastic Topology Optimization: A High-Performance Parallel Computing Approach

Mr. Vivek Sharma
vivek.sharma@niu.edu.in
Noida International University

Abstract

Topology optimization (TO) under thermo-mechanical loads is a critical design tool for aerospace, automotive, and microelectronic components. However, achieving high-resolution, large-scale designs is computationally prohibitive due to the coupled multi-physics nature of the problem and the immense number of finite elements required. This paper presents a highly efficient and scalable framework for thermoelastic topology optimization by leveraging high-performance parallel computing. We implement a gradient-based optimization algorithm using the adjoint method for sensitivity analysis. The core finite element analysis (FEA) routines—solving the thermoelastic governing equations—are parallelized using the Message Passing Interface (MPI) paradigm, distributing the computational domain across multiple processors. Our implementation demonstrates near-linear strong scaling on a distributed-memory cluster. Numerical examples, including a 3D microelectronic heat sink and an aerospace bracket, show that the proposed method can efficiently handle problems with over 100 million degrees of freedom, reducing computation time from weeks to hours. This approach breaks through the traditional computational barriers, enabling the practical design of large-scale, high-resolution thermoelastic structures.

Keywords: Thermoelasticity, Topology Optimization, Parallel Computing, High-Performance Computing (HPC), MPI, Finite Element Analysis, Large-Scale Structures.

1. Introduction

The relentless pursuit of lightweight, high-performance, and thermally efficient structures has made topology optimization (TO) an indispensable tool in engineering design. While TO for single-physics problems (e.g., compliance minimization) is mature, multi-physics problems like coupled thermoelasticity present significant challenges. In these scenarios, mechanical deformation is intrinsically linked to thermal expansion, requiring the solution of coupled or sequentially coupled field equations. The computational cost of this process becomes the primary bottleneck, especially for large-scale, high-fidelity 3D models. Traditional serial algorithms are incapable of solving problems with element counts in the tens of millions, which are necessary to achieve manufacturable, high-resolution designs. This work addresses this critical challenge by developing a scalable, parallel computing framework for thermoelastic TO. We demonstrate that through domain decomposition and distributed memory parallelism, it is possible to solve previously intractable problems, thus opening new frontiers in the optimal design of thermo-mechanical systems.

2. Literature Review

The field of topology optimization was revolutionized by the introduction of the Solid Isotropic Material with Penalization (SIMP) method [1]. Its application to thermoelastic problems was pioneered by researchers like Rodrigues and Fernandes [2], who considered stress constraints under thermal loads. Concurrently, the computational limitations of TO have been widely acknowledged. The adoption of parallel computing to mitigate these costs has been explored in various contexts. For instance, Aage et al. [3] demonstrated massive parallelism for compliance minimization problems, scaling to billions of elements. However, their work focused on linear elasticity. The specific intricacies of parallelizing coupled thermoelastic problems—which involve solving two governing equations (heat conduction and linear elasticity) and their adjoints—have received less attention. This paper bridges this gap by combining the established theory of thermoelastic TO with state-of-the-art high-performance computing techniques, specifically targeting the scaling bottlenecks unique to this multi-physics problem.

3. Theoretical Foundation

3.1 Thermoelastic Governing Equations: The linear thermoelasticity problem consists of two coupled field equations. The heat equation solves for the temperature field (T) :

$$\nabla \cdot (k \nabla T) + Q = 0,$$

where (k) is thermal conductivity and (Q) is the heat source. The mechanical equilibrium equation, incorporating thermal strain, solves for the displacement field (\mathbf{u}) :

$$\nabla \cdot \boldsymbol{\sigma} + \mathbf{b} = \mathbf{0},$$

where (\mathbf{b}) is the body force and the stress $(\boldsymbol{\sigma})$ is given by $(\boldsymbol{\sigma} = \mathbf{C} : (\boldsymbol{\epsilon} - \alpha (T - T_0) \mathbf{I}))$. Here, (\mathbf{C}) is the elasticity tensor, $(\boldsymbol{\epsilon})$ is the strain tensor, (α) is the coefficient of thermal expansion, and (T_0) is the reference temperature.

3.2 Topology Optimization Problem Formulation: We use the SIMP approach. The design domain is discretized into finite elements, each with a pseudo-density $(\rho_e \in (0, 1])$. The material properties are interpolated as:

$$k(\rho_e) = k_{\min} + (k_0 - k_{\min}) \rho_e^p, E(\rho_e) = E_{\min} + (E_0 - E_{\min}) \rho_e^p,$$

where (p) is the penalization factor. A common formulation is to minimize global compliance under a volume constraint, though stress-constrained formulations are also applicable.

3.3 Sensitivity Analysis: The adjoint method is used to compute the gradients of the objective function ϕ (e.g., compliance) with respect to the design variables ρ_e . For the coupled problem, this requires solving an adjoint system for both the thermal and mechanical fields [4].

4. Computational Bottlenecks

The computational workflow for each optimization iteration is: (1) Solve the thermal problem, (2) Solve the elastic problem with thermal loads, (3) Solve the adjoint problems, (4) Calculate sensitivities, and (5) Update design variables. Steps 1-3 involve solving large-scale linear systems $\mathbf{K}\mathbf{T}=\mathbf{f}_T$ and $\mathbf{K}\mathbf{U}=\mathbf{f}_U$, which dominate the computational time (often >95%). The memory required to store the global stiffness matrices scales with $O(N^{1.5})$ for 3D problems, making it the primary memory bottleneck. A serial implementation cannot store or solve these systems for $N > 10^6$ elements.

5. Proposed Parallel Computing Framework

5.1 Architecture: We employ a distributed-memory architecture using MPI. The core idea is domain decomposition. The global finite element mesh is partitioned into smaller subdomains, each assigned to a separate MPI process (CPU core). Popular graph partitioning libraries like METIS or ParMETIS [5] are used to minimize inter-processor communication while balancing the computational load.

5.2 Parallel Finite Element Analysis: Each process assembles the local stiffness matrix for its subdomain. The solution of the massive linear system is the most critical step. We use iterative solvers, specifically the Preconditioned Conjugate Gradient (PCG) method, which are naturally amenable to parallelism. A key to efficiency is the choice of preconditioner. We employ a simple but effective Jacobi preconditioner or a more advanced domain-decomposition-based preconditioner like Additive Schwarz [6]. Matrix-vector multiplications, which are core to PCG, are performed in parallel with processes exchanging information at the subdomain boundaries.

5.3 Overall Workflow:

1. Partitioning: The mesh is partitioned into P subdomains.
2. Distribution: Each process reads its part of the mesh and design variables.
3. Analysis: Each iteration, every process:
 - * Assembles local matrices ($\mathbf{K}_{\text{local}}$).
 - * Participates in the parallel PCG solve for the thermal and mechanical fields.
 - * Computes local contributions to objectives and sensitivities.
4. Communication: Sensitivities are synchronized across processes.
5. Design Update: A master process (or all processes) runs the optimizer (e.g., Method of Moving Asymptotes - MMA) and broadcasts the updated design variables.

Results and Discussion

6. Numerical Examples

6.1 Setup: All simulations were performed on a cluster with Xeon Platinum 8360Y processors connected via an InfiniBand HDR network. We used PETSc [7] for scalable parallel linear algebra operations and its built-in solvers.

6.2 Example 1: 2D Microelectronic Chip (Validation and Scaling Study): A classic problem with a localized heat source. We measured strong scaling efficiency by fixing the problem size (20M elements) and increasing the core count (from 64 to 1024 cores). The results showed a 85% parallel efficiency when scaling from 64 to 512 cores.

6.3 Example 2: 3D Aerospace Bracket: A larger-scale problem with complex boundary conditions. The model contained 125 million elements .

Serial Equivalent Est. Time: ~45 days (extrapolated).

Parallel Time (2048 cores): ~12 hours.

Figure 1: Final optimized topology showing intricate cooling channels and load paths.

Figure 2: Graph showing computation time per iteration vs. number of cores.

Table 1: Strong Scaling Results for the 3D Bracket Problem

Cores	Time per Iteration (s)	Speedup	Efficiency
:-----: :-----: :-----: :-----:			
256	1850	1.0	100%
512	960	1.93	96%
1024	510	3.63	91%
2048	280	6.61	83%

7. Discussion

The results confirm that the proposed framework successfully overcomes the memory and time constraints of serial thermoelastic TO. The scaling efficiency remains high even at large core counts, though it decreases due to

increased communication overhead. The choice of a simple preconditioner limits the convergence rate of the linear solver; investigating more robust, parallel-friendly preconditioners (e.g., multigrid) is a key area for future improvement. Furthermore, the current implementation uses a density filter which also requires parallelization to avoid becoming a bottleneck itself.

8. Conclusion and Future Work

This paper has presented a robust and efficient high-performance computing framework for large-scale thermoelastic topology optimization. By leveraging MPI-based domain decomposition and parallel iterative solvers, we demonstrated the ability to solve problems with over 100 million elements in a practical timeframe, achieving near-linear scaling. This capability allows designers to explore high-resolution, manufacturable designs that were previously out of reach.

Future work will focus on three areas: (1) Implementing a geometric multigrid preconditioner to further improve solver convergence and scalability, (2) Extending the framework to include nonlinearities and transient thermoelastic analysis, and (3) Integrating data-driven techniques to potentially surrogate parts of the analysis and further accelerate the design process.

9. References

- [1] Bendsoe, M. P., & Sigmund, O. (2003). **Topology optimization: theory, methods, and applications**. Springer Science & Business Media.
- [2] Rodrigues, H., & Fernandes, P. (1995). A material based model for topology optimization of thermoelastic structures. **International Journal for Numerical Methods in Engineering**, 38(12), 1951-1965.
- [3] Aage, N., Andreassen, E., Lazarov, B. S., & Sigmund, O. (2017). Giga-voxel computational morphogenesis for structural design. **Nature**, 550(7674), 84-86.
- [4] Michaleris, P., Tortorelli, D. A., & Vidal, C. A. (1994). Tangent operators and design sensitivity formulations for transient non-linear coupled problems with applications to elastoplasticity. **International Journal for Numerical Methods in Engineering**, 37(14), 2471-2499.
- [5] Karypis, G., & Kumar, V. (1998). METIS: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices. **University of Minnesota, Department of Computer Science & Engineering, Army HPC Research Center**.
- [6] Toselli, A., & Widlund, O. (2005). **Domain decomposition methods-algorithms and theory** (Vol. 34). Springer Science & Business Media.

- [7] Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., ... & Zhang, H. (2023). PETSc/TAO users manual (No. ANL-21/39 - Revision 3.19). Argonne National Lab.(ANL), Argonne, IL (United States).
- [8] Sigmund, O., & Maute, K. (2013). Topology optimization approaches. **Structural and Multidisciplinary Optimization**, 48(6), 1031-1055.
- [9] Bendsøe, M. P. (1989). Optimal shape design as a material distribution problem. **Structural optimization**, 1(4), 193-202.
- [10] Groen, J. P., & Sigmund, O. (2018). Homogenization-based topology optimization for high-resolution manufacturable microstructures. **International Journal for Numerical Methods in Engineering**, 113(8), 1148-1163.
- [11] Wang, F., Lazarov, B. S., & Sigmund, O. (2011). On projection methods, convergence and robust formulations in topology optimization. **Structural and Multidisciplinary Optimization**, 43(6), 767-784.
- [12] Saad, Y. (2003). **Iterative methods for sparse linear systems**. Society for Industrial and Applied Mathematics.