GPU-accelerated matrix-free topology optimization

A Delissen, J Wu, M Möller, AM Aragón, M Langelaar

Delft University of Technology

Motivation

Topology optimization of TU Delft's Flying-V airplane is a great computational challenge. For a fine-scale design resolution of a few millimeters, the large aircraft (65m span) requires billions of finite elements (13000×11000×2600). We propose a matrix-free approach in which the matrix-vector product is assembled on the fly on the GPU, avoiding explicit storage of the system matrix, which would otherwise exceed the memory capacity of even large supercomputers.

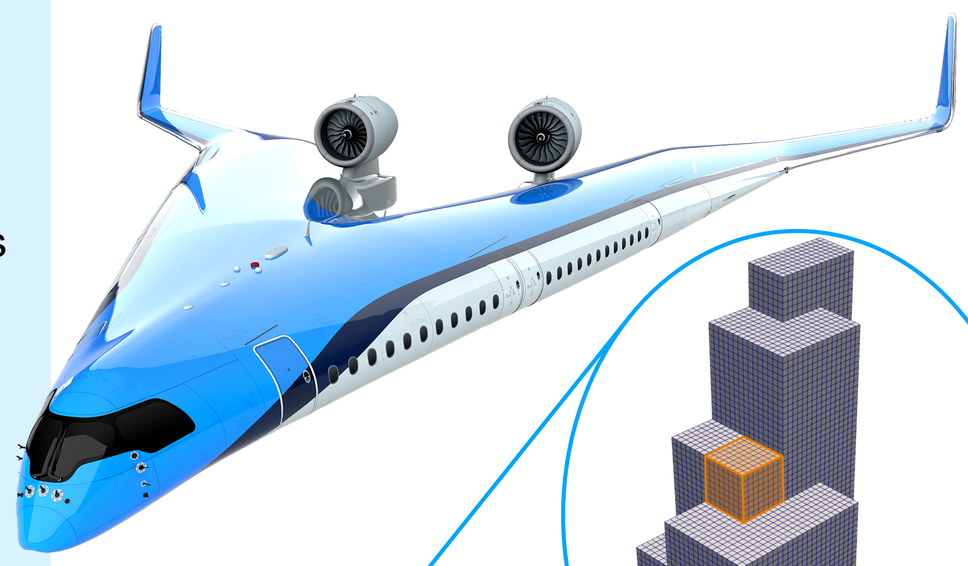

Finite element domain

Cube-shaped trilinear elements (voxels) enable the use of a singleton finite element matrix. The domain is built up using sectors of 8×8×8 elements, omitting sectors that are completely void. This allows for contiguous memory access and leads to highly efficient GPU utilization.

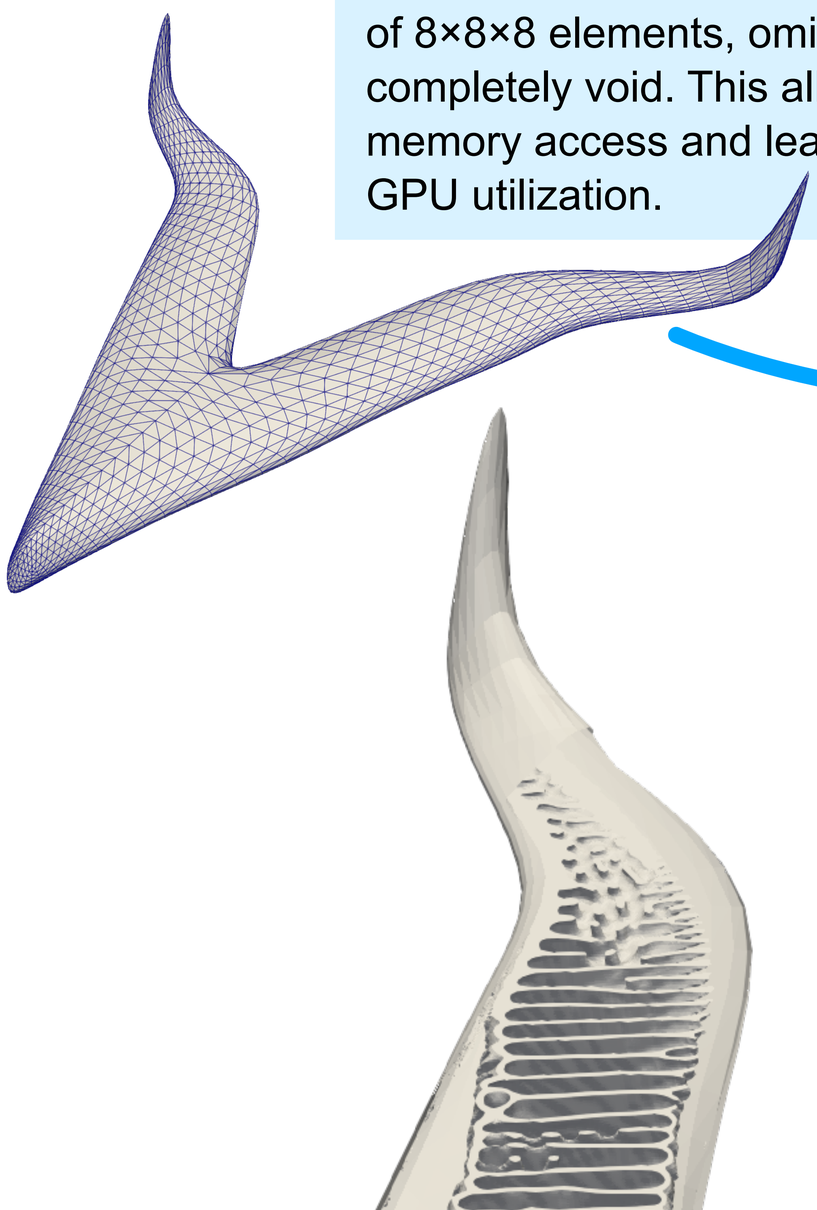

Topology optimization

The densities of voxels within the airplane surface are optimized. An optimal distribution of material for one or more given loadcase(s) is generated using an iterative process. In this case, 1.4 million voxels are meshed, requiring about 20 seconds per design iteration on a single V100s GPU.

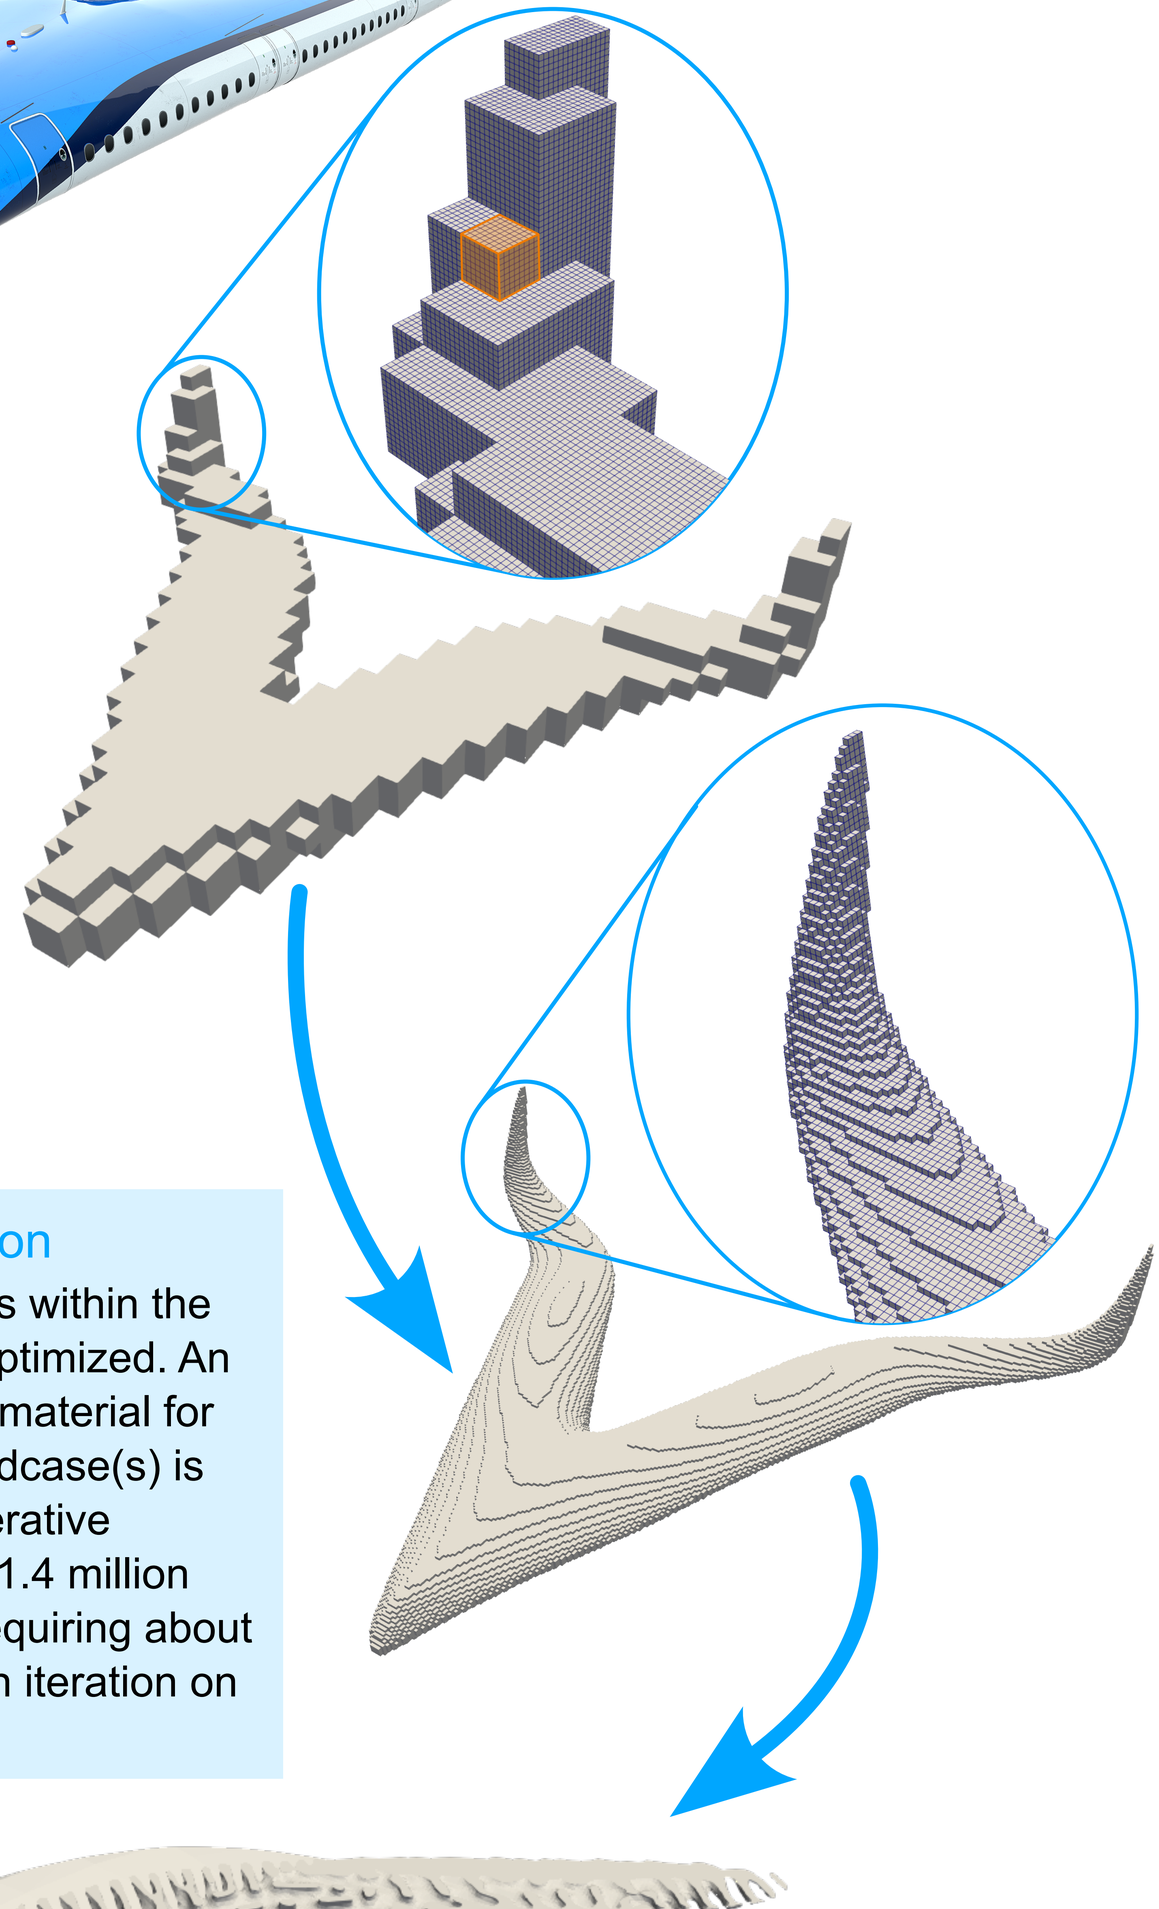

Conclusions

- Small memory footprint of matrix-free matrix-vector product.
- Efficient use of GPU with contiguous memory access.
- Ability to optimize irregular shapes without redundant finite-element calculations in void areas.
- Fast design iterations, with potential for even larger problems with multigrid preconditioning and multiple GPUs.