

 Open access • Journal Article • DOI:10.1038/NATURE23911

Giga-voxel computational morphogenesis for structural design — [Source link](#)

Niels Aage, Erik Andreassen, Boyan Stefanov Lazarov, Ole Sigmund

Published on: 04 Oct 2017 - Nature (Nature Publishing Group)

Topics: Topology optimization

Related papers:

- [Generating optimal topologies in structural design using a homogenization method](#)
- [The method of moving asymptotes—a new method for structural optimization](#)
- [Topology Optimization: Theory, Methods, and Applications](#)
- [Optimal shape design as a material distribution problem](#)
- [Structural optimization using sensitivity analysis and a level-set method](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/giga-voxel-computational-morphogenesis-for-structural-design-4fs8r63d2g>

Giga-voxel computational morphogenesis for structural design

DOI:
[10.1038/nature23911](https://doi.org/10.1038/nature23911)

Document Version
Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Aage, N., Andreassen, E., Lazarov, B. S., & Sigmund, O. (2017). Giga-voxel computational morphogenesis for structural design. *Nature*, 550(7674), 84-86. <https://doi.org/10.1038/nature23911>

Published in:
Nature

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Giga-voxel computational morphogenesis for structural design

Niels Aage¹, Erik Andreassen¹, Boyan S. Lazarov¹ and Ole Sigmund¹

¹Department of Mechanical Engineering, Technical University of Denmark,
Nils Koppels Allé, Building 404, 2800 Kgs. Lyngby, Denmark

Intelligent distribution of material determines performance and cost of industrial products from hearing aids to automobiles and airplanes. Historically, intuition and insight have driven mechanical design evolution, recently assisted by Computer Aided Design (CAD) approaches. A compelling synthesis approach allowing for unrestricted design freedom, called topology optimization, is showing great promise with regards to weight savings but its applicability has so far been limited to component design or simple structures (1,2). Here, we report on a supercomputer-based morphogenesis tool that allows giga-voxel resolution – more than two orders of magnitude higher than previously reported. Such resolutions provide insight into optimal structures, hitherto unachievable due to substantial challenges in the upscaling of existing modelling and optimization frameworks. The suggested morphogenesis process is generally applicable, not only to mechanical design, but also to flow systems (3), antennas (4), nano-optics (5) and micro-systems (6,7). As an illustrative example, the procedure is applied to full-scale airplane wing morphogenesis. An optimized full aircraft wing structure, as illustrated in Fig. 1, exhibits unprecedented structural detail at several length scales; intriguingly, it shows remarkable similarities to naturally occurring bone structures in, for example, bird beaks. An estimate of the weight savings predict a 2-5% mass reduction, which can be translated into a reduction in fuel consumption in the order of 40-200 tons of jet fuel a year for the aircraft example.

Computational morphogenesis is the engineering art of determining the best possible shapes and material distributions for prescribed structural objectives. A very common goal is to minimize structural weight subject to constraints on deflections and mechanical stresses. Similar goals are drivers for animal and plant evolution. Optimal loading capacity to weight ratio ensures efficient use of limited resources (plants) and minimizes energy consumption for walking and running (animals). Whereas nature has evolved efficient structures through tedious genetic evolution, engineering demands fast solutions that also take manufacturing limitations into account. Topology optimization (1,2) works by redistributing material within a predetermined design domain, not unlike natural bone-growth processes. Here, the design domain is defined as the geometric volume where material can be distributed, for example, the space above the building lot for a high-rise building or the internal part of a car, minus the space reserved for passengers, engine and other geometrical entities. Through an iterative, deterministic process that involves modeling of the physical system, gradient computations and mathematical programming-based design updates, structural material is gradually redistributed towards the optimal design as illustrated in Fig. 2a and Extended Data Fig. 1.




Figure 1: Optimized wing structure. *Birds-eye views of the result of the giga-voxel morphogenesis process applied to full-scale airplane wing design after 400 steps of the morphogenesis procedure. Fixed upper skins have been removed to reveal the internal details. Note that, apart from the thin fixed outer skin, no a priori assumptions were made on the internal geometry of the wing structure. Hence, all the intricate details such as curved spars, truss and wall structures that can be observed have appeared spontaneously as a result of the optimization process.*

Traditional stiffness-based topology optimization of components is now very well developed and is used routinely in all major mechanical engineering industries, including automotive (8) and aerospace (9). The automotive industry typically achieves weight savings of 20-40% on structural parts, reduced development time as well as elimination of needs to build prototypes. Despite its successes, applicability of the topology optimization approach is still limited to design of components or smaller structures. The major obstacle in applying it to large structures is resolution. Just as the visual quality of a TV-screen is limited by the number of pixels, the geometric quality and capability of the topology optimization process is limited by the number of voxels (3D equivalent of pixels) used to parameterize the design problem. Current state-of-the-art implementations limit voxel numbers to a few millions and hence render applications to detailed high-rise building, oil-rig or large airplane design impossible. This provokes the question: “Will giga-voxel resolution computational morphogenesis lead to radical design changes?” If so, significant weight-savings and reduced environmental impact will follow.

We investigate this question using high performance computing and demonstrate the potential of giga-scale morphogenesis on a full-scale wing design problem for a Boeing 777 type aircraft shown

in Fig. 1; Extended Data Figs. 2 and 3 as well as Supplementary images 1 and 2 (high-resolution and zoomable). The extension to solve problems with more than one billion voxels required overcoming several non-trivial challenges. These involved efficient solution of huge linear problems, as well as tailoring of optimization algorithms; data transfer and visualization (see Methods). Furthermore, obtaining the optimized designs presented here (for now) requires access to massive computational resources with run times of 1-5 days on 8,000 CPUs.

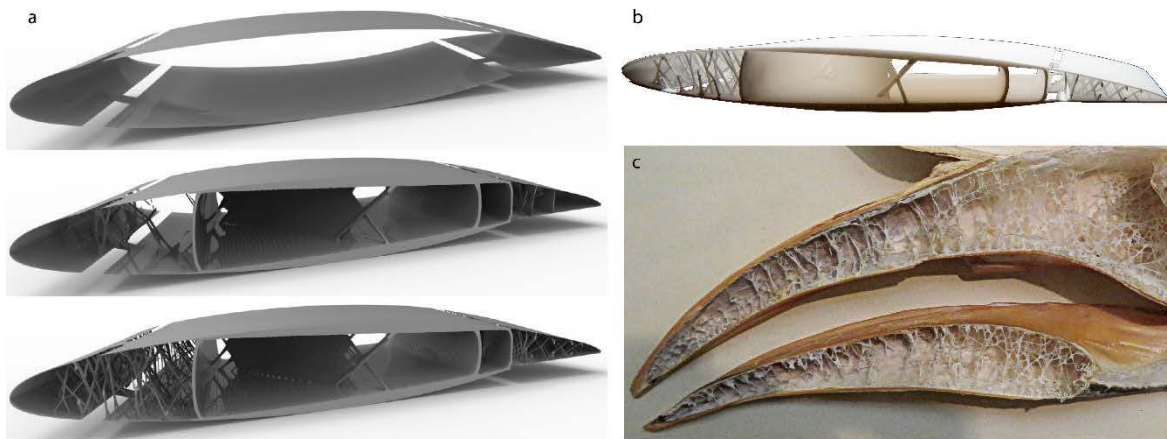


Figure 2: Morphogenesis evolution, 3D print and comparison to the hornbill bird beak. *a*, Steps of the morphogenesis process of the full wing structure demonstrated for a slice around the third connector from the wing root (see Extended Data Fig. 2 for precise location). The top figure corresponds to the initial design, where only the outer airfoil skin is given. The middle figure shows a snap shot during the morphogenesis process, ending in the final design shown at the bottom. *b*, A piece of the wing printed using additive manufacturing reveals additional detail and insight into the intricate but very regular structure of the optimized wings. The dimensions of the wing section are $300 \times 100 \times 45 \text{ mm}^3$. *c*, Hornbill bird beak with internal structure similar to the optimized airplane wing. Picture taken by Jessica Miller-Camp in NHMUK exhibits.

The case study considers the design of the load-carrying internal structure for a full airplane wing based on the publicly available NASA Common Research Model (CRM) (10, 11). This wing model is similar to a Boeing 777 and has a (half) wing span of 27 m and an inner and outer chord of 12m and 1.5m, respectively. The fixed-shape wing is discretized by 1.1 billion voxels (corresponding to a mesh of $1,216 \times 3,456 \times 256$ elements) with a maximum element size of 0.8 cm at the root, which is reasonable for a wing structure this size (12). From the NASA research model we have included two aerodynamic load scenarios corresponding to angles of attack of 0 and 4 degrees at Mach 0.85 and an altitude of 10,700 m. The outer layer of voxels is fixed to be solid, representing the airfoil outer skin and hence material can be freely distributed in the internal part of the wing with no further restrictions, see Fig. 2a(top). The optimized design considering a simplified full wing structure including flaps is shown in Fig. 1. The top-skin has been visually removed in order to reveal the optimized interior structures (the wing without the top-skin removed is shown in

Extended Data Figs. 2). Remarkably, all internal structural details have appeared spontaneously as a result of the morphogenetic process. An immediately obvious and striking result is the extreme amount of intriguing details and features that come out of the giga-voxel design resolution studies. It is also clear that expected structural features do not occur; the structure does not show standard straight wing box features made up of two spars, straight ribs combined with reinforced bottom and top skins. Instead, upon closer inspection and study of cross-sectional plots (Fig. 2a), spars of the wing box reveal to be curved to optimize torsional stiffness (the optimal pure torsion-rod cross-section is circular). Furthermore, the irregular wing box is reinforced by intriguing curved as well as diagonally placed ribs, which are not seen in conventional airplane wing designs. A conservative estimate of the weight savings achievable is 2-5%, corresponding to 200-500 kgs per wing. Converting this weight reduction into fuel savings, indicates potential saving of 40 to 200 tons of fuel per year is possible. The analysis behind these predictions can be found in the Method section available online. The unconventional structural details provide new insight and will help aircraft designers in developing lighter and stronger wing designs in the future.

The outcome of the optimization process depends strongly on the choice of loading cases. To illustrate this aspect we investigate the influence of varying loading scenarios in Extended Data Fig. 3. Optimizing for single aerodynamic loads result in open-walled wing boxes, see Extended Data Figs. 3a and b, and hence designs that have very low torsional stiffness. On the other hand, optimizing for two (or more) load situations results in closed and torsionally stiff wing boxes, see Extended Data Figs. 3c and d. In practice, engineers will first optimize for what is expected to be the most critical load cases. Subsequently, the resulting wing structure is checked for all design load cases and if it turns out that some load cases, not considered originally, are critical for the new design, they will be included in a new optimization run.

Comparing the presented results with the state-of-the-art in engineering design, we remind that currently available methods and software have limited design resolutions of approximately five million voxels and hence a feature size of 5 cm for the considered wing example. This is too large to capture typical wing features and hence, design studies based on coarse discretizations produce blurred and unclear designs with no spar or rib type formations. This issue is discussed in previous works on full-scale wing structure optimization, e.g. (9), (13), (14), (15). Ref. (15) attributes the lack of appearance of classical features such as ribs and spars either to: inappropriate physics; too few load cases and incorrect boundary conditions; or that the optimization problem formulation itself is incomplete. In contrast, based on our game-changing high-resolution study, we suggest that the real problem stems from lack of resolution. With a giga-voxel framework we can now solve engineering design challenges that hitherto were considered insurmountable.

The optimized wing designs from this first giga-voxel resolution case study are subject to limitations and manufacturing challenges. In the analyses we only included linear elasticity and have left anisotropic and non-linear materials, dynamic effects, buckling, aero-elasticity, etc. for future studies. Also, the fabrication of the complex design is currently infeasible, although one could imagine that large scale additive manufacturing will be a possibility in the future, which we illustrate by the small section print shown in Fig. 2b. Nevertheless, the designs are of huge interest

to structural engineers, who, for the considered case study, will get inspiration for future simplified designs, especially based on the intriguing curved and diagonal ribs seen in the optimized designs. Surely, the optimizer is giving the design engineer something to think about!

As a last remark to the considered case study, it is striking to see how the optimized designs resemble natural bone structure like the Hornbill bird beak structure shown in Fig. 2c. This is not a coincidence. As Wolff's law states that animal bone adapts to applied loads (16), computational morphogenesis works in a similar way. Both structures have closed outer shells (respectively for aerodynamic and food containment purposes) and need to be stiff towards longitudinal bending and applied surface pressure, hence resulting in similarly looking internal reinforcement structures. In other words; the outcome of the computational morphogenesis approach can be compared to the outcome of nature's evolutionary design processes, but with a major difference - it happens in a small fraction of the time due to the efficient, deterministic optimization approach applied.

The methodology presented is directly applicable to similar morphogenesis problems in other engineering disciplines as well as in architecture and industrial design. Structures relevant to this approach are ships, offshore structures, wind turbine blades and towers, power masts, high-rise buildings, bridges and other complex built structures. Giga-voxel resolution capability is expected to revolutionize design outcomes and yield new insight in optimal design. Computing time and resources reported in this paper provide a severe bottleneck and will, for the time being, limit the use to early design stages in research intensive (large) companies and academia. However, procedures for reducing the number of degrees of freedom by model reduction (17, 18), reducing the number of iterations, and increasing the granularity by deeper hierarchy (19), are being developed and are expected to result in significant reduction in computing requirements in the near future. Additionally, improvements are needed in peripheral computational tools such as meshing algorithms and visualization, as discussed in the Methods section.

Acknowledgements:

This work was funded by the Villum Foundation through the NextTop project and a PRACE grant TopWING giving access to the Curie supercomputer (GENCI@CEA, France). Access to, and efficient support from, the technical staff at Curie is highly appreciated. We also acknowledge access to and support from the Visualization Cluster at Copenhagen University through Troels Haugbølle and Åke Norlund, as well as discussions with Andy Horsewell and Jon Juel Thomsen from the Technical University of Denmark.

Author contributions:

N.A. contributed with the original idea, method development, implementation, supercomputing, visualization, renderings and manuscript preparation.

E.A. contributed with the original idea, method development, implementation, visualization and manuscript preparation.

B.S.L. contributed with mesh mapping and manuscript editing.

O.S. contributed with the original idea, method development, analytical studies and manuscript preparation.

Author information:

The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to Niels Aage (naage@mek.dtu.dk).

METHODS

General. Amongst a multitude of topology optimization methodologies available, we work with the so-called density approach. This section first describes the basic methodology and secondly, the details of extending this approach to the giga-voxel resolution regime.

Topology optimization. The structural design problem considered in this work, i.e. stiffness maximization, is solved based on an assumption of static, linear elastic behavior. Hence, the governing physics correspond to the linear Navier-Lamé equations with appropriate boundary conditions. The partial differential equation is solved by the finite element method, using 8-node hexahedral elements, which results in a linear system of equations

$$\mathbf{K}(\mathbf{r})\mathbf{u} = \mathbf{F}$$

Here $\mathbf{K}(\mathbf{r})$ is the stiffness matrix with properties depending on spatial coordinates \mathbf{r} , \mathbf{u} the displacement vector and \mathbf{F} the load vector. The stiffness dependence on the spatial coordinate is introduced to facilitate the structural optimization procedure presented next. That is, by smoothly interpolating between a reference solid (here aluminum with Young's modulus $E_{solid} = 70$ GPa) and a very weak material ($E_{void} = 10^{-9}E_{solid}$) using the so-called SIMP (1, 2) (Solid Isotropic Material with Penalization) model, we are able to use gradient-based optimization methods to solve the stiffness maximization problem. In the following we assign a single density variable to each element in the mesh and state the SIMP interpolation scheme as

$$E(\rho) = E_{void} + \rho^p(E_{solid} - E_{void})$$

Here $\rho \in [0,1]$ is a continuous, relative density and $p \geq 1$ is a penalization parameter that enforces almost discrete final designs with ρ being either 0 or 1 through penalization of the stiffness for intermediate densities. The final ingredient, before stating the mathematical optimization problem, is to introduce compliance (C) as the objective function for structural design. Compliance is the work done by the external forces, i.e. $C = \mathbf{F}^T \mathbf{u}$, and is inversely proportional to structural stiffness. Hence, to maximize structural stiffness we pose the following minimization problem

$$\begin{aligned}
& \min_{\mathbf{x} \in \mathbb{R}^n} \phi = \sum_i^L \mathbf{F}_i^T \mathbf{u}_i \\
s. t. & \quad \mathbf{K}(\mathbf{x}) \mathbf{u}_i = \mathbf{F}_i, \quad i = 1 \dots L \\
& \quad \frac{V(\mathbf{x})}{V^*} - 1 \leq 0, \\
& \quad 0 \leq x_j \leq 1, \quad j = 1 \dots n
\end{aligned} \tag{1}$$

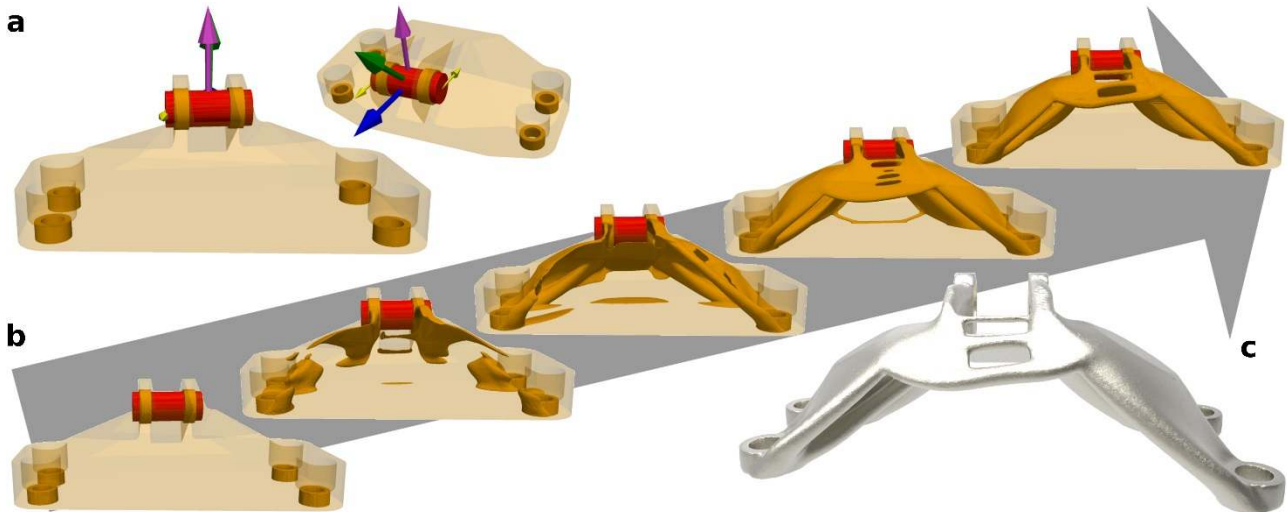
Here \mathbf{x} is a vector of mathematical design variables with length corresponding to the number of design elements and ϕ is a sum of multiple compliances arising from L loading scenarios. The three constraints are as follows; the first one ensures mechanical equilibrium, the second poses a limit on the amount of available material with $V(\mathbf{x})$ being the volume of the current structure and V^* the total amount of available material and the last statement constitutes box constraints on the design variables. Due to inherent issues with numerical stability and mesh dependence (2), we use a filtering technique to relate the mathematical design variables x to the physical density ρ (20,21). The filter operation is a convolution type operator that can be stated for the discretized model as

$$\rho_e = \frac{\sum_{i \in N_e} w(r_i) v_i x_i}{\sum_{i \in N_e} w(r_i) v_i}$$

Here $N_e = \{i \mid \|r_i - r_e\| \leq R\}$ is a neighborhood set with R being the filter radius, r_i is the spatial (center) coordinate of element i , $w(r_i) = R - \|r_i - r_e\|$ the weighting function and v_i the element volume. The sensitivities of the objective function are determined efficiently without additional analysis using the adjoint method and the filtering operation is included by the chain rule (22). The optimization problem is solved using a nested approach and an optimality criteria algorithm as described in (2, 23). Due to the massive increase in computational power over the past decades, the design problem in Eq. 1 can now be solved interactively even on hand-held devices such as smartphones or tablets (download from www.topopt.dtu.dk) for small numbers of n , e.g. up to several thousands of design variables.

Example problem. To demonstrate the iterative procedure of the topology optimization method we present the solution to a simple part-design problem in Extended Data Fig. 1. We consider the design of a jet-engine bracket used in the 2013 GrabCAD challenge. Extended Data Fig. 1a shows the design domain together with supports and a number of applied mechanical loading cases. Extended Data Fig. 1b shows snapshots of the material distribution during the 200 step morphogenesis process, starting from a blank design.

An animation of the iteration process can be found as Supplementary video 1. The optimized design, obtained by topology optimization, saves an additional 10% in weight compared to the winner of the GrabCAD challenge that had 638 world-wide participants. Interestingly, even though we only optimized for compliance, the resulting design easily satisfies the strength requirement set in the completion. This is partly due to the filtering process that not only enforces minimum length scale, but also avoids sharp corners and thereby stress concentrations. The example demonstrates that compliance-based topology optimization generally, but not always, also results in structures with favorable strength response. Remarkably, but not surprisingly, the lack of initial geometric



Extended Data Figure 1: Schematic of the morphogenesis process. *Design of a jet engine bracket by morphogenesis. The design domain and the four load scenarios are shown in a) and b) shows snapshots from the design evolution. c) Rendering of the final titanium bracket design.*


constraints and bindings to CAD-parameterizations result in the morphogenesis process giving rather organic looking structures that provide optimal force paths from applied forces to the supports, with features reminiscent of human and animal bone structure.

Meshing for giga-voxel resolution. Generating a body fitted mesh with more than one billion finite elements is not a trivial task for standard mesh tools and hence one has to resort to alternative procedures. In this work we have used a structured, logical mesh which is then mapped, slice-by-slice, to the geometry of the NASA common research model. We note that a shell-element discretization of the outer shell would increase accuracy of the model; however, based on experience and tests of various element types, we do not expect that a mixed solid/shell formulation will change the resulting morphologies in any significant way.

Ensuring structural details. Due to the non-convexity of the stiffness penalized optimization problem in Eq. 1, any gradient based solution method is likely to end in local minima. In order to ensure that the produced designs are of high quality, i.e. strong local minima, we therefore resort to a continuation strategy for the penalization parameter in the SIMP interpolation. This is motivated by two aspects. First, we note that Eq. 1 with $p = 1$ poses a convex problem. However, without stiffness penalization it will not result in pure 0-1 solutions and thus intermediate densities will dominate the optimized design. Secondly, it is of key interest that this study produces optimized designs with a high level of structural details. Starting the optimization procedure with e.g. $p = 3$ will quickly force the design into a 0-1 structure, but with very little and crude detail. Both of these issues are alleviated by a continuation strategy where the penalization parameter is slowly raised in steps of 0.25 from 1 to 4, distributed over a total of 400 design cycles.


Giga-scale solution procedure. By far the most costly and challenging part of the topology optimization process is the solution of the huge linear equation systems originating from the finely discretized finite element problem. For a one billion element discretization, the linear systems have

more than three billion unknowns. Despite sparsity, they cannot be solved using direct solvers due to storage and memory requirements and require iterative approaches that make use of efficient preconditioners. Construction of a powerful and robust linear solver and preconditioner is hard for several reasons. First of all, the equation system is ill-conditioned due to high stiffness contrast. Secondly, depending on the design problem, high aspect ratios of the design domain and ill-formed elements may challenge convergence further. Thus, it is experienced that solution strategies that work well for coarser problems are inadequate when the problem is scaled up (24). The solution methodology proposed next is based on the large-scale topology optimization framework presented in (25).



Extended Data Figure 2: Wing design model problem. *Rendering of full wing model without removal of top skins. All load cases use symmetry conditions at the root (orange color) with only the top-most layer of elements (the skin) being fully fixed (red color). The bottom right inserts show the pressure coefficients of the two aerodynamic load cases (adapted from Fig. 10 in ref. (11)) and the top left insert indicates the position and direction of the simulated engine loads for the third load case. The section highlighted with blue marks the cut-out region used in Figs. 2a and 2b.*

In our work we have applied Krylov subspace equivalent methods as the overall tool for both solver and preconditioning steps. These are chosen due to their explicit dependence on (sparse) matrix-vector products and hence, very efficient parallel computer implementation. The key to a robust and efficient linear solver lies in the construction of a good preconditioner, i.e. one that is fast to compute and fast to apply. This is an art and requires long development times, often not reported in articles (26). The present methodology is built as an add-on to the numerical library PETSc (27). The result is a very robust multi-level preconditioner strategy that uses Galerkin projected multigrid with successive over-relaxation preconditioned Chebyshev smoothers down to a coarse level of approximately 6 million degrees of freedom using four geometric multigrid levels. The resulting coarse grid problem is then solved with an algebraic multigrid preconditioned conjugate gradient method. As the outer solver we use the flexible generalized minimum residual approach (28) due to the slightly varying preconditioner arising from the smoothers. The ultra-large scale problems considered require massive computer storage capability and visualization is also a challenge when the designer wishes to inspect the results of the morphogenesis process.

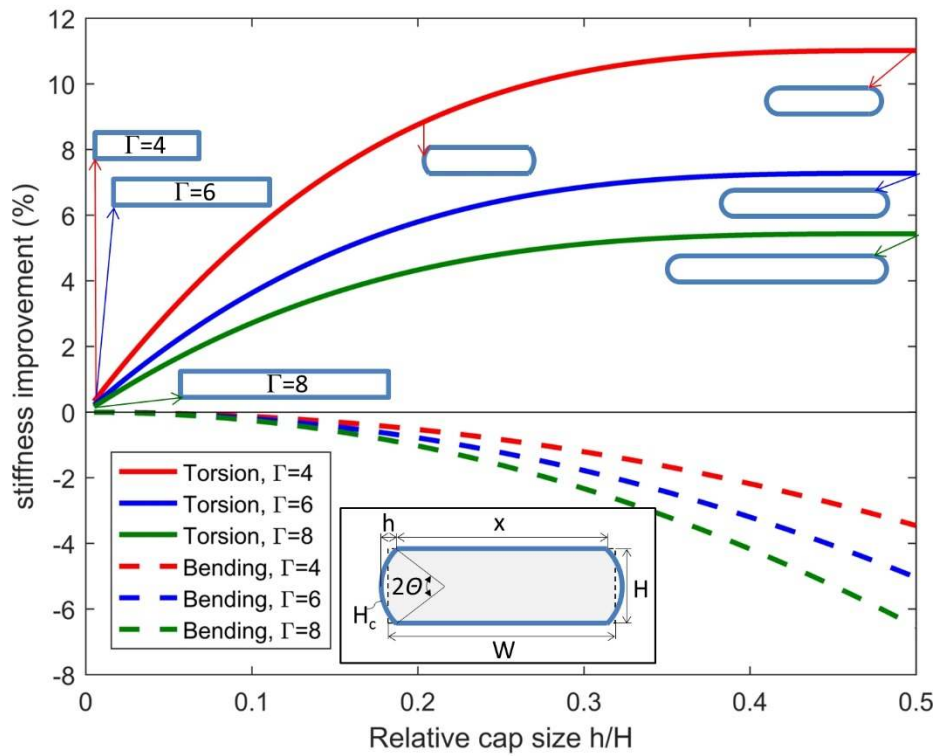


Extended Data Figure 3: Influence of multiple load cases. *Top view of the common research model with optimized wing box structures for varying loading cases. In the center, a standard wing box design with spars and ribs and optimized designs for; a) a single load case at 0 degree incidence, b) a single load case at 4 degree incidence, c) a two load case design for 0 and 4 degree incidence and d) a three load case design including engine load and two aerodynamic loads at 0 and 4 degree incidence. The load cases are shown in Extended Data Fig. 2.*

Computations. The giga-scale optimization problem, and the presented solution strategy, requires a massive amount of computational resources. Therefore all results presented in the work have been generated on the Curie Supercomputer (GENCI@CEA) in France. The optimal amount of processors capable of solving the problem was found by numerical experiments and memory requirements to be around 8,000 processors.


Estimate of potential weight savings. The wing study presented is qualitative. To come up with an exact figure of weight savings would require more detailed studies and tens of thousands of engineering hours for one specific aircraft. However, some simpler, partly analytical studies can be performed to get an estimate of potential weight savings. The topology optimization process resulted in two major geometrical findings. First, wing box spars are rounded instead of straight and second, the few appearing wing ribs represent odd geometries represented by curved and diagonally placed ribs.

First, we perform an analytical study to estimate the stiffness improvement (or weight saving) achievable by rounded instead of straight wing spars. Here the wing box is assumed to be straight, with constant wall thickness and rectangular cross section with aspect ratio $\Gamma=L/H$ (for definition of parameters see insert in Extended Data Fig. 4). The torsional stiffness of the wing box is proportional to the ratio of the squared cross sectional area and its perimeter (this is why a circular cross section represents the optimal torsion rod). The torsional stiffness can hence be improved by substituting the straight spars with circular ones. Assuming constant perimeter the relative torsional stiffness improvement for a cross section with circular end caps of relative width h/H can be found



Extended Data Figure 4: Curved spars. *Stiffness improvement in percent as function of curved spar dimensions ($h/H=0$ means straight spar and $h/H=1/2$ means half-circular spar). Red, blue and green lines correspond to torsion boxes with aspect ratios 4, 6 and 8, respectively. Solid lines indicate increased torsion stiffness and dashed lines (decreased) bending stiffness. Insert shows parameterization of the wing box cross section.*

analytically, resulting in a simply derivable but complicated expression. The resulting stiffness improvement for three different wing box aspect ratios is shown in Extended Data Fig. 4. It is seen that the improvement is largest for smaller aspect ratios (where the torsional stiffness difference between a circular and a square cross section is the largest). For all aspect ratios, cross sections with half circular ($h/H=1/2$) end caps are the stiffest. The improved torsional stiffness comes at the cost of decreased bending stiffness. The curved shear webs (spars) are less efficient in carrying shear loads. This time, again assuming constant material volume, a circular end cap results in a longer and thinner shear web resulting in a significant stiffness decrease. The analytical study for bending assumes a thin walled prismatic wing box with height to span length ratio of 20. The optimal spar curvature for the full wing structure thus represents a compromise between increase in torsional stiffness and decrease in bending resistance. Estimating cap sizes from the wing design studies to be $h/H=0.12$, and the wing box aspect ratio to be around $\Gamma=6$ (see cross-sectional plot in Fig. 2a(bottom) and Fig. 2b), expected torsional stiffness improvement or similar weight saving is 4% (see Extended Data Fig. 4). At the same time the loss in bending stiffness is only 0.5%. Depending on actual distribution between bending and torsional forces, the increase in stiffness or corresponding weight saving may be 1-4%. Considering that a full airplane wing may weigh 10 tons [29] and that 100-200 kg fuel may be saved per kg saved weight per year for a commercial airplane [30], this is a significant number – i.e. even a small percentage may have a huge influence.



Extended Data Figure 5: Non-traditional ribs. *Simplified wing box models used to estimate stiffness gain from non-traditional wing rib geometries. Both images show the optimized wing box for two aerodynamic load cases, c.f. Extended Data Fig. 3c, in grey. The optimized design is overlaid with blue (left): wing box with conventional straight ribs and red (right): wing box with unconventional ribs. Thicknesses of ribs are tailored to ensure equal mass.*

Second, we perform a numerical study on a crudely simplified wing box model in order to estimate the stiffness improvements achieved by curved and diagonal ribs placements compared to conventional straight ribs. As a basis for the study, we use the optimized wing box shown in Extended Data Fig. 3c as inspiration for an engineering interpretation. The result of this post processing step is shown in Extended Data Fig. 5, where it should be noted that the two reinforced wing box models are fitted to have the the same mass. The study is performed on a mesh consisting of 141 million elements (i.e. 430 million dofs), which is fully capable of capturing the details of the post processed wing box designs, and the numerical experiment is conducted using 640 cores on a small cluster. To quantify the performance improvement of the optimized rib/spar design compared to the conventional rib/spar design, we subject both models to the two aerodynamical load scenarios as described and depicted in Extended Data Fig. 2, i.e. the same two load cases used in the initial topology optimization. The numerical results are collected in Extended Data Table 1. From the results in Extended Data Table 1, we observe that even a crude interpretation of the optimized wing box design leads to a performance improvement of 0.8-0.9% for the two load scenarios.

Load scenario:	Conventional design [J]	Optimized Design [J]	Improvement
Load case 1:	511479	506756	0.92%
Load case 2:	15112593	14987090	0.83%

Extended Data Table 1. Performance of non-traditional vs. traditional rib structures. *Table showing the compliances for two aerodynamic load cases for the post processed designs presented in Extended Data. Fig. 5. The improvements obtained by the curved ribs are stated in percent.*

Finally, we can summarize the possible weight savings identified through the study of two of the main features observed in the optimized wing structure. We recall that the weight of both wings is

in the order of 20 tons, and that the presented work indicates that a total weight reduction of 2-5% is possible using both the curved spars and the novel rib layout. With regard to weight reduction, this corresponds to a total saving between 400kg and 1,000kg, corresponding to an estimated reduction in fuel consumption of 40 to 200 tons a year [29,30].

Data and code availability. The basis code used for the presented work is publically available at https://github.com/topopt/TopOpt_in_PETSc (25) and the results, i.e. the optimized designs represented by STL-files, are available from the corresponding author on reasonable request.

References:

1. Bendsøe, M. P., Kikuchi, N., Generating optimal topologies in structural design using a homogenization method, *Computer Methods in Applied Mechanics and Engineering* **71**, 197-224 (1988).
2. Bendsøe, M. P., Sigmund, O., *Topology Optimization - Theory, Methods and Applications* Springer Verlag, Berlin Heidelberg (2004).
3. Alexandersen, J., Sigmund, O., Aage, N., Large scale three-dimensional topology optimisation of heat sinks cooled by natural convection, *International Journal of Heat and Mass Transfer* **100**, 876-891 (2016).
4. Nomura, T. , Sato, K., Taguchi, K., Kashiwa, T., Nishiwaki, S., Structural topology optimization for the design of broadband dielectric resonator antennas using the finite difference time domain technique, *International Journal for Numerical Methods in Engineering* **71**, 1261–1296 (2007).
5. Jensen, J. S., Sigmund, O., Topology optimization for nano-photonics, *Laser & Photonics Reviews* **5**, 308-321 (2011).
6. Sardan, O., Eichhorn, V.; Petersen, D. H.; Fatikow, S.; Sigmund, O., Bøggild, P., Rapid prototyping of nanotube-based devices using topology-optimized microgrippers, *Nanotechnology*, **19**, 495-503 (2008)
7. Dühring, M., Mortensen, N.A., Sigmund, O., Plasmonic versus dielectric enhancement in thin-film solar cells, *Applied Physics Letters*, **100**, 211914 (2012).
8. Cavazzuti, M., Baldini, A., Bertocchi, E., Costi, D., Torricelli, E., Moruzzi, P., High performance automotive chassis design: a topology optimization based approach, *Structural and Multidisciplinary Optimization* **44**, 45–56 (2011).
9. Zhu, J., Zhang, W., Xia, L., Topology optimization in aircraft and aerospace structures design, *Archives of Computational Methods in Engineering*, **23**, 4, 1–28 (2015).
10. Vassberg, J. , DeHaan, M., Rivers, S., Wahls, R., Development of a common research model for applied CFD validation studies, *AIAA Paper* 6919 (2008).
11. Bell, J. H., Pressure-Sensitive Paint Measurements on the NASA Common Research Model in the NASA 11-ft Transonic Wind Tunnel, *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* 1128 (2011).
12. Kennedy, G. & Martins, J., A comparison of metallic and composite aircraft wings using aerostructural design optimization, *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 5475 (2012)
13. Rao, J., Kiran, S., Kamesh, J., Padmanabhan, M., Chandra, S., Topology Optimization of Aircraft Wing, *Journal of Aerospace Science and Technologies* **61**, 402-414 (2009).
14. Chin, T., Kennedy, G., Large-Scale Compliance-Minimization and Buckling Topology Optimization of the Undeformed Common Research Model Wing, *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 0939 (2016).

15. Stanford, B., Dunning, P., Optimal topology of aircraft rib and spar structures under aeroelastic loads, *Journal of Aircraft* **52**, 1298-1311 (2015).
16. Wolff, J. The Law of Bone Remodeling (translation of the German 1892 edition), *Berlin, Heidelberg, New York: Springer* (1986)
17. Hesthaven, J., Rozza, G., Stamm, B., Certified Reduced Basis Methods for Parametrized Partial Differential Equations, Springer (2015).
18. Groen, J., Sigmund, O. Homogenization-based topology optimization for high-resolution manufacturable micro-structures, *International Journal for Numerical Methods in Engineering*, online first (2017).
19. Alexandersen, J., Lazarov, B., Topology optimisation of manufacturable microstructural details without length scale separation using a spectral coarse basis preconditioner, *Computer Methods in Applied Mechanics and Engineering* **290**, 156-182 (2015).
20. Bourdin, B. Filters in topology optimization, *International Journal for Numerical Methods in Engineering*, **50**, 2143-2158 (2001)
21. Bruns, T. E. & Tortorelli, D. A., Topology optimization of non-linear elastic structures and compliant mechanisms, *Computer Methods in Applied Mechanics and Engineering*, **190**, 3443-3459 (2001)
22. Tortorelli, D. A., Michaleris, P., Design sensitivity analysis: overview and review, *Inverse Problems in Engineering* **1**:71-105 (1994).
23. Sigmund, O., A 99 line topology optimization code written in Matlab, *Structural and Multidisciplinary Optimization* **21**, 120-127 (2001). MATLAB code available online at: www.topopt.dtu.dk.
24. Evgrafov, A., Rupp, C. J., Maute, K., Dunn, M. L., Large-scale parallel topology optimization using a dual-primal substructuring solver, *Structural and Multidisciplinary Optimization* **36**, 329-345 (2008).
25. Aage, N., Andreassen, E., Lazarov, B.S., Topology optimization using PETSc: An easy-to-use, fully parallel, open source topology optimization framework, *Structural and Multidisciplinary Optimization* **51**, 565-572 (2015).
26. Balay, S., et al., PETSc Users Manual, Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory (2015).
27. Amir, O., Aage, N., Lazarov, B.S., On multigrid-CG for efficient topology optimization, *Structural and Multidisciplinary Optimization* **49**, 815-829 (2014).
28. Saad, Y., Iterative Methods for Sparse Linear Systems, SIAM (2003).
29. Kenway, G. K. W., Kennedy, G. J., & Martins, J. R. R. A. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA Journal*, **52**(5), 935-951 (2014).
30. Helms, H., & Lambrecht, U. The potential contribution of light-weighting to reduce transport energy consumption. *International Journal of Life Cycle Assessment*, **12**, 58-64 (2007).

Supplementary Data:

A video showing the iteration history for the jet-engine bracket design problems (Supplementary video 1) and two high resolution and zoomable images (Supplementary image 1 and 2) of the design in Fig. 1.