Implementing a structural continuity constraint and a halting method for the topology optimization of energy absorbers

Published in:
Structural and Multidisciplinary Optimization

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
Link to publication record in Queen's University Belfast Research Portal

Implementing a structural continuity constraint and a halting method for the topology optimization of energy absorbers

Daniel Stojanov · Brian G. Falzon · Xinhua Wu · Wenyi Yan

Abstract This study investigates topology optimization of energy absorbing structures in which material damage is accounted for in the optimization process. The optimization objective is to design the lightest structures that are able to absorb the required mechanical energy. A structural continuity constraint check is introduced that is able to detect when no feasible load path remains in the finite element model, usually as a result of large scale fracture. This assures that designs do not fail when loaded under the conditions prescribed in the design requirements. This continuity constraint check is automated and requires no intervention from the analyst once the optimization process is initiated. Consequently, the optimization algorithm proceeds towards evolving an energy absorbing structure with the minimum structural mass that is not susceptible to global structural failure. A method is also introduced to determine when the optimization process should halt. The method identifies when the optimization method has plateaued and is no longer likely to provide improved designs if continued for further iterations. This provides the designer with a rational method to determine the necessary time to run the optimization and avoid wasting computational resources on unnecessary iterations. A case study is presented to demonstrate the use of this method.

D. Stojanov
Department of Mechanical and Aerospace Engineering,
Monash University, Clayton 3800, Australia.

B. G. Falzon
School of Mechanical and Aerospace Engineering,
Queen's University Belfast, Belfast, UK

X. Wu
Department of Materials Engineering,
Monash University, Clayton 3800, Australia.

W. Yan
Department of Materials Engineering,
Monash University, Clayton 3800, Australia.
Tel.: +61 3 9902 0113
Fax.: +61 3 9905 1825
E-mail: wenyi.yan@monash.edu

1 Introduction

A common goal for load-carrying structures is to make them as light as possible. For an energy absorbing structure there is an additional requirement that it absorb the loads specified by the design conditions. In the case of typical metallic structures the energy from a dynamic load is primarily absorbed through plastic deformation of the material of that structure. Too much material deformation will lead to fracture and failure. When designing energy absorbing structures that are efficient with respect to their mass it is necessary that they have enough mass to ensure that plastic deformation can occur without catastrophic fracture and failure, but not with more mass than necessary to meet the design requirements.

Common types of structures used in energy absorbers include thin-walled tubes, shells, honeycomb materials, and foams (Lu and Yu (2003)). The design and optimization of energy absorbing structures has received considerable research interest in recent years. These include the optimization of truss structures for energy absorption (Pedersen (2003)); axially loaded tubes (Avalle et al (2002)), where the range between the lowest and highest crushing forces experienced were minimised by optimizing the parameters that defined the tube geometry; and, functionally graded foams (Cui et al (2009)) capable of minimising peak stresses by spreading the load over a larger time interval.

Topological optimization, which seeks to optimize the distribution of material with a given set of design constraints, to meet one (or more) objectives, has been applied to maximising stiffness (Chu et al (1996)); optimizing the stiffness of a mechanically controlled structure (Ou and Kikuchi (1996)); maximising heat conduction (Li et al (2004)); maximising the resonant frequency (Du and Olhoff (2007)); and, maximising the critical load of buckling structures (Neves et al (1995)).

Topology optimization accounting for geometric and material non-linearity has been used for the optimization of truss (Pedersen (2004)) and continuum structures (Mayer et al (1996); Jung and Gea (2004); Huang et al (2007)) for improved crashworthiness. This includes continuum structures subjected to short duration transient loads simulated using an explicit finite element (FE) solver (Forsberg and Larsgunnar (2007)). There has also been use of topology optimization methods for energy absorbing structures made of hybrid materials (Jung and Gea (2006)). An algorithm to halt the optimization process can make topology optimization methods more useful by supplementing/replacing the analyst's judgment with an algorithmic process to determine the point in the optimization process after which further optimization will not be useful toward generating an improved design. This problem is also related to finding the lowest feasible volume ratio[1] for a design and is often not dealt with in the literature; instead, a final volume ratio is often selected arbitrarily at the beginning of the optimization. The mass of the final structure is centrally important to the primary goal of the analyst and it is therefore inadequate to simply pick the final volume ratio arbitrarily before the optimization process has begun. The optimization method presented will seek to find the lowest volume ratio that still meets the requirement that it avoid failure. An automated method should decide to halt when further optimization will not be useful because it will not be likely to produce significant further improvements. Strategies addressing when the optimization should halt have included running the algorithm without a halting strategy, then selecting from the most suitable design that was created during the optimization process (Xie and Steven (1993)). Most often the final volume ratio has been pre-selected at the beginning of an analysis, then the algorithm allowed to run until that volume ratio is reached, or once reached, to continue an arbitrary number of iterations at this volume ratio to achieve an improved result. Sometimes this relationship is specified indirectly by an upper limit on a performance criterion, which leads to the algorithm halting when it is exceeded. One criterion is the maximum stress allowed in any element (Duysinx and Bendsøe (1998)). More sophisticated performance criteria have been presented in the literature to rank the results of an optimization process. These have included the method of extended optimality (Rozvany et al (2002)), in which the design was measured by the ratio between the progress of the designs toward improving the optimization objective, and the volume ratio. Another example was a performance index (Querin (1997)), which was also a single scalar quantity that assessed an optimized design against a performance criterion. A criterion that has been used with the Bi-Directional Evolutionary Structural Optimization (BESO) method is the BESO con-

vergence criterion (Huang et al (2006)). This measured how quickly the algorithm had converged between iterations and halted once this rate decreased below a pre-selected value.

Energy absorbing structures may be loaded to the plastic region, or beyond, to the point of material failure. To properly model and design these structures requires incorporating material damage into an analysis. This type of material damage is often simulated in a finite element (FE) analysis using either the method of element erosion, or tied nodes with failure. Both of these methods have been used to simulate the damage of fragments striking aircraft fuselages (Knight-Jr. et al (2000); Ambur et al (2001)). Element erosion has also been used to simulate the penetration of projectiles, through metal plates, during ballistic impact events (Børvik et al (2003b); Guo et al (2003)); to simulate the behaviour of tensile test specimens loaded to failure (Børvik et al (2003a)); to study high velocity impact of rigid projectiles on beams (Teng and Wierzbicki (2005)); and, to study perforation from rotor fragments during a catastrophic turbine failure (Stamper and Hale (2008)). When deployed in a real world application an energy absorbing structure would fracture if it were poorly designed and unable to absorb and transfer the energy of a load. As described, it is desired to find a structure that would absorb the energy of a load that would not experience any failure that caused it to be unsuitable for service. It is therefore necessary to detect these fractures when they occur in simulations. This would make it possible to determine whether a potential design would successfully absorb the energy of a load when put into service.

This paper presents a method to optimize the topology of energy absorbing structures. The method is based on BESO with an added constraint that the structure maintains its integrity after absorbing the energy of an external transient load. For this paper "maintaining integrity" will mean a structure in which there continues to be a load path between the supports and applied loads. Small amounts of damage are tolerated, but not amounts sufficient to break this load path. The topology optimization of energy absorbing structures is a topic that has previously been addressed. This includes optimizing structures that deform non-linearly and which experience plastic deformation. The topic addressed in this paper is different to optimizing energy absorbing structures alone. A different problem to this is one in which the reality of material damage is addressed. The very mechanism by which an energy absorbing structure completes its task is by deforming plastically. Too much of this deformation leads to failure. It is necessary when applying these optimization methods that they ensure that the structure be fit for service, while also completing its task with the minimum mass. The solution to be presented allows for some plastic deformation, including some material erosion, but not sufficient to lead to a critical fracture of the structure. It allows the designer to specify critical

---

[1] The ratio of the volume of a part/design over the volume of the entire design domain. It is described here using the symbol Ψ.

regions and to ensure that no fracture takes place between these critical regions. While these conditions are met, the optimization method also seeks to achieve a design that is the most efficient in the mass of material used to meet the design requirements. The mass efficiency is compromised if it is found necessary to meet requirements, but it is desired that this compromise be the minimum necessary. Finally, a halting method method is specified. This is a method the optimization algorithm can use to halt the optimization procedure at a point at which further optimization effort is likely to only be of minimal use. This works with the volume ratio history, rather than a performance metric.

This paper begins with an overview in Section 2 of the method used to optimize energy absorbing structures with a structural continuity constraint. This includes an overview in Section 2.1 of the BESO method. Section 3 gives a description of the continuity constraint checking algorithm and how it was incorporated to work with the BESO method. This is followed by Section 4 which describes a halting strategy and how it relates to the constraint checking algorithm. Section 5 describes a case study in which this optimization method was applied. The material model used for the finite element model is also described. Finally, case study results are discussed in Section 6.

2 Structural topology optimization of energy absorbing structures with a material damage model

The optimization task addressed in this paper was to design a structural component that would be of minimum weight and meet the requirement of absorbing the energy of a transient mechanical load without gross structural failure. It is noted that stiffness was not a design objective for this work and so was not optimized for during the optimization process.

When a transient, high velocity load is applied to an energy absorbing structure there is an amount of energy that will be absorbed by this structure as it responds to the applied load. A load of sufficient energy will cause sufficient damage to the structure that it will fracture and fail. A load of energy less than this sufficient amount will be successfully absorbed and the structure can be described as having successfully absorbed the energy of the load applied to it. The quantity of energy applied to a structure, which it needs to be able to absorb to successfully meet its design requirements, is given by, Θ. The maximum amount that a structure could absorb, beyond which fracture occurs, is given by, Θ_{max}. The magnitude of these quantities is a complex result of the geometry; the strain rate induced by the load; boundary conditions; and, location of applied loads. What is of interest is that whether measured quantitatively or not, the consequence is that when $\Theta > \Theta_{max}$ the structure will fail (experience a fracture breaking the load paths between the boundary and the applied loads). This method has the analyst tag

critical locations on the structure. These elements constitute a set of elements labelled C. If at least one load path exists between the elements in the set tagged by the analyst then the structure is considered successful and meets the constraints. In this way small amounts of damage that do not destroy the load paths are acceptable.

The optimization method used a finite element mesh to represent a 2D domain. It selected a subset of the n elements in this domain to be the design for a given iteration. Any given design could be described by a vector, $\mathbf{X} = (x_1 \dots x_i \dots x_n)$, where $x_i \in \{0, 1\}$, for a domain of n elements. Here 0 and 1 represented the inclusion or exclusion of element, i, in the design. The optimization task was to

$$\text{minimise } f(\mathbf{X}); \tag{1}$$
$$\text{subject to } \Theta \leq \Theta_{max},$$

where $f(\mathbf{X})$ is the mass of the structure.

A flow chart of the method is shown in Figure 1. The initialisation steps, including reading analyst-specified parameters, and the finite element geometry, were completed in step 1. Steps 2 to 4 were adapted from previous work, available in the literature, of applying the BESO method to optimize energy absorbing structures (Huang et al (2006, 2007)). The constraint checking algorithm corresponded to steps 5 to 11. After the FE analysis and optimization completed for each iteration, the new design was created at step 9 and the optimization repeated if the halting condition was not yet met.

For clarity, when an element was included in the design ($x_i = 1$ for element i) it is described in this paper as active and inactive/deactivated when not included ($x_i = 0$). When simulating a transient load using FE analysis, an element may be deleted and removed from the FE model for the rest of that iteration's FE analysis. This is a mechanism to simulate material damage and was discussed above and in Section 5.1.2 below. Here the terms *inactive* and *deleted* refer to two specific and different ideas that should not be confused; the deletion of an element is unrelated to the deactivation of an element, which occurred when the optimization method was generating a new structural design.

2.1 BESO optimization to create high energy density structures

The BESO process used in this paper is an iterative method to optimize the topology of a structure (Huang and Xie (2010); Huang et al (2006)). A finite element model was analysed at each iteration, corresponding to step 2 in Figure 1. From these results a sensitivity number, α, was calculated using data from that iteration's FE analysis. This corresponds to step 3. The sensitivity number was a measure of the corresponding element's influence on the objective function. The calculation of this sensitivity number depends on the optimization
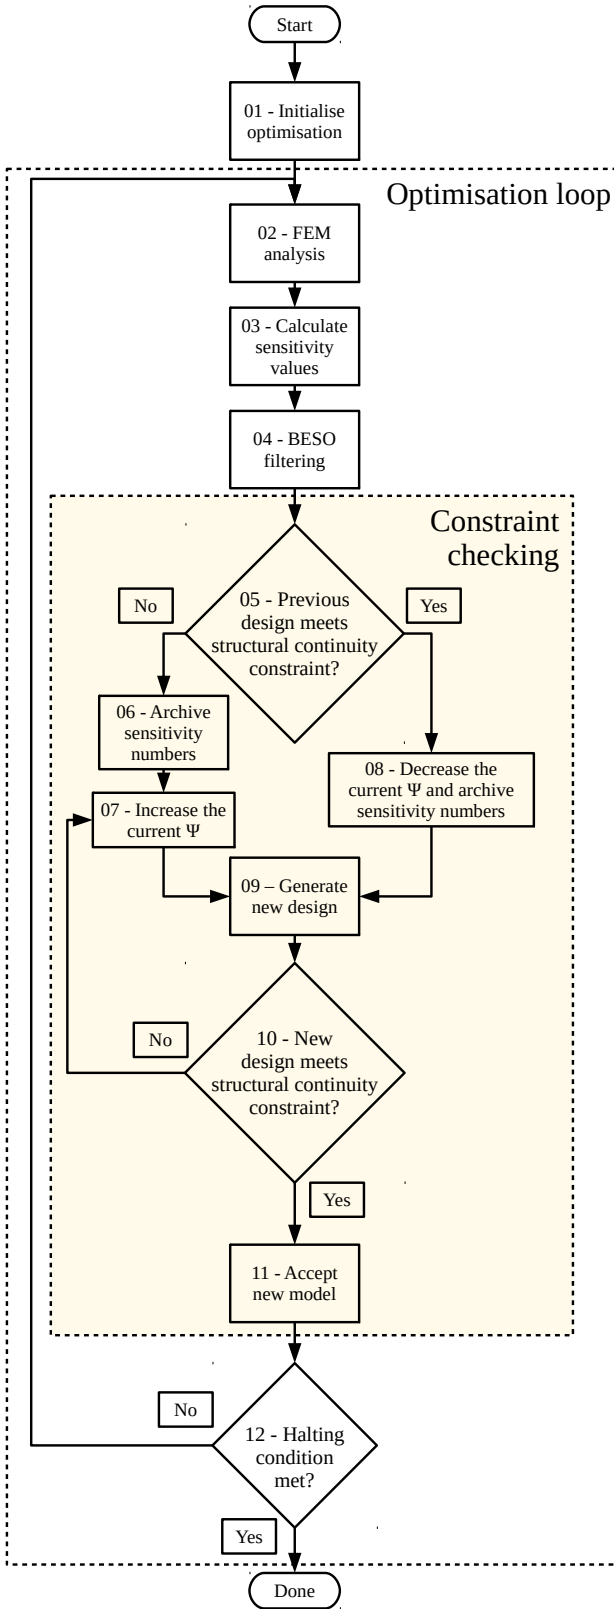
Fig. 1: Flow chart of the optimization process.

task for which the structure is to be optimized. For this analysis the optimization task was to increase the strain energy density of the structure. For an element, j, the sensitivity number was calculated using

$$\alpha_j = \frac{V_j}{V} - \frac{E_j}{E} \qquad (2)$$

where V and V_j were the total volume of the part and the volume of element j. Similarly E and E_j corresponded to the total plastic energy absorbed by the structure and the energy absorbed by element j (Huang et al (2007)). Use of Eq. (2) to calculate the sensitivity number caused the BESO method to increase the strain energy density of the structure. This is the desired result if what is being sought is a structure that absorbs energy as efficiently (with regard to structural mass) as possible.

After the sensitivity numbers were calculated they were filtered using a smoothing algorithm, at step 4. The filtering addressed numerical problems that can develop when using BESO and other topology optimization methods. These include mesh dependency and checkerboarding and have been discussed in the literature (Sigmund and Petersson (1998)). Sensitivity numbers were also calculated for elements that were previously deactivated, i.e. not part of the design at the current iteration, to provide a basis for their reactivation, hence enabling the bi-directionality of the optimization method (Huang and Xie (2007)). For each node, elements within an analyst-specified distance, r, of that node were identified. Node sensitivity numbers, β_k, were then calculated using (Huang et al (2007))

$$\beta_k = \frac{[V_1 \ ... \ V_i \ ... \ V_M] [\alpha_1 \ ... \ \alpha_i \ ... \ \alpha_M]^{\mathrm{T}}}{[V_1 \ ... \ V_i \ ... \ V_M] [1 \ ... \ 1]^{\mathrm{T}}} \qquad (3)$$

for, M, elements within distance, r, of node, k. Element i had volume, V_i, and sensitivity, α_i.

A new element sensitivity number was calculated for each element using a procedure similar to that described for calculating nodal sensitivity numbers. Here the same parameter r was used to collect all N nodes within a distance r of the centroid of each element. Then, for element i, the new, filtered, sensitivity numbers, γ_i, were calculated as

$$\gamma_i = \frac{[\omega_{r1} \ ... \ \omega_{rk} \ ... \ \omega_{rN}] [\beta_1 \ ... \ \beta_k \ ... \ \beta_N]^{\mathrm{T}}}{[\omega_{r1} \ ... \ \omega_{rk} \ ... \ \omega_{rN}] [1 \ ... \ 1]^{\mathrm{T}}} \qquad (4)$$

where r_{ik} was the distance between element i and node k and $\omega_{rk} = \omega(r_{ik}) = r - r_{ik}$.

Unless a constraint is used, like that described below in this paper, the volume of the structure will steadily decrease over successive iterations. Elements may go though optimization iterations during which they are deactivated, then reactivated after being inactive (absent) from the designs of previous optimization iterations. Elements are prioritised by their γ value with a lower γ leading to removal over elements with larger values.

3 Constraint checking and enforcement

The volume of the structure was controlled using a volume ratio, $\{\Psi \mid 0 \leq \Psi \leq 1\}$. It measured the current volume of the design, as a proportion of the sum of the volume of all elements in the design domain. For a material of constant density the mass ratio equalled the volume ratio. Therefore, minimising the volume of a structure was equivalent to minimising its mass. The decision to increase or decrease Ψ was determined by a series of checks against the structural continuity constraint. The current Ψ was decreased at the end of an iteration where the design constraints were met and otherwise increased.

Load paths exist in a structure between the location of boundary conditions and the applied loads. During a finite element simulation some elements in the model may be deleted in simulating material damage using the element erosion method. The continuity constraint required that a load path continued to exist after the energy of the applied load was absorbed.

The example given in Figure 2 shows a simple finite element mesh of quadrilateral elements. There are 20 elements in the entire design domain. These elements constitute the set, D. A subset of these, shown shaded in Figure 2a, is active (i.e. selected by the algorithm to be the design of the structure at this iteration of the optimization). This is the set T_i for iteration i, with $D \backslash T_i$ all elements that are not part of the design. Only elements in T_i are a part of the finite element analysis at iteration i. Figure 2b is the graph representing the connectivity relationship between the active elements of the mesh. Edges are unique and undirected, making Figure 2b a simple graph.

Finite element analysis was used to test the design at each iteration of the optimization. This corresponded to step 2 of Figure 1. The finite element analysis results may show some elements had been damaged. Two examples of different possible results from the FE analysis for this illustrative example are shown in Figure 3 and Figure 4. In both cases some elements received sufficient plastic strain that they were deleted during the analysis. Such elements are shown in red (darker shade). Their corresponding graphs in Figures 3b and 4b are subsets of the graph in Figure 2b. Before the analysis, elements 5 and 16 were selected by the designer as the ends of the load path across the structure. These two elements constitute the set C; this is the set C described in Section 2. During any of the structural continuity constraint checks all elements in C were checked to ensure that there was a continuous path of active, non-deleted elements connecting the entire set. It is easily seen that the result in Figure 3b has elements 5 and 16 on the same component, satisfying the structural continuity constraint check and allowing the design to be accepted. The result in Figure 4 fails this check and is outside of the prescribed constraints. In this case elements 5 and 16 lie on two different components, causing
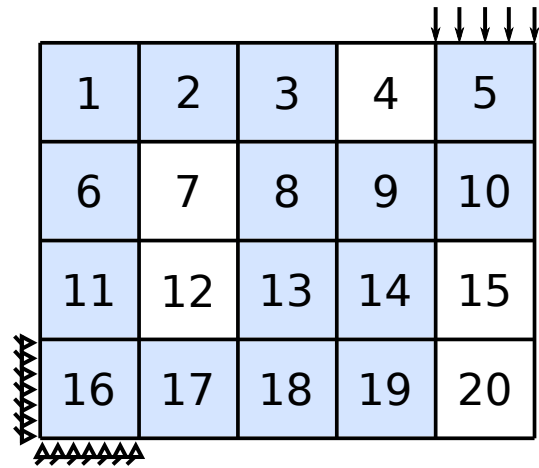
(a) Finite element mesh with elements in design T_i highlighted.
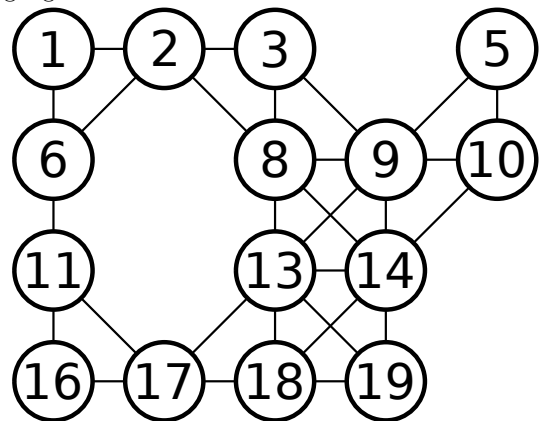
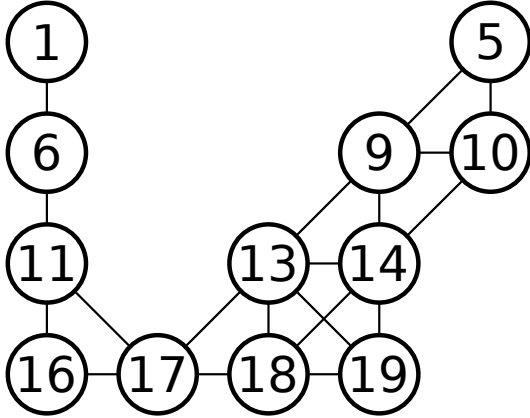(b) The network corresponding to the model of design T_i.

Fig. 2: The mesh and network of the design at T_i.

the design to be rejected. This check was performed in step 5 of the process shown in Figure 1.

In this illustrative example the evolutionary rate, Ω, was set to 0.05, giving a change of 5% of the volume of the entire design domain, between iterations. After the FE results of Figure 3 the algorithm calculated a sensitivity number, γ_i, for each of the elements, using Eq. (4). The next task was to generate a new design for iteration $i + 1$. After passing the first constraint check the total volume ratio was decreased to $\Psi_{i+1} = \Psi_i - \Omega$. For a domain of n total elements with a volume ratio Ψ, there were $\lfloor n \times \Psi \rfloor$ elements selected for the design of the structure, with elements prioritised by the value γ calculated for that element. Figure 3a shows that elements 4 and 7 were both inactive at i, but were likely to have large values of γ as they were near elements that were deleted. Elements are deleted because they experience large strains. From Eq. (2) such elements will have large sensitivity numbers and will be prioritised over other elements when ranked and a new design is generated. Conversely, some elements would need to be removed to meet the smaller volume ratio. If all other inactive elements at i, and elements 14, 18, and 19 had the lowest γ they would become inactive at the next design. This leads to a proposed design that would be-
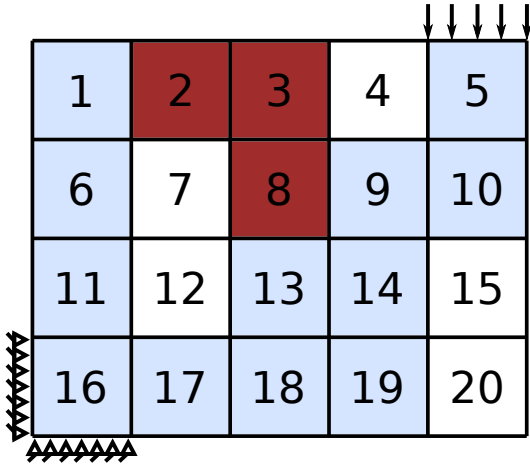
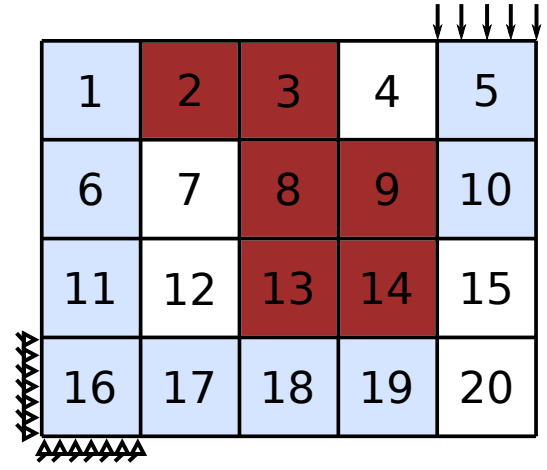(a) Model at T_i after simulation experiencing light damage.

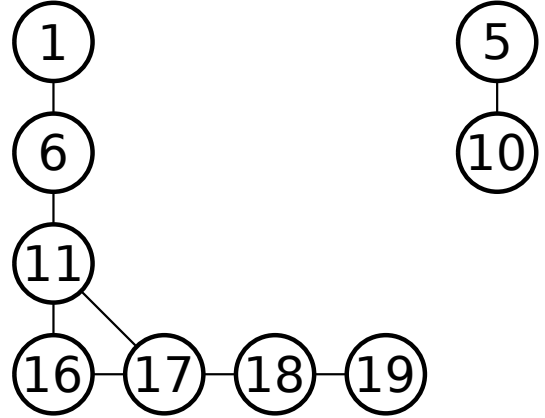(b) The network corresponding to the model of design T_i after light damage.

Fig. 3: The mesh and network of the design at T_i with light damage.

(a) Model at T_i after simulation experiencing heavy damage.

(b) The network corresponding to the model of design T_i after light damage.

Fig. 4: The mesh and network of the design at T_i with heavy damage.

come T_{i+1}(proposed). This design would need to pass the second continuity constraint, corresponding to step 9 of Figure 1. If T_{i+1}(proposed) passed the second check it would become T_{i+1}; otherwise, Ψ would be increased by Ω, a new design created, and this process repeated until the second check also passed. The loop between steps 7 and 9 will always produce a successful design except for the case in which it is not possible to satisfy the continuity constraint using the entire domain D in the design. This occurs when D does not contain enough elements to create a connected path between the members of C. Such a problem is impossible to solve and no optimization process would ever provide a valid solution.

The result in Figure 4 shows that the structure is unsuitable for these design requirements. This causes an increase of Ψ and the generation of a new design. The values of γ used must be those from the last *successful* design. The distribution of plastic strain changes with fracture. The load duration is shorter, leading to less energy being transferred to the structure. When a fracture destroys a load path, the energy of the load is no longer transferred through this load path to other areas of the structure. For this reason only sensitivity num-

bers of successful designs can be used. The sensitivity numbers of successful designs were therefore archived. If a design failed the first structural continuity constraint check, the archived sensitivity numbers were used in place of those calculated from the most recent iteration. Once the check at step 9 passed, the next FE analysis took place unless the halting method decided the optimization had completed.

The algorithm used to perform the constraint check required the generation of an associativity matrix. The associativity matrix, $[\mathbf{A}]$, corresponding to the network shown in Figure 5 is

$$A = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{pmatrix} 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix} \end{matrix}. \tag{5}$$

The labels correspond to the vertices of the graph. The graph is undirected and therefore edges connect vertices in both directions, leading to a symmetric matrix. Connected components were found using a published
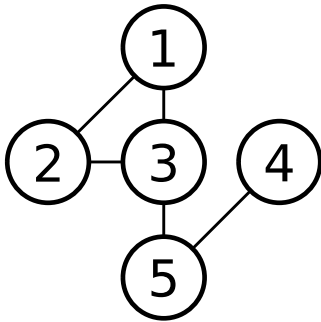
Fig. 5: Example graph for a structure with small number of elements.

Fig. 6: Two phases in an ideal optimization.

algorithm (Pearce (2005)) implemented in the Python SciPy (Jones et al (2001–)) library. The library function[2] accepts a matrix of the type shown in Eq. (5) and returns a mapping that takes an element number and returns the component on which that element is present. This mapping can be checked for each element in the model. If every member of C is on the same component the continuity constraint condition is met. Subsequent work has shown that this method is computationally too expensive for 3D models with a larger number of elements and more graph edges connecting neighbouring vertices. A solution that has been implemented in later work but was not used in the particular examples of this paper is to simply represent a graph data structure and complete a depth-first or breadth-first search through the graph beginning at an arbitrary element in the set C and continuing until all connected vertices are exhausted, or all remaining elements in C are visited. Further improvements, such as the use of the Euclidean metric to speed up searches, as is the case of the A* algorithm, was seen by the author as unnecessary given that the linear time, $O(n)$ to the number of vertices, to conduct this search was much faster than the other much more non-linear computations undertaken during an iteration of an optimization loop, such as the explicit finite element analysis of the candidate structure.

The continuity constraint checking method can be used to determine whether different locations within a finite element model are connected. In the context of structural topology optimization it provides verification that any design generated is connected between the pre-selected regions and in the context of optimization of structures in which material damage is modelled by element deletion, it allows for the verification that the simulation showed that the design was able to resist fracturing between those critical locations. It is fundamentally tied to the graph mapped to the discreet geometry of the finite element mesh.

Similar work for optimizing dual phase materials (Andreasen et al (2014)) for a mechanical damping objective used a different method to ensure connectivity in a structure. The work solved an additional conductiv-

ity optimization problem, different to the damping optimization problem that was the design task, in which a minimum conductivity was specified and which needed to be enforced by the optimization procedure. Searching through a graph is a straightforward and established problem of computer science and discreet geometry. The method described in this paper is a much simpler solution that runs an algorithm that completes in linear time to determine whether all elements in C are connected.
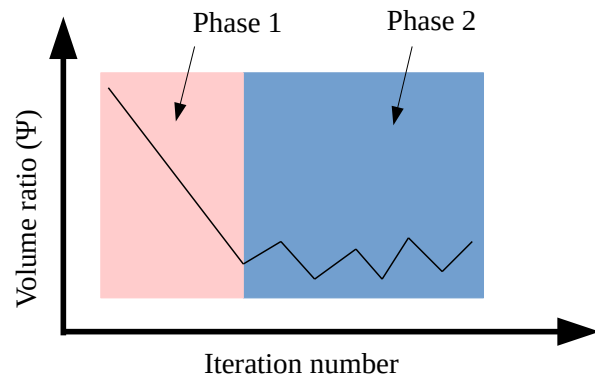
There is always a fixed number of neighbours for any element. The size of this upper maximum possible neighbours depends on the type of elements used. For example, quadrilateral elements can have up to eight neighbours. Any row or column in Eq. (5) can therefore only have up to eight positive values. As the mesh density increases and the size of the matrix increases, the proportion of matrix elements that have 0 values increases. Therefore for this application a sparse matrix representation of the associativity matrix should be used.

4 Halting method

The constraint checking algorithm controlled the volume fraction, Ψ, for the structure but did not provide a mechanism for determining when the optimization should halt. The general behaviour of this optimization method, when beginning with an initial design that had all elements in D active, can be described as occurring in two phases. This is illustrated in Figure 6. There is an initial phase in which most iterations successfully absorb the transient load. The volume ratio history between iterations will be almost monotonically decreasing. As the lowest permissible volume ratio is reached, the structure will begin to fail more frequently. In this second phase the volume ratio history remains unchanged, within a small range of fluctuation. The method used here to automatically determine when an algorithm should halt assumes this overall behaviour.

The volume ratio history of Figure 6 is a gross idealisation and in reality a higher level of complexity in the evolution of Ψ is observed. Some iterations will fail
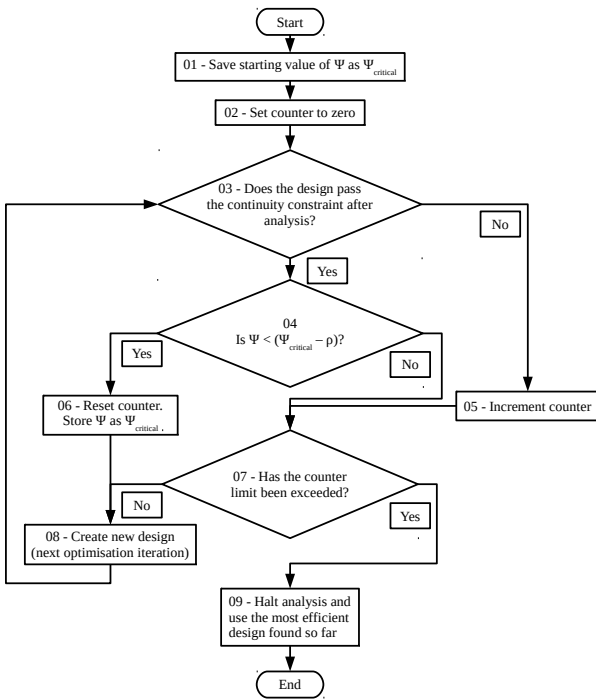
---

[2] Full function name, using version 0.13.2 of the SciPy library is "scipy.sparse.csgraph.connected_components".

Fig. 7: Flow chart of the halting method.

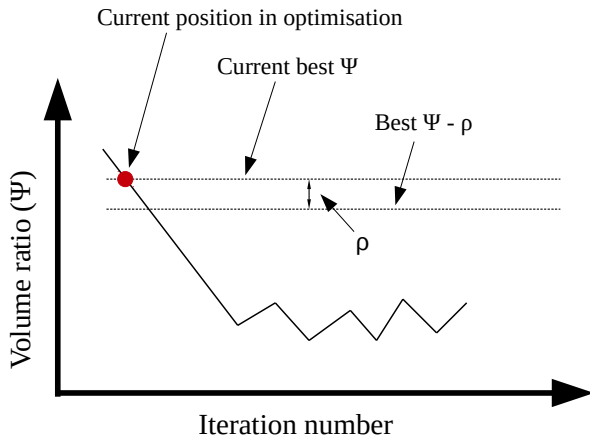Fig. 8: An ideal optimization in the first phase.

Fig. 9: An ideal optimization in the second phase.

the structural continuity constraint check during phase 1. In this phase these failures are transient and the process continues with most iterations producing designs that satisfy the structural continuity constraint. Simply halting after the first failed design is not likely to produce a highly optimized topology. In the second phase a critical volume ratio is reached. The structure will always be unable to meet its design goals, i.e. fail the structural continuity constraint, after Ψ is decreased below this critical value. The volume ratio history for an optimization will show frequent cycles as Ψ is decreased, causing a failure against the continuity constraint; this is followed by an increase in Ψ, a successful design, and another decrease in Ψ to repeat the cycle.

A counter tracked these cycles, halting the optimization when that counter reached a value prescribed by the designer before the optimization began. This pre-
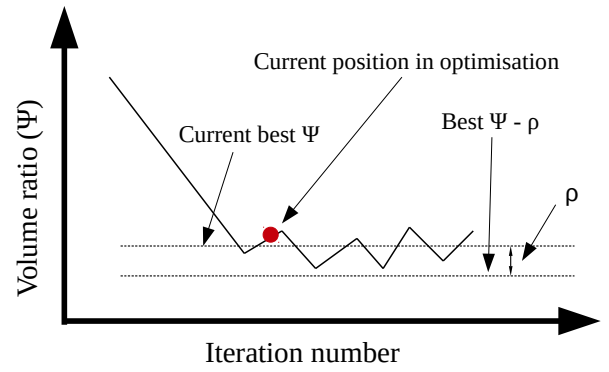
scribed limit, ξ, was a parameter specified to the optimization by the designer, before the analysis began. At any point in the optimization history the smallest value of Ψ for any structure that had satisfied the constraint check was stored as Ψ_{critical}. The initial value of Ψ_{critical} used was the value of Ψ at the beginning of the optimization. For the work in this paper this initial value was always 1.0. The counter halting the optimization was initially zero. Both of these bootstrapping operations are shown at steps 1 and 2 of Figure 7. Figures 8 and 9 show the progress of an ideal optimization at two stages. The first stage is shown in Figure 8. Two horizontal lines correspond to Ψ_{critical}, the higher of the lines, and $\Psi_{critical} - \rho$, the other. The scalar parameter, ρ, was provided by the designer before the optimization. This was the second of the two parameters used by the designer to control the automated process of halting. Together ξ and ρ would determine when the optimization would halt. At each iteration a design would be created that would either pass or fail the first structural continuity constraint check. This occurred at step 5 of Figure 1, corresponding also to step 3 of Figure 7. If the design failed it caused the counter to increment at step 5 of Figure 7. If Ψ, for a successful design, had fallen below the line of $\Psi_{critical} - \rho$ it would reset the counter to zero and update $\Psi_{critical}$. This was checked at step 4 and if found to have occurred would trigger a reset at step 6. The counter would be set to zero and Ψ_{critical} updated to the lowest Ψ of any design that had passed the structural continuity constraint after analysis. The result of this is shown in Figure 9 where a steady decrease of Ψ caused this floor to be pushed lower. The optimization is now in phase two and the volume ratio stays approximately steady. The floor cannot be reached because further decreases of Ψ lead to failed designs, which lead to increases of Ψ. At step 7 the counter was checked and found to have reached its limit. In doing so it caused the optimization to halt and the task was complete.

It has been discussed in literature, e.g. (Huang et al (2006)), that the most efficient design found using the BESO, or other similar methods, will not necessarily be the global optimum. The halting method allowed for a rational, automated, approach to searching for a design

but limiting the search once a local optimum had likely been reached. There are transitions between spaces of different local optima. These occur, as for the regular BESO method, when there is a topology change, often coinciding with a failure of the continuity check. There are also changes that are coincident only with different geometries, rather than topology. These failures will usually correspond to a search toward a different local optimum. In this regard the complete method will complete after searching through several spaces of different local optima. This convergence speed can be controlled by the selection of ρ and the counter limit. Generally, selecting larger values for these parameters will lead to a slower convergence speed.

5 Case study

The case study used to demonstrate this optimization process is shown in Figure 10. The 2D model was similar to a previously published example geometry (Huang et al (2007)). It was a vertically symmetric model of size $L \times W = 0.2m^2$. The model had two supports at distances between $D_1 = 0.15m$ and $D_2 = 0.22m$ from the nearest edge. A section thickness of $0.01m$ was used. The round, rigid impactor shown at the top of Figure 10 had an initial vertical velocity in the downward direction. The optimization exercise was repeated 27 times using different parameters for the impactor mass, impactor radius, and initial velocity. Case study 14 employed median values of all parameters investigated.

A value of $r = 0.01$ was used in this analysis for the user selected parameter, r, described in Section 2.1. The initial Ψ was 1.0 with $\Omega = 0.01$. The BESO algorithm does allow an initial value of $\Psi < 1.0$ (Huang et al (2006)), however this was not used for this study. The structure was defined using symmetry to analyse only one half of the model in Abaqus/Explicit with $10,240$ CPS4R elements. A convergence study of the mesh showed that there was convergence in the kinetic energy and energy absorbed by the structure, using a lower mesh density.
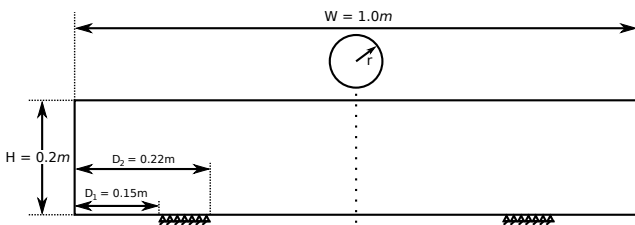

Fig. 10: Model dimensions

5.1 Material model

The material modelled in this analysis was aluminium alloy Al 2024T3/T351. Elastic material behaviour was modelled using a Young's modulus of, $E = 73GPa$,

and Poisson's ratio of, $\nu = 0.33$. The remaining constants used are listed in Table 1. Plastic material properties were modelled using the Johnson-Cook plasticity model (Johnson and Cook (1983)). Material damage initiation (see below) was modelled using the Johnson-Cook damage model (Johnson and Cook (1985); Dassault Systèmes (2011); Kay et al (2007)) and damage evolution was also accounted for.

5.1.1 Johnson-Cook plasticity model

The Johnson-Cook plasticity model was used to model the plastic material behaviour for this analysis. The model gave a yield stress, σ_y, as a function of strain (equivalent plastic strain), $\bar{\varepsilon}_{pl}$, strain-rate, $\dot{\bar{\varepsilon}}_{pl}$, and material temperature, T, i.e. $\sigma_y\left(\bar{\varepsilon}_{pl}, \dot{\bar{\varepsilon}}_{pl}, T\right)$. This function was defined by

$$\sigma_y\left(\bar{\varepsilon}_{pl}, \dot{\bar{\varepsilon}}_{pl}, T\right) = \left[A + B\left(\bar{\varepsilon}_{pl}\right)^n\right]\left[1 + C\ln\left(\dot{\bar{\varepsilon}}_{pl}^{*}\right)\right]\left[1 - (T^*)^m\right]. \quad (6)$$

The equivalent plastic strain rate, $\dot{\bar{\varepsilon}}_{pl}$, was regularised against a reference strain rate. The value 1.0 or a quasi-static strain rate is usually used with this material model. This then gave $\dot{\bar{\varepsilon}}_{pl}^* \equiv \dot{\bar{\varepsilon}}_{pl}/\dot{\bar{\varepsilon}}_{pl_0}$. Similarly, the variable, T^*, was regularised using

$$T^* \equiv \begin{cases} 0 & \text{for } T < T_0 \\ \frac{T-T_0}{T_m-T_0} & \text{for } T_0 > T_m \\ 1 & \text{for } T > T_m \end{cases}.$$

The subscript, m, denotes the melting temperature and the subscript 0 denotes the reference temperature, usually taken as room temperature in Kelvin.

5.1.2 Damage initiation and damage evolution in Abaqus

Modelling damage in Abaqus/Explicit involves defining the conditions for damage initiation and damage evolution. Figure 11 shows a simplified one-dimensional stress-strain response of a ductile material loaded to failure. As the strain increased the material transitioned from initially elastic during 1-2, to plastic during 2-3. Damage initiated at 3, followed by material degradation, 3-5, until complete failure at 5. The damage initiation material model specified the location of 3 and the damage evolution model described the deviation of the material model from the theoretical plastic response. The path 3-5 is the path actually taken by the material and 3-4 is the path of the theoretical plastic material if damage were not accounted for.

5.1.2.1 Damage Initiation
A change in the material response due to material damage occurs after position 3, in Figure 11. Between 2-3 a state variable, the criterion for damage initiation, $\{\omega_D \mid 0 \le \omega_D \le 1\}$, was incremented until it reached its critical value. This critical
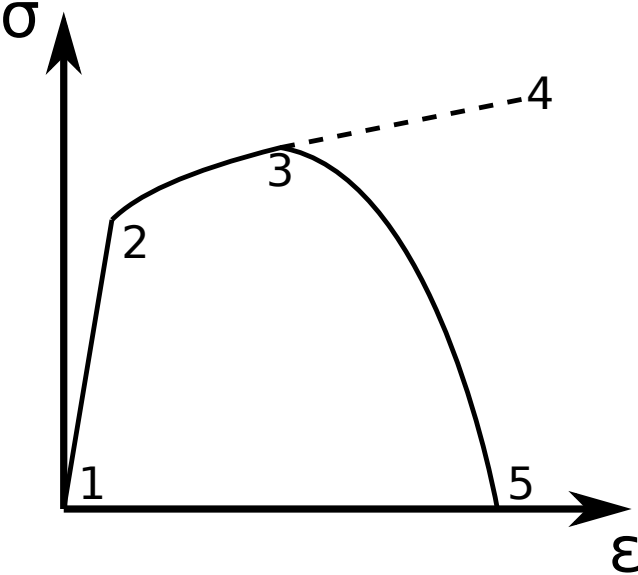
Fig. 11: Schematic stress-strain response for a damaged, ductile material.

Table 1: Material constants for the models used in this analysis.

Johnson-cook plasticity model				
A	B	n	m	C
$369 \times 10^6 MPa$	$684 \times 10^6 MPa$	0.73	1.7	0.0083
Johnson-cook damage model				
D_1	D_2	D_3	D_4	D_5
0.31	0.045	-1.7	0.005	0
Reference values				
T_{melt}		T_{room}		$\dot{\varepsilon}_0$
775K		294K		1
Displacement at failure ($\bar{u}^{pl}$)				
$2.0 \times 10^{-5} m$				

characteristic length of the element, was the distance of the diagonal across that element.

6 Results and discussion

The optimization method was tested by applying it to a case study in which a projectile impacted on a structure. This two dimensional model considered variations in the mass, radius, and initial velocity of the impactor. Their values and the initial result data are given in Table 2. Figure 13 shows a linear relationship between the energy of the projectile and the mass of the optimal structure that resulted from the optimization method. Material in the structure, when deformed plastically, will absorb energy from the projectile. If there is a minimum volume ratio that any final structure must have, to absorb all the projectile's energy, then an optimization method would seek to find a successful design that is as near to this value as possible. This minimum mass would be larger for a load that transfers more energy onto the structure. This is the relationship shown by Figure 13 where the final mass varied with the projectile energy.

To illustrate the relationship between the amount of projectile energy and the final volume ratio a selection of examples is shown. Their points on the graph in Figure 13 are highlighted with a circle, diamond, triangle, and square; corresponding to cases 1, 14, 24, and 27. The final results are shown in Figure 12 and Figure 14a. Both the graphed results in Figure 13 and the resulting designs themselves show that there is a very close relationship between the final mass and the amount of energy of the load, with greater energy leading to a larger mass ratio for the final result. In these examples all structures shared the same topology except for case 27 which featured several voids. This was the trend observed, that more complex topologies tended to occur at higher amounts of absorbed energy, which also corresponded to higher volume ratios. Voids tending to disappear as the volume ratio was decreased. The results in Table 2 show that the only designs with voids were those corresponding all from the two groups with the

value is typically set to 1.0, as it was for the simulations in this paper. After this value was reached the material model transitioned from damage initiation to damage evolution. The damage initiation criterion, ω_D, was incremented at each timestep by

$$\Delta\omega_D = \frac{\Delta\bar{\varepsilon}^{pl}}{\bar{\varepsilon}_D^{pl}\left(\mu, \dot{\bar{\varepsilon}}^{pl}\right)} \quad (7)$$

where $\Delta\bar{\varepsilon}^{pl}$ was the equivalent plastic strain in an element during a timestep. The value $\bar{\varepsilon}_D^{pl}$ was the strain at which damage initiated. This was a function of the stress triaxiality, μ, and the equivalent plastic strain rate at the current timestep, $\dot{\bar{\varepsilon}}^{pl}$. It was recalculated at each timestep (Dassault Systèmes (2011)). The function, $\bar{\varepsilon}_D^{pl}\left(\mu, \dot{\bar{\varepsilon}}^{pl}\right)$, used in this analysis, was based on the Johnson-Cook damage criterion

$$\bar{\varepsilon}_D^{pl}\left(\mu, \dot{\bar{\varepsilon}}^{pl}\right) =$$
$$[D_1 + D_2 \exp\left(-D_3\mu\right)]\left[1 + D_4 \ln\left(\dot{\bar{\varepsilon}}_{pl}^*\right)\right][1 + D_5 T^*] \quad (8)$$

where $\dot{\bar{\varepsilon}}_{pl}^* = \dot{\bar{\varepsilon}}^{pl}/\varepsilon_0^{pl}$. Here $\dot{\bar{\varepsilon}}^{pl}$ was the strain rate at a given increment and ε_0^{pl} was a reference strain rate. A value of 1.0 or a quasi-static strain rate is usually used. The regularised temperature, T^*, was defined as for the Johnson-Cook plasticity model. The constants D_n for $n = 1 .. 5$ are obtained through test results for each material and fitted to observed data (Kay et al (2007)).

5.1.2.2 Damage Evolution When ω_D reached 1.0 the material response transitioned to section 3-5 of Figure 11. The overall energy absorption can be defined by specifying one of either the total strain or energy absorbed before failure; with one of the latter two parameters derived from the other. For this analysis a linear curve for 3-5 was used and a plastic displacement at failure, $\bar{u}_f^{pl} = 2.0 \times 10^{-5} m$. Here $\bar{u} = L\bar{\varepsilon}^{pl}$ and L, the

(a) Case study 1 - iteration 108.

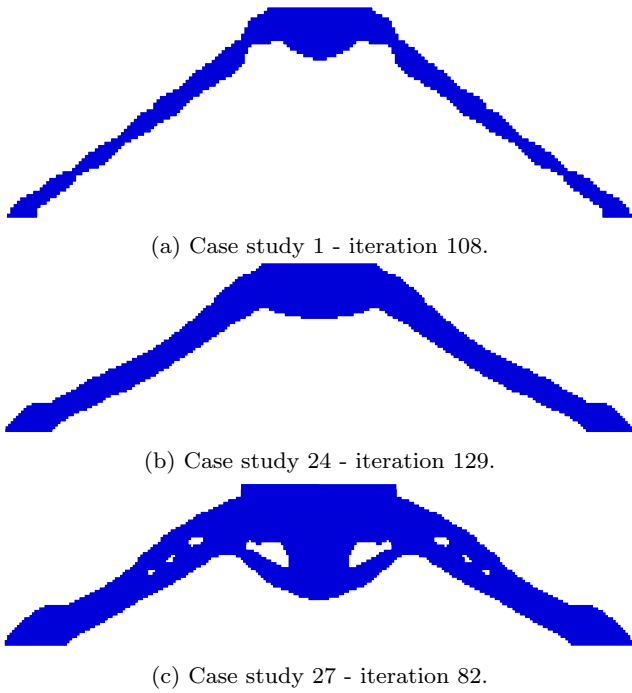(b) Case study 24 - iteration 129.

(c) Case study 27 - iteration 82.

Fig. 12: Resulting designs for case studies for applied loads of different magnitudes of energy.
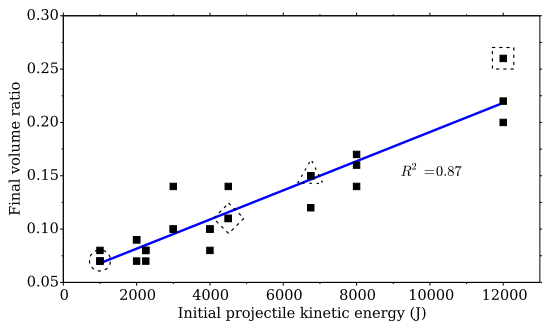
Fig. 13: Final volume ratios with initial impactor kinetic energy.

(a) Design meeting continuity constraint – iteration 123.

(b) Design not meeting continuity constraint – iteration 124.

Fig. 14: Examples of continuity constraint satisfying and failing designs generated during case study 14 – damaged elements orange (lighter colour).

(a) Distribution of plastic energy before optimization of case study 14.

(b) Plastic energy absorbed
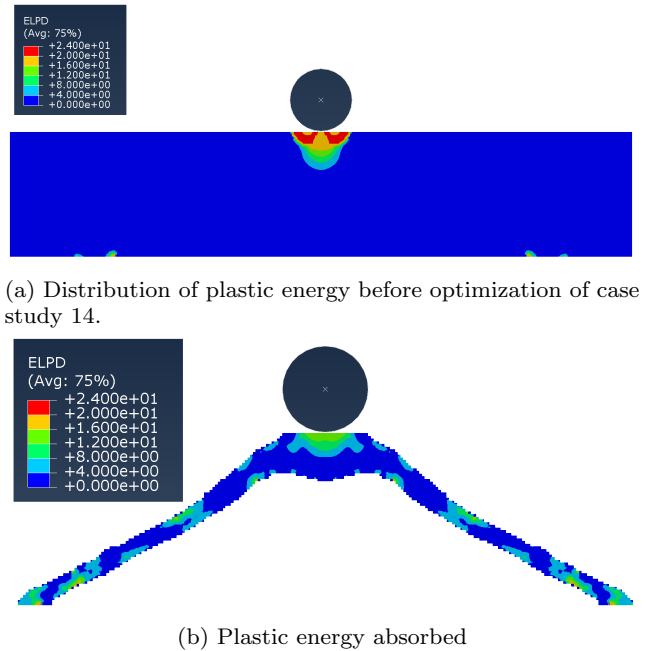
Fig. 15: Plastic energy absorption before and after optimization for case study 14.

largest amount of absorbed energy and case 19, which also featured one void.

Case study 14 had all three parameters at the median level. Optimization results for this case study are shown in Figures 14 to 16. The final result in Figure 14a shows the geometry without displacement and with the deleted elements highlighted. This is an example of a successful design that was assessed by the continuity constraint checking algorithm. Despite experiencing a small number of deleted elements, the structure was still able to contain the loading, successfully absorbed the projectile energy, and there was still a load path between the impact location and the supports after the energy was absorbed. Compare this with Figure 14b where the unsuccessful design experienced fractures breaking the load paths to the supports.

The distribution of plastic energy for case study 14, before optimization, is shown in Figure 15a. Almost all plastic energy is absorbed either near the impact location or near the supports. The distribution

after optimization is shown in Figure 15b. Here the plastic energy is absorbed more evenly throughout the structure. This is confirmed quantitatively in Figure 16 where the distribution of elements by plastic energy absorbed is shown in a pair of histograms. In the first, corresponding to the model shown in Figure 15a, the largest subset of elements is of those elements absorbing the least energy. This is two orders of magnitude larger than the next subset (note the broken vertical axis). Most elements absorbed little or no energy, with the size of subsets decreasing with energy absorption for those elements. At the far end there was a sharp increase in the histogram. This jump corresponded to deleted elements. Although the loading history for ele-

Fig. 16: Histograms of plastic energy absorbed before and after optimization for case study 14.

Fig. 17: Number of elements deleted during optimization of case study 14.

Fig. 18: Volume ratio history and energy absorbed for case study 14.

ments may vary, there was an approximate maximum plastic energy that an element would absorb before it was deleted from the finite element model. A deleted element would not 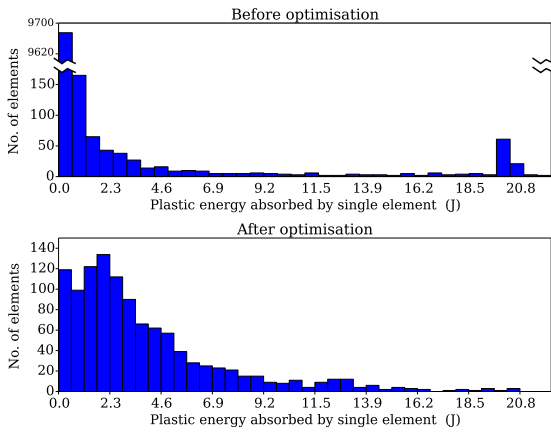absorb further energy, leaving a peak at the far end of the first histogram. This contrasts with the other histogram, corresponding to the optimized design. More elements dissipated larger amounts of plastic energy. There was also a shift away from the far end. There were fewer elements that experienced large enough strains—i.e. dissipated enough energy through plastic deformation—to be deleted. This shift in the distribution of plastic energy occurred together with a decrease in the stiffness of the structure. This is visible in Figure 17 where the number of elements deleted during the progress of the optimization, is shown. At selected itations the corresponding design is shown, with the undisplaced geometry outlined and the displaced geometry shown shaded. There was a sharp drop in the number of elements deleted between approximately 75 and 100 iterations into the optimization. The structure changed from small to large displacements within a small number of optimization iterations. By increasing the displacement over which the energy of the transient load was absorbed, it caused the structure to absorb more energy through a global response. A global response allows for more of the structure to be involved in absorbing the energy of the load, leading to the shift shown in Figure 16. This is visible in Figure 15a where much of the plastic energy is shown to be absorbed near the impact location, compared with Figure 15b. Also in Figure 18 is the volume ratio history for case study 14, shown in the blue/broken line, along with the energy absorbed by the structure, shown using the green/unbroken line. While the volume ratio is decreasing in stage 1, the energy absorbed increases slightly. This is a result of a decrease in the stiffness in the structure, allowing it to absorb slightly more energy from the projectile without it ricocheting away. Once stage 2 some designs are failures (marked on the green/unbroken line). These will often absorb far less energy as the structure fails before all of the projectile
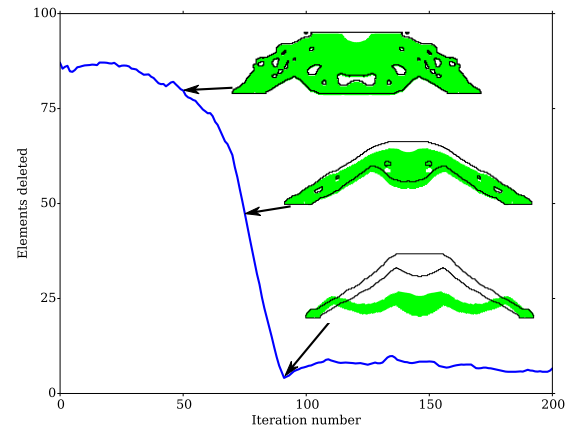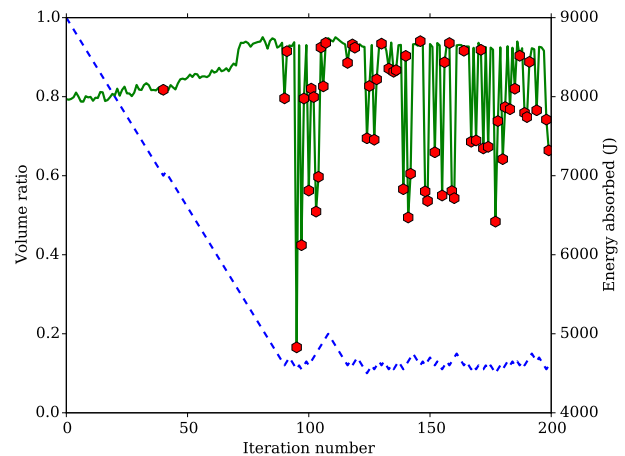
energy is absorbed. These are often at the bottom of a sharp trough before the design is recovered and the energy absorption recovers again.

The increased displacement occurred as plastic hinges began to form during simulation in the FE analysis step. The distribution of plastic energy dissipation of the final design in Figure 15 before and after optimization shows that the depth of plastic energy dissipation near the impact location is much greater before than after the optimization. There is a self-reinforcing process where this buckling allows the energy to be dissipated away from the impact location and to a shallower depth within the structure. This allows for a thinner structure that allows for greater buckling. It shows that for this combination of load, material, and boundary conditions the sensitivity number calculation in Eq. (2) will tend to generate thin structures with large amounts of buckling if stiffness is not a focus of the optimization method.

There were two conditions that were identified which, if they emerged, would lead to an optimization that was inefficient in the time taken to reach a final solution or in the final volume ratio of the structure.

They both related to selection of the elements in set C. If elements within this set were deleted during a finite element analysis, the algorithm would decide the component had fractured. Figure 19 shows an example from test case 17 in which a crack would form in line with the direction of the applied load. This crack would disappear in later iterations. Such a crack was acceptable under the structural continuity constraint if it did not fracture the component. However, the algorithm that calculated the structure's status against this constraint gave the status as violating for iterations in which this crack appeared. Figure 20 shows the volume ratio history over the course of two different optimizations of test case 17, together with test case 14 for comparison. During the first attempt an element in C was in the path of the crack. After this optimization the element was moved to a nearby location just outside of the path of this crack. The plot shows the difference in the volume ratio history caused by the design being incorrectly classified for some iterations, leading to unnecessary increases in Ψ. It also shows that despite this, the algorithm was robust enough to still reach an optimized design, albeit requiring more optimization time than when the analysis was repeated. The solution was to move the particular element in C to a nearby location just outside of the crack path. A comparison of elements in C against the list of elements deleted during an analysis could identify when this event has occurred and could be used to display warning messages to the analyst to address this if necessary. This example illustrates that the selection of elements in C need not strictly include those elements on which loads and boundary conditions were applied to be of practical use. Instead elements nearby to those on which loads were applied were selected because of the advantage this provided for the analysis.

Once an unsuitable design is encountered the constraint checking method will cause the structure to increase in volume until a suitable design is re-encountered. This causes the optimization to branch into a different space with a different local optimum toward which the optimization then heads as it begins to again reduce structural mass, converging toward a different solution. Experience from the case study showed that such an even causes a jump in volume ratios during phase 2 which leads to a small, transient increase in Ψ that then decreases as the optimization moves back toward the new local optimum. In the context of the optimization method this would cause a resetting of the counter limit and the optimization continues as usual. From the case study this was only of practical importance when too large a value of Ω was selected for the optimization. Otherwise smaller values of Ω tended to lead to more stable optimizations. This is illustrated in Figure 20.

If an element were selected and later optimized away it would similarly have caused the algorithm implementing the structural continuity constraint to incorrectly classify designs as failing against the constraint

(a) Crack formation in case study 17 at iteration 91.

(b) Crack not forming for case study 17 at iteration 92.
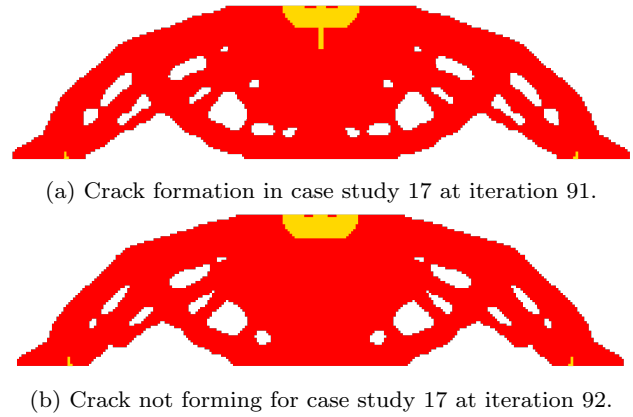
Fig. 19: Intermittent cracks formed during optimization of case study 17.
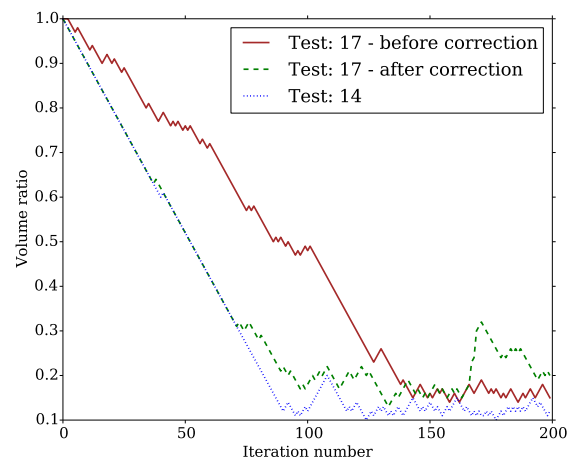
Fig. 20: Volume ratio history for case study 17.

when the structure had in fact not fractured. The most straightforward solution was to select elements in C near to regions of relatively higher sensitivity numbers. For this analysis elements with large amounts of plastic strain energy density also had large sensitivity numbers. Selecting elements in C with large values of γ would avoid this problem. An automated solution, not used in this paper, would be to check that the element is deactivated and automatically relocated to a nearby, suitable location, that still enforced the connectivity constraint. Figure 15a shows that for the analysis of an energy absorbing structure the location of applied loads and supports are the areas likely to have high sensitivity values. For this reason this problem is more likely to emerge if the analyst selects elements in C that are not near to, or the elements themselves, that have loads and boundary conditions applied on their boundaries.

The case study was repeated across a range of different parameters and repeated 27 times with results shown in Table 2. Each of these 27 tests held all conditions the same except for three variables which were changed between the tests. To show the generality of the method used to enforce the continuity constraint, test

case 14 was modified in the position of the boundary conditions and loads. These cases were no longer symmetrical. This necessitated modelling the interaction without the advantage of symmetry to speed processing, however, the results were equivalent to the faster case and no changes were made to the optimization process. In the first of these modified cases, shown in Figure 21a, all aspects of the optimization process were unchanged. The model was modified only by the addition of a single extra impactor to the right, at a position of $0.75m$, or 3/4 of the width of the structure. The second modified case, shown in Figure 21b changed the boundary conditions creating a cantilever structure. No changes were made to the loads from the example in Figure 21a.

It is noted that although loads and boundary conditions had been discussed in connection to the continuity constraint check in Figure 2, this was only made in the context of how the method can be best applied by the analyst in order to achieve the desired results. The load and boundary conditions are not used by the algorithm implementing the constraint check and so their change will not affect the results obtained for a given geometry and element set C (see Figure 3). The location of elements constituting C are shown in the diagrams as crosses. For the example shown in Figure 21a no change was made to C. The constraint still required that the span between the two supports kept a load path after loading. In the second case, shown in Figure 21b the change in boundary conditions resulted in a condition in which the most suitable geometry would likely have been one in which the bottom surface, near the supports in the previous example, would have been removed by the optimization method. The critical load path existed between the support on the left of the structure and the two locations at which loads were applied. These positions are shown in Figure 21b. They were the new locations making up the set C. Except for this change in the set C the rest of the optimization proceeded unchanged, including use of the halting method to halt after 3 cycles.

Results from the optimization are shown in Figure 22. What is observable is that despite changes in boundary conditions and loads there was no problem for the constraint checking algorithm to enforce the connectivity constraint. This can be expected, given that boundary conditions and loads are not directly used in any calculation of the algorithm and continued to execute as expected. Results are shown at iterations 105 and 111, as selected by the halting method.

It is noted that for the example shown in Figure 21b with results in Figure 22b the set C was changed from that used in all other cases discussed. If static loads are used, the supports and loads are applied to the surface/boundary and are available in the description of the model used by the finite element analysis. This information could be used to automatically populate C with the elements on which loads and boundaries are applied, removing the need for human tagging of

(a) Case with one additional impactor.

(b) Case with cantilever and additional impactor.
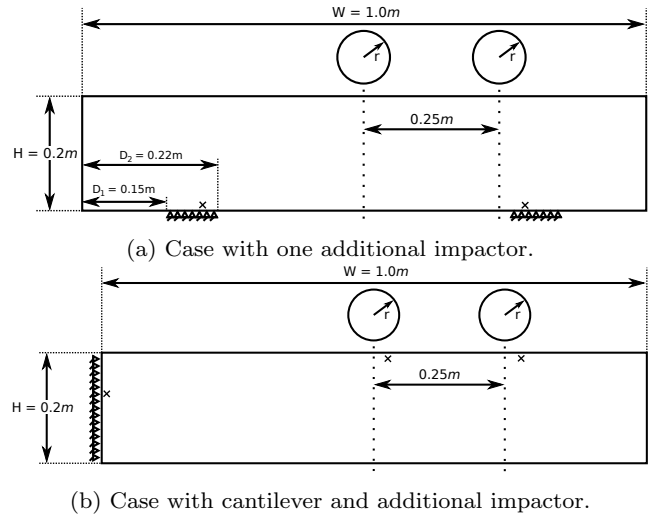
Fig. 21: Diagrams of additional cases with modified loads and boundary conditions.

(a) Case with one additional impactor.

(b) Case with cantilever and additional impactor.
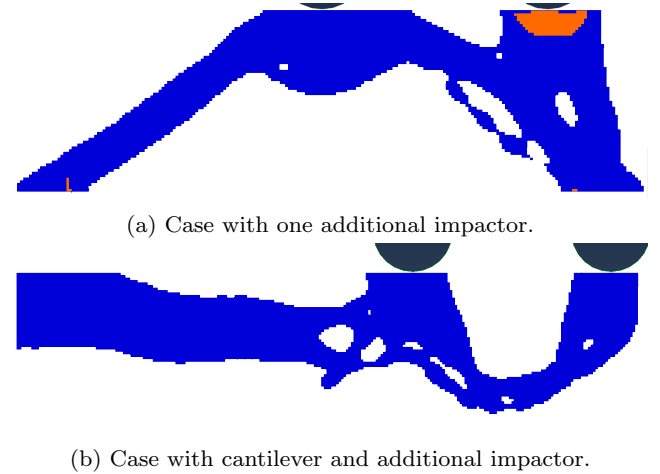
Fig. 22: Results from additional cases with modified loads and boundary conditions.

elements in C, however this approach was not used in this work.

For this analysis all test cases were allowed to run for 200 iterations. For each case the lowest successful volume ratio reached was compared against the halting method described in Section 4 above. The results are shown in Table 3. This table includes two groups of results. The first is the volume ratio of the most efficient design and the iteration at which it was first encountered during the optimization. This group of columns is headed with "*Best design*". The other group includes the best design reached when using the halting method and the iteration at which halting took place. These are shown for the case when the method would halt the optimization after 2, 3, 4, and 10 failures of the continuity constraint condition, without resetting the counter (described in Section 4 above), before halting. This allowed for a comparison between the halting method, repeated using several user-defined values for the counter limit, against a longer optimization of 200 iterations. The iteration at which transition to phase 2 of the optimization

Table 2: Results of 27 case studies.

Test no.	Mass (kg)	Velocity (ms-1)	Radius (m)	Final volume ratio	Elements deleted at final Ψ	No. of voids in topology	Last iteration no.
1	0.2	100	0.025	0.07	4	0	108
2	0.2	100	0.050	0.08	4	0	135
3	0.2	100	0.075	0.07	4	0	93
4	0.2	150	0.025	0.07	12	0	100
5	0.2	150	0.050	0.08	7	0	175
6	0.2	150	0.075	0.08	5	0	147
7	0.2	200	0.025	0.08	59	0	125
8	0.2	200	0.050	0.10	7	0	128
9	0.2	200	0.075	0.10	14	0	113
10	0.4	100	0.025	0.07	43	0	169
11	0.4	100	0.050	0.09	9	0	122
12	0.4	100	0.075	0.09	2	0	108
13	0.4	150	0.025	0.14	69	0	155
14	0.4	150	0.050	0.11	4	0	123
15	0.4	150	0.075	0.11	8	0	194
16	0.4	200	0.025	0.16	152	1	98
17	0.4	200	0.050	0.14	96	1	157
18	0.4	200	0.075	0.17	13	1	169
19	0.6	100	0.025	0.14	21	1	104
20	0.6	100	0.050	0.10	19	0	122
21	0.6	100	0.075	0.10	23	0	166
22	0.6	150	0.025	0.12	81	0	123
23	0.6	150	0.050	0.15	22	0	159
24	0.6	150	0.075	0.15	20	0	129
25	0.6	200	0.025	0.20	241	1	104
26	0.6	200	0.050	0.22	140	1	192
27	0.6	200	0.075	0.26	15	4	82

Table 3: Final volume ratios using the halting algorithm.

Test case	Best design		Best design seen using halting strategy												
			2x	Ratio	Halted	3x	Ratio	Halted	4x	Ratio	Halted	10x	Ratio	Halted
1	0.07	108	0.16	2.29	97	0.07	1.00	134	0.07	1.00	136	0.07	1.00	161
2	0.08	118	0.11	1.38	93	0.11	1.38	94	0.11	1.38	100	0.09	1.13	130
3	0.07	93	0.07	1.00	96	0.07	1.00	100	0.07	1.00	104	0.07	1.00	123
4	0.07	100	0.07	1.00	104	0.07	1.00	107	0.07	1.00	109	0.07	1.00	122
5	0.08	128	0.1	1.25	107	0.1	1.25	109	0.09	1.13	115	0.09	1.13	132
6	0.08	111	0.1	1.25	95	0.1	1.25	97	0.1	1.25	103	0.09	1.13	120
7	0.08	94	0.09	1.13	118	0.09	1.13	119	0.08	1.00	127	0.08	1.00	140
8	0.09	129	0.11	1.22	105	0.11	1.22	107	0.11	1.22	109	0.1	1.11	144
9	0.1	103	0.18	1.80	86	0.18	1.80	89	0.1	1.00	117	0.1	1.00	133
10	0.19	98	0.2	1.05	101	0.2	1.05	111	0.2	1.05	113	0.2	1.05	156
11	0.09	109	0.11	1.22	99	0.1	1.11	106	0.1	1.11	115	0.09	1.00	131
12	0.09	104	0.09	1.00	118	0.09	1.00	120	0.09	1.00	126	0.09	1.00	143
13	0.14	89	0.14	1.00	162	0.14	1.00	164	0.14	1.00	166	0.14	1.00	187
14	0.11	95	0.12	1.09	98	0.12	1.09	100	0.12	1.09	101	0.12	1.09	107
15	0.11	144	0.13	1.18	94	0.13	1.18	95	0.13	1.18	97	0.13	1.18	113
16	0.16	90	0.17	1.06	95	0.16	1.00	101	0.16	1.00	102	0.16	1.00	117
17	0.14	132	0.18	1.29	101	0.18	1.29	103	0.18	1.29	145	0.18	1.29	119
18	0.17	169	0.24	1.41	95	0.22	1.29	104	0.22	1.29	106	0.21	1.24	116
19	0.14	88	0.14	1.00	113	0.14	1.00	115	0.14	1.00	118	0.14	1.00	150
20	0.1	100	0.11	1.10	109	0.11	1.10	111	0.1	1.00	124	0.1	1.00	139
21	0.1	125	0.44	4.40	59	0.16	1.60	95	0.16	1.60	96	0.11	1.10	147
22	0.12	108	0.19	1.58	95	0.13	1.08	111	0.13	1.08	112	0.12	1.00	130
23	0.15	113	0.17	1.13	107	0.16	1.07	116	0.16	1.07	117	0.16	1.07	128
24	0.15	125	0.19	1.27	98	0.19	1.27	99	0.19	1.27	100	0.17	1.13	117
25	0.2	96	0.21	1.05	98	0.21	1.05	101	0.2	1.00	106	0.2	1.00	117
26	0.22	128	0.3	1.36	75	0.28	1.27	82	0.28	1.27	105	0.27	1.23	96
27	0.25	83	0.25	1.00	86	0.25	1.00	87	0.25	1.00	88	0.25	1.00	100

Table 4: Changes in Ψ waiting for three instead of two failed models.

Test case	Reduction in Ψ	Additional iterations
1	9%	37
10	0%	10
21	28%	36
22	6%	16

Table 5: Changes in Ψ waiting for four instead of three failed models.

Test case	Reduction in Ψ	Additional iterations
9	8%	28
20	1%	13

occurred varied, but was approximately less than 100 iterations. This allowed for comparison against results in which at least half of the optimization was at or near the critical volume ratio.

Increasing the counter limit would cause a longer search, which may have found a better design, at the cost of increased computation. A larger limit required more failures against the continuity constraint condition before halting. In changing from two to three repeated failures of the constraint condition, most test cases extended their optimization runs for fewer than ten iterations before halting. Four extended the optimization run by ten or more. These are shown in Table 4. Three of these four case studies also arrived at improved designs. The exception was test case 10 which experienced ten further iterations. In this case the increase of 1 to the counter limit only led to large increases in the computation time for case studies that also produced improved designs. A similar situation held when going from three to four failures for the counter limit. These results are shown in Table 5. Here two test cases experienced more than ten additional iterations. These also experienced an improvement in reducing the volume of the best successful design. The change from four to ten corresponded to significant increases in the number of iterations without significant changes to the mass of the best designs. There appeared to be a range, toward the lower end of the interval between four and ten after which significant increases in optimization time did not correspond to finding significantly more efficient designs.

There is a relationship between the size of the increment/decrement by which material is added or removed from the model and the volatility in the volume ratio history during phase 2 of the optimization. This effect is illustrated in Figure 23 where the optimization of test case 14 was repeated for different values of the evolutionary rate, Ω. Here the optimization was allowed to run for 200 iterations, as for the test cases with results in Table 3, except for $\Omega = 0.0025$ where the evolutionary rate was so low that it required significantly further iterations to complete. There is an approximate

relationship that a larger value of Ω will lead to greater variability in phase two of the optimization.

Increasing Ω corresponded to steeper decreases in volume ratio in the first phase of the optimization. Halting at the most suitable iteration is an easier problem to solve for smaller adjustment ratios. This would suggest that there is an ideal compromise between the extended time the optimization remains in phase 1, and the increased difficulty that an optimization method would have to reach a converged solution. Previous results on optimizing structures for maximum stiffness (Chu et al (1997)) and for even stress distributions (Abolbashari and Keshavarzmanesh (2006)) have found that the final result for 2D structures is not significantly diminished with increasing removal ratios, however, results for the former were all had similar shapes whereas the latter produced similarly performing results with different final shapes. The experience from this paper is that a value of $\Omega = 2$ produced large variations within phase 2 of the optimization.

It is already established and discussed in literature that the BESO optimization method, by simultaneously adding and removing material, can provide suitable solutions to some non-convex problems. Results from this case study show that different designs are obtained at the same volume ratio during the course of an optimization. With regards to the convexity of this optimization problem, the algorithm is exploring a space outside that which would have been explored using simply the BESO method with a fixed volume ratio. The first suitable solution found before phase two begins is not necessarily the best found after further optimization iterations. Results in Table 3 illustrate during the optimization where improved solutions were reached after revisiting a particular volume ratio. In fact, some designs were demonstrated to be suitable at volume ratios lower than other failing designs at higher volume ratios.

If a larger counter limit than necessary is chosen, it will cause the optimization to remain in phase two for a longer time without producing successful designs of volume ratio lower than those already encountered. Therefore the value of ρ selected must be larger than the amount of variation experienced during the regular cycling of stage two, but less than the range of Ψ during stage 1. In the case of this case study the variation was between 2-10% for the range of cycles in stage 2 and 90% for the range of Ψ in stage 1. Therefore a value of ρ between these two values (between 0.1 and 0.9) would have been suitable to identify the transition to stage 2. This is also illustrated in Figure 20. The value of ρ should be at the lower end of that range, without exceeding the lower limit. This avoids misclassifying some failed designs in stage 1 as occurring in stage 2, although, as shown in Figure 20, this transition between stages 1 and 2 can often be continuous and this is managed by the selection of the counter limit.

This halting method worked by monitoring the variation in Ψ and halting once this variation behaved in a manner characteristic of a structure unlikely to yield
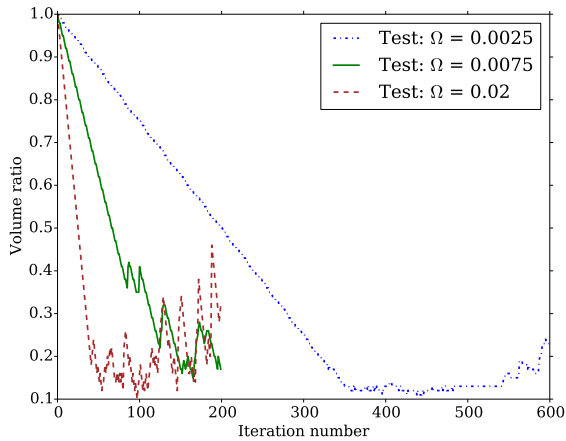
Fig. 23: Volume ratio history for different adjustment ratios.

further results. It was reasoned that a small number of small cycles in Ψ were behaviour indicative of this condition, with *small* defined by the analyst. Previous methods for halting have been presented in the literature. The BESO convergence criterion (Huang et al (2006)) will transition between an initial phase of changing Ψ to a second phase in which Ψ remains fixed. It then measures changes in performance between iterations and halts when this performance falls below a threshold value. Changes in designs during this second phase were subtle, with performance and visual comparisons (see (Huang et al (2006))) showing very small changes between iterations. In comparison, in this research, the value of Ψ was varied up to the point when the optimization halted. The BESO convergence criterion compared designs at the same Ψ. Any change in performance would be much more closely related to the change in volume ratio rather than how close or far the optimization was from converging to any solution. In place of the performance index of the BESO method, a similar quantity described by the concept of extended optimality (Rozvany et al (2002)) can be used. Here the topology and volume fraction were simultaneously varied to try to find the design that minimised/maximised the ratio between the objective function and the volume fraction.

A halting method would necessarily need to ensure that Ψ was not set higher than necessary, but that there was sufficient material in a structure such that $\Theta > \Theta_{max}$. It is therefore necessary that any halting method work ensuring that the correct value of Ψ has been reached, which meets these two conditions. Other methods, such as those just described seek instead to halt once the optimization has converged and that the result is the most efficient structure. Whereas this type of efficiency is desired, for real structures put into service the need to also meet practical performance requirements is important. The structure that is the theoretical most efficient, but fractures during operation, is not the most suitable. The efficiency of the design is

unrelated to the constraint condition. A very high efficiency tells nothing about whether the design will be able to meet the design requirements. The value of Ψ will vary near to a particular value at a point during the optimization.

The method used in this work was simple and straightforward, making it easy to understand and implement. Whereas it may be the case that this can be improved upon, it is an example of a halting method that relates its operation to the volume ratio history and to the continuity constraint. It halts for a design successful against the constraint and at a point at which further optimization efforts are unlikely to continue to be very fruitful. It does require the use of two constants, but default values should be suitable if a designer does not want/need to specify them.

Although not explored here, the previously mentioned methods and performance measuring statistics could be used in a second optimization phase in which the halting method presented in this research triggers a change between phases rather than a complete halting. A caveat to this approach is that the constraint must still be respected, leading to a possibility that designs when at a constant Ψ may still be generated which violate the constraint condition, in which case this will still need to be managed. Examples (Huang et al (2006)) have shown that optimizing a structure at a fixed volume ratio can provide improvements to the final optimization results.

7 Conclusion

A method was presented for the topology optimization of mechanical structures that absorb energy and in which material damage is used to formulate the optimization task. It used the BESO method augmented with a continuity constraint condition to ensure designs successfully absorbed energy from loads without global failure. This was confirmed by results in which the mass of test case designs varied in proportion with the energy that needed to be absorbed, as well as visual observation of some example models in which the algorithm correctly differentiated between successful and unsuccessful designs. A method to allow the optimization algorithm to determine when to halt was also presented. The halting method was shown to be able to halt the optimization algorithm once it had reached a volume ratio close to the critical volume ratio. This occurred after the optimization had transitioned to its second phase. Guides on how an analyst can use these algorithms in the design process were also presented along with the motivation and evidence for that advice.

8 Acknowledgement

References

Abolbashari MH, Keshavarzmanesh S (2006) On various aspects of application of the evolutionary structural optimization method for 2D and 3D continuum structures. Finite Elements in Analysis and Design 42(6):478–491, URL http://www.sciencedirect.com/science/article/pii/S0168874X05001216

Ambur DR, Jaunky N, Lawson RE, Knight Jr NF (2001) Numerical simulations for high–energy impact of thin plates. International Journal of Impact Engineering 25(7):683–702, URL http://www.sciencedirect.com/science/article/pii/S0734743X00000737

Andreasen CS, Andreassen E, Jensen JS, Sigmund O (2014) On the realization of the bulk modulus bounds for two-phase viscoelastic composites. Journal of the Mechanics and Physics of Solids 63(0):228–241, URL http://www.sciencedirect.com/science/article/pii/S0022509613001828

Avalle M, Chiandussi G, Belingardi G (2002) Design optimization by response surface methodology: application to crashworthiness design of vehicle structures. Structural and Multidisciplinary Optimization 24(4):325–332, URL http://dx.doi.org/10.1007/s00158-002-0243-x

Børvik T, Hopperstad OS, Berstad T (2003a) On the influence of stress triaxiality and strain rate on the behaviour of a structural steel. Part II. Numerical study. European Journal of Mechanics - A/Solids 22(1):15–32, URL http://www.sciencedirect.com/science/article/pii/S0997753802000050

Børvik T, Hopperstad OS, Langseth M, Malo KA (2003b) Effect of target thickness in blunt projectile penetration of Weldox 460 E steel plates. International Journal of Impact Engineering 28(4):413–464, URL http://www.sciencedirect.com/science/article/pii/S0734743X02000726

Chu DN, Xie YM, Hira A, Steven GP (1996) Evolutionary structural optimization for problems with stiffness constraints. Finite Elements in Analysis and Design 21(4):239–251, URL http://www.sciencedirect.com/science/article/pii/0168874X9500043S

Chu DN, Xie YM, Hira A, Steven GP (1997) On various aspects of evolutionary structural optimization for problems with stiffness constraints. Finite Elements in Analysis and Design 24(4):197–212, URL http://www.sciencedirect.com/science/article/pii/S0168874X96000492

Cui L, Kiernan S, Gilchrist MD (2009) Designing the energy absorption capacity of functionally graded foam materials. Materials Science and Engineering: A 507(12):215–225, URL http://www.sciencedirect.com/science/article/pii/S0921509308013646, j.msea.2008.12.011

Dassault Systèmes (2011) Abaqus Analysis User's Manual

Du J, Olhoff N (2007) Topological design of freely vibrating continuum structures for maximum values of simple and multiple eigenfrequencies and frequency gaps. Structural and Multidisciplinary Optimization 34(2):91–110, URL http://dx.doi.org/10.1007/s00158-007-0101-y

Duysinx P, Bendsøe MP (1998) Topology optimization of continuum structures with local stress constraints. International Journal for Numerical Methods in Engineering 43(8):1453–1478, URL http://dx.doi.org/10.1002/(SICI)1097-0207(19981230)43:8<1453::AID-NME480>3.0.CO;2-2

Forsberg J, Larsgunnar N (2007) Topology optimization in crashworthiness design. Structural and Multidisciplinary Optimization 33(1):1–12, URL http://dx.doi.org/10.1007/s00158-006-0040-z

Guo J, Shi G, Wang Y, Lu C (2003) Efficient modeling of panel-like structures in perforation simulations. Computers & Structures 81(1):1–8, URL http://www.sciencedirect.com/science/article/pii/S0045794902003899

Huang X, Xie YM (2007) Convergent and mesh-independent solutions for the bi-directional evolutionary structural optimization method. Finite Elements in Analysis and Design 43(14):1039–1049, URL http://www.sciencedirect.com/science/article/pii/S0168874X07000777

Huang X, Xie YM (2010) Evolutionary Topology Optimization of Continuum Structures: Methods and Applications. John Wiley & Sons, Ltd., West Sussex, UK

Huang X, Xie YM, Burry MC (2006) A new algorithm for Bi-Directional Evolutionary Structural Optimization. JSME international journal Series C, Mechanical systems, machine elements and manufacturing 49(4):1091–1099, URL http://ci.nii.ac.jp/naid/110005850896/en/

Huang X, Xie YM, Lu G (2007) Topology optimization of energy-absorbing structures. International Journal of Crashworthiness 12(6):663–675, URL http://www.tandfonline.com/doi/abs/10.1080/13588260701497862, 10.1080/13588260701497862, http://www.tandfonline.com/doi/pdf/10.1080/13588260701497862

Johnson GR, Cook WH (1983) A constitutive model and data for metals subjected to large strains, high strain rates and high temperatures. In: Proc. 7th International Symposium on Ballistics, The Hague, The Netherlands, pp 541–547

Johnson GR, Cook WH (1985) Fracture characteristics of three metals subjected to various strains, strain rates, temperatures and pressures. Engineering Fracture Mechanics 21(1):31–48, URL http://www.sciencedirect.com/science/article/pii/0013794485900529

Jones E, Oliphant T, et al PP (2001–) SciPy: Open source scientific tools for Python. Http://www.scipy.org/

Jung D, Gea HC (2004) Topology optimization of nonlinear structures. Finite Elements in Analysis and Design 40(11):1417–1427, URL http://www.sciencedirect.com/science/article/pii/S0168874X03002117

Jung D, Gea HC (2006) Design of an energy-absorbing structure using topology optimization with a multimaterial model. Structural and Multidisciplinary Optimization 32(3):251–257, URL http://dx.doi.org/10.1007/s00158-006-0011-4

Kay G, Goto D, Couch R (2007) Statistical testing and material model characterization of aluminum and titanium for transport airplane rotor burst fragment shielding. Technical Report DOT/FAA/AR-07/26, Lawrence Livermore National Laboratory, Livermore, California

Knight-Jr NF, Jaunky N, Lawson RE, Ambur DR (2000) Penetration simulation for uncontained engine debris impact on fuselage-like panels using LS-DYNA. Finite Elements in Analysis and Design 36(2):99–133, URL http://www.sciencedirect.com/science/article/pii/S0168874X00000111

Li Q, Steven GP, Xie YM, Querin OM (2004) Evolutionary topology optimization for temperature reduction of heat conducting fields. International Journal of Heat and Mass Transfer 47(23):5071–5083, URL http://www.sciencedirect.com/science/article/pii/S0017931004002352

Lu G, Yu T (2003) Energy Absorption of Structures and Materials. Woodhead Publishing, Cambridge, UK, URL http://www.knovel.com/web/portal/browse/display?_EXT_KNOVEL_DISPLAY_bookid=2524\&VerticalID=0

Mayer RR, Kikuchi N, Scott RA (1996) Application of topological optimization techniques to structural crashworthiness. International Journal for Numerical

Methods in Engineering 39(8):1383–1403, URL `http://dx.doi.org/10.1002/(SICI)1097-0207(19960430)39:8<1383::AID-NME909>3.0.CO;2-3`

Neves MM, Rodrigues H, Guedes JM (1995) Generalized topology design of structures with a buckling load criterion. Structural optimization 10(2):71–78, URL `http://dx.doi.org/10.1007/BF01743533`

Ou JS, Kikuchi N (1996) Optimal design of controlled structures. Structural optimization 11(1-2):19–28, URL `http://dx.doi.org/10.1007/BF01279651`

Pearce DJ (2005) An improved algorithm for finding the strongly connected components of a directed graph. Tech. rep., Victoria University, http://homepages.mcs.vuw.ac.nz/∼djp/files/P05.pdf

Pedersen CBW (2003) Topology optimization design of crushed 2D-frames for desired energy absorption history. Structural and Multidisciplinary Optimization 25(5-6):368–382, URL `http://dx.doi.org/10.1007/s00158-003-0282-y`

Pedersen CBW (2004) Crashworthiness design of transient frame structures using topology optimization. Computer Methods in Applied Mechanics and Engineering 193(68):653–678, URL `http://www.sciencedirect.com/science/article/pii/S0045782503005826`

Querin OM (1997) Evolutionary structural optimisation: Stress based formulation and implementation

Rozvany GIN, Querin OM, Gaspar Z, Pomezanski V (2002) Extended optimality in topology design. Structural and Multidisciplinary Optimization 24(3):257–261, URL `http://dx.doi.org/10.1007/s00158-002-0235-x`

Sigmund O, Petersson J (1998) Numerical instabilities in topology optimization: A survey on procedures dealing with checkerboards, mesh-dependencies and local minima. Structural and Multidisciplinary Optimization 16(1):68–75, URL `http://dx.doi.org/10.1007/BF01214002`

Stamper E, Hale S (2008) The use of LS-DYNA models to predict containment of disk burst fragments. In: 10th International LS-DYNA users conference, Livermore Software Technology Corp., Michigan USA

Teng X, Wierzbicki T (2005) Numerical study on crack propagation in high velocity perforation. Computers & Structures 83(12-13):989–1004, URL `http://www.sciencedirect.com/science/article/pii/S0045794905000490`

Xie YM, Steven GP (1993) A simple evolutionary procedure for structural optimization. Computers & Structures 49(5):885–896, URL `http://www.sciencedirect.com/science/article/pii/004579499390035C`