Multiple objective optimisation of composite sandwich structures for rail vehicle floor panels

Craig W. Hudson, Joe J. Carruthers*, A. Mark Robinson

NewRail, Newcastle University, Newcastle upon Tyne, NE1 7RU, England, UK

ABSTRACT

This paper describes the application of an ant colony optimisation (ACO) algorithm to the multiple objective optimisation of a rail vehicle floor sandwich panel. The ACO algorithm was used to search a design space that was defined by sandwich theory and a material database in order to identify constructions that were optimal with respect to low mass and low cost. A broad range of mass and cost optimal sandwich material designs were identified successfully. These provided mass savings of up to 60% compared to existing plywood-based flooring systems, although mass savings above 40% had an associated cost premium.

Keywords:
Sandwich material
Ant colony optimisation (ACO)
Rail vehicle
Floor panel

1. Introduction

Within the rail industry, lightweighting is becoming an increasingly important topic. Recent studies (e.g. [1]) have indicated that rail vehicles have generally become heavier over the last thirty years. Whilst these increases in vehicle mass can often be attributed to enhanced passenger environments (e.g. the provision of air-conditioning, improved accessibility, crashworthiness, etc.), there are clearly undesirable side-effects of heavier trains. Everything else being equal, a heavier vehicle will consume more energy in operation than a lighter one, thereby making it more costly to run. Increased energy consumption also implies a likelihood of higher CO_2 emissions at some point in the energy supply chain. Furthermore, heavier vehicles are likely to cause more damage to the track, thereby resulting

* Corresponding author. Tel.:+44 1246 281 634.
E-mail address: joe.carruthers@ncl.ac.uk (J.J. Carruthers).

in higher costs for infrastructure maintenance and renewal. In some countries, heavier vehicles also attract higher track access charges for operators.

A recent investigation [2] by a cross-industry consortium of rail vehicle manufacturers examined some of the issues surrounding the increased use of lightweight materials in metro vehicles. As part of this work, a number of applications were identified that were considered to have a high potential for lightweighting through material substitution. One such application was interior floor panels.

A typical six-car metro vehicle will have around 250 m^2 of flooring material as part of its interior (Fig. 1). This is likely to weigh a total of around 4 tonnes, thereby representing a significant lightweighting opportunity. In terms of functionality, the most fundamental requirement of a floor construction is that it is capable of supporting the loads induced by passengers without excessive deflection or failure. Additionally, floor constructions must also provide the required level of insulation. It can be seen from Fig. 2 that current interior floor constructions are often quite complex multi-material assemblies employing woods, inserts, elastomers and insulative materials. Is there a material configuration that would provide a lighter solution at a competitive cost?

Given the combined requirements of high stiffness, low weight and good insulation, it seemed interesting to investigate the concept of a sandwich design. Sandwich materials, consisting of two thin, stiff facings separated by a low density core, can be used to produce structures that are both light and flexurally rigid. However, the optimisation of sandwich structures is not straightforward. This is because there are typically multiple design variables, objectives and constraints. For example, for a given application, what is the "best" material to use for the facings? And for the core? How thick should the facings and core be? And how does one resolve the potential conflict between a construction that is optimised for low weight and one that is optimised for low cost?

Previous work on the optimisation of sandwich structures has approached the subject from a number of different perspectives. Steeves & Fleck [3] followed the approach of Gibson & Ashby [4] to produce failure mechanism maps for sandwich beams. However the study was restricted to a single objective problem (minimum mass in three-point bend) for which a characteristic mass index was minimised for a given load, material and geometry. More recently, Pflug & Verpoest [5] extended the well known Ashby [6] material selection chart method to include multiple objective sandwich problems. However, whilst such Ashby-based methods have been used to accommodate multiple objectives and even identify Pareto-optimal solution sets [7], their general approach, alongside other analytical techniques, is somewhat contrary to the direction taken in this paper. This is because they generally rely on narrowing down an exhaustive set of material options so that a decision can be made between a manageable few. In this paper, the aim was to keep the range of material combinations deliberately large to allow any potentially new or non-obvious solutions to be discovered. For this reason, *metaheuristic algorithms*, which use a population of search agents to identify optimal solutions within a design space, provided the basis for the investigation.

An early application that used metaheuristic algorithms for sandwich design is described by Bassetti et al. [8]. A genetic algorithm was used to optimise an insulating sandwich panel for a truck. However, only a single objective (minimum mass) optimisation was considered into which different stiffness or strength criteria could be integrated. Furthermore, the optimisation did not consider the material composition or lay-up of the laminated composite sandwich facings, which would otherwise have significantly increased the magnitude of the search space.

In a comparative study of different metaheuristic algorithms for sandwich design, Hudson [9] found that the ant colony optimisation (ACO) algorithm was more competitive than other techniques such as particle swarm optimisation and simulated annealing. For this reason, ACO will be the focus of this paper. The ACO algorithm aims to mimic the behaviour of ants searching for food. As each "ant" searches the design space, it leaves behind a pheromone trail that fellow ants can follow. The better solutions tend to build up a stronger pheromone scent as more ants are drawn towards them, whilst evaporation of the pheromone after each iteration of the algorithm discourages poorer solutions from being followed. Abachizadeh & Tahani [10] used ACO to examine the optimisation of a simply supported composite laminate. The optimisation considered two material options (carbon-epoxy and glass-epoxy), as well as the orientation angle of individual laminae. Furthermore, two design objectives were taken into account – low cost and high natural frequency. However, the number of possible alternative solutions was relatively small and only a single amalgamated objective function, with stiffness constraints, was optimised. This led to the identification of a single optimum design, rather than a broader set of Pareto-optimal solutions. Similarly, Aymerich & Serra [11] investigated the stacking sequence of a composite laminate with strength constraints. But again, it was only a single objective optimisation (to maximise buckling strength) in which one optimum solution was the intended goal.

This paper describes the application of an ACO algorithm to multiple objective sandwich material design. Using a metro vehicle floor panel as a case study, it describes how the ACO algorithm can be used to search a design space that is defined by a large material database and sandwich theory in order to identify Pareto-optimal sandwich constructions.

2. Problem definition

An interior flooring arrangement similar to that depicted in Fig. 2 will be considered. It consists of a series of sandwich floor panels supported by an underlying timber framework (Fig. 3). The optimisation will consider both the construction of the sandwich floor panels and the spacing of the supporting timber joists. So there will be a trade-off between having a more substantial supporting framework and less structural panels, or having larger supporting spans and stiffer panels. The main (exterior) structural floor, which is part of the vehicle bodyshell structure, is not considered in the analysis. In sections 2.1 to 2.4 that follow, the optimisation problem is defined in terms of the objectives, the variables, the constraints and the governing physical equations.

2.1. Objectives of the optimisation

The objective of the optimisation was to identify those sandwich flooring assemblies that were optimal for both low mass and low cost. Multiple objective handling in optimisation studies is not straightforward and the topic has been considered extensively elsewhere [12]. However, a brief outline of the fundamental methodology for retrieving optimal solutions from multiple objective problems, based upon "non-dominance", is described here.

Usually in multiple objective optimisations, there is no single optimal solution. Instead, a series of solutions exist that each contain an element of optimality. Consider, for example, an ordinary (non-sandwich) beam of fixed dimensions. Suppose that there was a requirement to optimise the mass of this beam subject to a certain minimum stiffness. If the beam material was the only variable, the optimisation would be trivial. The material with the lowest density that still met the required stiffness would be selected. Similarly, if the sole objective was to minimise the cost of the beam, the optimal material would be the cheapest option.

However, if the objective was instead to optimise both the mass AND the cost of the beam subject to a certain minimum stiffness, the situation becomes less clear. This is because it is unlikely that the material that produces the lightest solution would also provide the cheapest solution. Instead, when both objectives are considered, a trade-off boundary between mass and cost is formed. The result is a set of solutions which, when all objectives are considered, show some degree of optimal quality. The optimal solutions within this set are "non-dominated" with respect to all the other known solutions and are called the Pareto-optimal set. A non-dominated solution is one which, when compared to any other solution, has at least one superior objective value, or is equal across all objective values.

For the floor system considered here, the objective was to find the Pareto set of sandwich constructions that are optimal for both low mass and low cost. The two objective functions to be minimised therefore are given in Eq. (1) and (2).

Mass objective function:

$$M = \{\rho_{f1}t_{f1} + \rho_{f2}t_{f2} + \rho_c t_c\} + \left\{\frac{m_s}{lb}\left[l + (b - b_s)\right]\right\} \tag{1}$$

where M is the total mass per unit area of the sandwich panel and its supports, ρ is density, t is thickness, l is the length of the sandwich panel, b is the width of the sandwich panel, m_s is the mass per unit length of the supporting timbers, b_s is the width of the supporting timbers, and subscripts *f1*, *f2*, *c* and *s* pertain to the upper sandwich facing, lower sandwich facing, sandwich core and timber supports respectively. So the first curly bracketed term in Eq. (1) represents the mass contribution of the sandwich panel, and the second curly bracketed term represents the mass contribution of the supporting timber framework.

Cost objective function:

$$C = \{\rho_{f1}t_{f1}c_{f1} + \rho_{f2}t_{f2}c_{f2} + \rho_c t_c c_c\} + \left\{\frac{m_s c_s}{lb}\left[l + (b - b_s)\right]\right\} \tag{2}$$

Where C is the total cost per unit area of the sandwich panel and its supports, and c is the cost per unit mass of an individual component in the system.

2.2. Design variables

Table 1 summarises the main design variables. These are the parameters that the ant colony optimisation algorithm sought to obtain optimal values for.

For the facing materials, the optimisation algorithm was provided with a range of options to choose from including various aluminiums, steels, fibre-reinforced polymers and wood products. Furthermore, for the fibre-reinforced polymer facings, the algorithm could select between a range of fibre and matrix materials, as well as specifying the fibre volume fraction, the number of plies in the laminate, and the orientation angle of each ply (0^o, $+45^o$, -45^o or 90^o). Similarly, a number of core material options were available, including a variety of polymer foams, honeycombs and balsa woods of different densities. In total, there were 40 different core and facing materials for the algorithm to choose from. This material database, when coupled with the fibre-reinforced polymer laminate design options, provided a very large number of potential sandwich material combinations.

The upper limits on the sandwich floor span (i.e. the spacing between the underlying timber supports) were defined by a typical maximum panel size that can be manufactured in an industrial press.

2.3. Design constraints

Clearly, for the optimisation algorithm to be useful, it must be capable of discriminating between those sandwich constructions that are fit-for-purpose and those that are not. This fitness-for-purpose was defined by a number of design constraints or requirements that any prospective sandwich must satisfy. The constraints employed for the sandwich floor application were as follows:

- The sandwich must be sufficiently stiff, i.e. it must not deflect excessively under passenger loading. The limiting deflection was set at a maximum of 1 mm under a distributed load of 6000 N/m^2.
- The sandwich must provide sufficient thermal insulation. The maximum allowable thermal conductance of the sandwich was set at 0.0025 W/K, which is equivalent to the performance that might be expected from a conventional non-sandwich floor construction consisting of a 20 mm plywood panel with 30 mm of glass wool insulation.
- The upper facing must be sufficiently resilient to high localised loadings (e.g. heeled shoes). This aspect was arbitrarily handled by stipulating that the product of the upper facing Young's modulus and the upper facing thickness should be greater than 100 MN/m.
- The maximum allowable sandwich thickness was set at 20 mm. Again, for equivalence with a typical existing plywood panel.
- The maximum allowable panel dimension was set at 2.5 m x 1.5 m – the dimensions of a typical industrial panel press.
- The sandwich must not fail under passenger loading. The failure modes considered for the sandwich included tensile and compressive failure of the facings due to bending, shear and compressive failure of the core, and wrinkling of the facings.

The supporting timber joists were also assumed to be constant in terms of their material and geometry and were therefore constrained. They had a mass per unit length (m_s) of 0.9 kg/m and a panel-supporting width (b_s) of 100 mm.

2.4. Governing equations

The properties and performance of the sandwich materials and their constituents were estimated using analytical "textbook" solutions.

The fibre-reinforced polymer facing stiffness properties were estimated using classical laminate theory. This is well described in many standard texts (e.g. Gibson [13], Matthews & Rawlins [14]). To simplify the laminate equations, only balanced, symmetric laminates were considered, although orthotropic constructions were permitted.

The mechanics of the sandwich panels were estimated using sandwich plate theory, as described, for example, by Allen [15] and Zenkert [16]. Each facing was considered separately, so that the upper sandwich facing could be of a different material and thickness to the lower sandwich facing. The analytical expression employed for panel deflection, w_{max}, assumed that a given section of sandwich was simply-supported around its periphery (as a worst case boundary condition from a deflection perspective). The governing equation was [15]:

$$w_{\max} = \frac{qb^4}{D_x}\beta_1 \tag{3}$$

where q is the uniformly distributed load, β_1 is a sandwich coefficient [15], and D_x is the sandwich flexural rigidity in the x-direction (parallel to the length, l, of the panel) given by:

$$D_x = d^2\left(\frac{1}{E_{x,f1}t_{f1}} + \frac{1}{E_{x,f2}t_{f2}}\right)^{-1} \tag{4}$$

where d is the distance between centrelines of opposing facings and E_x is the Young's modulus in the x direction. The stiffness expression in Eq. (4) is applicable for sandwich panels with orthotropic facings of unequal thickness and different materials.

With respect to failure prediction, the tensile and compressive stresses due to bending in the facings, and the shear and compressive stresses in the core were compared against the respective material strengths. For the isotropic facings, the von Mises failure criterion was used. For the fibre-reinforced polymer facings, first ply failure was estimated using the Tsai-Hill criterion. Local facing wrinkling was also considered using the expression provided by Zenkert [16]:

$$\sigma_{wrinkling} = \frac{\sqrt[3]{E_{x,f1}E_cG_c}}{2} \tag{5}$$

where G is the shear modulus. The critical wrinkling stress, $\sigma_{wrinkling}$, was compared against the facing compressive stress to determine the onset of this mode of failure. Eq. (5) was specifically used to check for wrinkling of the upper ($f1$) facing in the x direction. A similar check was applied for the y direction (parallel to the width, b, of the sandwich).

Whilst more complex and accurate methods of predicting the behaviour of sandwich panels are available, their use would not have fundamentally altered the manner in which the optimisation was performed. It would just have required the substitution of one sandwich design algorithm for another within the optimisation process. For the purposes of this study, the textbook analytical solutions were considered sufficient for the purposes of demonstrating the application of ant colony optimisation for sandwich design.

3. Implementation of the ant colony optimisation (ACO) algorithm

The basic procedure adopted for implementing the ACO algorithm was as follows:

1. The population, or "colony" of ants, starting in their nest, set out to explore the design space. Each ant followed a certain path that was defined by the variables that controlled the problem (Table 1). This was implemented as a series of sequential decisions for each ant: which "facing thickness" path to follow; which "facing material" path to follow; etc. By taking a series of such decisions, each ant's path through the design space was defined. For the first iteration, these decisions were taken at random. For subsequent iterations, the propensity of an ant to pick a particular path was influenced by the pheromone levels left by previous ants (see points 4 – 6 below).
2. When each ant reached the end of its path (i.e. decisions had been made about all the required variables), the values of the corresponding objective functions for each ant were calculated using Eqs. (1) and (2). A check was also made for each path to ensure that no constraints had been violated, i.e. that the resulting sandwich was sufficiently stiff, not likely to fail, etc.
3. The objective values of all viable ants were compared to identify the best solutions. This was done on the basis of non-dominance. The identified best solutions were held in a separate global best repository (the Pareto-optimal set).
4. Except for the first iteration, the pheromone levels from the previous iteration were reduced (through "evaporation").
5. Variables that have been visited in the previous iteration had their pheromone levels increased in proportion to the number of ants that visited.
6. Extra pheromone was also given to all variables in the Pareto-optimal set in proportion to the number of times they had been visited.

7. The ants returned to their nest and steps 1-7 were repeated until a stopping criterion was satisfied (in this case, a fixed number of iterations).

The importance of parameter selection when setting-up an ACO has been highlighted elsewhere [17]. For this study, suitable algorithm parameters were selected on an observational basis and from the results of previous work [18]. The key parameters employed were as follows:

- Number of ants = 10.
- Number of iterations = 200,000.
- Maximum size of Pareto-optimal set = 50.
- Evaporation rate = 0.1 (i.e. during each iteration, the pheromone level for each variable reduces naturally by 90%).
- $\alpha = \beta = 0.5$ (parameters controlling the pheromone levels of currently popular paths and Pareto-optimal solutions respectively).

4. Results and discussion

The sequential graphs in Figure 4 illustrate the dynamic evolution of the ACO over an increasing number of iterations of the algorithm. Each graph shows both the position of the individual ants during the given iteration, and the current non-dominated Pareto-optimal solutions. In Figure 4a, the initial (random) distribution of calculated objective functions is shown. During the early stages of the optimisation (the first 10 iterations, Figure 4b), rapid progression was observed, with the Pareto-optimal solutions showing marked improvements between successive iterations as they moved towards the low mass and low cost regions of the design space. After around 500 iterations (Figure 4c), incremental improvements to existing best solutions had become smaller, and a larger and more diverse set of Pareto-optimal solutions had been identified. As the number of iterations continued to increase, changes to the Pareto-optimal set became less and less significant, with few improvements beyond 100,000 iterations. The final distribution, after 200,000 iterations, is shown in Figure 4d.

After 200,000 iterations, the ACO had identified a total of 32 non-dominated Pareto-optimal solutions (those plotted in Figure 4d). For validation purposes, a random sample of these solutions were verified manually using the governing equations in order to confirm that the algorithm had performed reliably. A pleasingly broad range of optimal material solutions had been found including extruded polystyrene and polymethacrylimide cores of various densities, and a wide variety of facing materials: carbon fibre-reinforced phenolics, steel and stainless steel for the upper facing; carbon and glass fibre-reinforced phenolics, aluminium, plywood and hardboard for the lower facing. Furthermore, for the fibre-reinforced materials, a range of fibre volume fractions and lay-ups were identified. In terms of the support geometry, the maximum longitudinal span of 2.4 m was preferred in all cases, but for the transverse spans an optimal range of 0.4 – 0.52 m was suggested.

Table 2 summarises the optimised design variables for three representative Pareto-optimal solutions – a low mass option, a low cost option and an intermediate option. The savings in mass and cost are in comparison to a typical existing 2.5 m x 0.5 m x 20 mm plywood / 30 mm glass wool construction with a mass of 12.7 kg/m^2 and a cost of 15 €/m^2 (including timber supports). It can be seen that, from a cost perspective, only the "low cost" option is cheaper than an equivalent plywood panel. Furthermore, this design also provides a 37% mass saving. The lighter "intermediate" and "low mass" optimal solutions were both more expensive than plywood, although their weight savings were also higher at 40% and 53% respectively. However, it should be noted that lightweight designs are likely to provide additional cost savings over and above those associated with materials. For example, an integrated, self-insulating sandwich might have lower installation costs than a separate plywood / glass wool insulation system. There will also be through-life operational cost savings associated with the use of lighter materials. For a single six-car metro vehicle, the estimated annual operational cost saving associated with a 53% reduction in flooring mass would be around 10,000 € [2]. Clearly, for a fleet of vehicles over a 40 year life, such operational cost savings would be very significant.

Finally, an important point to note is that whilst the ACO algorithm attempts to identify the non-dominated (i.e. "best") set of optimal solutions, it cannot be absolutely known that the set generated does indeed match the *true* Pareto-optimal set to the problem. However, by using a large number of iterations (200,000), and by running the simulation multiple times from different random starting positions, a reasonable level of confidence in the results can be obtained.

5. Conclusions

An ant colony optimisation (ACO) algorithm has been applied to the design of sandwich panels for a rail vehicle interior flooring application in which multiple objectives of low mass and low cost were considered. The problem definition and the associated implementation of the algorithm allowed considerable freedom in the choice of both materials and geometry subject to certain constraints associated with fitness-for-purpose.

A broad range of optimal solutions were identified by the ACO. These included sandwich constructions that provided a significant (approximately 40%) saving in both mass and cost compared to the plywood panels that are currently used, as well as designs that provided more significant mass savings (of over 50%), albeit at a cost premium.

References

[1] Ford R. Transport mass. Institution of Mechanical Engineers Seminar: Weight Saving and Structural Integrity of Rail Vehicles, Derby, UK, 25 September 2007.

[2] Carruthers JJ, Calomfirescu M, Ghys P, Prockat J. The application of a systematic approach to material selection for the lightweighting of metro vehicles. J Rail Rapid Transit 2009;223:427-437.

[3] Steeves CA, Fleck NA. Material selection in sandwich beam construction. Scripta Materialia 2004; 50:1335-1339.

[4] Gibson LJ, Ashby MF. Cellular solids: structure and properties. Oxford: Pergamon Press, 1988.

[5] Pflug J, Verpoest, I. Sandwich materials selection charts. J Sandwich Mater Struct 2006;8:407-421.

[6] Ashby MF. Materials selection in mechanical design. London: Butterworth Heinemann, 2005.

[7] Ashby MF. Multi-objective optimization in material design and selection. Acta Materialia 2000;48:359-369.

[8] Bassetti D, Brechet Y, Heiberg G, Lingorski I, Pechambert P. Genetic algorithm and performance indices applied to optimal design of sandwich structures. In: Vautrin A, editor. Mechanics of Sandwich Structures, Proceedings of the EUROMECH 360 Colloquium, Saint-Étienne, France, 13-15 May 1997.

[9] Hudson CW. Population-based techniques for the multiple objective optimisation of composite sandwich materials and structures. PhD thesis, Newcastle University, Department of Mechanical and Systems Engineering; 2009.

[10] Abachizadeh M, Tahani M. An ant colony optimization approach to multi-objective optimal design of symmetric hybrid laminates for maximum fundamental frequency and minimum cost. Structural and Multidisciplinary Optimisation 2009;37:367-376.

[11] Aymerich F, Serra M. Optimisation of laminate stacking sequence for maximum buckling load using the ant colony optimisation (ACO) metaheuristic. Comp A 2008;39:262-272.

[12] Deb K. Multi-objective optimization using evolutionary algorithms. Chichester: Wiley, 2001.

[13] Gibson RF, Principles of composite material mechanics. New York: McGraw-Hill, 1994.

[14] Matthews FL, Rawlins RD. Composite materials: engineering and science. London: Chapman & Hall, 1994.

[15] Allen HG. Analysis and design of structural sandwich panels. Oxford: Pergamon Press, 1969.

[16] Zenkert D. An introduction to sandwich construction. EMAS Publishing, 1995.

[17] Wong KY, Komarudin. Parameter tuning for ant colony optimization: a review. In: Global Links for Human Development, Proceedings of the International Conference on Computer and Communication Engineering (ICCCE08), Kuala Lumpa, Malaysia, 13 – 15 May 2008.

[18] Hudson CW, Carruthers JJ, Robinson AM. Application of particle swarm optimisation to sandwich material design. Plas Rub Comp 2009:38;106-110.

Table 1
Sandwich flooring design variables

Variable	Range	Notes
Facing thickness	0.5 - 5 mm.	Upper and lower facings can have different thicknesses. Different discrete thicknesses within the stated range were permitted for different materials to reflect availability.
Facing material	Selected from a material database.	Upper and lower facings can be of different materials. For fibre-reinforced polymer facings, there are further variables relating to the laminate construction (fibre material, matrix material, fibre volume fraction, and the orientation angle of each ply).
Core material	Selected from a material database.	A range of different densities were available for each core material option.
Span (spacing between timber supports)	0.1 – 2.4 m (longitudinal) 0.1 – 1.4 m (transverse).	The support span can be different in the longitudinal and transverse directions.

Table 2

Representative Pareto-optimal solutions for the metro vehicle floor panels

	Low mass design	Low cost design	Intermediate design
Upper facing material	Carbon fibre- reinforced phenolic	Steel	Steel
Upper facing lay-up	$[0^o/90^o]_s$	-	-
Upper facing fibre volume fraction	0.5	-	-
Upper facing thickness, t_{f1}	1 mm	0.5 mm	0.5 mm
Core material	Extruded polystyrene	Extruded polystyrene	Extruded polystyrene
Core density	40 kg/m^3	40 kg/m^3	45 kg/m^3
Lower facing material	Plywood	Glass fibre-reinforced phenolic	Carbon fibre-reinforced phenolic
Lower facing lay-up	-	$[90^o/90^o]$	$[90^o/90^o]$
Lower facing fibre volume fraction	-	0.35	0.3
Lower facing thickness, t_{f2}	3 mm	0.5 mm	0.5 mm
Longitudinal span	2.40 m	2.40 m	2.40 m
Transverse span	0.42m	0.40 m	0.50 m
M (kg/m^2)	6.0 (53 % reduction)	8.0 (37% reduction)	7.6 (40% reduction)
C (€/m^2)	29 (93% increase)	9 (40% reduction)	17 (13% increase)

Fig. 1. Typical floor panels in a metro vehicle interior.

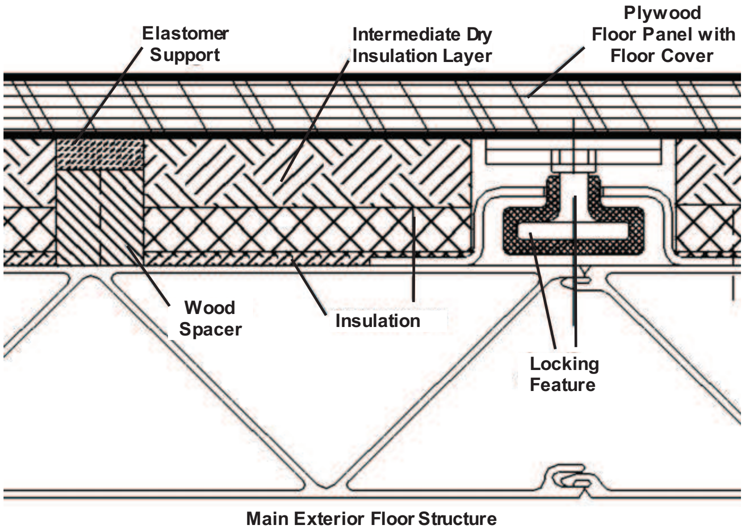

Fig. 2. A cross-section of a typical current interior floor construction employing an assembly of different materials.

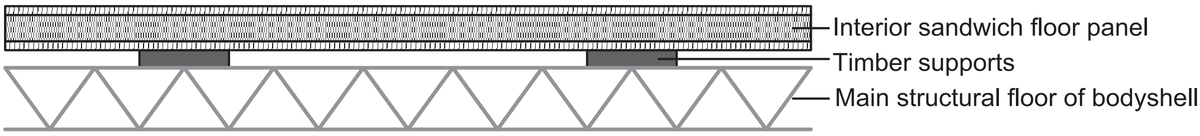

Fig. 3. A cross-section of the assumed configuration of the sandwich flooring.

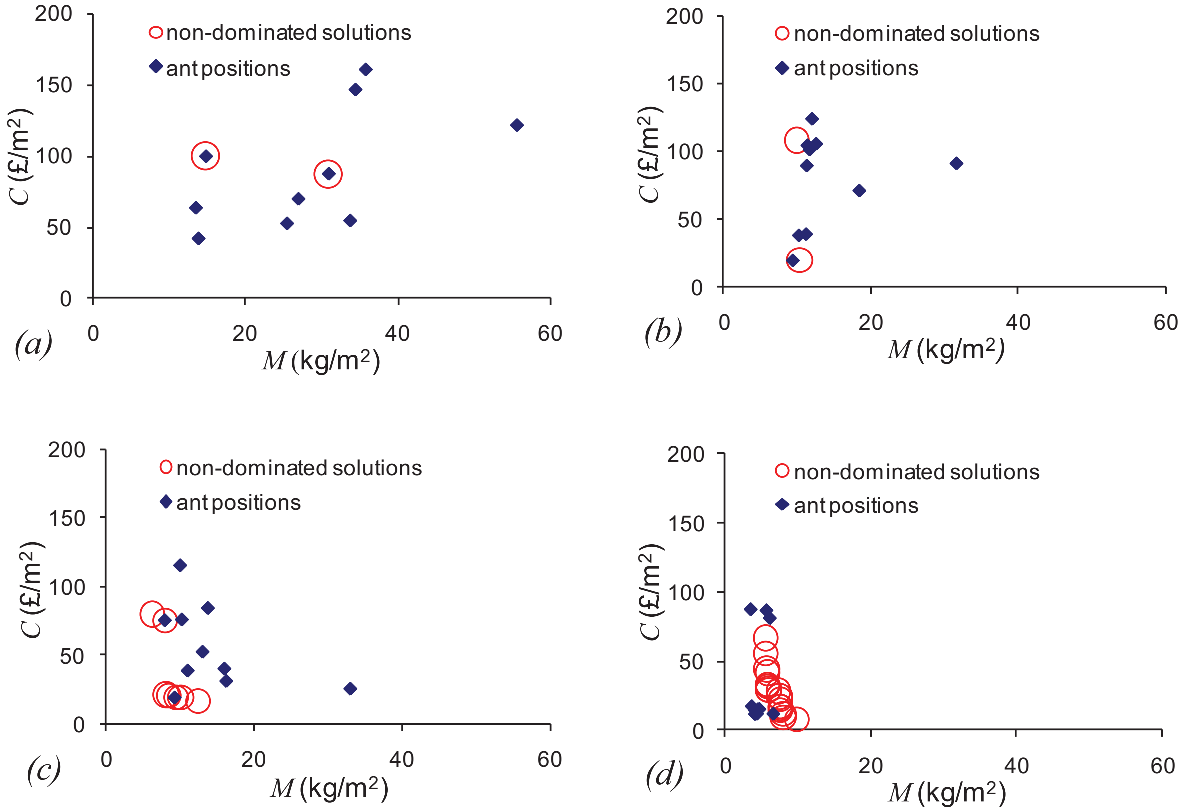

Fig. 4. The progression of the ant colony optimisation after (a) 1 iteration, (b) 10 iterations, (c) 500 iterations and (d) 200,000 iterations.