

Automatic Optimised Design of Umbilicals

Problem presented by

David Fogg

Technip Umbilicals Ltd



ESGI116 was jointly hosted by
Durham University
Smith Institute for Industrial Mathematics and System Engineering



Report author

Jacqueline Christmas (University of Exeter)

Executive Summary

A control umbilical is an assemblage of different cables and pipes, bound together into a single unit for robustness and flexibility. An umbilical design generally comprises a cross-section layout that shows the positions of the specified components within the construction, along with inert filler that provides mechanical protection and stability. The positioning of the components and filler affect the properties of the finished product and the ease and cost of manufacture.

We have investigated possible methods for automating the design process, including how to generate feasible solutions, how to compare the relative quality of different designs, how to determine how similar two designs are, and how an automated system could make use of existing expertise.

All of these elements would contribute to an automated design system

Version 1.0
May 11, 2016
v+55 pages

Contributors

Niall Bootland (University of Oxford)
Ferran Brosa Planella (University of Oxford)
Jacqueline Christmas (University of Exeter)
Davin Lunz (University of Oxford)
Jeff Dewynne (University of Oxford)
Nabil Fadai (University of Oxford)
Artur Gower (University of Manchester)
Peter Hicks (University of Aberdeen)
Vasyl Kovalchuk (Polish Academy of Sciences)
Andrew Lacey (Heriot-Watt University)
Marie-Louise Lackner (Technische Universität Wien)
Sara Lee (Ruprecht-Karls-Universität Heidelberg)
John Ockendon (University of Oxford)
Victoria Pereira (University of Oxford)
Bernard Piette (University of Durham)
Richard Purvis (University of East Anglia)
Thomas Roy (University of Oxford)
Tamsin Spelman (University of Cambridge)
Robert Timms (University of East Anglia)

Contents

1	Problem statement	1
2	Overview of work	3
2.1	Some possible optimization options	4
3	Generating feasible solutions: part 1	6
3.1	Describing the problem mathematically	6
3.1.1	Estimating cost	6
3.1.2	Constraints	8
3.2	General considerations	9
3.2.1	Two assembly process and radial symmetry	9
3.2.2	Choosing an initial cable configuration	10
3.2.3	Measuring similar cable configurations	10
3.3	Results	10
3.4	Conclusions	11
4	Generating feasible solutions: part 2	14
4.1	Single-objective optimisation	14
4.2	Multi-objective optimisation	17
4.3	Conclusions	18
5	Filler Placement	21
5.1	Optimal placement of the fillers	21
5.1.1	Initial guess	22
5.2	Example solution	22
5.3	Further work	22
6	Maximising variety in the starting sets	24
7	Generating feasible solutions: part 3	28
7.1	Classification of the main frames	28
7.2	Example	28
7.3	Adding an auxiliary frame for a given main frame	29
7.4	Uniqueness of solutions	31
7.5	Automatic or semi-automatic umbilical design creator	31
7.6	Additional classification of centre pieces	31
8	Assessing the similarity of different solutions	33
8.1	Classification of umbilical designs with the help of contact graphs	35
8.2	Formal definition of contact graphs	35
8.3	Possible limitations of this approach and suggestions how to resolve these	38
8.3.1	Incorporating distances between components	38
8.3.2	Understanding graph isomorphism	41
8.3.3	Multi-layered designs	42

8.4	Numerical Implementation	42
8.4.1	Numerical Scheme	43
8.4.2	Numerical Results	43
8.4.3	Producing a Manageable Number of Configurations	44
8.5	Conclusions	45
9	Mechanical response of the Umbilical	46
9.1	Linear elastic model	46
9.2	Possible alternative approaches	48
10	Learning from the experts	51
10.1	Learning from the archive	51
10.2	Learning from the human expert	52
11	Conclusions	54
	References	55

1 Problem statement

- (1.1) A subsea control umbilical is defined as an assemblage of electrical and fibre optic cables, and fluid conduits bound together for robustness and flexibility. In offshore oil- and gas-field developments, umbilicals are critical components for the production of hydrocarbons, providing electrical and hydraulic power, control and signals to and from the subsea components in the system, along with a means of supplying injection chemicals for flow assurance purposes or as a conduit for gas transportation.
- (1.2) Technip Umbilicals designs and manufactures bespoke umbilicals to clients' individual specifications. A design generally comprises a cross-section layout that shows the positions of the specified components within the construction, along with inert filler that provides mechanical protection and stability. The positioning of the components and filler affect the properties of the finished product and the ease and cost of manufacture. The process of cross-section design is currently performed manually, using AutoCAD. Figure 1 shows five different designs for the same set of components, where each design is optimal in some respect. There is no known solution that is optimal in all respects.
- (1.3) A large amount of judgement is applied to the designs during preparation and review, based on calculated properties, experience, manufacturing considerations, knowledge of costs, etc., which makes judgement by non-experts difficult and the design process time-consuming.

Requirement


Based on the requested functional specification, produce an optimised umbilical design automatically. The design will achieve required mechanical properties, primarily strength and linear weight targets, and minimise

- the finished diameter of the umbilical
- manufactured cost
- manufacturing risk


- (1.4) An evolution of the system might accommodate learning from feedback, or learning from previous successful designs, even where those designs did not have exactly the required properties or components. It is anticipated that an automated design tool might generate designs outside the usual rules of thumb applied by human designers, but which are more optimal in some sense.

Components:


- 2 x Quad Power Cables (black quatrefoils)
- 5 x 19.05mm ID Tubes (magenta circles)
- 3 x 12.7mm ID Tubes (red circles)
- various fillers, as necessary (black shading)




(a) smallest diameter and lowest cost



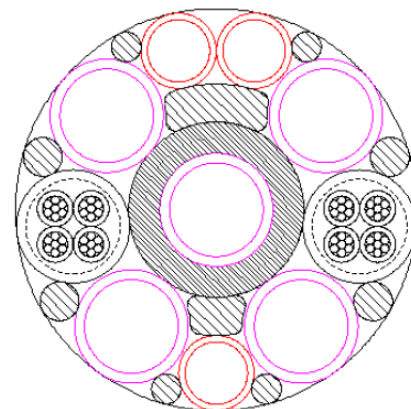
(b) best mechanical protection



(c) no contact between metal tubes



(d) no tube covering required



(e) most efficient; possible manufacturing risk

Figure 1: A number of different designs for a single component set. There is no single optimal solution; each design is optimal in some way, as indicated by the sub-captions.

2 Overview of work

- (2.1) With such a large group working on the problem, it was possible to split up into teams to tackle different aspects of it. These different aspects are described in different sections of this report, as follows.
- (2.2) The first challenge is to devise a system that is able automatically to generate feasible designs for the umbilical. Feasible designs contain the correct components, relatively densely packed together into a circular cross-section, with no components overlapping. For any given list of components to be included in a single umbilical (other than trivial cases with only a very small number of components), there is likely to be a very large, indeed possibly infinite, number of designs. Even if we could automatically generate all possible designs, the amount of time it would take to generate and then evaluate each of them would far exceed the time available, even with the most powerful computers. Instead we have concentrated on methods based on the “good enough, fast enough” principle. Three different solutions to this problem are described. Section 3 uses mathematical optimisation of a single objective (i.e. it seeks designs which minimise a single numerical value, calculated from the design layout). Taking methods developed under the umbrella of Artificial Intelligence, section 4 describes the use of *genetic algorithms* to optimise a single objective (section 4.1) or multiple objectives simultaneously (section 4.2). Section 7 describes a method inspired by the orbits of electrons within atoms.
- (2.3) The fundamental difference between components and filler is that the size and number of components are fixed, whereas those of filler may vary (subject to some minimum size). With this in mind, the group approached the overall design process in two stages: first place the components, and then fill the gaps. Section 5 describes a method of adding filler to a component design.
- (2.4) The component/filler placement methods described in sections 3 and 4 rely on starting from some random position, usually a set of randomly-generated designs, that are then combined and/or adjusted to make them better in some way. Variety in this starting set is key to the search for good designs. Complete randomness may not be the best option; for example, it may result in starting designs that are just rotations of the same layout. Section 6 investigates a method for maximising the variety of the starting set, while section 8 looks at methods for determining how structurally similar two designs are.
- (2.5) In order to find *good* solutions to this problem, we need to consider how to measure “goodness”. Section 2.1, below, summarises a number of different considerations in this respect. For many of these the mathematical means of calculating a value to be optimised is relatively straightforward and various methods are described elsewhere in the report. However, there are significant challenges in assessing the mechanical properties of a given design, particularly since umbilicals must withstand the stresses of different physi-

cal environments, such as (a) the dynamic compression and bending loads of deployment from a ship into the sea, (b) the (relatively) static environment of lying on the sea-bed, and (c) the dynamic bending stresses of connecting from the sea-bed to some fixed or floating structure. Section 9 investigates the first of these, looking particularly at the deformation of an umbilical as it is squeezed; this is a key physical process as the umbilical is gripped whilst being deployed.

- (2.6) Finally, in section 10 there is a discussion about how we might use Technip’s existing expertise to drive umbilical design, based on their archive of previous designs and the knowledge of the human designer.

2.1 Some possible optimization options

- (2.7) A number of different measures that are to be optimised have been discussed in the work group. The following list is a starting point, and elements of this list are discussed in each section of the report. It should be noted that these automated optimisation systems are only able to optimise values for which a numerical value can be calculated. A particular challenge would be to quantify, for example, something like “manufacturing risk”.

- Minimize the overall radius of the final design.
- Minimize the cost of the final design (including manufacturing costs), perhaps by minimising
 - the number of manufacturing passes, where each pass wraps a number of components into a single super-component for inclusion into a future pass,
 - the overall number of fillers,
 - the number of different types of filler, and
 - the number of adjacent circular fillers (this is equivalent to minimizing number of fillers with irregular shapes and number of extra layers imposed on the main components).
- Maximize the mechanical strength of the final design and its safety:
 - no touching between some components (e.g., steel tubes),
 - more vulnerable components place towards the centre of the design, and
 - homogenize the density in the cross-section of the design.
- Maximize the elastical properties of the final design (for installation and storage):
 - minimize the MBR (Minimal Bending Radius) of the final design.

- Maximize the geometrical properties of the final design, perhaps resulting in an easier process of manufacturing, of a more aesthetic design:
 - maximize a symmetry of the allocation of the main components and/or fillers, or
 - choose the shape and density of the fillers as similar as possible to the shape and density of the main components
- Minimize the manufacturing risk.

3 Generating feasible solutions: part 1

- (3.1) Looking at the suggested designs given by Technip, we can see that there seems to be no limit to the number of positions that the cables can occupy within the umbilical. Since we can add any number of fillers of any size, there is not a finite number of cable configurations, so it is not possible to just check all viable cable configurations one by one. Instead, we need to allow, mathematically, these cables to be placed anywhere as long as they do not overlap or break any other given restriction. Naturally this leads us to formulate the problem as an optimisation problem with continuous variables.

3.1 Describing the problem mathematically

- (3.2) Here we translate from English to mathematics the several desirable properties of an optimised umbilical design.
- (3.3) We begin by specifying an Euclidean (x, y) coordinate system, where $(0, 0)$ is the centre of the umbilical. We will refer to all the possible conduits, cables and tubes as just cables. For simplicity, say we want the umbilical to contain two steel cables and two quad cables with radius R^S and R^Q , respectively. We describe the position of the centre of the steel cables by $\mathbf{X}_1^S = (x_1^S, y_1^S)$ and $\mathbf{X}_2^S = (x_2^S, y_2^S)$, and the centre of quad cables by $\mathbf{X}_1^Q = (x_1^Q, y_1^Q)$ and $\mathbf{X}_2^Q = (x_2^Q, y_2^Q)$. See figure 2 for an illustration.
- (3.4) We can now translate the desirable features of an umbilical to simple mathematical expressions. Figure 3 shows how the different costs described below influence the optimised design.

3.1.1 Estimating cost

- (3.5) We begin by assuming the cost of fabrication is proportional to the area of the umbilical¹, so that there is an added cost of

$$C_{area}R^2, \quad (1)$$

where R is the radius of the umbilical and C_{area} is the cost per unit area.

- (3.6) In certain situations, it is undesirable for the steel cables to touch. We can assume that when the steel cables do touch, they wear away more quickly and this incurs an added long-term cost C_{steel} . A simple, smooth function that only adds cost when the steel cables touch is

$$C_{steel} f\left((2R^S)^2 - \|\mathbf{X}_1^S - \mathbf{X}_2^S\|^2\right), \quad (2)$$

¹The cost of having a two-assembly design is not considered, but could have easily been included.




Figure 2: the position of a cable is described in terms of the vector to its centre and its radius.

where $\|(x, y) - (w, z)\| := \sqrt{(x - w)^2 + (y - z)^2}$, $f(\cdot)$ is a sigmoid function such as $f(t) = (1 + e^{-20t})^{-1}$, and we used a power of two so that this cost is a smooth function of its variables (an important feature in gradient based nonlinear optimisation).

- (3.7) We can also include a design rule of thumb: stiffer cables (i.e. steel cables) should be well-distributed throughout the umbilical. We can again assume that a long term cost C_{stiff} is incurred if the stiffer cables are badly distributed. Say that steel and quad cable have a stiffness of S_{steel} and S_{quad} . We can use principals of physics such as centre of mass or centre of inertia (when applying torque to the boundary of the umbilical), but with the stiffness's S_{steel} and S_{quad} instead of the masses. The simplest of these is the centre of stiffness:

$$C_{stiff} \|S_{steel} \mathbf{X}_1^S + S_{steel} \mathbf{X}_2^S + S_{quad} \mathbf{X}_1^Q + S_{quad} \mathbf{X}_2^Q\|^2. \quad (3)$$

The above formula gives zero when the centre of stiffness is at the centre of the umbilical, otherwise a cost is incurred.

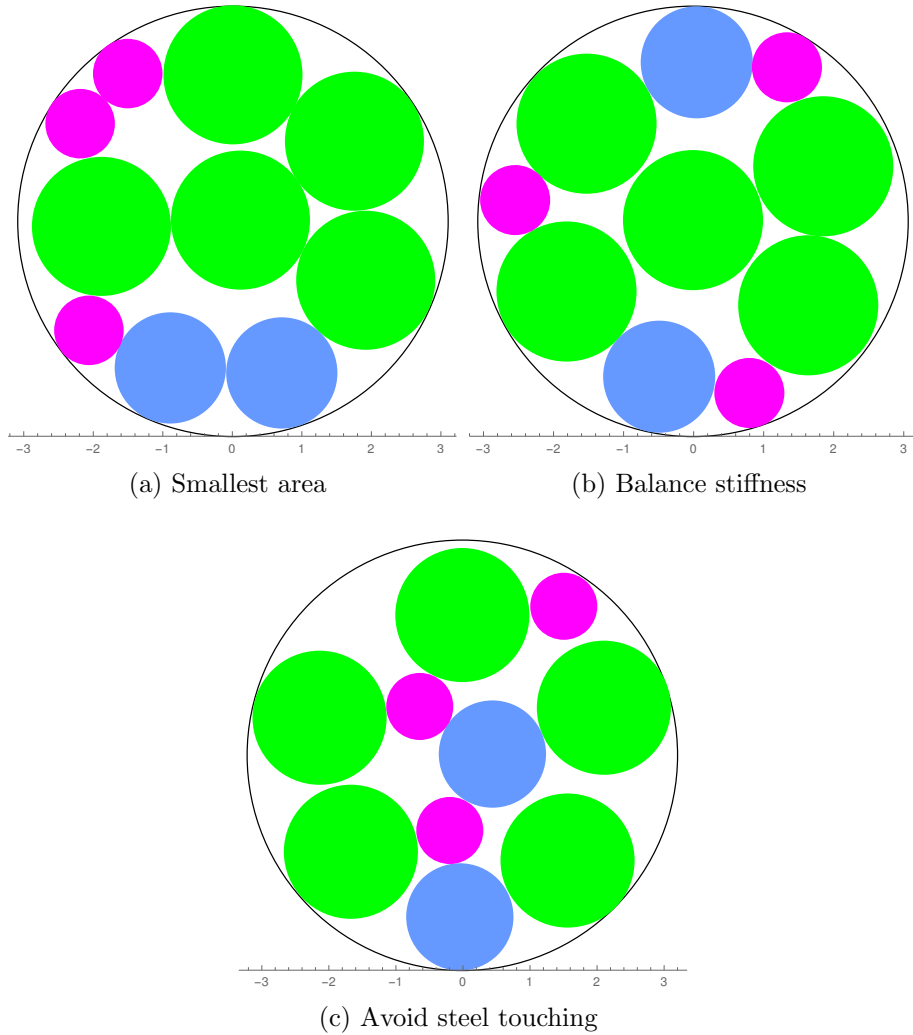


Figure 3: Shows how the results of an optimisation method change when adding extra cost functions. (a) only considers the cost of the area ($C_{area} = 1$ and $C_{steel} = C_{stiff} = 0$); (b) also includes the cost of balancing the stiffness ($C_{area} = C_{stiff} = 1$ and $C_{steel} = 0$); (c) further includes the cost of steel cables touching ($C_{area} = C_{stiff} = C_{steel} = 1$). Note that due to their lack of radial symmetry *these results can not be manufactured*, something which is corrected later and the improved results presented in section 3.3.

3.1.2 Constraints

- (3.8) The cables can neither overlap, nor can they be outside the umbilical casing. For the cables to not overlap translates to

$$\begin{aligned}
 \|\mathbf{X}_1^S - \mathbf{X}_2^S\|^2 - (2R^S)^2 &\geq 0, & \|\mathbf{X}_1^S - \mathbf{X}_1^Q\|^2 - (R^S + R^Q)^2 &\geq 0, \\
 \|\mathbf{X}_1^S - \mathbf{X}_2^Q\|^2 - (R^S + R^Q)^2 &\geq 0, & \|\mathbf{X}_2^S - \mathbf{X}_1^Q\|^2 - (R^S + R^Q)^2 &\geq 0, \\
 \|\mathbf{X}_2^S - \mathbf{X}_2^Q\|^2 - (R^S + R^Q)^2 &\geq 0, & \|\mathbf{X}_1^Q - \mathbf{X}_2^Q\|^2 - (2R^Q)^2 &\geq 0.
 \end{aligned} \quad (4)$$

A trick often used in optimisation is to turn a constraint into a cost penalisation. For this penalisation to be a smooth function of its variables, we again used the power of two in the constraint.

- (3.9) Forcing the cables to be within the umbilical casing leads to the following constraints:

$$\begin{aligned} (R - R^S)^2 - \|\mathbf{X}_1^S\|^2 &\geq 0, & (R - R^S)^2 - \|\mathbf{X}_2^S\|^2 &\geq 0, \\ (R - R^Q)^2 - \|\mathbf{X}_1^Q\|^2 &\geq 0, & (R - R^Q)^2 - \|\mathbf{X}_2^Q\|^2 &\geq 0. \end{aligned} \quad (5)$$

- (3.10) For a two-assembly design we need to add another constraint for the components added in the second assembly. Let R^0 be the radius of the first assembly, and assume we want to place both our steel and quad cables in the second assembly process, then

$$R - R^0 - 2R^S \geq 0, \quad R - R^0 - 2R^Q \geq 0, \quad (6)$$

$$\begin{aligned} (R^0 + R^S)^2 - \|\mathbf{X}_1^S\|^2 &\leq 0, & (R^0 + R^Q)^2 - \|\mathbf{X}_1^Q\|^2 &\leq 0, \\ (R^0 + R^S)^2 - \|\mathbf{X}_2^S\|^2 &\leq 0, & (R^0 + R^Q)^2 - \|\mathbf{X}_2^Q\|^2 &\leq 0. \end{aligned} \quad (7)$$

The first inequality (6) states that each of the remaining cables fit between the inner and outer layer. The following inequalities (7) stop any of the remaining cables from overlapping with the inner radius.

3.2 General considerations

- (3.11) There are many optimisation methods, packages and programmes available. Instead of describing one particular approach, we will outline some general considerations for all gradient based methods. Later, in Section 3.3, we show some results.

3.2.1 Two assembly process and radial symmetry

- (3.12) To actually build any design it needs to possess enough radial symmetry near the centre; see figure 3 for examples of designs that cannot be built. We can remedy this by splitting the method into two steps: first choose a small set of cables (with at least one cable) to go in the centre of the umbilical, then run the optimisation method just for these cables. For the next step let R^0 be the radius of the smallest circle that contains the cables so far. With this first set of cables fixed, now run the optimisation algorithm for the remain cables with the added constraints (6) and (7).
- (3.13) Splitting the method in two steps guarantees radial symmetry in the centre of the umbilical. If only one cable is placed in the centre then we can interpret

this as a one-assembly process, whereas if two or more cables are placed in the centre then the design is a two-assembly process.

3.2.2 Choosing an initial cable configuration

- (3.14) Putting all the constraints together can leave little room for the cables to move and rearrange. It is easy for an optimisation method to get stuck in one particular cable configuration (a local minimum). This is why it is important to have many different initial values for the method, that should ideally be well distributed. Then, the method can be run separately for each of these initial values and compare the results. We achieve this by generating a large set of random initial cable configurations, and then removing duplicate (or sufficiently similar) configurations.
- (3.15) In the next section we show in more detail how to determine whether two cable configurations are similar, as this technique is also useful when choosing between optimised designs.

3.2.3 Measuring similar cable configurations

- (3.16) How do we measure if two cable configurations are alike, in the sense that they would give umbilicals with the same mechanical properties?
- (3.17) Let all the positions of the centre of the cables be denoted by $\mathbf{X}_1, \mathbf{X}_2 \dots \mathbf{X}_N$. Create a matrix \mathbf{M} , where $M_{ij} = \mathbf{X}_i \cdot \mathbf{X}_j$ and $(x, y) \cdot (z, w) := xz + yw$. This way a rotation of all the \mathbf{X}_i 's will not affect \mathbf{M} .
- (3.18) Now the order in which we list the cables $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N$, should also not influence our measure. Note that swapping the order of two cables, for example swapping \mathbf{X}_1 with \mathbf{X}_2 , is equivalent to swapping \mathbf{M} for \mathbf{PMP} where \mathbf{P} is a row-switching matrix with the properties $\mathbf{PP} = \mathbf{I}$ and $\det \mathbf{P} = -1$. One way to eliminate the influence of the order is to apply several row-switching matrices until $(M_{11}, M_{22}, \dots, M_{NN})$, the diagonal of \mathbf{M} , is in decreasing order. Yet another way is to combine the components of \mathbf{M} to form quantities that are invariant to row-switching operations, such as the isotropic invariants

$$\mathcal{S}(\mathbf{M}) = (\text{tr}(\mathbf{M}), \text{tr}(\mathbf{M} \cdot \mathbf{M}), \dots, \text{tr}(\mathbf{M}^N)). \quad (8)$$

With the above, the matrices \mathbf{M}_1 and \mathbf{M}_2 of two cable configurations can be compared by computing the difference $\|\mathcal{S}(\mathbf{M}_1) - \mathcal{S}(\mathbf{M}_2)\|$.

3.3 Results

- (3.19) Here we present the results of a gradient based method implemented in

Mathematica 10, which is run with a large set of random initial cable configurations.

(3.20) To illustrate, we choose the following cables:

- 5 steel cables with radius $R^S = 10\text{cm}$ and stiffness $S_{steel} = 10\text{cm}^{-2}$
- 2 quad cables with radius $R^Q = 9\text{cm}$ and stiffness $S_{quad} = 2\text{cm}^{-2}$
- 3 optical cables with radius $R^O = 5\text{cm}$ and stiffness $S_{optical} = 1\text{cm}^{-2}$

(3.21) In section 3.1.1 we see that we have to choose three cost parameters C_{area} , C_{steel} and C_{stiff} . We pick two examples

$$\text{example 1: } C_{area} = 2C_{stiff} = \frac{1}{2}C_{steel}, \quad (9)$$

$$\text{example 2: } C_{area} = C_{stiff} \quad \text{and} \quad C_{steel} = 0, \quad (10)$$

We do not give a value for C_{area} as only the ratio between the costs will effect the resulting optimal designs.

(3.22) Example 1 means that it is most important that steel cables do not touch and least important to balance the stiffness; see figure 4 for the results. Example 2 gives equal importance to balancing the stiffness and reducing the area of the umbilical, while not considering whether the steel cables touch; see figure 5 for the results.

(3.23) Remember we have not included the cost of having a two-assembly process (two layers), nor do we consider any complications arising from adding fillers. The method presented is simply a proof of concept, but could be extended to include all these considerations that have left out and more.

3.4 Conclusions

(3.24) To summarise, this section of the report shows how to translate several features of an optimised umbilical into mathematics. We also proposed that we can calculate the optimised designs by using nonlinear optimisation methods. As a proof of concept, we used Mathematica's optimisation tools to generate the optimal designs shown in figures 4 and 5.

(3.25) Both the cost and stiffness parameters used in the method were not known, such as C_{area} and S_{steel} . To improve the method, these parameters can be roughly estimated and then the method can search within these estimates to produce optimised designs.

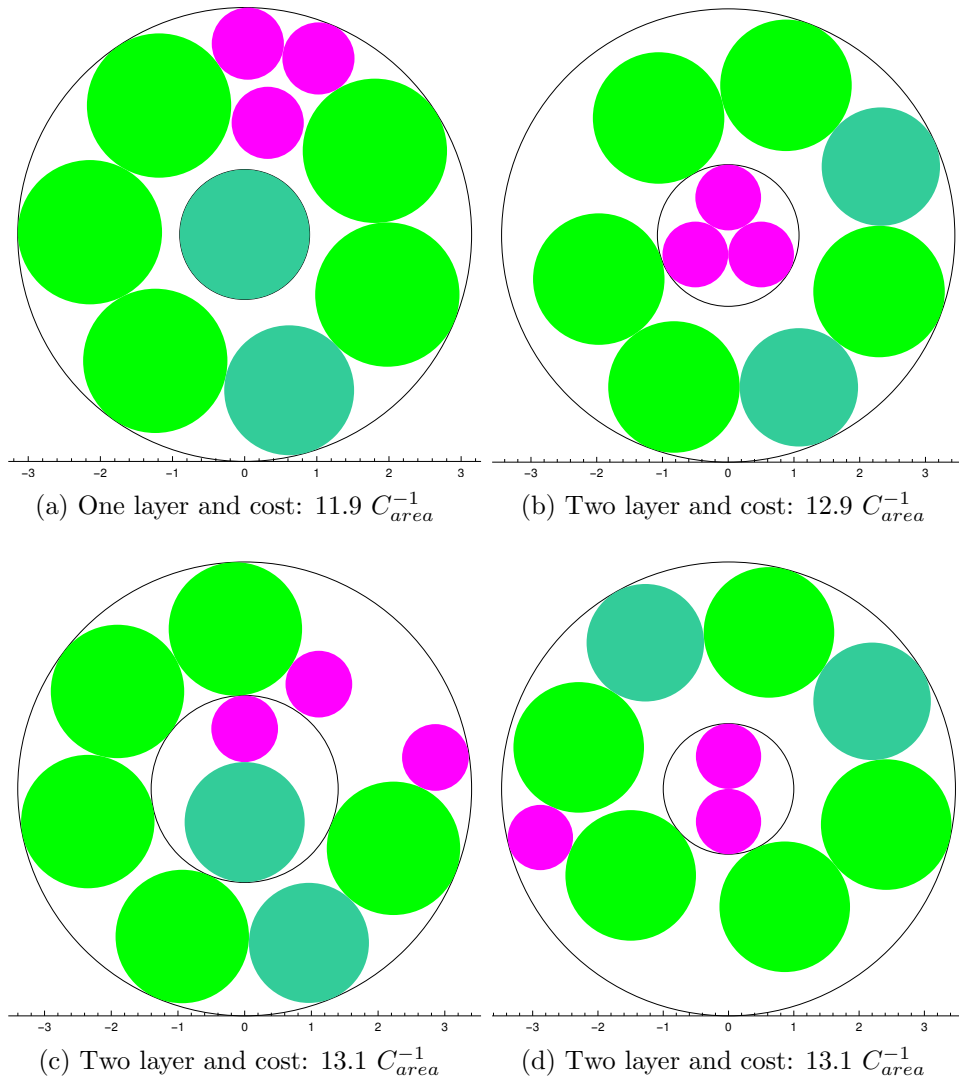


Figure 4: Shows the optimal designs for the cables specified just above equation (9), which gives the cost parameters. The large light green circles are the steel cables, the small purple circles are the optical cables and the mid-sized jade circles are the quad cables. Note how in general the steel cables avoid touching and are well distributed in the umbilical (due to the balance of stiffness).

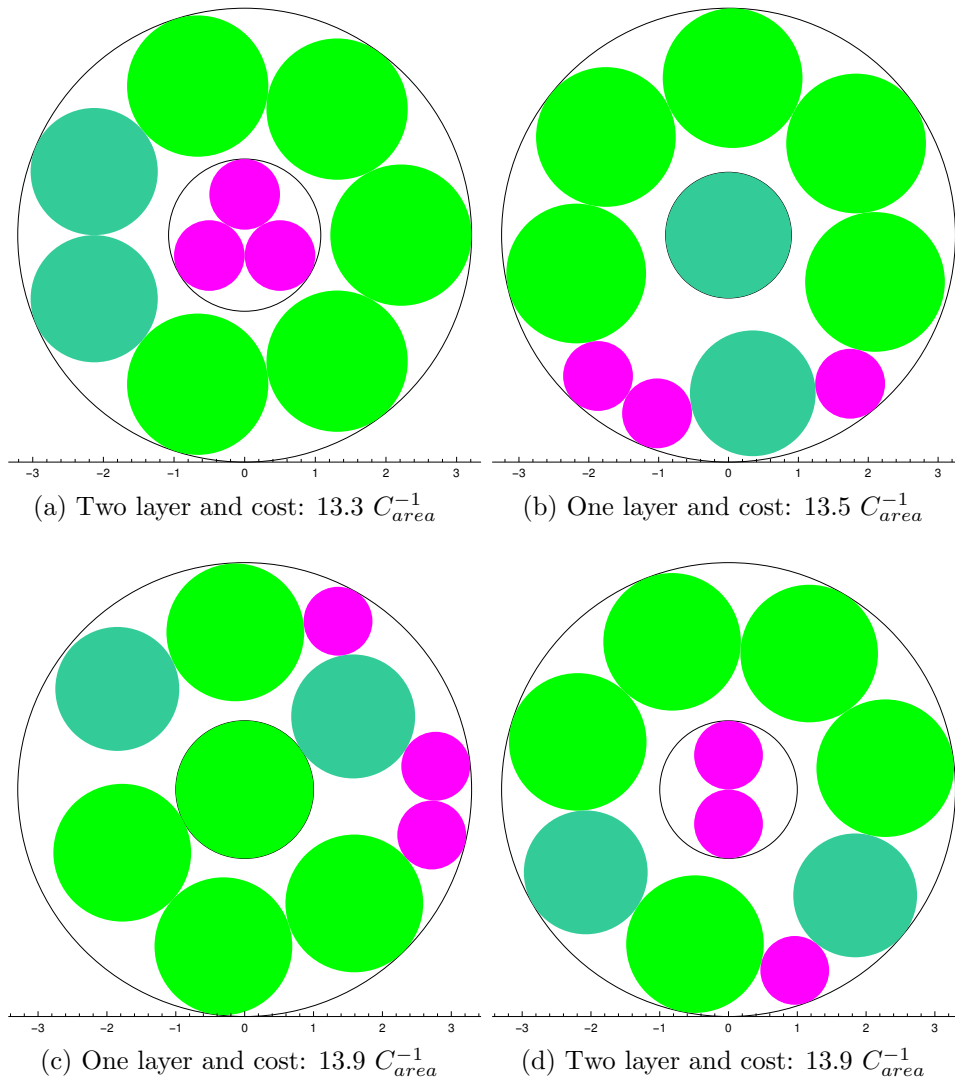


Figure 5: Shows the optimal designs for the cables specified just above equation (10), which gives the cost parameters. The large light green circles are the steel cables, the small purple circles are the optical cables and the mid-sized jade circles are the quad cables. Note how in general these umbilicals are more compact than those given in Figure 4 because steel cables are allowed to touch.

4 Generating feasible solutions: part 2

- (4.1) In this section we describe an alternative way of optimising umbilical design, using two different genetic algorithms (GA) [4] to *evolve* potential designs. A GA mimics the process of natural selection, maintaining a population of plausible solutions to the problem and then, through an iterative procedure, combining and mutating these solutions to improve their performance as measured by some *fitness function*. The fitness of a potential design (i.e. of one member of the population) provides a measure of how good the solution is and hence a means of comparing two potential solutions. The aim of the GA is thus to achieve the best possible fitness value. In this section, it makes sense to define a fitness function whose value we wish to minimise, since we wish to minimise, for example, the radius of the umbilical’s cross-section, so instead of *fitness function* we will refer to the more generic *objective function*.

4.1 Single-objective optimisation

- (4.2) Our goal was to consider a set of circles of different, fixed radii, and configure them in a single umbilical in a way that is “good”. We implemented a multiphase optimisation algorithm to automate the configuring, by formulating an objective function that expresses design ideas we want to satisfy (for example circles not overlapping in conjunction with a small bounding circle radius) and converging to an optimal solution with a genetic algorithm.
- (4.3) The objective function was designed to consider (and penalise) the following configurations:
- circle overlap
 - total bounding radius
 - steel cables (circles) touching
 - a tunable affinity or preference to be placed close to the centre of the umbilical

It is worth noting that the relative importance of each of these is easily changed, as well as adding further considerations discussed in the project outline, most importantly mechanical considerations, if these can be formalised prescriptively.

- (4.4) Considering the list above, the objective function, to be minimised, is defined as:
- the sum of pairwise overlaps
 - + the radius of the bounding circle
 - + steel touching penalty (distance within allowed separation)
 - + (weighted) distances from centre of the umbilical

where the relative weighting of these four contributions can be adjusted, as necessary.

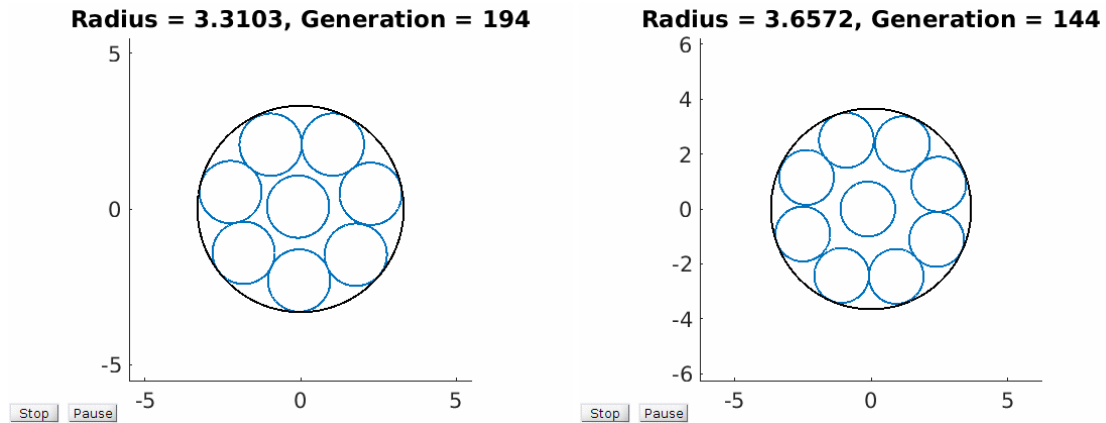


Figure 6: Screenshots of the resulting configurations of packing identically sized balls, essentially achieving the theoretically optimum designs.

- (4.5) We also note some implementation details of the genetic algorithm. We used a population size of around 2000, and evolved the population through a few hundred generations. Of course this can be extended with a more careful implementation and more computing power. The two key features in the genetic algorithm are the mutations introduced into the population, which allow us to traverse the search space, and crossovers, which prescribe how to propagate the population from one generation to the next.
- (4.6) At first we considered the standard mutations, swapping the circle centres in various ways. To allow motion we added random translations (all scaled to be somewhat invariant to the radius sizes). We also considered more geometrically rich mutations, such as hops (allowing a circle to reflect its position in the line connecting the centres of two other, randomly chosen, circles), plane reflections (randomly choosing a line, and reflecting the centres of all those circles above the line) and centring around the centre of mass.
- (4.7) Some mutations are better than others at different stages. For example, slightly larger translations are good at the beginning to leave the initial (random) configuration and make large steps toward converging to feasible configurations. However towards the end of the evolution when the configuration is in a local minimum, we are less interested in large translations and perform smaller movements to encourage a smaller scale settling. This annealing proves helpful. In general, the style and size of mutations can be tailored to be vary with time.
- (4.8) The crossover functions used were fairly standard. We used both single-parent breeding, as well as two-parent, where the children take some centres from one parent and the rest from another.
- (4.9) A useful testing benchmark is the case of identical circles, for which the optimal packing is known. With no other considerations, our algorithm

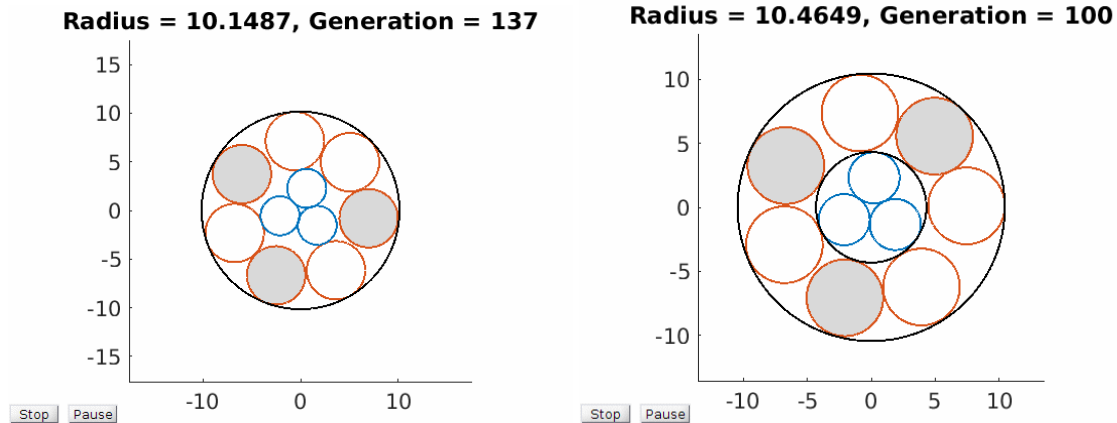


Figure 7: Screenshots of the resulting configurations of two simulations. On the left we ran a single phase, placing importance on the smaller circles being in the centre, whilst on the right we achieved a similar result with two phases. The grey shading shows steel circles that cannot touch.

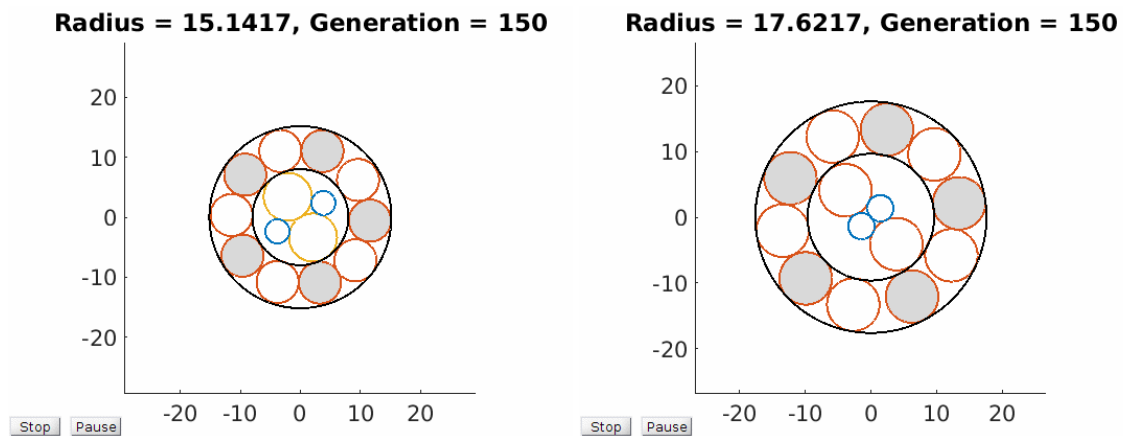


Figure 8: Screenshots of the resulting configurations of two separate simulations. On the left we gave more importance to the larger circles being in the centre, whilst on the right we gave more importance to the smaller circles being in the centre. The grey shading shows steel circles that cannot touch.

should converge to the optimal packing. See figure 6 for examples.

- (4.10) We call the algorithm multiphase because we allow the user to choose separate phases to run the algorithm. The first phase settles with some configuration around the origin. We then encircle this with a bounding circle. The next phase configures the circles around, that is, outside the outermost bounding circle. And we can continue as such for an arbitrary number of phases. See figure 7.
- (4.11) Initial conditions are also important, and may lead to a broader set of final configurations. In these experiments we simply placed the circles at random at the start of the evolution, but other ideas are worth considering. For

example, placing the centres equispaced on a large ring would allow a rich array of initial configurations, and admits a systematic and discrete way of enumerating potential initial conditions. Section 6 of this report considers different methods for generating starting configurations.

- (4.12) Ideally, we envisage the algorithm running many times, with many initial conditions. There were discussions about automatically identifying similar configurations using methods from graph theory (see section 8), so the output would select only those that were distinct. By running with different objective functions penalising different things (e.g. mechanical properties, geometric properties, etc.) it is possible to achieve a variety of potential configurations. For example, see figure 8.

4.2 Multi-objective optimisation

- (4.13) In section 4.1, the objective function combines all the different costs and penalties into a single, weighted value. The weightings must be specified up front and a single run of the algorithm searches for the single best design for that particular selection. The reason for the weightings is that there is no single umbilical design that is optimal in all senses.
- (4.14) An alternative form of GA takes into consideration the fact that different designs may be optimal in different ways, but that no one design is optimal in all ways. A common trade-off in two objectives is cost versus quality; cost to be minimised and quality to be maximised. The *multi*-objective optimising GA maintains a population of potential solutions that represent the trade-off between two or three objectives; more than three is generally referred to as *many*-objective optimisation. This trade-off set, usually called the *Pareto* set, is the set of potential designs that are better than all the others in the set in at least one objective, but not in all objectives.
- (4.15) The Pareto set output by this type of GA presents the user with a range of designs, from which they may choose the best according to their particular requirements, without them having to define the weightings of each objective in advance. Seeing the form of the Pareto set informs this decision-making process. An example will make this clearer.
- (4.16) Consider designs that must both minimise the radius of the bounding circle and produce designs that are as symmetrical as possible. For some sets of components (for example, a single component), there may be a single design that is optimal in both objectives, but in many cases increasing the symmetry is likely to expand the bounding circle.
- (4.17) The calculation of the radius of the bounding circle is straightforward. The symmetry measure has been formulated as the distance of the centre of mass of the components from the centre of the umbilical, and hence smaller values




Figure 9: A plot of bounding radius versus symmetry (both to be minimised) for the Pareto set resulting from a single run of the GA. Note the obvious discontinuity in the centre of the plot, where a very small increase in the radius of the bounding circle gives rise to a significant improvement in symmetry.

mean more symmetry. Thus both objectives are to be minimised.

(4.18) Figure 9 shows a plot of bounding radius versus symmetry for the Pareto set resulting from a single run of the GA. There is an obvious discontinuity in the centre of the plot, where a very small increase in the radius of the bounding circle gives rise to a significant improvement in symmetry. The decision-maker can choose which of these points, each of which represents one design, is most acceptable for a given requirement, with full understanding of the trade-offs available.

(4.19) Figure 10 shows two example solutions generated by the multi-objective GA. The one on the left has been selected from the region of the Pareto set that “prefers” the minimum bounding radius, while that on the right is selected from the region that “prefers” symmetry.

(4.20) In these examples, the red circles represent steel components. Incorporating a constraint that these should not be in contact with one another is straightforward. Figure 11 shows a design solution that is symmetrical, tightly-packed (i.e. the bounding circle is small), *and* the steel components are not in contact.

4.3 Conclusions

(4.21) During testing of the single-objective GA, a valuable test of its ability to find optimal solutions was to compare its minimum-radius solutions for sets of identically-sized components to the mathematically-proven op-




Figure 10: Two example solutions generated by the multi-objective GA. The one on the left has been selected from the region of the Pareto set that “prefers” the minimum bounding radius, while that on the right is selected from the region that “prefers” symmetry.




Figure 11: An example solution generated by the multi-objective GA, where the steel components (in red) are not touching.

tima. Wikipedia provides a useful summary of these optima (see https://en.wikipedia.org/wiki/Circle_packing_in_a_circle). The GA repeatedly found the optimal solutions for problems from 1 to 13 components, where 13 is the largest number for which the true optimum has been proven.

(4.22) The single-objective GA requires that all the different measures that we would like to optimise be combined into a single value, requiring the user

to decide beforehand what the relative importance of each of the measures is. The multi-objective GA removes this necessity, and evolves a range of solutions. The result is a trade-off between the different measures, from which the user may select make an informed selection of their preferred design.

5 Filler Placement

- (5.1) Once an optimal solution for the placement of all the required components has been found, it remains to fill the remaining space in the umbilical with filler material. In practice, lengths of filler material are wound together with the required components during umbilical manufacture such that the maximum unfilled gap in the umbilical cross section is less than some prescribed threshold. These filler components can have any shaped cross section, but often circular fillers are favoured in the interest of cost. The aim is to meet the design specification (which may include criteria related to strength, bending etc.) in some optimal way (e.g. at the least cost). In the following we focus on meeting the objective of filling some prescribed fraction of surface area using the fewest number of filling components. Of course, the approach could be extended to include different, or multiple, objective functions.

5.1 Optimal placement of the fillers

- (5.2) In order to solve the filler placement problem we set the following optimisation problem.

$$\begin{aligned}
 & \max \varphi, \\
 \text{s.t. } & \|\mathbf{X}_i - \mathbf{X}_j\| \geq r_i + r_j, \\
 & \|\mathbf{X}_i - \mathbf{Y}_k\| \geq r_i + \rho_k, \\
 & \|\mathbf{X}_i\| \leq R, \\
 & r_i \geq r_{min},
 \end{aligned} \tag{11}$$

where the objective function φ is the filling factor defined as

$$\varphi = \frac{1}{R^2} \left(\sum_i r_i^2 + \sum_k \rho_k^2 \right). \tag{12}$$

The variables of the problem are the position of the centre of the fillers \mathbf{X}_i and the radii of the fillers r_i . The parameters of the optimisation problem are the configuration of the cables (required components) given by the centre of the cables \mathbf{Y}_k and their radii ρ_k , a minimum radius of the fillers r_{min} and the radius of the external cable R .

- (5.3) The first two constraints enforce no overlapping of elements (the first, filler-filler; the second, cable-filler). The third constraint enforces the fillers to fit into the umbilical and the fourth constraint enforces the fillers to have a minimum radius. The optimisation problem was coded in MATLAB using the non-linear constrained optimisation routine `fmincon`.
- (5.4) It is not known in advance how many fillers will be required to achieve the prescribed filling factor φ , so the algorithm starts with one filler and finds the optimal configuration. If the filling factor does not reach the set threshold

the algorithm adds another filler and optimises the result again. This process continues until one of the stopping criteria is satisfied. This can be either by achieving the required filling factor, by exceeding the maximum number of fillers set by the user, or by not finding a feasible solution to start with.

5.1.1 Initial guess

- (5.5) The first approach for the initial guess in the optimisation routine was a random configuration of the fillers. After a few trials it was found that the outcome was very sensitive to the initial guess and that starting from a random configuration led to clearly non-optimal solutions most of the time. To improve the algorithm the approach was changed so that for each number of fillers N a number of initial guesses were chosen from the feasible region. The optimisation routine was completed several times and the best solution was selected. The optimal solution for N fillers was then used to provide an initial guess for the $N + 1$ filler problem, with the position of the additional cable being selected at random. As before, a number of initial guesses for the random placement were made, and the best was chosen. This process the repeated until the filling criteria had been met, or the maximum number of fillers had been placed. It was also found that reducing the radius of the filler to the minimum radius r_{min} in the initial condition provided improved solutions. Reducing the radius prevented a large filler from becoming “stuck” in a large gap where it may be favourable to place two filler components at a later stage in the optimisation algorithm.

5.2 Example solution

- (5.6) Given an optimal configuration of required components, the algorithm was used to fill some required threshold. Figure 12 shows a filled two-layer umbilical. The required components, shown in blue, are the result of a genetic algorithm approach, which was filled by the above optimisation routine. The solution looks remarkably similar to an example umbilical provided to the study group (see figure 1a), and provides some justification that the algorithm provides solutions that are not only optimal (in the sense of filling a given space with the fewest number of fillers) but can also be produced by manufacturing methods currently in use.

5.3 Further work

- (5.7) The algorithm outlined above could be extended to include additional objectives and criteria. For example, during the study group an alternative algorithm was devised which aimed to fill the cable in such a way that the centre of mass was as close to the centre of the umbilical as possible. In order




Figure 12: A filled two-layer umbilical. The required components are depicted as blue circles, and the filler components are depicted as red circles.

to deal with fillers of arbitrary cross section an example approach might be to develop an algorithm which detects if there are a large number of circular fillers in close proximity to one another, and suggests to the user that such collections of circular fillers may be replaced with a single filler of arbitrary cross section.

6 Maximising variety in the starting sets

- (6.1) Whichever optimisation scheme is to be applied, the choice of initial guess is an important consideration. Since the problem is inherently non-convex, with many local minima present in the solution space, we must search based on many initial guesses. To well sample the solution space we select initial guesses at random. However, this can be complemented with heuristic techniques to seek more preferable choices of initial guess.
- (6.2) What must be chosen to prescribe an initial guess is the centre of each cable within the umbilical. If we suppose that we wish to minimise the radius of the umbilical then a random initial guess can be given by specifying that the centres of the cables are independent and identically distributed normally around the origin. Such an initial guess is likely to include overlapping of the cables, a constraint which must be made tight when optimising. As an alternative, we can construct a good feasible initial guess, having no overlap, using a heuristic argument based on packing the cables together.
- (6.3) The basic approach is to add the cables one by one, fitting them into a gap between two cables. Figure 13 demonstrates this positioning. To start, we arrange the first two cables so that they are touching, this leaves two gaps on either side of the point of touching to add a third cable. By keeping track of the gaps, we continue trying to add cables into the gaps. If the cable will not fit, we try the next gap, until we find a place it can fit. This process continues until all cables have been added. We note that the order in which the cables are added is important for the initial guess constructed.
- (6.4) In order to apply this heuristic we are required to find the position a cable should take to fit into a gap between two cables. Here we relax the assumption that the first two cables are touching to allow for the more general case of a gap between two nearby cables. This geometric problem can be solved as follows, aided using quantities shown in Figure 14. Suppose the centres of the two existing cables are given by X_1 and X_2 with radii r_1 and r_2 respectively.




Figure 13: Fitting a new cable (blue) into the gap between two existing cables (red).

- (6.5) Each time we add a new cable, touching two existing cables, we create four new gaps, being either side of the two touching points. We can maintain this list of gaps throughout the procedure ready for trying new cables. To test if a new cable fits in a gap, say positioned at X_{new} , we calculate the distances from this new centre to all existing cable centres X_i and check that it is greater than or equal to the radii of the two cables being considered. That is, for each centre of an existing cable X_i we check that

$$\sqrt{(x_{new} - x_i)^2 + (y_{new} - y_i)^2} \geq r_{new} + r_i, \quad (21)$$

where r_{new} is the radius of the new cable. Note that since we calculate the distance to the centre of each existing cable we can identify any nearby cables which, while not touching, may still provide a gap to slot in a further cable. Hence we can operate within the more general framework shown in Figure 14.

- (6.6) Now that we have an understanding of the geometry of placing new cables between the gaps of existing cables there are two choices to make in constructing an initial guess. Firstly, in what order do we add in the cables and, secondly, in what order do we try the gaps between the existing cables. To characterise the gaps we define the *gap size* to be the distance between the centres of the cables to which the gap belongs, namely c as defined in (15). If the cables are touching this is simply the sum of the two radii of the cables.
- (6.7) To see that the two choices given above are important we give an example of a bad combination of choices. If we add cables in ascending order of size (radius) and try gaps in decreasing order of size then we can obtain a long and slender arrangement of the cables, giving a large bounding radius for the umbilical. As an example we consider a set of cables having radii 1, 2, 3, ... 10. Figure 15 plots the resulting arrangement of the cables illustrating its undesirable characteristics. The bounding radius of this configuration is 30.31.
- (6.8) While we could add the cables in a random order and also try gaps randomly to give a variety of initial guess configurations, a good configuration can be gained by ordering the cables and gaps both in descending order; the resulting configuration, using the same set of radii as Figure 15, is given in Figure 16 and gives a bounding radius of 23.41. In contrast using 100 trials of a fully randomised version gave an average bounding radius of 27.31 with the minimum being 24.16, higher than that using the descending orderings. However, the order of adding the cables appears to be the most important consideration. If we add them in decreasing order of size but try gaps at random then in 100 trials the mean bounding radius drops to 25.06. We also note that in this set of trials the minimum was found to be 22.53 and so improves the configuration of Figure 16 when only decreasing orderings are used. Overall, using such randomisation we can gain a set of initial guesses to input into an optimisation routine.




Figure 15: A bad arrangement of cables with bounding radius 30.31.




Figure 16: A good arrangement of cables with bounding radius 23.41.

- (6.9) We stress that in the heuristic arguments presented above there is no explicit optimisation involved, only a set of rules which aim to provide a reasonably good initial guess (or set of initial guesses, if randomisation is included) for an optimisation routine. However, we should be careful that when using an optimisation routine we do not allow it only latch onto the best initial guess and be unable to move away to find better solutions. It is also not clear that using only the, possibly randomised, heuristics here provides an adequate sampling of the whole space. More investigation is required to deduce whether such a ‘good’ set of initial guesses, as might be derived by the above heuristics, will likely also give a ‘good’ outcome from a given optimisation routine.

7 Generating feasible solutions: part 3

- (7.1) In order to make an automated system that is able to generate all possible designs which are in infinite number of combinations of components and fillers and their shapes, it is essential to categorize the group within an efficient structure to narrow down the choices which leads to the optimised one. We use the analogy from the orbits of electrons within an atom to place components in orbital layers about the centre of the umbilical.
- (7.2) The intuition for a combinatorial algorithm for multi-layer cable placement, based on orbits:
- Introduce orbits where components lie
 - Adjust the radii of orbits such that umbilical radius is minimised
 - Previous geometric and symmetry constraints continue to hold

Figure 17 shows, on the left, an initial assignment of components to two different orbits, and, on the right, the shrinking of these orbits to achieve a design.

7.1 Classification of the main frames

- (7.3) The maximum number of orbits is equal to the number of components, i.e. N . Describing the allocations of the components on the orbits and all possible orderings of components on the same orbit results in a representative $(N + 1) \times N$ matrix for each main frame.

7.2 Example

- (7.4) For example consider 3 small tubes, 5 large tubes and 2 cables, i.e. $N = 10$:






Figure 17: Orbital placement: (left) assign the components to orbits and then (right) shrink the orbits to squeeze them together.

categorize their type by its radius. Even if two different types of components have same radius, categorize them with different r_i and r_j .



Then place the components in the order one wants on each orbit.



In this example, Orbit 0 is empty and so do Orbit 3 – 10. Now we are ready to write the allocation matrix:

$$M = \begin{bmatrix} 0 & 0 & & \dots & & & & & & 0 \\ 1 & 1 & 1 & 0 & & \dots & & & & 0 \\ 3 & 2 & 2 & 2 & 3 & 2 & 2 & 0 & 0 & 0 \\ 0 & & & \dots & & & & & & 0 \\ \vdots & & & & \ddots & & & & & \vdots \\ 0 & & & \dots & & & & & & 0 \end{bmatrix}$$

where each row represents the orbits and the entries are type of components. For this example there is no component on Orbit 0. Thus the first row is empty. There are three kinds of components and all three of first kind marked as r_1 or 1 in matrix entry lie on Orbit 1. Other entries in the first row are filled with zeros. This continues for all $(N + 1) \times N$ entries which represents the allocation matrix for this main frame.


7.3 Adding an auxiliary frame for a given main frame

(7.5) Optionally we are introducing the additional auxiliary impenetrable orbits between the existing orbits. Then each impenetrable orbit can be taken as a complete umbilical and has freedom to have independent assembly processes for the final design.


(7.6) Then we merge main frames consists only main components and auxiliary frames consists the main components and fillers) into a final design. With

the frame we continually scaling distances between orbits and components on the same orbits as well as rotating the orbits independently and moving components on the same orbits in the way that their ordering still remains the same. In this way we ensure that we stay in the same main frame configuration. This process can be parametrized with the help of some set of parameters (which can be used later for the multivariate optimization purposes).

- (7.7) Finally, we are enveloping the whole structure with the circle (confining it in the thermoplastic covering) and inserting the fillers (preferably circular ones) into a given main frame to produce a final design (a solution) ensuring that the left-out free space (not filled either with main components or fillers) is less or equal than some critical value ($S_{\text{free}} \leq S_{\text{crit}}$) which adds the flexibility of the cord but no displacement of components.



The space enclosed by three small circles is small enough to satisfies $S_{\text{free}} \leq S_{\text{crit}}$.



As the size of circles increases, the enclosed space also increases until it is big enough to insert a given component or needs a filler to fix the position of circles. Then it violates the condition and $S_{\text{free}} > S_{\text{crit}}$.

- (7.8) Main purposes of adding filters:
- fixing the component in the places where we want them to be
 - giving the design desired mechanical properties, e.g., desired MBR, degree of overall elasticity, strength, homogeneity of density/weight in the cross-section, degree of resistance to hydrostatic pressure, etc.
 - ensuring some additional symmetry of the solution (in shape, geometrical distribution of components, etc.)

- (7.9) We begin the design with circular fillers. However, at the very end, the option to merge the set of adjacent circular fillers into more costly fillers of irregular shapes or extra layer covering of the main components.

7.4 Uniqueness of solutions

- (7.10) Each final design (each solution) can be uniquely represented with the help of:
- allocation of main components on orbits
 - angular ordering of components on the same orbits
 - adjacency graph associated with the given main frame.

7.5 Automatic or semi-automatic umbilical design creator

- (7.11) An user-friendly automatic or semi-automatic umbilical design creator can be made by questioning the user on each step to design on each orbit. By mechanical or design purpose the user can decide first the centre piece: which cable to place or leave it empty. This design for the centre piece is the additional classification which will be explained in detail later. Then it moves to the next orbit to ask which components to place to generate the possible orders and rotation of allocation of components regarding the symmetry or more mechanical requirement. In this way the possible combinations of allocations of components narrows down in order to follow the steps explained previously.

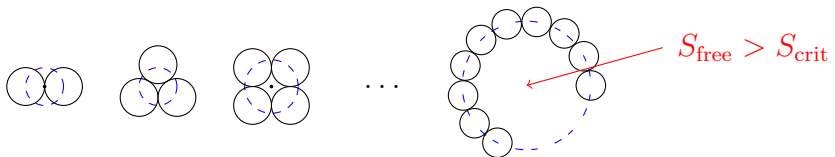
7.6 Additional classification of centre pieces

- (7.12) We can provide some classification of the central designs (the predefined design of the lowest orbits, i.e., giving an allocation of main components on the zeroth, first, and maybe second orbits) in order to maximally diversify the initial conditions of numerical algorithms for optimization process. Not only this classification will help to categorize from the initial step, it will help to ensure that the algorithm considers all possible combinations. Define some conditions for cutting-off options (e.g., $S_{\text{free}} \leq S_{\text{crit}}$). Predefine the final design by given some packing trends in the arrangement of main components (e.g., smaller components inside, larger components outside or vice versa; start with small elements, then place larger ones, and finally again small, etc.).
- (7.13) Here are some examples of the centre pieces.

Case 1: When Orbit is is filled, only one possibility is a component is placed on Orbit 0.

Orbit 0 

Case 2: In case Orbit 0 is empty, we can place the components as long as $S_{\text{free}} \leq S_{\text{crit}}$.

Orbit 1 

The last possible case is to place all given components on Orbit 1. However, it might violate the condition as indicated.

8 Assessing the similarity of different solutions

- (8.1) As shown in figure 1, designers can come up with multiple, structurally different designs for the same umbilical. If one design is merely a rotation of a previous design, this would not be presented as a different option. But it is much harder for a computer to identify these similarities. Here we will define an umbilical based on which components are adjacent to each other and use this to identify umbilical designs which are structurally different. We show proof of concept and discuss further improvements including: incorporating distance between components and limiting the number of different structures.
- (8.2) The goal of this part of the project was to establish and implement a mathematical framework for the classification of umbilical designs in order to be able to say when two umbilical designs can be considered as the same.
- (8.3) *Why is it necessary to be able to say that two designs are the same?* Given a set of components that should be contained in an umbilical, we want to be able to automatically generate and present different possible design solutions that are optimal with respect to different goals. However, if we run algorithms various times in order to generate a multitude of umbilical designs, many of the solutions will be essentially the same from the point of view of the geometry of the umbilical. For example, consider the two designs depicted in figure 18. The design on the right hand side can be obtained from the one on the left hand side by first flipping the umbilical over a horizontal axis passing through the centre of the umbilical and then performing a rotation. For us humans it is rather straightforward to see that these designs are the same because we are very good at detecting symmetry and other geometrical properties. However, it is a priori not clear how a computer programme that only knows the shapes and positions of the involved components should be able to tell that these designs do not need to be distinguished. It is algorithmically not possible to check all symmetry operations, all rotations etc., this would lead to a combinatorial explosion of the problem. The goal of this part of the project was thus to find an automated way to determine whether two designs are the same in order to reduce a multitude of possible optimal solutions to those which really differ from each other.
- (8.4) *What industrial advantages does this bring?* The algorithms we have suggested in this report for finding optimal designs run multiple times producing many “optimal” solutions. We then identify which of these meet the mathematical criteria best. Instead, we can consider all those optimal solutions, use the framework presented in the following to break them down into a handful of options which are structurally different and present these designs as possibilities. Then a final design choice, which can depend on project specific requirements, would be made by a designer. This allows for flexibility in design properties which vary between umbilical projects. Alternatively, other design criteria, which might be computationally more difficult




Figure 18: Two umbilical designs that should be considered as the same, the only difference being that the central clutch of three components plus three fillers has been rotated with respect to the rest of the design.

to characterise or mathematically harder to quantify, could then be applied to choose between the remaining designs.

- (8.5) *What do we mean by saying that two umbilical designs are “the same”?* In the following, saying that two designs can be considered as the same means that they correspond to the same geometrical set-up and can be obtained from one another by central or axial symmetry operations, by rotations, or very small displacements of the components that do not effect the properties of the umbilical (mechanical properties, production costs, production risks aso.).
- (8.6) *What do umbilical designs look like?* In an optimal umbilical design, components and fillers will lie very close to each other (fillers will always be used to fill up gaps as soon as they are large enough. If gaps between components are necessary, e.g., since steel tubes should not be touching, this will be ensured with the help of fillers or other components). In general, designs are very compact for various reasons (stability, costs, manufacturing risks etc.) and every component and filler will be touching at least two other elements of the umbilical². This leads us to the conclusion that the geometric set-up of an umbilical design is highly characterized by the information which elements are touching which other elements and which type of component-component and component-filler adjacencies occur.

²For simplicity, we assume that all components and fillers are circular. Incorporating other filler shapes in this model could however be done without any trouble.

8.1 Classification of umbilical designs with the help of contact graphs

- (8.7) *Our main idea:* Every umbilical design is associated with its coloured *contact graph*—a graph that stores all the relevant information about its geometrical set-up. An example of such contact graphs can be found in figure 19: once for the case that only components are considered and once for the case that both components and fillers are considered. In the following section we will formalize this notion.
- (8.8) Many of the ideas presented in the following are standard graph theoretic concepts or inspired by such concepts. As a general reference, we refer to Diestel’s *Graph theory* [2] that is used as a standard textbook in modern graph theory.

8.2 Formal definition of contact graphs

- (8.9) Let us start with a formal definition of graphs³:

Definition 1 (Graph). A **graph** is an ordered pair $G = (V, E)$ consisting of a set V of **vertices** or **nodes** together with a set E of **edges** which are 2-element subsets of V .

- (8.10) That is, an edge consists of two distinct vertices and edges are not directed (there is no difference between the edge connecting vertex v with vertex w and the edge connecting w with v). Moreover, multiple edges between the same pair of vertices cannot occur. In the following, we will consider *coloured graphs*, i.e., graphs in which every node is given a specific color.
- (8.11) We now define a particular type of graph, a so-called contact graph, that we associate with umbilical designs.

Definition 2 (Definition and construction of the contact graph). Let D be an umbilical design which is specified by the number of components⁴ of different types and their positions in the plane. We construct the *contact graph* $G(D)$ corresponding to the design D as follows:

1. We associate a node in the graph to every component of the design. Depending on the type of the component, nodes are drawn in different colors (one colour each per type of component).

³In the literature, such graphs are referred to as simple graphs. In the following we will however only use the term *graphs*.

⁴For simplicity, we do not distinguish between components and fillers in this definition; we consider fillers to be components of the type “filler”. As can be seen in figure 19, including fillers in the contact graph can be done in a straightforward way.




Figure 19: An example of an umbilical design (left), the contact graph of its components (middle) and the contact graph of both its components and its fillers (right).

2. We connect two nodes with each other by drawing an edge if the corresponding components touch each other in the design.

(8.12) For example, consider the umbilical design depicted on the left hand side of figure 19. We have three different components: quad cables correspond to black nodes, large steel tubes to pink nodes and smaller steel tubes to red nodes. In the center of the design are the three smaller steel tubes and they are grouped in such a way that every one of the tubes touches the other two. This configuration is reflected by the triangle of red vertices in the center of the contact graphs. As shown on the right hand side of the figure, one can also represent fillers (here, in gray) with the help of nodes and construct the contact graphs of both components and fillers.

(8.13) Contact graphs are very particular graphs with strong graph theoretic properties: they are *planar*. This means that they can be drawn in the plane in such a way that no edges intersect⁵ Moreover, contact graphs of both components and fillers are always *connected* which means that there is a path between every pair of vertices (there are no unreachable vertices).

(8.14) If we want to specify the graph G_m represented in the middle of figure 19 as in Definition 1, this could for instance be done as follows:

$$\begin{aligned}
 G_m &= (V, E), \\
 V &= \{r_1, r_2, r_3, b_1, b_2, p_1, p_2, p_3, p_4, p_5\} \text{ and} \\
 E &= \{\{r_1, r_2\}, \{r_1, r_3\}, \{r_1, b_1\}, \{r_2, r_3\}, \\
 &\quad \{r_2, b_2\}, \{r_2, p_2\}, \{r_3, p_3\}, \{r_3, p_4\}\},
 \end{aligned}$$

where the labelling of the nodes has been done as represented on the left hand side of figure 20.

⁵A simple example of a graph that is not planar is K_5 , the *complete graph* on 5 vertices, that is the graph consisting of 5 vertices in which every vertex is connected to every other vertex.

- (8.15) This representation of graphs with the help of the two sets V and E is not particularly handy, especially not as a data structure handling contact graphs in an algorithmic context. Let us therefore present the following compact and clear representation of (contact) graphs with the use of adjacency matrices.

Definition 3 (Adjacency matrix). Let $G = (V, E)$ be a graph with n vertices and let us assume that $V = \{v_1, v_2, \dots, v_n\}$. Then the **adjacency matrix** of G , denoted by $A(G) = (a_{i,j})_{1 \leq i, j \leq n}$ is the symmetric $n \times n$ matrix in which the entry $a_{i,j}$ is equal to 1 if there is an edge between v_i and v_j and 0 otherwise.

- (8.16) Using the same labelling of vertices as before (see figure 20), we group vertices by the type of component they represent and obtain the following adjacency matrix for G_m :

	r_1	r_2	r_3	b_1	b_2	p_1	p_2	p_3	p_4	p_5
r_1	0	1	1	1	0	0	0	0	0	0
r_2	1	0	1	0	1	0	1	0	0	0
r_3	1	1	0	0	0	0	0	1	1	0
b_1	1	0	0	0	0	0	0	0	0	0
b_2	0	1	0	0	0	0	0	0	0	0
p_1	0	0	0	0	0	0	0	0	0	0
p_2	0	1	0	0	0	0	0	0	0	0
p_3	0	0	1	0	0	0	0	0	0	0
p_4	0	0	1	0	0	0	0	0	0	0
p_5	0	0	0	0	0	0	0	0	0	0

Note that $A(G)$ is always a symmetric matrix ($a_{i,j} = a_{j,i}$ since there is an edge between v_i and v_j if and only if there is an edge between v_j and v_i) and that the diagonal elements are always zero ($a_{i,i} = 0$ since edges always consist of two distinct vertices; components do not “touch” themselves). Thus we only need to store $n \cdot (n - 1)/2$ values in order to fully describe $A(G)$.

- (8.17) Now that we have found a way to represent the geometrical set-up of umbilical designs it is time to define when two designs can be considered as the same. This is quite straightforward: Two umbilical designs are the same if they have the same contact graph or, to be more precise, if their contact graphs are isomorphic:

Definition 4 (Isomorphism of (coloured) graphs). Let $G = (V, E)$ and $H = (W, F)$ be two graphs. We say that G is **isomorphic** to H , in symbols $G \simeq H$ if there exists a bijection

$$f : V \rightarrow W$$

such that any two vertices u and v of G are adjacent in G if and only if $f(u)$ and $f(v)$ are adjacent in H . If the graphs G and H are coloured the bijection f additionally has to satisfy that $f(v)$ has the same color as v for all vertices $v \in V$.




Figure 20: Three isomorphic contact graphs.

- (8.18) For example, consider the graphs depicted in figure 20: The graph on the left-hand side is the contact graph of the design depicted on the left hand-side of figure 18 and the one on the right corresponds to the design on the right-hand side of the same figure. These two graphs are isomorphic since the function f that maps r_i to R_i , b_i to B_i and p_i to P_i fulfills all the requirements of Definition 4. The graph depicted in the middle of figure 20 simply shows an intermediate step of the bijection: flipping the left-hand graph horizontally and relabelling according to the bijection or, alternatively, rotating the right-hand graph by approximately 30 degrees in anti-clockwise direction.

8.3 Possible limitations of this approach and suggestions how to resolve these

8.3.1 Incorporating distances between components

- (8.19) One possible issue with the idea of classifying umbilical designs by identifying them with their contact graphs is the following. Contact graphs do not record the actual distances between components, but merely whether components touch each other not. Consider for example the four umbilical designs D_1 , D_2 , D_3 and D_4 consisting of one large and four smaller components that are depicted in figure 21. All four designs have the same contact graph G that is depicted on the left-hand side of the figure. However, we might not want to consider all four of these designs as the same—they do not all have the same amount of symmetry and are very likely not to have the same mechanical properties. To be more precise, we will most likely want to consider designs D_1 and D_4 as the same, but designs D_2 and D_3 are too different from D_1 to consider them as the same.
- (8.20) In the following we will see how this issue could be resolved by introducing **scaled distance matrices** and using them instead of adjacency matrices of contact graphs associated with umbilical designs. Let us however first observe that a situation as depicted in figure 21 is not likely to occur in a complete umbilical design consisting of a set of prescribed components and a multitude of fillers that have been added to increase stability, improve




Figure 21: The four (partial) designs depicted above all have the same contact graph G even though we might not want to consider them as equivalent.

mechanical properties and create a circular shape of the design. Indeed, as remarked earlier, an umbilical design is highly characterized by its compactness and the fact that components and fillers will all be in contact with at least two other components or fillers. Thus, such sparse contact graphs with as few edges as the graph G in the example, will most likely only appear in partial umbilical designs where fillers have not yet been placed or in intermediate steps of algorithms that generate designs.

Definition 5. Let D be an umbilical design with n components⁶. Further, let x_1, \dots, x_n be the x -coordinates of the components in the plane, let y_1, \dots, y_n be the y -coordinates and r_1, \dots, r_n the radii of the respective components. Then the **scaled distance matrix** $S(D)$ is the $n \times n$ -matrix with entries $s_{i,j}$ defined as follows:

$$s_{i,j} = \frac{\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}}{r_i + r_j} \text{ for all } 1 \leq i, j \leq n.$$

(8.21) Note that the scaling factor $1/(r_i + r_j)$ leads to the fact that $s_{i,j} = 1$ if and only if the i -th and the j -th component are touching each other, i.e., the corresponding entry in the adjacency matrix $A(G(D))$ is also equal to 1. Moreover, entries on the diagonal will always be equal to 0. All other entries in the matrix $S(D)$ are larger than 1 and how far away from 1 the value $s_{i,j}$ is indicates how far away the i -th and the j -th component are from touching each other.

(8.22) The four designs in figure 21 are characterized by the (x, y) coordinates (rounded to one decimal):

	(x_1, y_1)	(x_2, y_2)	(x_3, y_3)	(x_4, y_4)	(x_5, y_5)
D_1	(0, 0)	(3, 0)	(0, 3)	(-3, 0)	(0, -3)
D_2	(0, 0)	(3, 0)	(-1.5, 2.6)	(-3, 0)	(-1.5, -2.6)
D_3	(0, 0)	(3, 0)	(1.5, 3.6)	(-1.5, 2.6)	(-3, 0)
D_4	(0, 0)	(3, 0.2)	(0.2, 3)	(-3, 0)	(-0.2, -3)

⁶For simplicity, we do not distinguish between fillers and components here. Moreover, we assume that there are no two different components that are of the same size. Incorporating these distinctions could be done very easily by considering a coloured version of scaled distance matrices.

Moreover, it holds that $r_1 = 2$ and $r_i = 1$ for $2 \leq i \leq 5$. Thus we obtain the following four scaled distance matrices (entries are rounded to two decimals):

$$\begin{aligned}
 S(D_1) &= \begin{array}{c|ccccc} & 1 & 2 & 3 & 4 & 5 \\ \hline 1 & 0 & 1 & 1 & 1 & 1 \\ \hline 2 & 1 & 0 & 2.12 & 3 & 2.12 \\ 3 & 1 & 2.12 & 0 & 2.12 & 3 \\ 4 & 1 & 3 & 2.12 & 0 & 2.12 \\ 5 & 1 & 2.12 & 3 & 2.12 & 0 \end{array}, \\
 S(D_2) &= \begin{array}{c|ccccc} & 1 & 2 & 3 & 4 & 5 \\ \hline 1 & 0 & 1 & 1 & 1 & 1 \\ \hline 2 & 1 & 0 & 2.6 & 3 & 2.6 \\ 3 & 1 & 2.6 & 0 & 1.5 & 2.6 \\ 4 & 1 & 3 & 1.5 & 0 & 1.5 \\ 5 & 1 & 2.6 & 2.6 & 1.5 & 0 \end{array}, \\
 S(D_3) &= \begin{array}{c|ccccc} & 1 & 2 & 3 & 4 & 5 \\ \hline 1 & 0 & 1 & 1 & 1 & 1 \\ \hline 2 & 1 & 0 & 2.12 & 3 & 2.12 \\ 3 & 1 & 2.12 & 0 & 2.12 & 3 \\ 4 & 1 & 3 & 2.12 & 0 & 2.12 \\ 5 & 1 & 2.12 & 3 & 2.12 & 0 \end{array} \text{ and} \\
 S(D_4) &= \begin{array}{c|ccccc} & 1 & 2 & 3 & 4 & 5 \\ \hline 1 & 0 & 1 & 1 & 1 & 1 \\ \hline 2 & 1 & 0 & 2.6 & 3 & 2.6 \\ 3 & 1 & 2.6 & 0 & 1.5 & 2.6 \\ 4 & 1 & 3 & 1.5 & 0 & 1.5 \\ 5 & 1 & 2.6 & 2.6 & 1.5 & 0 \end{array}
 \end{aligned}$$

(8.23) Using this concept of scaled distance matrices we can now define a **metric** on the set of all umbilical designs with identical contact graphs that allows us to say when two umbilical designs should be considered the same:

Definition 6. Let D and E be two umbilical designs with isomorphic contact graphs and let their respective scaled distance matrices be $S(D) = S = (s_{i,j})$ and $S(E) = T = (t_{i,j})$ ⁷. Moreover, let $\tau > 0$ be a fixed threshold. Then the distance between D and E is:

$$d(D, E) = \max_{i,j} (|s_{i,j} - t_{i,j}|)$$

and D and E are considered as the same if $d(E, D) < \tau$.

⁷This means that $s_{i,j} = 1$ if and only if $t_{i,j} = 1$.

(8.24) For our example, we obtain the following distances:

	D_1	D_2	D_3	D_4
D_1	0	0.62	0.88	0.15
D_2	0.62	0	1.1	0.69
D_3	0.88	1.1	0	0.74
D_4	0.15	0.69	0.74	0

(8.25) So, if one sets the threshold value $\tau = 0.2$ the two designs D_1 and D_4 are considered as the same but all other pairs of designs need to be considered as different.

(8.26) Note that it would be worthwhile investigating other possible definitions of the distance $d(D, E)$. For instance, one could consider the **Frobenius norm** of the matrix $(|s_{i,j} - t_{i,j}|)_{i,j}$. Also note that choosing an appropriate value for τ will highly depend on the (number of) involved components and the desired properties of the specific umbilical design.

8.3.2 Understanding graph isomorphism

(8.27) Another issue with this approach—that is perhaps more of a theoretical nature—is that graph isomorphism is not well understood geometrically. Definition 20 is very clear from a formal graph-theoretical point of view but it says little about the geometry of isomorphic graphs. To what extent do isomorphic graphs look the same? Can it happen that two contact graphs are isomorphic even though their underlying umbilical designs are geometrically not the same?

(8.28) In order to tackle this question, it is important to note that contact graphs are not just arbitrary graphs: they are always connected and planar. There are two graph-theoretic results that are relevant in order to understand graph isomorphism for contact graphs:

1. The **Circle packing theorem**, also known as Koebe-Andreev-Thurston theorem. The following description is taken from https://en.wikipedia.org/wiki/Circle_packing_theorem: “For every connected planar graph G there is a **circle packing** in the plane whose intersection graph is (isomorphic to) G and this circle packing is unique up to Möbius transformations and reflections in lines. A circle packing is a connected collection of circles [...] whose interiors are disjoint. The [contact] graph [(also referred to as a coin graph)] of a circle packing is the graph having a vertex for each circle, and an edge for every pair of circles that are tangent [i.e. touching each other].”

Thus circle packings are highly related to umbilical designs and there contact graphs are nothing else than what we defined here a contact graph.




Figure 22: An example of how an umbilical design consisting of two layers (left-hand side) could be represented with the help of two contact graphs.

A scientific discussion of this theorem and other related graph theoretic statements was given by Sachs [7]. We also refer to the references therein. A history of circle packing theory that is more accessible to a non-specialized audience was published in the notices of the AMS [9].

2. **Fáry's theorem** [3] states that any planar graph can be drawn without crossings so that its edges are straight line segments.

It is certainly worthwhile analysing the implications of these two theorems in order to resolve this question.

8.3.3 Multi-layered designs

- (8.29) So far, we have only considered single-layered umbilical designs in this part of the project. However, optimal designs often involve several layers, i.e., are assembled in multiple stages and thus extending the classification of designs with the help of contact graphs to multi-layered designs is necessary. This could be done in the following way: An inner layer is considered as a vertex of a special kind (it is itself a graph), and components are connected to this vertex if they touch the outside border of the inner layer. For an example, consider the two-layered umbilical design depicted in figure 22 together with a representation of its contact graph. The central vertex labelled “c” is here a graph itself that consists of three component-nodes and of three filler-nodes.

8.4 Numerical Implementation

- (8.30) We will now demonstrate proof of concept through our numerical scheme implemented in Matlab using contact graphs to identify structurally different solutions.

8.4.1 Numerical Scheme

- (8.31) For each umbilical configuration we considered two elements as adjacent if the gap between them was smaller than ϵ_d , i.e., if $(x_1 - x_2)^2 + (y_1 - y_2)^2 - (r_1 + r_2)^2 < \epsilon_d$. The value ϵ_d is non-zero as numerical errors mean adjacent elements will always be slightly separated.
- (8.32) Using this adjacency definition we constructed the adjacency matrix (see Definition 3) for every umbilical by comparing the centres of every pair of elements. By swapping rows and columns we then used these matrices to determine which configurations had isomorphic graphs. It is still an unsolved problem whether checking graph isomorphism is an NP-complete problem or if it is solvable in polynomial time. However, contact graphs are always planar, and there are many algorithms for graph isomorphism of planar graphs which run in polynomial time. In particular all umbilical designs have < 100 components, which is a relatively small number of graph nodes, so this comparison can be done quickly.
- (8.33) We examined 200 “optimal” designs obtained from the genetic algorithm discussed in Section 4.2. We identified those designs that are structurally different and calculated the proportion of initial conditions which led to these structures (given to 2 decimal places). Two examples are given in this report: firstly 5 components of equal radius as displayed in figure 23 and secondly 9 components with 2 of radius one, 2 of radius two, 3 of radius three and 2 of radius 4 as shown in figure 24.

8.4.2 Numerical Results

- (8.34) The base case of 3 equally sized components was carried out and we identified that all designs had the same single optimal packing configuration up to rotation. Similarly the case of one large and five small components (similar to figure 21) was studied and the different configurations identified.
- (8.35) The case of 5 equally sized components (see figure 23) gave rise to four different structures. Either of the first two could be desirable for the umbilical dependant on other design requirements. The third structure might want to be considered separately as components are arrayed more evenly. Alternatively it might be desirable to group it with the very similar second option which can be achieved by allowing one difference in contact graphs. The fourth structure shows that error checking to ensure solutions have converged is important.
- (8.36) The different geometric structures of a more complicated umbilical were considered in figure 24. Initially, there were too many geometrically different optimal structures but there are a number of ways to reduce this number. How this can be done is discussed in the next section. Here in particular

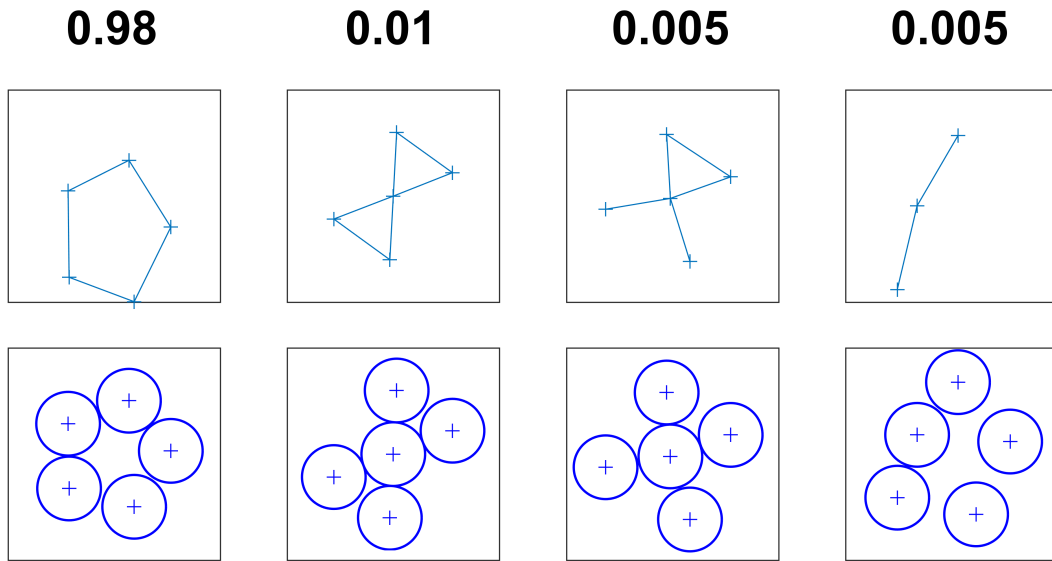


Figure 23: Geometrical shapes for the optimal packing of 5 equally sized components. Each column corresponds to one geometrical configuration. Each row in order corresponds to the proportion of designs which correspond to this configuration to 2dp), the contact graph of this configuration and an example of this configuration.




Figure 24: Geometrical shapes for the optimal packing of 2 components of radius one, 2 of radius two, 3 of radius three, and 2 of radius 4. Each column corresponds to one geometrical configuration. The first row shows the contact graph of this configuration and an example of this configuration is on the second row. Differences between contact graphs has been allowed here to reduce the number of cases.

we only considered cases structurally different if there were more than 5 differences in their contact graphs.

8.4.3 Producing a Manageable Number of Configurations

(8.37) For many components, particularly with a wide range of sizes, there are too many geometrically different optimal configurations. There are a number of ways to reduce this:

- Error check optimal solutions to ensure they have converged to a packed configuration
- Improve the choice for ϵ_d . We expect ϵ_d should be a function of the radii of the two components being compared for adjacency.
- Group solutions as structurally the same if there is up to a (small) specified number of differences in their contact graphs. This was used in figure 24 to limit the available configurations which otherwise gave most of the solutions as different. One disadvantage is that this parameter means some structures aren't considered. Hence ideally it is small or zero but should be chosen to balance the number of required output solutions. This is an easy way to limit the number of configurations while still distinguishing those which differ the most. But other filters based on solution symmetry etc. would be better. We expect this parameter should depend on the total number of components and the number of different types of component.
- Filter structures to ignore undesirable structures, such as very non-symmetric ones.

A detailed parameter study still needs to be carried out to determine the formulations the above parameters should take for any general umbilical. These parameters could also be dependent on the desired number of outputted structures.

8.5 Conclusions

- (8.38) We introduced the concept of contact graphs in order to classify umbilical designs which are structurally equivalent. Numerical implementation in Matlab allowed us to provide proof of concept. This method adds flexibility to optimisation criteria by offering a handful of options to choose from. We suggested further improvements including the use of distance matrices and parameter analysis.

9 Mechanical response of the Umbilical

- (9.1) An additional constraint on the umbilical design not yet discussed is the mechanical response. During installation the umbilical is held in high tension and compressed whilst being laid. There is a design limitation that only a limited amount of deformation of each component is allowed within the optimal design. Currently the response under this loading is established in part through experience of the designers, and then confirmed through fully three-dimensional numerical finite element models to ensure the required mechanical properties in the design are met. This is clearly resource intensive and not compatible with the optimisation methods discussed above, although of course can still be used to validate the mechanical properties of feasible designs. As an alternative, the group also considered whether a simpler analytical approach could offer an insight into the mechanical properties of the design, perhaps to offer additional constraints in the optimisation process.

9.1 Linear elastic model

- (9.2) The main progress made was considering a linear elastic model of individual components (either a steel annulus, or solid filler), subjected to arbitrary external loading (from adjacent components). The ambition is to then combine these individual components to get a measure for the overall behaviour of the entire design. For simplicity, we neglect the three-dimensional effects introduced by the helical nature of the real umbilical and restrict attention to two-dimensional cross-sections of a cylinder. We also consider only plane-strain, with deformations only occurring in the plane aligned with the umbilical cross section, with no deformations occurring along the umbilical. Although umbilicals are subjected to a wide range of bending and stretching during installation, it is thought that the response to the squeezing in the rollers is the most important feature to account for in the design process. Two articles considering related problems were found. Firstly, article [1] considers the behaviour of a composite layered cylinder, subjected to squeezing in plane-strain. Secondly, article [10] uses a similar technique to examine a squeezed thin annular disk. The derivation below draws upon both.
- (9.3) For a linearly elastic material in equilibrium, the radial (σ_{rr}), hoop ($\sigma_{\theta\theta}$) and shear ($\sigma_{r\theta}$) stresses satisfy in polar coordinates (r, θ):

$$\frac{\partial \sigma_{rr}}{\partial r} + \frac{1}{r} \frac{\partial \sigma_{r\theta}}{\partial \theta} + \frac{1}{r} (\sigma_{rr} - \sigma_{\theta\theta}) = 0, \quad (22)$$

$$\frac{1}{r} \frac{\partial \sigma_{\theta\theta}}{\partial \theta} + \frac{\partial \sigma_{r\theta}}{\partial r} + \frac{2}{r} \sigma_{r\theta} = 0. \quad (23)$$

In plane-strain the constitutive equations are given by

$$\sigma_{rr} = \frac{E}{(1+\nu)(1-2\nu)}(\nu\epsilon_{\theta\theta} + (1-\nu)\epsilon_{rr}), \quad (24)$$

$$\sigma_{r\theta} = \frac{E}{(1+\nu)}\epsilon_{r\theta}, \quad (25)$$

$$\sigma_{\theta\theta} = \frac{E}{(1+\nu)(1-2\nu)}(\nu\epsilon_{rr} + (1-\nu)\epsilon_{\theta\theta}), \quad (26)$$

where E is the Young's modulus and ν is Poisson's ratio. We write the radial and tangential displacements as $u = u(r, \theta)$ and $v = v(r, \theta)$, and use the usual strain-displacement relations

$$\epsilon_{rr} = \frac{\partial u}{\partial r}, \quad \epsilon_{\theta\theta} = \frac{u}{r} + \frac{1}{r} \frac{\partial v}{\partial \theta}, \quad \epsilon_{r\theta} = \frac{1}{2} \left(\frac{1}{r} \frac{\partial u}{\partial \theta} + \frac{\partial v}{\partial r} - \frac{v}{r} \right). \quad (27)$$

Combining these expression gives Navier's equations for electrostatics:

$$2(1-\nu) \left(\frac{\partial^2 u}{\partial r^2} + \frac{1}{r} \frac{\partial u}{\partial r} - \frac{u}{r^2} \right) + \frac{1-2\nu}{r^2} \frac{\partial^2 u}{\partial \theta^2} + \frac{1}{r} \frac{\partial^2 v}{\partial r \partial \theta} + \frac{4\nu-3}{r^2} \frac{\partial v}{\partial \theta} = 0, \quad (28)$$

$$(1-2\nu) \left(\frac{\partial^2 v}{\partial r^2} + \frac{1}{r} \frac{\partial v}{\partial r} - \frac{v}{r^2} \right) + \frac{2(1-\nu)}{r^2} \frac{\partial^2 v}{\partial \theta^2} + \frac{1}{r} \frac{\partial^2 u}{\partial r \partial \theta} - \frac{3-4\nu}{r^2} \frac{\partial u}{\partial \theta} = 0. \quad (29)$$

- (9.4) For the two problems considered (a solid disc and an annulus), the boundary conditions (tractions) on the outer radius $r = b$ are specified as

$$\sigma_{rr}(b, \theta) = -P(\theta), \quad \sigma_{r\theta}(b, \theta) = 0, \quad (30)$$

where $P(\theta)$ represents the external load per unit length. For the problem considered P will be zero for all θ , except in a small number of regions (which may be infinitesimally small), where a constant load is applied. To maintain equilibrium, the integral over θ of the applied loads is zero. Additionally for the solid disc we require the displacements to be non-singular as $r \rightarrow 0$, whereas for the annulus of inner radius $r = a$ we instead impose $\sigma_{rr}(a, \theta) = \sigma_{r\theta}(a, \theta) = 0$.

- (9.5) Given that u and v must be periodic, we expand the displacements in the form of Fourier series. For simplicity we restrict attention to cases where P can be represented by even functions but the method is readily extendible to arbitrary load distributions. Writing the displacements in the form

$$u(r, \theta) = \sum_{n=0}^{\infty} f_n(r) \cos(n\theta), \quad v(r, \theta) = \sum_{n=0}^{\infty} g_n(r) \sin(n\theta), \quad (31)$$

and writing the imposed load in the form

$$\sigma_{rr}(b, \theta) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n \cos(n\theta), \quad (32)$$

where the c_n are the appropriate Fourier coefficients. Substitution into the governing equations and solving the standard ordinary differential equations for both f and g we obtain

$$u(r, \theta) = A_0^1 r + \frac{A_0^2}{r} + \sum_{n=1}^{\infty} \sum_{j=1}^4 \frac{n(\lambda_j + 4\nu - 3)}{2\lambda_j(1 - \nu)(1 + n^2\alpha - \lambda_j^2)} A_n^j r^{\lambda_j} \cos(n\theta), \quad (33)$$

$$v(r, \theta) = \sum_{n=1}^{\infty} \sum_{j=1}^4 \frac{A_n^j r^{\lambda_j}}{\lambda_j} \sin(n\theta), \quad (34)$$

where $\alpha = (1 - 2\nu)(2(1 - \nu))^{-1}$ and $\lambda_1 = n + 1$, $\lambda_2 = n - 1$, $\lambda_3 = -n - 1$ and $\lambda_4 = -n + 1$. The unknown constants A_n^j are then fixed by the boundary conditions. For the symmetric case we consider here we assume $u(r, \theta) = u(r, \pi - \theta)$ and $v(r, \theta) = -v(r, \pi - \theta)$, which leaves $A_n^j = 0$ for n odd. For a solid cylinder, A_0^2 , A_n^3 and A_n^4 are taken to be zero to ensure the displacements remains finite. For an annulus all terms are retained.

- (9.6) The remaining unknown constants A_n^j are determined by imposing the boundary conditions on σ_{rr} and $\sigma_{r\theta}$ given by (24), (25) and (30). Substituting in the expressions for u and v and (32) leaves a simple linear system to solve for the unknown constants. In figure 25 we show typical solutions for the solid cylinder case with pinching at two and four locations. In figure 26 we present the the solution for the thin annulus subjected to the same loading. The imposed loading was of the form

$$\sigma_{rr} = \begin{cases} -P_0 & \text{for } |\theta - \frac{\pi}{2}| < 0.4 \text{ and } |\theta - \frac{3\pi}{2}| < 0.4 \\ 0 & \text{otherwise} \end{cases} \quad (35)$$

or

$$\sigma_{rr} = \begin{cases} -P_0 & \text{for } |\theta - \frac{n\pi}{2}| < 0.2 \quad n = 1, 2, 3, 4 \\ 0 & \text{otherwise} \end{cases} \quad (36)$$

with P_0 constant. Note these results are for arbitrary values of the parameters rather than for specific components, chosen to emphasise the deformations that can take place.

- (9.7) As mentioned above the solution is readily extendible to arbitrary loading. In principle, it is then possible to consider a combination of touching components. There is still a need to resolve the extent of the contact patches between components (Hertzian contact) under plane-strain but then the above method could feasibly be extended to consider multicomponent designs. The whole discussion above limits the study to linear elastic material. In reality, the design criteria allows for some (small) permanent deformation of the annulus. Thus there is a further plastic problem, originating on the inner radius, where the yield stress is exceeded.

9.2 Possible alternative approaches

- (9.8) An alternative approach to that detailed above would be to exploit the rela-




Figure 25: Deformation of a solid cylindrical component pinched at two locations (left) and four locations (right) with the load P as given in the text. The red line shows the undisturbed shape.

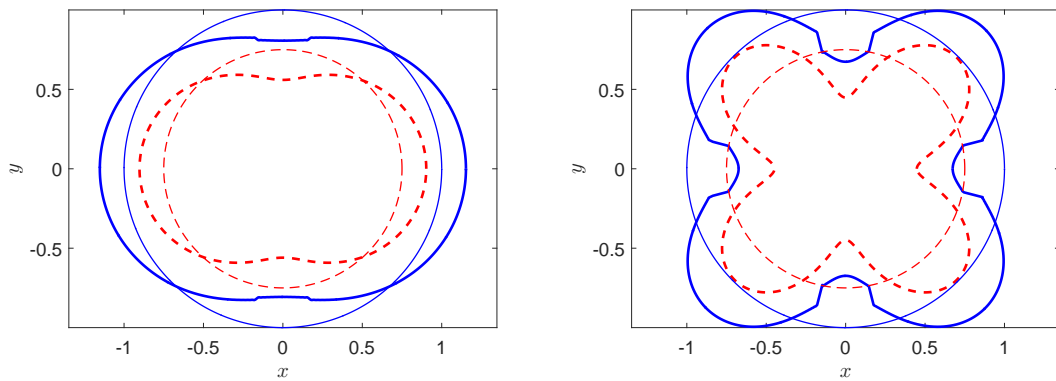


Figure 26: Deformation of a thin annulus component pinched at two locations (left) and four locations (right) with the load P as given in the text. The undeformed shape is also shown. The red line shows the inner radius, the blue the outer radius.

tive thickness of the steel pipes to describe the behaviour of the annular case using a nonlinear beam. Describing the deformation of the annulus using the arc-length s and angle $\Theta(s)$ between the centreline and x -axis. Following the derivation in [5], we find the Euler-Bernoulli beam equation

$$EI \frac{d^2 \Theta}{ds^2} + N_0 \cos \Theta - T_0 \sin \Theta = 0. \quad (37)$$

For a finite beam N_0 and T_0 are the normal and tangential forces applied at each end. In the current cylindrical case, N_0 and T_0 must be chosen to ensure Θ is periodic. Additionally N will include some source terms (perhaps delta-functions for simplicity) to model the pinching between the rollers. If the tube was just an arc of a circle, and T was prescribed at each end, then we would have a classical buckling problem, but now we will be interested in reaching the yield stress for a tube that is becoming oval in shape. For small enough displacements, it is proposed that an array of nearly touching tubes could be modelled as a network of springs whose moduli would come out of a linearization of (37).

- (9.9) The other possible approach considered was to attempt an homogenisation. Given the multi-component nature of the more complex umbilical design, it was felt that by considering averaged mechanical properties across the whole device, we could achieve a simplified problem, more compatible with the optimisation routine. Following the derivation for a 2D Cartesian geometry given in reference [5], some progress was made at extending this to a Cartesian array of circular components. The outline approach is to consider a single component, assumed embedded in a large array of components, under the assumption that although mechanical properties vary considerable over a local length coordinate, they only change gradually over a longer length scale. By adopting a multiple scales approach, the solution in the longer length scale can be found in terms of uniform properties homogenised over the local length scale. This approach showed some promise, although the contact between components still needs to be addressed as in section 9.1. The extension from a 2D infinite periodic array into the cylindrical geometry seen here does not seem a trivial task however.

10 Learning from the experts

- (10.1) So far we have concentrated on quantifying the known objectives of the problem so that the automated systems are able to determine the quality of a given design. While some objectives, such as the size of the bounding circle, are easy to measure, others are less clear-cut, and there remain a number of objectives for which quantification is a significant challenge, for example, manufacturing risk.
- (10.2) Technip has two sources of design expertise: an archive containing several years' worth of successful designs, and a human design expert. In this section we consider how these two different sources might be used to aid the automated design process.

10.1 Learning from the archive

- (10.3) Over the last 15 years or so, Technip has archived several hundred successful umbilical designs. Some of them were manufactured; the remainder are intermediate designs leading to a final manufacture, though no link is recorded between them. Altogether, this represents a set of good designs from which we may be able to measure how good a given computer-generated design is without having to fully understand and measure individual objectives. We can consider this as a statistical classification problem: given a potential design, how likely is it that it is a good one?
- (10.4) Usually such a system is trained by giving it a wide range of examples, covering both good and bad designs, from which it learns to extract or identify features that differentiate between the two classes. It can then extract/identify the same set of features from a previously unseen example and use them to predict which class the example comes from. In many cases, rather than just assigning it a class, the system provides a probability of membership of each class, e.g. "example X is 90% likely to be good design", from which the user may determine some measure of certainty.
- (10.5) In the case of the Technip archive we only have good design examples. This is very like the problem of assessing social media (e.g. Facebook) web pages on the basis of whether or not they have been "liked". A "like" may imply that a user considers the page to be good in some way (a *positive* example), but a lack of a "like" (an *unlabelled* example) does not imply that they consider the page to be bad; they may just not have seen it. So there is a body of literature investigating methods of classifying potential solutions into positive and negative, based only on the experience provided by positive examples. One-class [6], or PU (Positive-Unlabelled) (for example, [11]) learning algorithms work from a set of positive examples, augmented by a set of unlabelled examples that may be randomly generated.

(10.6) In order to train a classification model we need to extract the same set of design features from each example. The umbilical design archive records the following items about each design:

- a parts list,
- the physical properties of each of the parts,
- an AutoCAD design file, from which the position of each part and the shape and position of each filler can be extracted,
- the date of the design,
- the depth of water for which the design was made,
- whether the design was manufactured or not,
- whether the design was for installation in a dynamic or a static environment, and
- for a small selection of designs only (those that are sufficiently novel), results of full technical testing.

Within this set of good designs we might also suggest that newer designs are better than older ones, and manufactured designs are better than those that were not manufactured.

(10.7) Given that each design is bespoke, it is highly unlikely that a new design will be an exact copy of an existing one. Since they are likely to contain different numbers and types of parts and have different mechanical requirements, we must probably rely on features that describe each design as a whole rather than its individual parts. Examples of such features are: the radius of the finished product, its centre of mass, the number of components, the average density along a radius, and the dynamic/static environment classification.

10.2 Learning from the human expert

(10.8) The other source of significant design expertise is, of course, the human designer. This designer is able to intuitively assess a concept as complex as “manufacturing risk” without having to enumerate or measure all the different contributions to this risk. It is this lack of quantification that makes it hard to design an algorithm that can optimise it.

(10.9) Consider one of the algorithms required for self-driving cars: selecting an appropriate speed. “Stanley” was a standard road vehicle (a VW Touareg), altered so that it could drive autonomously or be driven by a human. It took part in the 2005 DARPA Grand Challenge race, which required it to drive a cross-country course over many different types and qualities of road surface. How do you design an algorithm that can select the fastest possible, safe speed for any given road surface and set of environmental conditions? Stanley’s designers opted to learn by observation (see [8, section 2.4.3]); the vehicle was human-driven over many miles and the algorithms learned how to extract information from which they could make their speed decisions. The

result was that Stanley famously won the race, and has laid the groundwork for the current work on autonomous vehicle design.

(10.10) So, how do we use the human to drive the design process in some similar way? Given two different designs, the designer is able to state a preference for one or other of them, and we might use this preference within the design algorithm. Incorporating the human's opinion has three advantages: (a) it removes the need to quantify something that might be unquantifiable, (b) it improves the optimisation of the umbilical designs, and (c) it, in effect, provides a new archive containing "preference" data. The latter could be used to drive a future version of the optimisation algorithm.

11 Conclusions

- (11.1) Given that there are likely to be too many potential solutions to a given umbilical design problem to find the overall best design (the *global* optimum) within an acceptable time-frame, we have investigated methods for generating feasible designs that are as good as possible (one of a number of *local* optima) given the time constraint.
- (11.2) We have shown a variety of methods for generating feasible designs that are optimum in some way, while respecting constraints such as the fact that the umbilical components are solid, physical things that cannot overlap. A subsequent optimisation step has been shown to add appropriate circular fillers to these designs.
- (11.3) We have considered a number of measures for comparing the relative qualities of different feasible designs, including the size of the bounding circle, symmetry, and the preference for certain types of components to be closer to the centre of the umbilical. In general, due to the time constraint imposed by the one-week work-group, these are relatively simple, geometric measures. A more complete solution would need to evaluate the physical properties of a design. In this area we have investigated the mechanical response of an umbilical as it is subjected to stretching and compression during the process of deployment.
- (11.4) It is highly beneficial to be able to assess how similar two designs are, and for some of the algorithms, to be able to generate a set of random-ish starting designs that are as dissimilar as possible. A good range of designs improves the chances of finding the global optimum solution, or at least a better local optimum.
- (11.5) We have suggested ways in which the expertise available within Technip could be used to drive or improve the design process, either by using their design archive, or by including the human designer in the process.
- (11.6) All of these elements would feature in a full design system. The system would take in the customer's specification. Using a range of different measures that can be calculated from a plausible design (including its mechanical properties) and the design archive, the system would generate pairs of designs that it considers to be "good" and present these to the human designer who would specify which of the pair is the preferred option. This preference is used by the system to provide an improved pair for the human to comment on. There would be several iterations of system-human interaction, until the human is satisfied with the quality. A separate process would try and learn from the pair-preferences what features of a design leads to improved human preference, and this would then be fed back into the design system.

References

- [1] Davison, T. and Wadley, H. (1994). Elastic response of a layered cylinder subjected to diametral loading. *Composites Engineering*, 4(10):995–1009.
- [2] Diestel, R. (2010). *Graph theory, Graduate Texts in Mathematics; 173*. Springer-Verlag Berlin and Heidelberg GmbH.
- [3] Fáry, I. (1948). On straight line representations of planar graphs. *Acta. Sci. Math.(Szeged)*, 11:229–233.
- [4] Goldberg, D. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Reading, MA, USA.
- [5] Howell, P., Kozyreff, G., and Ockendon, J. (2009). *Applied Solid Mechanics*. Cambridge University Press.
- [6] Moya, M. and Hush, D. (1996). Network constraints and multi-objective optimization for one-class classification. *Neural Networks*, 9(3):463–474.
- [7] Sachs, H. (1994). Coin graphs, polyhedra, and conformal mapping. *Discrete Mathematics*, 134(1):133–138.
- [8] Stavens, D. (2011). *Learning to drive: perception for autonomous cars (available from <http://purl.stanford.edu/pb661px9942>)*. PhD thesis, Stanford University,.
- [9] Stephenson, K. (2003). Circle packing: a mathematical tale. *Notices of the AMS*, 50(11):1376–1388.
- [10] Tokovyy, Y., Hung, K.-M., and Ma, C.-C. (2010). Determination of stresses and displacements in a thin annular disk subjected to diametral compression. *Journal of Mathematical Sciences*, 165(3):342–354.
- [11] Yu, H., Han, J., and Chang, K. C.-C. (2002). PEBL: positive example based learning for web page classification using SVM. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*.