

A Genetic Algorithm with a Variable-Length Genotype and Embryogeny for Microstructured Optical Fibre Design

Steven Manos

Optical Fibre Technology Centre & School of Physics
University of Sydney, NSW 2006, Australia
smanos@oftc.usyd.edu.au

Peter J. Bentley

Department of Computer Science
University College London, London WC1E 6BT, UK
p.bentley@cs.ucl.ac.uk

Leon Poladian

School of Mathematics and Statistics
University of Sydney, NSW 2006, Australia
leonp@maths.usyd.edu.au

Maryanne Large

Optical Fibre Technology Centre
University of Sydney, NSW 2006, Australia
m.large@oftc.usyd.edu.au

ABSTRACT

Microstructured optical fibres are a relatively recent advance in fibre technology which guide light by using arrays of air holes which run the length of the fibre. The internal microstructure of optical fibres can be altered to reshape and transform light for use in medical applications, sensing, long distance and local area network high bandwidth communications. Recent progress in the production of polymer fibres allows designs with complex microstructures consisting of hundreds of holes to be manufactured. In this paper we present a generative (embryogenic) representation which can produce symmetric fibre designs with a variable number of holes. The resulting genetic algorithm has the ability to search designs of varying complexity over time, allowing less or more complex designs to be evolved as required. Various aspects of this representation are discussed in light of the supporting genetic algorithm such as recombination of designs and the conversion of the variable length binary genotype to the phenotype (optical fibre structure). We include some single objective design results for a high-bandwidth optical fibre along with manufactured designs.

Categories and Subject Descriptors

I.2.0 [Applications and Expert Systems]: Medicine and Science

General Terms

Design

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'06, July 8–12, 2006, Seattle, Washington, USA.
Copyright 2006 ACM 1-59593-186-4/06/0007 ...\$5.00.




Figure 1: Examples of radially symmetric structures found in nature. A and B are images of *radiolarians*, small ocean-dwelling microbes whose skeletal structure is made from a silica-like material. C is the shell of common sea urchin, which exhibits radial 5-fold symmetry. Images A and B are from [3].

Keywords

Microstructured optical fibres, fibre design, constraint satisfaction, variable length genotype

1. INTRODUCTION

Traditional (solid cored) fibres are generally made of glass (silica) where the internal structure, or the refractive index profile is used to alter the propagation of light. The internal structure is generally circularly symmetric, resulting in a one dimensional parameterisation in terms of finding the most optimal profile given a particular application. A more recent advance in the last decade has been the invention of microstructured optical fibres (MOF) [1], where instead of introducing chemicals into the silica to alter its refractive index, small holes of air are used instead. This has opened up optical fibres to new application areas such as air-core guiding and ‘endlessly single mode’ fibres [2].

One of the major limitations, in terms of design possibilities, in silica MOF is the manufacturing technique used - stacking. Typically a series of hollow rods is stacked to form a *preform*, which is then heated in a furnace and stretched (drawn) down to an optical fibre, where the internal struc-




Figure 2: Examples of manufactured optical fibres. A is a silica microstructured fibre (or photonic crystal fibre). The hexagonal patterning can clearly be seen. B and C are polymer microstructure fibres. The central image is a square array of holes which is used for medical imaging purposes, and the right is a more complex polymer fibre with large variations in hole size and positioning. *Image A is courtesy of the University of Bath.*

ture of the preform is reproduced on a smaller scale. This stacking results in the array being hexagonal (Figure 2A). Drilling is possible, but is difficult and not practical for large-scale manufacture of silica fibres.

A more recent technological advance in microstructured fibres has been the development of microstructured polymer optical fibres (MPOF) [4]. Polymer offers the advantage that it can not only be stacked into a preform, but drilling, casting, extrusion and moulding can also be used to manufacture a fibre preform. As manufacturing technologies become more sophisticated, allowing arbitrary arrays of holes to be inserted into the polymer, the complexity of possible fibre designs expands dramatically.

Nature, through the processes of evolution and adaptation, has managed to learn to create symmetric structures which have an uncanny resemblance to microstructured fibres (Figures 1 and 2). Through a process of embryogeny, where the phenotype is formed by a growth phase, these natural systems self assemble into radially symmetric structures.

In this paper we present a binary Genetic Algorithm (GA) with a variable length representation which can generate symmetric microstructured fibre designs. The algorithm can evolve designs of arbitrary complexity, without any predefined limitations on the number of holes in a design. This means that for a given objective function, less complex or more complex designs can be evolved over time if required. Given the wide range of possible designs, the handling of constraints (for example, non-overlapping holes) becomes difficult. By using a growth phase for the genotype to phenotype mapping, these manufacturing constraints are automatically satisfied.

Section 2 introduces microstructured optical fibres along with some previous examples of fibre optimisation using evolutionary approaches. In Section 3 we give an overview of embryogenies and general design schemes. Section 4 details the genotype to phenotype growth scheme, the mutation and recombination operators and the GA as a whole. In Section 5 we discuss in more detail designing microstructured fibres in order to improve bandwidth, along with results and manufactured designs. We finish up with a discussion on other application areas and future work in Section 6.

2. MICROSTRUCTURED OPTICAL FIBRE DESIGN

Microstructured Optical Fibres (MOF) first emerged from the research work at the University of Bath in the early 1990's. These fibres exhibited some remarkable properties, such as the ability to guide light in air and also strongly alter the properties of light, such as dispersion and polarisation in relatively short lengths of fibre. As a result, research into all aspects of MOFs is a very active and growing research area.

Many examples of microstructured fibre design (both manual and automated) have so far focussed on hexagonal arrangements of holes. These types of structures arise from the manufacturing process of stacking capillaries to produce silica fibres. The regularity and fixed symmetry of such structures (Figure 2A) allows the simple parameterisation of such structures to form a small search space. Typically the size of the holes d and the distance between holes Λ are used to describe a hexagonal array.

The first published example that numerically optimised the properties of microstructured fibres is [5]. Here the authors used relatively simple designs which consisted of 1 or 2 rings of holes that could be parameterised by 2 or 3 real values to optimise the fibres for multiple characteristics of interest such as splice loss and dispersion. The authors in [6] compared the Simplex algorithm, Simulated Annealing and an Evolutionary Strategy in optimising a hexagonal array for mode field diameter and dispersion slope. The hexagonal array was fixed in complexity with 4 rings of holes, and the parameters Λ, d formed the genotype of an individual.

More recent examples include [7, 8] where the sizes of the holes in different rings of a hexagonal structure have been added as a degree of freedom, bringing the search space up to 7 variables. This allowed the fibre to exhibit more complex behaviour, where fibre characteristics such as single-modeness, dispersion control and non-linear effects were considered. The area of microstructured fibre design using evolutionary search and other automated algorithms is in its early stages, with a large scope for development.

3. GROWING THE PHENOTYPE: EMBRYOGENIES IN GENETIC ALGORITHMS

Generally speaking optimisation or design algorithms consist of three main parts [9]:

- The representation,
- The search algorithm,
- The design problem.

This structure is also reflected in the software architecture, allowing the easy interchange of problem definitions, representations or fitness measures.

In order to automate an optimisation or design process, an appropriate *representation* must be used which can describe solutions to the problem of interest. The representation must also be adapted for use in the *search algorithm*, which typically uses operations that can alter and recombine designs to form new and different designs. A binary representation and a genetic algorithm are used, which are discussed in the following sections. The design problem gives

us a method of judging designs relative to one another - a measure of fitness. Here we are comparing fibre designs based on a computationally efficient measure of the *bandwidth* of a fibre. The bandwidth is related to the maximum data rate that can be transmitted down an optical fibre, and is discussed in more detail in Section 5.1.

3.1 Representations: Optimisation -vs- Design

Most GAs typically use a representation where the conversion from genotype to phenotype is relatively straightforward. This *parametrisation* of a solution is typically used in *optimisation* problems as opposed to *design* problems. The complexity of the designs are relatively fixed because it is straightforward for the human designer of the representation to envisage all possible designs.

As the solution space grows, the interactions between various gene elements become far more complex. This can result in design requirements such as hard constraints becoming near impossible to solve. This is where the other types of representations, which can be grouped into more general *open-ended* design, are required. Rather than a one to one correspondence between genotype and phenotype, fundamental modules encoded by various genes can be re-used through concepts such as modularity, hierarchy and symmetry [9]. These types of representations are known as *artificial embryogenies* [10].

Using this approach, not only can the parameters of the design be altered, but the fundamental topology and complexity can evolve over time as well. These representations allow a more novel design process, with less pre-conceptions about the range of possible designs. Since fundamental units are re-used, mutations need only be made on one unit which are then propagated to all other units. The method of propagation and its associated parameters can also be changed, bringing global alterations (such as a change in symmetry or overall size scale) which would be near impossible using a directly parameterised representation.

In relation to our fibre design problem, the main global characteristic of interest is the rotational symmetry, which is explicitly encoded into the genotype. Each modular unit in the genotype then encodes an air hole. Through the process of growth, this genotype is converted into a symmetrical optical fibre design where constraints such as hole overlap are automatically satisfied. This process is discussed in detail in the next section.

4. THE GENETIC ALGORITHM

Genetic algorithms are search algorithms which operate on a population of designs, and through the processes of selection, recombination and mutation, superior designs are evolved. This section outlines the representation and embryogeny used to describe microstructured fibres, along with the corresponding genetic operators.

4.1 Binary Genotype

A binary genotype was chosen to represent the parameters of the fibre such as the symmetry n_{symm} , the hole positions x_i, y_i and hole sizes r_i . This allows the mutation and recombination operation to be relatively straightforward since we are just cutting and splicing binary strings. The fibres considered are made out of polymer and holes all contain air, so no information on the materials needs to be included




Figure 3: Overview of the variable length binary genotype used to represent the symmetry and holes of a microstructured optical fibre.

(although this can be added for more complex designs). The information that needs to be encoded is

- The symmetry of the fibre n_{symm} ,
- The position and size of each hole, x_i, y_i, r_i where $i \in [1, \dots, n_h]$ and n_h is the number of holes.

This is outlined in Figure 3. Although a non-binary genotype which uses a mix of integers and real-values could be implemented, this would then require unnecessary complexity in the recombination operation. The length of the n_{symm} genes is defined as b_s , the x_i, y_i genes have length b_{xy} and the hole diameter value has length b_r .

The variable length nature of the representation arises from the ability to add and delete holes from the genotype through mutation or recombining only particular holes or segments of the genotype. Through these operations, a constraint is placed to guarantee correctness, requiring the total length of the genotype b_t to be

$$b_t = b_s + n_h (b_{xy} + b_{xy} + b_r). \quad (1)$$

4.2 Growth process: Conversion of the binary genotype to an optical fibre

In this section we outline the process of genotype to phenotype conversion: decoding the binary genotype into numerical values, the symmeterisation of the microstructured fibre, followed by the growth of the holes to satisfy constraints.

4.2.1 Decoding the binary genotype

The first phase of the genotype to phenotype mapping involves decoding the genotype (Figure 3) into its constituent numerical values. The first b_s bits are decoded into an integer representing the symmetry n_{symm} . For the purpose of this study we use $b_s = 5$ resulting in a maximum $n_{\text{symm}} = 31$, but higher values are permitted. Values of 0 and 1 can arise, but these are mapped to the minimum allowed symmetry of $n_{\text{symm}} = 2$. Designs with $n_{\text{symm}} = 1$ also possible, but here we are interested in designs which exhibit at least two-fold symmetry (discussed further in Section 4.3).

Following this the n_h holes consisting of triplets x_i, y_i, r_i are decoded into hole positions and sizes. x_i, y_i are decoded directly into unsigned real values with first 6 bits representing the whole component and remaining 3 bits representing the fractional component. This results in a maximum position value of $63.875\mu\text{m}$. No scaling of these values takes place here since this represents the typical upper limits of core sizes manufactured.

The mapping of the r_i values to hole radii is treated a little differently. Because of the manufacturing method currently used for microstructured polymer fibres (where holes are

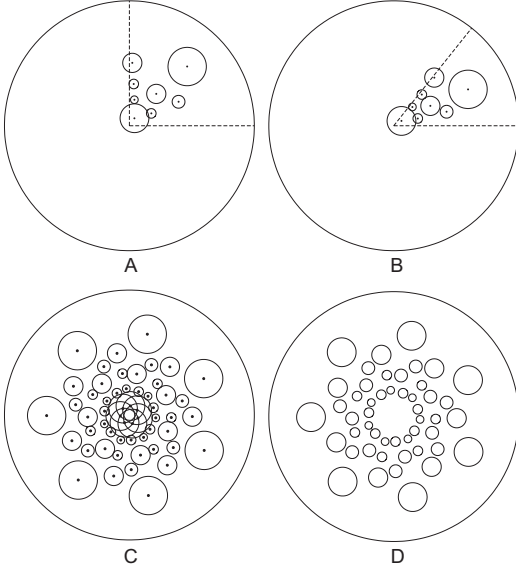


Figure 4: Genotype→phenotype conversion as outlined in section 4.2 of the text. Step A shows the fibre immediately after decoding from the binary genotype. Step B is the initial transformation into the fibres new symmetry (in this case $n_{\text{symm}} = 7$). Step C shows the result of copying the symmetry unit from step B rotationally around the fibre, and step D shows the result of growing the hole radii in order to satisfy hole overlapping constraints.

drilled into the preform) only particular hole sizes (drill bits) are available for use. The 3 bit encoding for r_i is converted into an integer value and the corresponding hole sizes are outlined in Table 1. In this listing, 8 sizes were chosen which give a good spread of hole sizes.

These decoded positions form the basic $n_{\text{symm}} = 4$ structure, an example of which is in Figure 4A.

4.2.2 Converting to a symmeterised fibre

Once we have the basic decoded structure which initially corresponds to $n_{\text{symm}} = 4$ (Figure 4A), this design is transformed into its new genotype-defined n_{symm} . The polar r_i, θ_i coordinates are evaluated for each hole, then they are transformed

$$\theta_i = \frac{4}{n_{\text{symm}}}\theta_i, i \in [1, \dots, n_h] \quad (2)$$

Table 1: Air hole encoding and sizes.

r_i value	Preform drill hole diameter (mm)	Fibre hole diameter (μm)
0	1.2	3.5
1	1.5	4.4
2	2.0	5.9
3	3.0	8.8
4	4.5	13.2
5	6.0	17.6
6	7.0	20.6
7	8.5	25.0

which corresponds to Step B in Figure 4. This group of holes, which we refer to as the *unit holes* are then copied radially $n_{\text{symm}} - 1$ times in steps of $\theta = 2\pi/n_{\text{symm}}$.

4.2.3 Growing the fibre design

As we can see in Figure 4C the holes are initially overlapping or touching after being converted into a complete symmetrical structure. The idea of growing the holes is to solve the hard manufacturing constraint where each hole must be separated a minimum wall thickness w_h (typically $0.5\mu\text{m}$) for structural stability.

Below we give an outline of the growth algorithm in pseudocode. Each hole in the fibre initially has defined:

- Genotype defined radius r_{Geno}
- Phenotypic hole radius r
- Growth value of true or false.

Initially all hole growths are set to true and all radii are set to $r = 0.0$.

```

loop over list of available hole sizes {
  get 'currentHoleSize' from list of available hole sizes
  if no holes are growing {
    break
  } else {
    // If any holes are allowed to grow, set their size to the
    // currentHoleSize

    loop over all holes in the fibre {
      if current fibre hole growth is 'true' {
        set current fibre hole radius to 'currentHoleSize'
      }
    }

    define list holesToStoprLimit;

    // Check to see if any holes have reached their own
    // genotype defined limits on radius

    loop over all holes in the fibre {
      if current fibre hole radius == genotype-defined radius {
        add hole to holesToStoprLimit;
      }
    }

    define list holesToStopOverlapLimit;

    // Do a pairwise check of all holes to check if any
    // overlap

    loop over all unique pairwise holes [H1,H2] in the fibre {
      if holes H1,H2 overlap {
        add hole H1 to holesToStopOverlapLimit
        add hole H2 to holesToStopOverlapLimit
      }
    }

    // Note that now the rules are applied in parallel
    // to be consistant

    loop over all holes in holesToStoprLimit {
      if hole growth is set to true {
        set hole growth to 'false'
        if using the first (smallest) currentHoleSize {
          set hole radius to 0
        } else {
          set hole radius to previous holeRadius
        }
      }
    }
  }
}

```



```

loop over all holes in holesToStopOverlapLimit {
  if hole growth is set to true {
    set hole growth to false
  }
}
}
}

```

This algorithm results in Figure 4D, where none of the holes are overlapping. Some of the holes in Figure 4 have maintained their original size, whereas others have been reduced in size because of surrounding features. Some of the holes, such as the large ones in the centre have not been expressed phenotypically at all because of their close proximity to one another.

Through the growth process of mapping the genotype to the phenotype, we can see this as a ‘correction’ to the phenotype [11], forming a *legal mapping* or phenotype. Since all individuals are phenotypically correct with the respect to constraints, it is expected that the overall performance of the GA is increased since no effort is wasted on individuals who do not satisfy the hard manufacturing constraints through penalty approaches or other methods.

These changes are not propagated back into the genotype. In other words, the genotype specifies the ‘potential’ size of holes, and the surrounding phenotypic features and the genotype itself decide the extent by which these are allowed to increase in size. A number of randomly generated designs are shown in Figure 5.

A phenotypic feature we have not explicitly coded for is a central air hole at the position $x = 0, y = 0$. This could be remedied by including a binary string of length b_r before the coding for n_{symm} which would represent the size of a hole at the point $x = 0, y = 0$. The symmetry copying operation would not be applied to this hole, and the inclusion of this whole would still maintain the global n_{symm} of the fibre structure.

4.3 Symmetry and Constraints

One integral property of the representation is the ability to easily alter the symmetry of designs, where we can convert a design from n_{symm} to any other symmetry. The symmetry approach also fits well with the numerical modelling algorithm used to evaluate the properties of the fibre [12], since it only needs to know the hole coordinates of one symmetry-unit of the microstructured fibre. In fact, the symmetry of the intensity field of the light which is guided down the fibre takes on the same symmetry as the structure itself. Figure 7 shows an example of a randomly generated fibre of $n_{\text{symm}} = 6$ which has then had its symmetry altered to $n_{\text{symm}} = 2, 3, \dots, 10$. In this work we consider designs with a symmetry of at least 2, for three main reasons. Firstly, symmetrical designs generally allow greater control of the optical characteristics. Secondly, they maintain a more regular structure when drawn down to an optical fibre. Finally, any form of symmetry can be exploited to significantly reduce the computational complexity when simulating the behaviour of light within the fibre.

In order to maintain symmetry through the growth process, the constraints themselves are also symmetric. The constraints applied to a unit cell are automatically valid and correct in any copy of the unit cell around the fibre. Using the growth technique, constraints are automatically satisfied, and some holes in the phenotype may not appear.




Figure 5: Randomly generated microstructured fibre designs, using the fixed hole sizes outline in Table 1. Each design varies in symmetry, overall size, the number of holes and complexity.

Although, they exist in the genotype, inhibited by surrounding holes that are in close proximity. Through recombination and mutation, these holes may later appear in the phenotype.

4.4 Recombination and Mutation

The problem of variable length crossover has been previously explored in various ways. Typically the genotypes tend to consist of fixed length blocks, resulting in a type of hierarchical representation of the phenotyps [13]. As a result, limitations are placed on how splicing takes place during variable length crossover to maintain valid genotypes.

In this work we have used single point crossover where two parents are used to produce a single child. The breeding pairs are chosen through a $k = 2$ tournament selection scheme.

Given the first parent with total length given in Eqn. 1 (Section 4.1) a recombination point is chosen randomly in the range $[1, b_t]$. In order for the child to maintain a correct length (Eqn. 1 must hold for an integer number of holes n_h), the recombination point on the second parent must be chosen carefully. Given Parent 2 with n_h holes, we choose a random integer $[1, n_h]$. The exact cut position within the Parent 2 binary code depends on the cut point selected for Parent 1 since we need to maintain the correct length. Some examples of the phenotypes of recombined individuals are shown in Figure 6.




Figure 6: Examples of recombination only (no mutation) from two randomly generated parents with different lengths to single children. The inheritance of features such as symmetry, groups of holes and single holes can be seen in the examples. For example, in the first recombination example a child with $n_{\text{symm}} = 4$ is formed which contains holes from both parents. In the second example Parent 1 is acting as a symmetry operator, almost turning the child into a copy of Parent 2 but with $n_{\text{symm}} = 3$ and the single holes from Parent 1. Characteristics from both parents is clearly evident in the final example.

The mutation operator uses different binary mutation rates depending on the gene segment. A mutation rate of $p_s = 1/2n_{\text{symm}}$ is used for the binary gene representing the symmetry, such that on average one bit within n_{symm} is flipped every two individuals. A similar argument follows for the position mutation rate $p_{xy} = 1/2b_{xy}$ and the hole radii binary rate $p_r = 1/2b_r$. Higher level (hierarchical) operators are used which can randomly delete holes from the genotype with probability p_d or add random holes with probability p_a . Currently values of $p_a = p_d = 0.001$ are used. These low rates are used since the recombination operator already has a strong tendency to generate designs with more or less holes. Finally, to avoid selection bias to holes which lie on the ends of the genotype (due to the recombination operator used), each individual has the gene locations of a randomly selected pair of holes swapped every generation. More work is planned on ascertaining which sort of operators are best to use, and what mutation rates are most appropriate for this representation.

5. SIMULATION AND RESULTS

Two of the most important performance measures for data transmission fibres relate to bandwidth and transmission light loss [14]. The accurate evaluation of these parameters can be computationally expensive, but for the purpose of these experiments we can make some approximations to demonstrate the feasibility of using this GA for generic, creative fibre design.

5.1 High bandwidth fibre design

Fibre loss, or confinement loss, refers to the rate of leakage of light out of the sides of the fibre. Light is reflected back as it hits the polymer-air interface of a hole, but it can leak through gaps between the holes and also tunnel through the smaller air holes. Light is supported in different ‘modes’ in the core of the fibre, and large core polymer fibres are typically multi-moded. In a multi-moded fibre, these modes travel at different velocities, causing light pulses to spread as they propagate down the fibre. Bandwidth can be approximated by considering the difference between the fastest and slowest propagating modes Δv . To increase bandwidth we need to minimise Δv , since adjacent light pulses need to be distinct and not overlap.

The most accurate method of evaluating bandwidth is to calculate *every single mode* supported by a fibre, which can take days to weeks of evaluations. Approximations can be made by only considering a few select modes, but even this is problematic. Loss, a very important characteristic, is difficult and computationally expensive to evaluate. Algorithms such as [12] do not behave like a ‘black box’, as one would hope in an automated GA setting, especially with the diversity of designs shown in Figure 5. Improving this algorithm for use in a GA is currently underway.

In multimode fibres, the optimal profile to increase bandwidth is a parabolic refractive index profile (Eqn. 3). This idea can be applied to microstructured fibres since a large array of holes can be designed to approximate a circularly symmetric parabolic profile $\bar{n}(r)$. Thus we can design microstructured fibres with a better bandwidth by evolving designs with an arrangement of holes giving a parabolic profile of the form

$$n(r) = 1.48 - \alpha r^2 \quad (3)$$

where 1.48 is the refractive index of polymer. The average index profile is evaluated by dividing the fibre structure into a series of equally spaced concentric rings. The average refractive index of each ring corresponds to the (total area of polymer)/(total area of air) within that ring. This can be clearly seen in Figure 8, where closely spaced large holes result in a low average refractive index. The objective is to minimise the mean difference between the profile in Eqn. 3 and the average index profile of the candidate microstructured fibre.

5.2 Discussion

The GA was run over a total of 10,000 generations with $N_p = 50$ parents and $N_c = 50$ children using a completely random starting population. The evolution of designs over time is shown in Figure 10.

The initial best design from the first generation is fairly unfit, but rapidly within 50 generations the GA begins to evolve designs which satisfy the requirement of a large flat area of $n = 1.0$ (flat region in Figure 8). As the algorithm proceeds it starts adding more, but smaller and more spaced out holes to the central area of the fibre (generations 200 to 3,500) to satisfy the requirement of a parabolic region. More towards generation 10,000 the GA evolves holes which sit on the outer edges of the core to further reduce the average index in this region. This design was fabricated as a preform (Figure 9A) and another example of a similar design drawn down to an optical fibre is shown in Figure 9B.

The designs with the highest fitness tended to have a symmetry of $n_{\text{symm}} = 14$. This number was not initially obvious in this setting, but relates to the core size requirement in Figure 8, the sizes of holes available in Table 1 and how they can pack together closely enough to achieve an area of low average refractive index at the outer edges of the core. Between different runs the GA never finds the *exact same* optimal microstructured fibre, but always finds designs which are very close in hole arrangement, the total number of holes, and symmetry. For example, three runs with different randomly seeded populations yielded very similar optimal designs with fitness values 0.0334, 0.348, and 0.315 after 2,000 generations. Finally, optimal designs tend to have nearly no unexpressed holes in the genotype. All the evolved high-fitness genotypes consist entirely of holes which are not inhibited by surrounding holes. Thus, the GA evolves designs where the genotype is very ‘efficient’ with respect to the embryogeny.

6. CONCLUSION

Two of the most important results presented in this paper were

- A representation which can generate microstructured fibre designs that automatically satisfy manufacturing constraints,
- A genetic algorithm which can evolve designs which were not fixed in symmetry and the number of holes.

Relatively early on the GA found that a symmetry of $n_{\text{symm}} = 14$ was the most optimal symmetry for the objective function. It also changed the complexity (the total number of holes) of the designs over time as it better achieved the design objective. It first formed the large holes around the edges of the core, refined the parabolic inner-core, then added features to improve the outer-core region. Another interesting feature was the GA quickly evolved to a small core size, as compared to the large cores which tended to emerge in random designs.

These are all features which were not possible to evolve in other microstructured fibre design algorithms which are simply parametrisations of designs with a fixed complexity and symmetry. One further advantage of this representation is the ability to change various attributes of the type of structure we evolve. For example, we may only want to evolve designs where each hole is the same size. This becomes important in manufacture, since any difference in holes sizes tends to lead to slight deformations as the preform is drawn down to an optical fibre. Other possibilities include considering multiple objectives simultaneously. Future work will primarily focus on further developing algorithms which can automatically evaluate important fibre properties such as light loss, allowing us to more accurately evaluate the fitness of these designs with respect to real world applications.

7. ACKNOWLEDGMENTS

The authors would like to acknowledge helpful discussions with members of the Microstructured Polymer Optical Fibre research group at the OFTC, and members of KanGAL, the Kanpur Genetic Algorithms Laboratory at the Indian Institute of Technology, Kanpur, India.

8. REFERENCES

- [1] J C Knight, T A Birks, P St J Russell, and D M Atkin. All-silica single-mode optical fibre with photonic crystal cladding. *Optics Letters*, 21(19):1547–1459, October 1996.
- [2] T A Birks, J C Knight, and P St J Russell. Endlessly single-mode photonic crystal fiber. *Optics Letters*, 22(13):961–963, July 1997.
- [3] D’Arcy Thompson. *On Growth and Form*. Cambridge, 1961.
- [4] M van Eijkelenborg, M Large, A Argyros, J Zagari, S. Manos, N A Issa, I M Bassett, S C Fleming, R C McPhedran, C M de Sterke, and N A P Nicorovici. Microstructured Polymer Optical Fibre. *Optics Express*, 9(7):319–327, September 2001.
- [5] M A van Eijkelenborg, L Poladian, and J Zagari. Optimising holey fibre characteristics. In *CLEO/Pacific Rim: Proceedings of the 4th Pacific Rim Conference on Lasers and Electro-Optics*, volume 1, pages 436–437, 2001.
- [6] S Manos, A Mitchell, M Lech, and L Poladian. Automatic synthesis of microstructured holey optical fibres using numerical optimisation. In *Proceedings of the 27th Australian Conference on Optical Fibre Technology, 8th–11th July, Darling Harbour, Sydney, Australia*, pages 47–49, 2002.
- [7] E Kerrinckx, L Bigot, M Douay, and Y Quiquempois. Photonic crystal fiber design by means of a Genetic Algorithm. *Optics Express*, 12(9):1990–1995, May 2004.
- [8] F Poletti, V Finazzi, T M Monro, Broderick, V Tse, and D J Richardson. Inverse design and fabrication tolerances of ultra-flattened dispersion holey fibers. *Optics Express*, 13(10):3728–3736, May 2005.
- [9] Peter J Bentley and David W Corne, editors. *Creative Evolutionary Systems*. Morgan Kaufmann, 2002.
- [10] Kenneth O. Stanley and Risto Miikkulainen. A Taxonomy for Artificial Embryogeny. *Artif. Life*, 9(2):93–130, 2003.
- [11] Tina Yu and Peter Bentley. Methods to evolve legal phenotypes. In *PPSN V: Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, pages 280–291, London, UK, 1998. Springer-Verlag.
- [12] Nader Issa and Leon Poladian. Vector wave expansion method for leaky modes of microstructured optical fibres. *Journal of Lightwave Technology*, 22(4):1005–1012, 2003.
- [13] Peter J. Bentley and Sanjeev Kumar. Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999), 13–17 July 1999, Orlando, Florida, USA*, pages 35–43, 1999.
- [14] Geoff Barton, Martijn van Eijkelenborg, Geoff Henry, Nader Issa, Karl-Friedrich Klein, Maryanne Large, Steve Manos, Wayne Padden, Wilson Pok, and Leon Poladian. Characteristics of multimode microstructured pof performance. In *Plastic Optical Fibre (POF) Conference, Seattle, USA*, pages 81–84, September 2003.

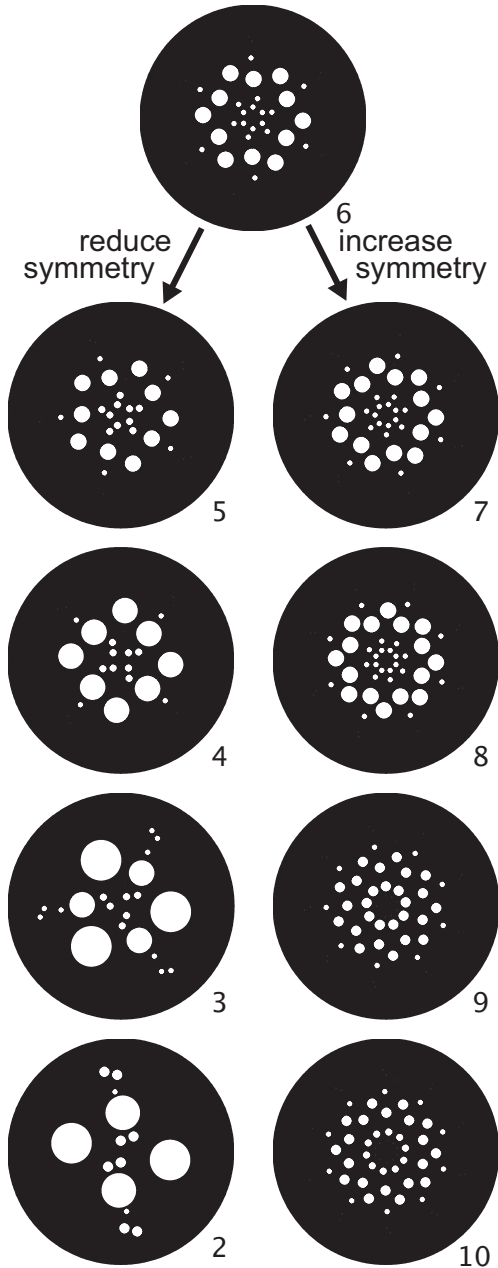


Figure 7: A randomly generated fibre with $n_{\text{symm}} = 6$ is converted to various symmetries. On the left it is reduced from 6 down to 2, and on the right increased from 6 to 10 (note that higher symmetries are possible). The influence of the growth phase is evident here with the reduction of hole sizes as the distance between neighbouring holes decreases as the symmetry is increased. Some holes disappear completely, whereas others maintain the same size through most symmetries.




Figure 8: A randomly generated microstructured fibre (left) and its averaged $n(r)$ profile (right). The solid line indicates the target profile, and the dashed line the average profile of the fibre structure. The objective value is 0.14.




Figure 9: A. Photograph of the drilled preform of the fibre design evolved in figure 10. The preform is 70mm in diameter, made from PMMA (polymethylmethacrylate) polymer and contains 126 holes with $n_{\text{symm}} = 14$. This is then drawn down to an optical fibre, reduced in diameter by a factor of approximately 340. Photograph B shows another evolved design which was found to be easier to draw to fibre, which had accurately maintained its structure during the draw. This fibre is $220\mu\text{m}$ ($1/5\text{mm}$) in diameter, which guides light down the central core region.




Figure 10: Evolution of a microstructured fibre with an average parabolic index profile. The best design of the selected generations are shown.