Citation for published version:

Sullivan, P & McCombie, P 2013, 'Optimisation of tidal power arrays using a genetic algorithm', *Proceedings of the Institution of Civil Engineers - Energy*, vol. 166, no. 1, pp. 19-28. <https://doi.org/10.1680/ener.12.00011>

DOI:

[10.1680/ener.12.00011](https://doi.org/10.1680/ener.12.00011)

Publication date:

2013

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

Publisher Rights

Unspecified

Permission is granted by ICE Publishing to print one copy for personal use. Any other use of these PDF files is subject to reprint fees. See: <http://www.icevirtuallibrary.com/content/journals>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

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

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Optimisation of tidal power arrays using a genetic algorithm

Sullivan and McCombie

ice | proceedings

Optimisation of tidal power arrays using a genetic algorithm

1 Paul Sullivan MEng
Civil Engineer, BP, London, UK

2 Paul McCombie MSc, DIC, PhD, CEng, FICE, FRSA
Senior Lecturer, BRE Centre for Innovative Construction Materials,
Department of Architecture and Civil Engineering, University of Bath, UK

Machines for generating electricity from tidal flows have seen substantial development in recent years, and studies have examined the issues that govern the positioning of devices in relation to each other. This is a complex problem, because the installation of a device can affect the flow both upstream and downstream; also the effects of multiple devices are not simply cumulative, but give rise to complex interactions. The complexity is greatly increased by the geometry of the tidal channel. Determination of an optimum arrangement of devices is a multi-objective problem, which lends itself to solution using genetic algorithms. This approach has been proposed for the design of wind farms; this paper uses a simplified analysis to investigate its potential for the optimisation of tidal power arrays.

Notation

A	swept blade area
C_P	power coefficient
C_T	thrust coefficient
CC	cable cost
CS	support structure cost
CT	turbine cost
D	turbine design life (h)
F	fitness
I	income per kWh produced
N	number of turbines required for an array
O	turbine operational time (%), output
P	turbine power (kWh), probability
R	turbine radius
T	turbine thrust
v	water velocity
ρ	water density

1. Introduction

Tidal power is expected to make an increasingly important contribution to renewable energy, owing to its high degree of predictability. Wind power is essentially unavailable during anticyclones, owing to low wind speeds. Yet the lowest temperatures and greatest heating demand are in anticyclonic weather in winter, whereas the highest temperatures and greatest cooling demands are during anticyclonic weather in summer. Even though good building design can minimise both these demands,

and the use of electricity for heating is very wasteful, increasing use of ground source heat pumps for heating and cooling in addition to the growing population and demand for electricity will place exceptional demands on the power generation infrastructure, leading to soaring emissions of carbon dioxide unless the electricity comes from renewables. Because wind power will not be available when these demands peak, and owing to its unreliable nature, alternative sources of renewable energy such as solar cells will be important. However, these demand high initial investment and make use of rare earth elements.

Although tides do not flow constantly, they have the great advantage over wind power of being almost totally reliable. Exceptional weather can affect the timing and speed of tidal flows, but come what may, the tide must flow as the earth rotates. Tidal flows reach peak velocity approximately four times per day, resulting from two ebbs and two flows, yielding a cyclical power output that is unlikely to coincide with the diurnal cycle of demand for electricity. This would be a problem for a single installation, but the timings of the peak flows vary so considerably around an island the size of Great Britain that a geographically distributed set of installations would between them feed a smoothed supply of power into a national grid, which would peak just twice per month at spring tides, rather than twice per day. Although this would only produce a baseline output corresponding to neap tides, the spring tide peaks are so reliable and predictable that the power supply system could be

managed to use this power efficiently (Royal Academy of Engineering, Institution of Engineering and Technology and Institution of Civil Engineers, 2012).

Although prototype tidal power generators have been isolated devices, major arrays are now being planned and built. Tidal turbines face similar issues to wind turbines in the way they remove energy from the surrounding fluid, reducing the energy available to another turbine. This interaction means that their relative positioning is governed by more than just standardised separation distances. Because the fastest tidal flows are found in narrow channels, which are often also shipping routes, the position of turbines requires careful optimisation to achieve maximum efficiency in the space available. However, the cost of installation and connection can be critical, and dependent upon location, so this too must be accounted for in the optimisation.

The problem of placing tidal turbines may be characterised as a need to make decisions about plan location, resulting in two degrees of freedom for each individual turbine, so the search space has a dimensionality of $2N$, where N is the number of turbines required for the array. The fitness of a solution to define the arrangement of N turbines must consider both the power output and the cost of installation. Additional requirements may include planning restrictions and access for operations and maintenance; these may be included within the algorithm or assessed objectively by an engineer following the optimisation presenting a number of similar solutions. It might seem that the optimum layout of turbines would be easily determined, but this might not be so, even in a large, open body of water with a uniform depth. In practice, locations favourable for tidal power are characterised by fast tidal flows resulting from confinement by coastlines and seabed, so initial conditions are anything but uniform; then the effect of the turbines extracting energy from the flow inevitably alters the flow.

2. Genetic algorithms

Genetic algorithms (GAs) are a search and optimisation method used to find solutions to complex problems that are sometimes impossible to solve manually. Other search methods exist, such as random searches and hill climbing. A random search would need to check a large proportion of all potential solutions to the problem, depending upon the complexity. A hill-climbing algorithm would attempt to find the direction of change within the search space that resulted in greatest improvement, and then continue to search in that direction. It is helpful to visualise the search space as a two-dimensional surface in a three-dimensional space, with the third dimension being the quality or 'fitness' of the solution for that plan location. If the high points are localised peaks, then a random search may never find a peak, whereas the hill-climbing algorithm could end up at the top of a relatively minor peak, which is nowhere

near the highest point either in plan or in fitness. This representation gives insight into the nature of the optimisation problem, but it is difficult to visualise the extent to which the complexity increases as the number of dimensions of the search space increases beyond two or three. GAs address these problems by working from multiple points in the search space in parallel. The consequent ability to find the global optimum rather than just a local optimum is their principal advantage compared with other search methods.

The concept of a GA was created by John Holland in 1975. Based on Darwinian principles, they enact an imitation of evolution and natural selection through a number of generations, until the optimum solution is found. GAs have been applied in numerous fields, mainly in engineering and science. The optimum solution may not even represent something to be made. GAs can be used to find any solution that can be defined and tested, but is difficult to find, such as the critical surface in slope stability analysis (McCombie and Wilkinson, 2002; Zolfaghari *et al.*, 2005). Although generic GA software is widely available, problem-specific algorithms are commonly used because they allow the encoding of the problem to be tailored closely to its specific features. This is a critical factor in the efficiency of the algorithm (Goldberg, 1989).

Although variations exist in GAs, they all begin with the initialisation of a set or 'population' of proposed solutions, which is then evolved through processes of reproduction, cross-over and mutation (Goldberg, 1989). Each solution is defined by a chromosome, which contains genes, which in turn are encoded parameters from which the solution is generated (Goldberg, 1989). The set of possible solutions may then be altered by altering the chromosomes from which it is generated. Each gene within a chromosome may represent an actual design value, or a rule used to generate a design value. A fitness value is calculated for the solution generated from each chromosome, which is an expression of how well that solution meets the defined aims of the design. The fitness may be based on a single value, such as weight of material in the case of a structure, or weighted combinations of values and logical tests.

A new generation of the population is produced by applying concepts from evolution. First, 'survival of the fittest' discards the weakest of the population, and the strongest are used to produce the next generation; the chances of each member being used are in proportion to its strength. A small number of the very best individuals (an elite) may pass unchanged to the next generation, whereas the remainder go through a breeding process in which pairs of chromosomes exchange some of their genes – this is crossover. This results in a new population that contains aspects of the better members of the previous generation; these may be combined in favourable or unfavourable ways. Over successive generations, genes that are more useful will come to

dominate the population. Depending on the complexity of the problem, and the ‘roughness’ of the search space, a large population may be needed to allow beneficial combinations of genes to arise. Similarly to the evolution of natural species, innovative solutions may only appear when a degree of random mutation is introduced, but too much mutation will interfere with the evolutionary process.

The most important consequence of this evolutionary approach to finding solutions is its ability to find multiple solutions, not just a single solution. If a number of significantly different solutions show promise, the process will persist with the evolution of those solutions in parallel, rather than attempt to converge on a ‘single right answer’, something that rarely exists in engineering. It is nevertheless necessary that the assessment of the fitness of a solution is broad and realistic rather than narrow and simplistic.

3. Wind farm optimisation

A comparison can be made between the optimisation of tidal arrays and wind farms. The main optimisation needed for both is related to the wake behaviour, site constraints, economics and environmental issues. The shape of the site is more critical for tidal arrays, because boundary conditions can drastically affect the flow. Mosetti *et al.* (1994) first attempted to apply a GA to optimise the positioning of turbines within a wind farm, by considering whether or not a turbine would be placed in a particular cell within a grid. Grady *et al.* (2005) subsequently demonstrated that this simple approach was in fact feasible, and Wan *et al.* (2009) used a real-coded GA to refine turbine placement within the cells, instead of using a finer grid. This showed negligible clear benefits, but Sedat *et al.* (2009) modified the methodology by omitting some cells to represent a real situation. This sequence of papers demonstrated the problem of trying to build on what had been done before, because it was in fact necessary to experiment with different approaches to defining the problem to make real progress. Turbine positions could be defined much more efficiently by integer values specifying their locations on a fine grid, rather than by applying a Boolean on/off to determine the placement of a turbine in every cell. Achieving an efficient algorithm is heavily dependent upon using an efficient coding to generate each instance, or design layout, for this problem (Goldberg, 1989). Although the work on wind turbine arrays considered power output, efficiency and cost, in various combinations, it did not address the more challenging issue of large, irregular sites, for which the optimum solution is not reasonably obvious to begin with.

4. Tidal turbines

Tidal turbines obstruct the water flow and present a semi-permeable obstruction that removes energy from the flow (Myers and Bahaj, 2007). They create a negative back effect on the current (Garrett and Cummins, 2008) as with wind power, and the rate of wake propagation and free stream

mixing relates to the interaction between the turbines (Blunden and Bahaj, 2007).

An indicative minimum spacing between turbines may be obtained by considering the wake behind the turbine for the calculation of the kinetic energy reduction (Bryden *et al.*, 2007; Myers and Bahaj, 2007). For the wind farm optimisations outlined above, the modelling method was to sum the energy deficits caused by the wake effects (Mosetti *et al.*, 1994), then calculate the resulting power. For tidal arrays, the principle is similar in that the wake will reduce the velocity, kinetic energy and hence the power output of the affected turbines. The turbine is extracting power from the flow, and the more turbines there are, and the narrower the channel, the more the turbines will impede the overall flow velocity. They will hence reduce the velocity of flow they are working in, as well as reducing the downstream velocity.

Myers and Bahaj (2010) explain that tidal wakes are commonly much longer than for wind turbines, because the wake is constrained by the seabed and water surface. Consequently, the extent to which interactions between turbines affect optimal arrangements is much greater too. Seabed roughness is also known to have an effect on the turbulence. As with wind power, tidal power output is proportional to the cube of the velocity (Myers and Bahaj, 2010) so the wake position should be carefully calculated, as it will have a large effect on the optimum positioning. For wind power the maximum available energy is often calculated using the Betz limit (Betz, 1966), but this is inappropriate for tidal power because of the constraints on the wake diffusion. Garrett and Cummins (2008) consider the actual limit for tidal power, showing that maximum power can be obtained from a given flow using surprisingly few turbines if they are carefully placed. This may be considerably less than the Betz limit, but the consequence is that the smaller reduction in flow velocity is less likely to produce significant changes in the marine environment.

5. Optimisation objective

For the purposes of this study, a simplified model is used to give an indication of the true complexity of the optimisation that would be required for a real project. Although differing objectives have been explored, such as maximum power output and investment efficiency, the objective used will be the maximum profit over the lifetime of the project. This will ignore economic factors such as the cost of capital and varying energy rates, although they may have a significant effect on the optimisation of a real project.

The fundamental data that must be generated from a chromosome defining a single tidal turbine are the coordinates of its plan location. A number of factors in the turbine placement may have a significant influence on the overall cost, and hence should affect the optimisation – for example, the nature of the sea bed and depth of water will affect installation cost and

type of support structure required. A substantial cost that will have a direct influence on a turbine layout is the cabling. For the present study, this is assessed using the simple proxy total cabling length. This may be determined from the data about the placement of a turbine if a gene is defined to indicate the adjacent turbine or power hub a turbine is connected to.

The optimisation function (i.e. fitness) of the algorithm is expressed below.

Fitness = lifetime array profit

$$= \sum_{i=1}^n [(P_i \times O_i \times D_i \times I) - (CT_i + CS_i + CC_i)]$$

where P is the turbine power (kWh), O is the turbine operational time (%), D is the turbine design life (h), I is the income per kWh produced, CT is the turbine cost, CS is the support structure cost, CC is the cable cost (for associated turbine), and n is the number of turbines in the array.

A different turbine or flow condition will require a different support structure, based on the axial thrust and hence moment resistance required. This may be significant where the peak obtainable power from a turbine is reached – for example, a flow rate of 3 m/s, resulting in a decrease in fitness if placed in higher flows, owing to the support structure costs. It may be noted that a full assessment would require the summation of energy output for both incoming and outgoing tides, but for the purposes of the demonstration presented here, only one direction of flow is considered.

5.1 Power

The power output of a turbine is related to the swept blade area (A), and the cube of the water velocity (v). Other factors included are the performance coefficient (C_P), and the density of the fluid (ρ) where $\rho = 1025 \text{ kg/m}^3$ for seawater.

$$P = \frac{1}{2} C_P \rho A v^3$$

5.2 Axial thrust

The axial force acting on the turbine is similar to the power equation, but proportional to the square of the velocity and a thrust coefficient (C_T).

$$T_{\max} = \frac{1}{2} C_T \rho A v^2$$

The moment acting on the support structure (from the turbine alone) is therefore $T_{\max} \times \text{hub height}$. This value allows the structural section to be chosen from a stored list.

The C_P and C_T values will be set at 0.45 and 0.90, respectively, for all turbine types, with the minimum operating water velocity being 0.7 m/s, and no increase in power being obtainable over

Item	Values	Cost: units
Turbine	Diameter: 8–24 m	1000–1850
Support structure	Moment capacity: 10–50 MNm	200–500
Cabling	Per metre	1.2

Table 1. Basis for costing

3 m/s. The values in Table 1 are used to define the allowable parameters and their associated cost.

6. Development of computer program

The work presented here explores the feasibility of using a GA to optimise the placement of tidal turbines in an array in a defined body of water. To achieve this requires a method for specifying turbine locations that can be encoded by a chromosome, a GA to optimise a population of chromosomes, and a means of generating a fitness value for an array of turbines, which must involve a hydraulic model – although that used for this study is very much a simplification.

6.1 Genetic algorithm

The GA works by following a number of steps, iterated over a number of generations. These have been outlined in the literature cited above, and the process is listed below.

- Initialise population of individuals through random generation.
- Perform fitness function on all individuals (perform analysis of fluid model and costs).
- Select parents for mating, based on fitness (greater profitability results in higher probability of selection for crossover).
- Produce new individuals using the crossover process.
- Perform mutation on newly created individuals.
- Repeat for total number of generations.

Methods such as selective pairing, which chooses the best two parents of four (Kaveh and Bondarabady, 2003), have been reviewed, but not used, in order to limit the risk of reducing diversity in the complex problem. Elitism is used, however, to allow a small number of the very best individuals through to the next generation without alteration, in order to prevent the potential loss of strong genetic code (Zitzler *et al.*, 2000).

The probability of selection increases proportionally with increasing fitness using the following ratio

$$\text{Probability of selection} = \frac{\text{individual fitness}}{\text{sum of population's fitness}}$$

$$P_i = \frac{F_i}{\sum_{j=1}^n F_j}$$

R	R	θ	θ	N	N	x	x	y	y	N	N	x	x	y	y	N	N	x	x	y	y	N	N	x	x	y	y
Turbines				Member 1						Member 2						Member 3						Member 4					

Figure 1. Representation of chromosome used

6.2 Turbine siting

The initial development of the program used a simple wake deficit model to test for fitness, in a similar way to Mosetti *et al.* (1994). This also followed the established and obvious practice of defining turbine positions individually, but the resulting optimisations did not stabilise easily if the scale of the problems examined was increased. It was found that successful local arrangements of turbines were too easily disrupted, so instead a system was developed that allowed groups to be defined, as well as individual turbines. This led to a chromosome that contains a series of definitions of groups, in each of which the location of turbines is defined relative to the group's origin, as shown in Figure 1. This allows the movement of a single turbine or the movement of a whole group. A solution then defines which groups are placed where.

The same group definition can be used any number of times in the solution, and it can be used within another group. Because a group can contain groups, each of which can contain groups containing groups, the members of which could include individual turbines or other groups, a sequence of arrays of turbines can be generated that can become infinitely recursive. The recursion would automatically be stopped, however, when a turbine location is no longer in the sea, or otherwise falls outside a defined area. A limit can, in any event, be placed on the level of recursion. This is illustrated in Figure 2, in which the 'Extract group' routine can be seen to be recursive.

Figure 1 shows the genes used within the chromosome. It can be seen that the turbines are stored as a reference to the

Extract solution

For each member in solution:

Extract x and y coordinates

If member is group then: *Extract group*, passing through coordinates

If member is turbine then: *Extract turbine*, passing through coordinates

Extract group

For each member in solution:

Extract deltax and deltay coordinates

Apply deltax, deltay to received coordinates from parent routine (cumulative)

If member is group then: *Extract group*, passing through coordinates

If member is turbine then: *Extract turbine*, passing through coordinates

Extract turbine

If turbine is within site and does not overlap another, add turbine to turbine list with type, coordinates, radius, angle

Figure 2. Generation of turbines from the chromosome, by way of groups

stored list with a parameter for their placement angle. Each member then makes a reference (N) to either a turbine or a group and contains corresponding x and y coordinates.

6.3 Model

The simple wake deficit model can only work for uniform flow in one direction, whereas optimal locations for tidal turbines are where velocities are increased owing to the confining effect of headlands. Because the turbines themselves effectively confine the flow further, the fitness of a solution can only be assessed by taking into account the flow of water in a more sophisticated way. This work comprises an investigation of the usefulness of a GA for optimising a tidal array, rather than the generation of an actual solution for a real site, so it was decided that a simplistic two-dimensional model of water flow would suffice for this initial investigation. Greater accuracy and the added complexity of modelling flow in three dimensions would considerably increase the solution time. This two-dimensional model is not intended to provide accurate flow regimes, but does allow the introduction of some complexity of flow conditions attributable to headlands and the downstream effect of turbine placement. For application to a real problem this model would require modification to take into account the upstream, regional flow and three-dimensional effects – this would also allow the placement depth to be optimised.

The site has been split into a number of cells. The simplified model is based on the total volume of water entering each cell from each of the four edges over a given time period. This is assumed to exit the cell equally in all four directions. Where a site boundary exists, the water is reflected back into the cell and forced to exit through a free edge, increasing the volume and thus flow in that particular direction. Referring to Figure 3, the input through each side of the cell is the output of the adjacent cell. The total input (I_T) is therefore the sum of the inputs

$$I_T = \sum_{i=1}^4 I_i; \quad O_i = \frac{I_T}{4}$$

where I is input and O is output.

Figure 3. Flow model

This process is repeated for each cell and iterated until the flow stabilises. This simple iterative model provides sufficient values for the demonstration of the suitability of the GA.

6.4 Program

The overall structure of the program is shown in Figure 4, and the user interface in Figure 5. The method for setting up the program is as follows.

- User enters site details – defining points, turbine position accuracy, cell size for accuracy of flow calculation.
- Cost details are input – turbines, support structures, cabling.
- GA and group options are edited if needed.
- User starts the generation process – automatic site and cell creation, evolution.
- Program outputs the site and cells of the best solution.

Once the generation process has started, the program displays the best solution found in each iteration (Figure 6). This allows the changes to be seen, and assists in understanding how and why the solution is changing. The user can click on a particular turbine or cell to see information such as water velocity, direction, turbine power output, size and position. A number of resizable views are available, and certain layers of detail can be added or removed, such as the cells, water flow, cabling and a fitness graph, which charts the evolution of the fitness of the solution.

Options allow the optimisation method to be changed to see the maximum power obtainable, maximum efficiency or maximum profit. The program can be paused or restarted should the user change certain options or want to look at the current solution.

An example of a near-optimal solution is shown in Figure 7. The darkness of the shading indicates water velocity, the direction of which is shown by the short lines. It should be noted that in this case a mid-blue colour represents 2 m/s, and white 1.2 m/s. A headland is shown on the south side of the area, around which the velocity increases. A single turbine in the SW corner (T0) takes the energy from the beginning of this increase, and the velocity does not increase significantly again until nearly half-way across the view, reaching its maximum as it passes the headland. No turbine is placed at this point, however, because there is an upper velocity limit at which the turbines can operate, and a higher velocity calls for a more expensive support structure with no gain in power and significant reduction in output from downstream turbines. Instead, a turbine (T1) is placed where the flow starts to spread out again (attributable to mixing of the increased velocity with adjacent cells). The next turbine (T2) is placed a little behind and to one side of this, where it can pick up the faster flows that have gone past the turbine immediately east of the headland. Being stepped

Figure 4. Program overview (note: Rand indicates random number)

back allows the increased flow to affect a larger number of turbines. The near-optimal cable routes are also shown, with cable length minimised. Connection of turbines using cabling has been automatically determined during the process because

Figure 5. Program screenshot. The main area shows a plan of the site, with the dark lines showing turbine positions, the short lines showing flow direction, and the darkness of the shading indicating velocity. The graph below the plan view shows the evolution of the fitness

an unconnected turbine would yield no power output and is thus a weaker solution. The position of the remaining turbines is not yet quite optimal, as bringing the northerly row (T7–T10) further west in line with the turbines to the south would reduce the cable length. This demonstrates a feature of the GA approach. A near-optimal solution can be found from an unmanageably large number of possibilities quite quickly, but it is often more effective to make a final adjustment with the use of a local optimisation technique or input of a knowledgeable engineer. Repeating the same scenario a number of times yields very similar results, up until the point discussed in Figure 7. This demonstrates the repeatability in outcome but reinforces the limitation at the final stage of optimisation. By again considering the search space as a three-dimensional landscape, the more specific and unique the solution is (i.e. the sharper the peak), the less reliable or repeatable the GA, and other search methods, will become.

It should be noted that this layout is optimised for flow in one direction only. In a real tidal situation the flow would of course occur in alternating directions, not necessarily with similar velocities. It can be seen that this headland is not symmetrical, and the flow patterns would certainly be asymmetrical. To determine an overall optimum was beyond the scope of this work, because it would require considerably longer to calculate a fitness value for each case. However, in such a situation it can be seen that the optimum arrangement may be far from obvious, and the use of an evolutionary technique as described may be the most practical method to find an optimum.

Figure 6. Optimisation of turbines in eastwards flow past a headland: (a) initial, random solution ($i = 0$); (b) turbines start aligning to optimum flow positions ($i = 87$)

Figure 8 shows a graph of the associated array costs throughout the evolutionary process and reinforces what is shown in Figure 7. It can be seen that with increasing turbine numbers (reflected by jumps in cost) comes an increased power output

Figure 7. Optimised arrangement of turbines in eastwards flow past a headland ($i = 2000$)

(reflected by electricity income). The addition of a new turbine, or the movement of an existing one into faster flows, evidently has priority over the increased cabling costs. When the addition or move has been made, the cables can be seen to readjust; this is reflected by the reduction in cable costs soon after a turbine movement. This prioritisation of objectives is a consequence of the weightings given in the fitness calculation, where weighting is simply the defined cost.

These examples were run on a single-core 2.4 GHz computer and took 22 min to complete 2000 iterations. Approximately 85% of the computational time is for the iterative calculation of the flow regime for each member of the population (100) for each generation (2000). This model is therefore run approximately 200 000 times and is the governing factor in the computational resources required. Clearly, determining the smallest population size and number of generations required to find the optimum will drastically improve the speed of the GA. It should, however, be noted that once the GA is running from generation 0 it can be left unattended until completion for the user to review the results.

The limitation here is the number of analysis calls required to assess each individual. This is an issue that would be apparent for engineering problems. The GA could therefore be combined with hill climbing to speed up the process at the end (where the GA seems least efficient).

Figure 8. Evolution of optimisation parameters through successive generations

One method to allow engineering input towards the end would be to allow the user to pause and edit certain aspects of the solution, as proposed by Ceranic and Fryer (2000); this may, however, bias the results incorrectly if the optimum solution is substantially different from that expected by the engineer.

7. Conclusions

The work presented here has demonstrated the ability of a GA to find successful arrangements of tidal turbines, taking into account a number of the complicating factors that influence the cost and performance of an array. The long wake behind a tidal turbine and upstream effect due to the confinement of the flows where tidal power is most efficiently exploited make optimisation of tidal arrays intrinsically more difficult than optimisation of wind turbine arrays.

Taking into account the actual shape of the sea bed, and a more rigorous analysis of water flows and their variations with time, will increase considerably the time required to assess the fitness of each potential solution. This would be the case whether or not an evolutionary algorithm is being used to find the solution, but the algorithm has the advantage that it can work automatically, night and day, to assess a population of solutions, evolving those solutions towards an optimum, without the need for human supervision.

The automation of the process, and the non-sequential nature of the algorithm, allowing multiple different solutions to evolve alongside each other, permits the emergence and exploration of subtle interplays between different factors, which may not have been imagined by a design engineer. The design of a tidal array is a problem that increases in complexity very rapidly when the representation of the problem is made more realistic. The cabling costing, investigated as an example of an additional input into the assessment of a solution, required the addition of

a single gene and a simple calculation. This required no update of the algorithm and represents the ease of adding further objectives once the algorithm has been defined. It is thus well suited to the particular strengths of a GA for finding an optimised solution, and this work has demonstrated that this potential can be realised.

REFERENCES

- Betz A (1966) *Introduction to the Theory of Flow Machines*, translated by Randall DG. Pergamon, Oxford, UK.
- Blunden L and Bahaj A (2007) Tidal energy resource assessment for tidal stream generators. *Proceedings of the Institution of Mechanical Engineers, Part A – Journal of Power and Energy* **221**(2): 137–146.
- Bryden I, Couch S, Owen A and Melville G (2007) Tidal energy resource assessment for tidal stream generators. *Proceedings of the Institution of Mechanical Engineers, Part A – Journal of Power and Energy* **221**(2): 125–135.
- Ceranic B and Fryer C (2000) An application of a real number micro-genetic algorithm to the shape optimisation of in-plane stressed plate elements. In *Computational Engineering Using Metaphors From Nature* (Topping BHV (ed.)). Civil-Comp Press, Edinburgh, UK.
- Garrett C and Cummins P (2008) Limits to tidal current power. *Renewable Energy* **33**(11): 2485–2490.
- Goldberg DE (1989) *Genetic Algorithms in Search, Optimisation, and Machine Learning*. Addison-Wesley.
- Grady SA, Hussaini MY and Abdullah MM (2005) Placement of wind turbines using genetic algorithms. *Renewable Energy* **30**(2): 259–270.
- Kaveh A and Bondarabady HAR (2003) A hybrid graph-genetic method for domain decomposition. *Finite Elements in Analysis and Design* **39**(13): 10.
- McCombie P and Wilkinson P (2002) The use of the simple genetic algorithm in finding the critical factor of safety in

-
- slope stability analysis. *Computers and Geotechnics* **29**(8): 699–714.
- Mosetti G, Poloni C and Diviacco B (1994) Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm. *Journal of Wind Engineering and Industrial Aerodynamics* **51**(1): 105–116.
- Myers L and Bahaj AS (2007) Wake studies of a 1/30th scale horizontal axis marine current turbine. *Ocean Engineering* **34**(6): 758–762.
- Myers LE and Bahaj AS (2010) Experimental analysis of the flow field around horizontal axis tidal turbines by use of scale mesh disk rotor simulators. *Ocean Engineering* **37**(2–3): 218–227.
- Royal Academy of Engineering, Institution of Engineering and Technology and Institution of Civil Engineers (2012) *The Future of Marine Renewables in the UK*. Memorandum submitted by Engineering the Future Alliance (MAR 37), response from the Engineering the Future alliance to the Energy and Climate Change Committee, see <http://www.publications.parliament.uk/pa/cm201012/cmselect/cmenergy/writev/marine/mar37.htm> (accessed 20/11/2012).
- Sedat S, Ozgu T, Murat T and Unal C. (2009) Optimal positioning of wind turbines on Gökçeada using multi-objective genetic algorithm. *Wind Energy* **13**(4): 297–306.
- Wan C, Wang J, Yang G and Zhang X (2009) Optimal siting of wind turbines using real-coded genetic algorithms. *Proceedings of EWEC 2009, Marseille, France*.
- Zitzler E, Deb K and Thiele L (2000) Comparison of multiobjective evolutionary algorithms: empirical results. *Evolutionary Computation* **8**(2): 173–195.
- Zolfaghari AR, Heath AC and McCombie PF (2005) Simple genetic algorithm search for critical non-circular failure surface in slope stability analysis. *Computers and Geotechnics* **32**(3): 139–152.

WHAT DO YOU THINK?

To discuss this paper, please email up to 500 words to the editor at journals@ice.org.uk. Your contribution will be forwarded to the author(s) for a reply and, if considered appropriate by the editorial panel, will be published as a discussion in a future issue of the journal.

Proceedings journals rely entirely on contributions sent in by civil engineering professionals, academics and students. Papers should be 2000–5000 words long (briefing papers should be 1000–2000 words long), with adequate illustrations and references. You can submit your paper online via www.icevirtuallibrary.com/content/journals, where you will also find detailed author guidelines.