4 Entropy-Based Method and its Application to Field and Laboratory Lagoons

Entropy-based methods are an alternative to conventional modelling, that have been used in the past to predict equilibrium river profiles and planforms (Leopold and Langbein, 1962; Rodríguez-Iturbe et al., 1992; Molnár and Ramírez, 2002; Huang et al., 2004). Entropy principles suggest that an open system, where matter and energy can enter and exit the system will attain minimum entropy production at equilibrium (Prigogine, 1967). For the calculation of equilibrium morphologies it is more beneficial to consider the energy of the morphology rather than its entropy value. The entropy of a system can be related to energy production. If the energy dissipation of a system is minimised, the energy available to increase the entropy of the system is limited and entropy production reaches a minimum (Yang, 1994; Phillips, 2002). This then provides the rationale for the method presented in this Chapter.

In the previous Chapter, it was shown that for the simplified lagoon system with unidirectional flow, energy dissipation was indeed at a minimum when a stable equilibrium morphology formed. This was the result expected according to the theoretical background discussed in Chapter Two, Section 2.2.4. Hergarten and Neugebauer (2001) obtained a similar result in their study of drainage networks. One aim of the research presented in this thesis was to develop a method based on entropy principles that could be applied to determine the equilibrium morphology of a sandy lagoon system. The developed method is discussed in this Chapter. It is not time-step-based, but rather used optimisation to predict stable equilibrium morphologies directly, based on entropy principles.

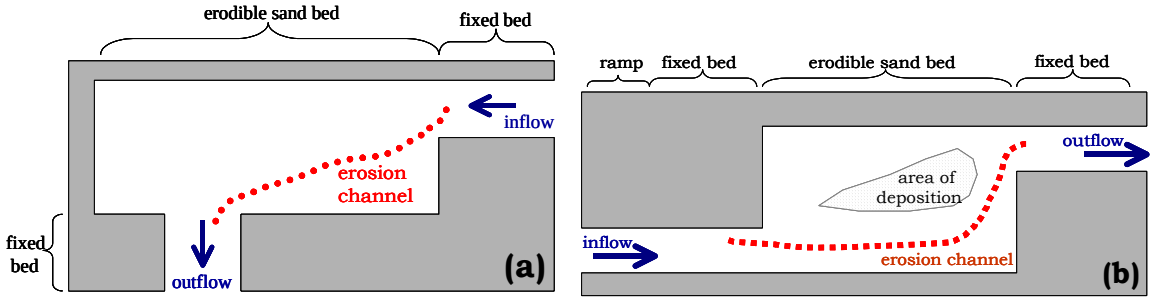

Figure 4.1 Model setup for (a) field sized lagoon and (b) the laboratory scale case study.

This Chapter discusses the development of an entropy-based model, outlining the method used and the application of the model to a lagoon type situation (as

shown in Figure 4.1 (a)) with either unidirectional flow or flow reversal, as well as a small laboratory scale example (as shown in Figure 4.1 (b)).

4.1 Method Description

In this Section the general method applied in the case of a unidirectional lagoon is outlined. Further modifications are made to this method when it is applied to different situations, but these modifications are outlined in the latter sections of this Chapter.

This research extends work undertaken in river morphology, where local and global energy dissipation is minimised to determine the dendritic river patterns that form in a catchment, by Molnár and Ramírez (2002) and Rodríguez-Iturbe et al. (1992) amongst others. Like these researchers, an objective function has been formulated where energy dissipation is minimised both globally and locally. This investigation differs from previous research as it examines the detailed morphology, rather than a large-scale planform pattern prediction. It fills a gap in current modelling methods, extending previous river formulations and applying them to coastal, complex environments, as well as making detailed comparisons between the new model predictions, conventional models and laboratory findings.

In the case of a unidirectional lagoon, with water entering a wide sandy area through a narrow fixed bed channel and exiting through a similar channel, an objective function has been developed as discussed below.

This research considers energy that is dissipated globally by the water as it passes from the narrow entrance to the narrow exit of the lagoon, and locally by applying penalties. A penalty is included in the objective function calculations to dissuade the optimisation search routine from choosing morphologies where the bed level of an area has been reduced to a depth with a velocity less than the threshold velocity. Similarly a penalty is applied to avoid morphologies where sediment must be removed, as deposited sand has reached a height where the velocity moving over the deposited mound is greater than the threshold velocity and sediment transport will occur. This is similar in some respects to the critical velocity rule of the cellular automata models used in Chapter Three. However, in the entropy-based model there is no actual movement of sediment, so morphological snapshots of the potential equilibrium bedforms must be penalised.

In order to derive equilibrium morphologies without stepping through time, randomly generated snapshot morphologies were assessed through an objective function based on the product of global energy calculations and penalties of morphologies with velocities that were above or below critical when areas of deposition or erosion were considered respectively. This objective function was then utilised in a number of global optimisation routines, used to minimise the objective function value. The general methodology is displayed on the flow chart in Figure 4.2.

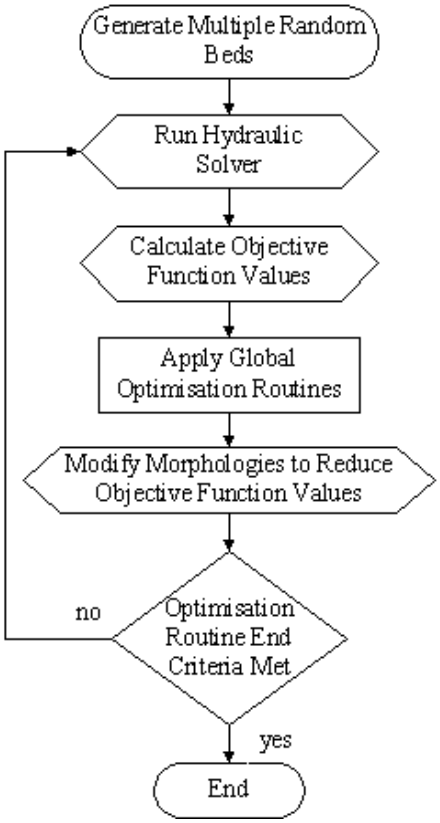

Figure 4.2 General methodology of optimisation routines.

Firstly, in order to calculate the energy dissipation due to the water movement, (the global energy dissipation of the present system), the average specific energy at the entrance to the lagoon is determined and compared to the average specific energy of the water leaving the lagoon (see Eqs. 4.1 to 4.3).

Velocities and water elevations are calculated using a rectangular staggered grid finite difference hydrodynamic solver. The same solver used in the self-organisation models discussed in the previous Chapter and explained in detail in Chapter Three, Section 3.3. The positions of the energy calculation points are shown in Figure 4.3.

$$E_e = \frac{u^2 + v^2}{2g} + \eta \qquad (4.1)$$

$$\Delta E = \left(E_{entrance} - E_{exit}\right) \qquad (4.2)$$

$$E_s = \psi \Delta E \qquad (4.3)$$

where E_e is the average energy in the entrance channel or exit channel in metres;

u and v are the velocity vector components in the x and y direction respectively;

η is the surface elevation with respect to a datum;

ΔE is the energy difference between the entrance and exit channels in metres;

ψ is a scalar factor; and,

E_s is the global energy contribution to the objective function.

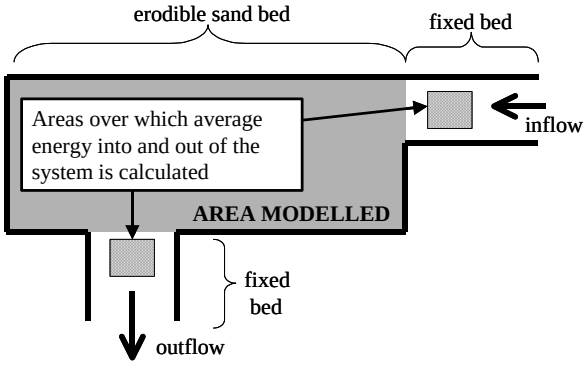

Figure 4.3 Sketch of areas where global energy dissipation of the system are calculated.

A scaling factor was used to allow similar magnitudes of all components of the objective function value calculations. This is purely to aid in the efficiency of the global search methods to identify near-global minima. The optimisation methods would still be able to find a near global minimum without the use of scaling factors, but they would first try to minimise the larger part of the objective function, followed by the smaller part. This was observed to be the case during sensitivity analyses. Making the components have similar magnitudes allows the optimisation module to minimise the two components simultaneously, thereby reducing the computation time. In order to determine suitable scaling factors, sensitivity analyses were undertaken (as discussed in Section 4.3.4) considering objective function value size if each component was optimised independently.

Secondly, a penalty was applied related to local points on the snapshot of the morphology that have been eroded to form overly deep holes or deposited to form unrealistically high mounds in situations where the velocity is less than or

greater than that required for sediment transport to occur. This is calculated by a comparison of the present velocities to threshold velocity values based on the van Rijn (1984) equation (Eq. 3.1, Section 3.1). It is similar in application to that discussed by Wright et al. (1973), where they utilised a threshold velocity to determine whether the tidal delta had reached equilibrium and sediment transport was minimised.

In the entropy-based method developed in this thesis, areas where erosion has occurred relative to an average baseline were firstly considered. If the velocity of the eroded area is less than critical, the bed elevation is lower than the threshold for sediment transport. This means that more sediment has been eroded from the bed than is necessary under the average conditions considered. In this situation a penalty based on the velocity of the morphological snapshot is applied (P_e) (see Eq. 4.4).

$$P_e = \beta\left(\sum_{\text{all } i,j} u_{crit} - \overrightarrow{u_{ij}}\right) \qquad \text{for } \overrightarrow{u_{ij}} < u_{crit} \text{ and } \left(z_{ij} - z_0\right) < 0 \qquad (4.4)$$

$$P_d = \beta\left(\sum_{\text{all } i,j} \overrightarrow{u_{ij}} - u_{crit}\right) \qquad \text{for } \overrightarrow{u_{ij}} > u_{crit} \text{ and } \left(z_{ij} - z_0\right) > 0 \qquad (4.5)$$

$$OF = E_s \times P_d \times P_e \qquad (4.6)$$

Where P_e and P_d are excess velocity sediment eroding and depositing penalties respectively. They must be greater than 1, else the total value is taken as 1, to prevent the penalty falsely reducing the objective function;

z_{ij} and z_0 are the predicted and baseline elevations respectively; and,

OF is the overall objective function value.

Next areas where deposition has occurred relative to the baseline were considered. If the snapshot velocity is greater than critical, the bed elevation produces a velocity higher than the threshold for sediment transport. This means that the velocity can still erode the bed and the system is not at equilibrium. The bed still needs to be eroded to move the system towards a favourable morphology that minimises energy dissipation. The penalty (P_d) for this additional erosion requirement (or excess deposition) is based on the snapshot velocity and is included in the objective function as the sum of the excess velocities above the critical velocity (see Eq. 4.5).

A scaling factor (β) is applied to bring all values into similar magnitudes to facilitate solution efficiency. An overall objective function value is calculated by the product of the local penalties and the global energy value (see Eq. 4.6). The factors ψ and β are determined based on knowledge of the order of magnitude of

the components in the objective function value calculations. They are also related to the number of significant figures required for accuracy and the minimum values obtained if each component is optimised separately. In the unidirectional lagoon case study they have been assigned values of 1000.0 and 10.0 respectively based on a sensitivity analysis.

Though deposition is predicted to some extent, the objective function has been designed with a bias towards predicting erosive environments. It was expected that the model would be able to determine channel formation. However, the direction of deposition is not taken into account as there is no time considered in the model, so any deposition is only taken account of if it is big enough to have an effect on the energy dissipation. There may be some accumulation of sediment in areas where the global energy dissipation is not increased or the velocity does not increase above a critical value for the initiation of sediment transport, but this accumulation is not determined directly.

In order to find equilibrium lagoon morphologies global search methods are used, namely a combined genetic algorithm (GA) and simulated annealing (SA) algorithm, to minimise the overall objective function value. These methods were chosen after an initial sensitivity and comparison study was undertaken. This study compared the predictive capabilities of the two methods in the determination of flow fields around a plate for an example where a finite difference model is generally employed. Their description, comparison and application to the flow field problem are outlined in the following Section.

4.2 Global Search Method Characterisation

A study was undertaken to determine a suitable global search method, able to deal with a solution space involving a large number of decision variables. In the situations where it was required, such as a unidirectional lagoon, the number of decision variables was upwards of 265. A test case was devised to test different methods, in which flow fields were solved around a plate in a channel. This test case is shown in Figure 4.4.

In this type of problem, computational fluid dynamic models are generally used. These models are effective, and provide a comparison for the global search methods. The optimisation routines were used to minimise the residuals of the finite difference equations. Three different search methods were compared. Firstly a binary GA (BGA) and a real GA (RGA) were compared to each other. Secondly an SA was tested. Finally, a combination of the RGA and the SA was found to give the best results.

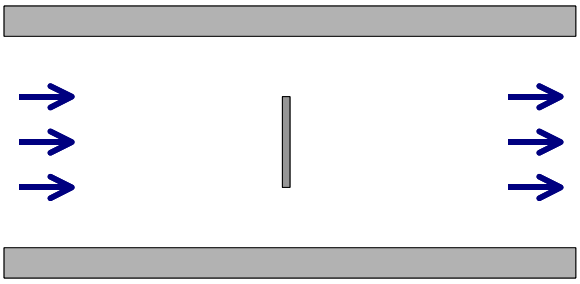

Figure 4.4 Solving flow around a plate in a flume using optimisation methods.

4.2.1 Method

To compare the effectiveness of BGAs, RGAs and SA, the flow around a plate in a channel was solved for 6 by 6 and 18 by 6 rectangular central staggered grids using finite difference and solving for continuity and momentum. This severe example was chosen to test the limits of the method. Random flow fields were generated and used to calculate objective function values, based on the residuals of continuity and momentum calculations.

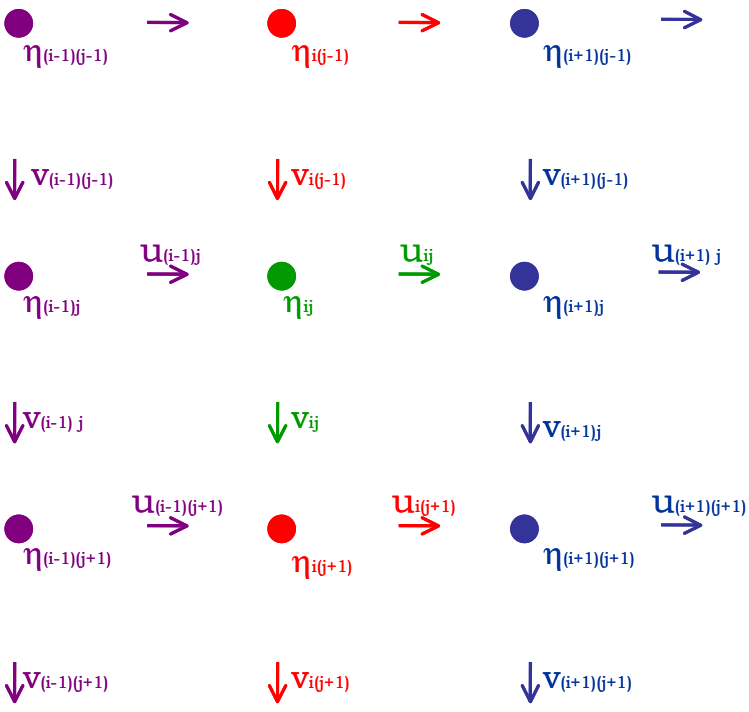

Figure 4.5 Staggered central finite difference grid system.

The different optimisation methods were evaluated on the minimisation of their objective functions, and the run time taken. A standard finite difference hydrodynamic solver was also used to determine flow fields. The results of this standard model were compared quantitatively to the best result obtained with the optimisation model using Relative Mean Absolute Error (RMAE), the method

preferred by van Rijn et al. (2003) to evaluate the performance of velocity calculations in models. Figure 4.5 shows an example of the type of grid used to calculate the objective function (Eq. 4.13), where v and u are velocities and η is water surface elevation.

The objective function minimised the squares of the errors from computing continuity and momentum about the whole grid. Eqs. 4.7-4.13 show the Navier-Stokes steady-state equations used for the centre of the grid. These equations were modified for the edges of the grid, so that they were only half centred instead of completely centred over an η point.

$$\text{Continuity}: D\frac{\partial u}{\partial x}+D\frac{\partial v}{\partial y}=0 \qquad (4.7)$$

$$\left(H+\eta_{ij}\right)\left(\frac{u_{ij}-u_{(i-1)j}}{\Delta x}\right)+\left(H+\eta_{ij}\right)\left(\frac{v_{ij}-v_{i(j-1)}}{\Delta y}\right)=\text{error}_{\text{continuity}} \qquad (4.8)$$

$$\text{Momentum}_{\text{x}}: u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}+g\frac{\partial \eta}{\partial x}-\frac{1}{\rho D}\left(4\rho\frac{C_f}{\pi}u^2\right)=0 \qquad (4.9)$$

$$\left(\frac{u_{ij}+u_{(i-1)j}}{2}\right)\left(\frac{u_{ij}-u_{(i-1)j}}{\Delta x}\right)+\left(\frac{v_{ij}+v_{i(j-1)}}{2}\right)\left(\frac{\left(\frac{u_{i(j+1)}+u_{(i-1)(j+1)}}{2}\right)-\left(\frac{u_{i(j-1)}+u_{(i-1)(j-1)}}{2}\right)}{2\Delta y}\right) \qquad (4.10)$$

$$+g\left(\frac{\eta_{(i+1)j}-\eta_{(i-1)j}}{2\Delta x}\right)-\frac{1}{\rho\left(H+\eta_{ij}\right)}\left(4\rho\frac{C_f}{\pi}\left(\frac{u_{ij}+u_{(i-1)j}}{2}\right)^2\right)=\text{error}_{\text{momentum}_{\text{x}}}$$

$$\text{Momentum}_{\text{y}}: u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}+g\frac{\partial \eta}{\partial y}-\frac{1}{\rho D}\left(4\rho\frac{C_f}{\pi}v^2\right)=0 \qquad (4.11)$$

$$\left(\frac{u_{ij}+u_{(i-1)j}}{2}\right)\left(\frac{\left(\frac{v_{(i+1)j}+v_{(i+1)(j-1)}}{2}\right)-\left(\frac{v_{(i-1)j}+v_{(i-1)(j-1)}}{2}\right)}{2\Delta x}\right)+\left(\frac{v_{ij}+v_{i(j-1)}}{2}\right)\left(\frac{v_{ij}-v_{i(j-1)}}{\Delta y}\right) \qquad (4.12)$$

$$+g\left(\frac{\eta_{i(j+1)}-\eta_{i(j-1)}}{2\Delta y}\right)-\frac{1}{\rho\left(H+\eta_{ij}\right)}\left(4\rho\frac{C_f}{\pi}\left(\frac{v_{ij}+v_{i(j-1)}}{2}\right)^2\right)=\text{error}_{\text{momentum}_{\text{y}}}$$

$$\text{Objective Function}: \text{minimise}: \left(\text{error}_{\text{continuity}}\right)^2+\left(\text{error}_{\text{momentum}_{\text{x}}}\right)^2+\left(\text{error}_{\text{momentum}_{\text{y}}}\right)^2 \qquad (4.13)$$

where x and y are the two axes that define the solution area

u and v are the two components of velocity in the x and y directions respectively.

η is the surface elevation

D is the local water depth

A number of different optimisation methods were trialled to minimise the objective function. The following sections describe the RGAs, BGAs and SAs used and give comparisons between them.

4.2.2 Real Genetic Algorithm Description

Real genetic algorithms are a type of global optimisation technique that is described in this Section.

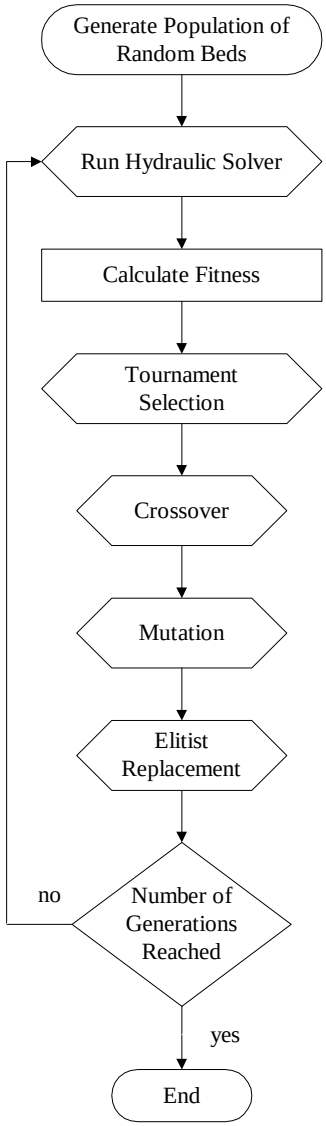

Figure 4.6 Real genetic algorithm methodology.

An RGA involves the use of a single real number to represent each decision variable. In contrast to a BGA, this means that the length of each population member is considerably reduced, as each decision variable is represented by a single number. This also means that the precision of the answer found using a

RGA is greater than that of a BGA, as there is no discretisation. RGAs have been found to be more accurate and computationally efficient than BGAs by a number of researchers (e.g. Zalzala and Fleming, 1997; Yoon and Shoemaker, 2001).

The RGA utilises tournament selection, as can be seen in the flow chart in Figure 4.6. Tournament selection involves the comparison of two different morphologies, paired at random from the population. The member that has the greater fitness (smaller objective function value) is moved to the mating pool, where crossover and mutation can take place.

Crossover involves the random selection of pairs from the mating pool. If a random number between zero and one is less than the probability of crossover, some of the genetic information (the bed elevation) is swapped between the pair. Numerous different crossover types were trialled in the developed RGA including average, one-point, two-point, uniform and average-uniform.

Average crossover is performed by cutting the population members taking part in the crossover in a number of places, and replacing the first section in one member with an average value between the decision variables of each of the two members. In the second section, the other member has its decision variables replaced by average values and so on. This is illustrated in Figure 4.7.

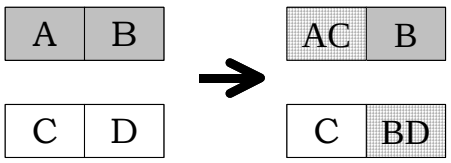

Figure 4.7 Average crossover, where A, B, C and D are four different elevation values, and AC and BD are averages of these values.

In one-point crossover, a random point is chosen. The two population members chosen to participate in crossover are cut at this point and rejoined, with the fellow member in the pairs second half of its string being joined to the first half of the other member, and vice versa. This is illustrated in Figure 4.8.

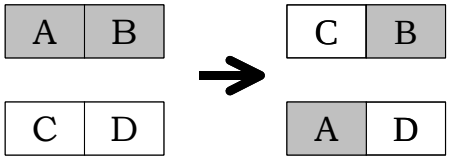

Figure 4.8 One-point crossover, where A, B, C and D are four different elevation values.

Two-point crossover is similar to one point, but instead of one cut, two cuts are made and the information contained between the two cuts is swapped between the two population members.

If average-uniform crossover occurs, a mask cell is created to consist of randomly ordered integers ranging from zero to six. If a zero is encountered, the decision variable in the position remains unchanged, if a one is encountered, the value of the decision variable of the first population member is replaced by the average of the two original values, but the decision variable value of the second population member remains the same. If a two is encountered, then the second population member is the one where the decision variable value is replaced by the average. If a mask cell value of three occurs, both population member decision variables are replaced with the average value, while a mask cell value of four means that a straight swap of values occurs between the two members without any averaging. A mask cell value of five or six swaps the first members decision variable with the second members and places the average value in the second population or vice verse respectively.

After crossover has occurred, the mutation process is undertaken. A separate random number between zero and one is chosen for each member in the mating pool. If this number is less than the probability of mutation, a single decision variable (bed elevation) is chosen and modified to a new random value within the prescribed limits. Numerous different mutation types were evaluated including random, random-adjacency, neighbour, true-adjacency and a type suggested by Wang and Zheng (1998).

Random mutation involves a decision variable being replaced by a totally new random value that is within the limits imposed on that variable. In adjacency mutation the decision variable value is modified by a step size, either randomly chosen or prescribed based on system knowledge (random adjacency or true adjacency respectively). This is a more controlled modification. The formula used to determine the mutation value of the decision variable is given in Eq. (4.14).

$$\text{child}(\text{gene}) = \text{parent}(\text{gene}) + \text{SS1}(2 * \text{RND} - 1) \qquad (4.14)$$

where child(gene) is the mutated bit value;
parent(gene) is the original decision variable value;
SS1 is a random or a predetermined step size; and
RND is a random number between 0 and 1.

Wang and Zheng (1998) mutation is described by Eq. (4.15), which was implemented if the newly generated number fell within the upper and lower limits of the decision variable, otherwise random mutation was utilised in its place.

$$\text{child}(\text{gene}) = \text{parent}(\text{gene}) + 2(\text{RND} - 0.5)(\text{UL} - \text{LL}) \qquad (4.15)$$

where UL and LL are the upper and lower limits of each decision variable respectively.

At the end of each generation, when crossover and mutation had occurred, elitist selection was implemented. This involved replacing the worst population member in the current generation with the best population member from the previous generation. This allowed good genetic information to be carried through all the calculations, making its information available in each gene pool, without the risk of losing it by a poorly made crossover or mutation.

The GA is stopped after a specified number of generations, which can be reduced or increased depending on the results obtained.

4.2.3 Binary Genetic Algorithm Description

The BGA uses similar reproduction mechanisms to the RGA. It includes tournament selection, crossover, mutation and elitist selection. It uses a population string of binary coded numbers, instead of real numbers. Each decision variable is represented by a chromosome made up of ones and zeros. On average each chromosome was five bits in length, when a discretisation interval of 0.1 was specified.

In order to calculate objective function values, the binary strings had to be converted into real numbers which could then be used to calculate continuity and momentum residuals. However, the rest of the reproduction operations were carried out on the binary strings.

One-point, two-point and uniform cross-over were available for use in population modifications of the BGA. These have already been described in the previous Section (Section 4.2.2).

Uniform mutation was the type of mutation chosen to modify binary populations. Uniform mutation occurred in each bit of the chromosome if the random number assigned to that bit, was less than the probability of mutation. If uniform mutation occurred, the value was flipped, so that a bit of value one was replaced by a value of zero and vice versa.

Further information on GAs can be found in Goldberg (1989) and Zalzala and Fleming (1997).

4.2.4 Binary and Real Genetic Algorithm Comparison

Initially, a number of tests were undertaken to determine the best crossover and mutation types to use for either a real or binary GA, as well as the most suitable population, generation and probability of mutation or crossover values. These were chosen using comparisons of objective function values. A table showing comparison values can be found in Appendix B. The test situation used in this first instance was flow around a plate in a 6 grid cells by 6 grid cells channel. This size was chosen as it took a relatively small amount of computation time, allowing rigorous sensitivity analyses to be undertaken. It was found that the best combination for the RGA, which allowed a near global minimum to be found in a reasonable computation time, was average crossover with a probability of 0.8, random mutation with a probability of 0.01, a population of 400 and 2000 generations. In contrast, the best BGA configuration was found to be uniform crossover with a probability of 0.7, uniform mutation with a probability of 0.001, when a similar population and generation size to the RGA was used. The resultant velocity path around a plate in a short channel, using either a RGA or BGA is shown in Figure 4.9. It was found that the RGA was able to outperform the BGA in terms of objective function value, with values recorded of 1.36 and 8.89 respectively, as can be seen in Figure 4.10. This is because the RGA does not discretise the answer so it is able to get closer to the global minimum as the values it is able to utilise do not come from a limited set.

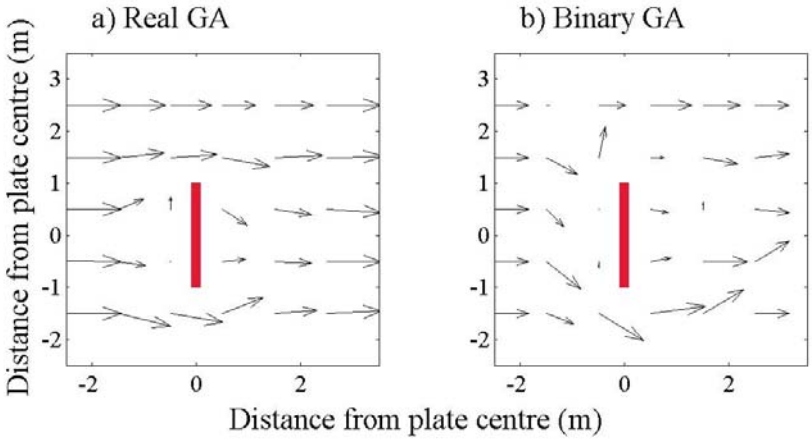

Figure 4.9 Velocity pattern around a plate – comparison of (a) RGA and (b) BGA results in a short channel.

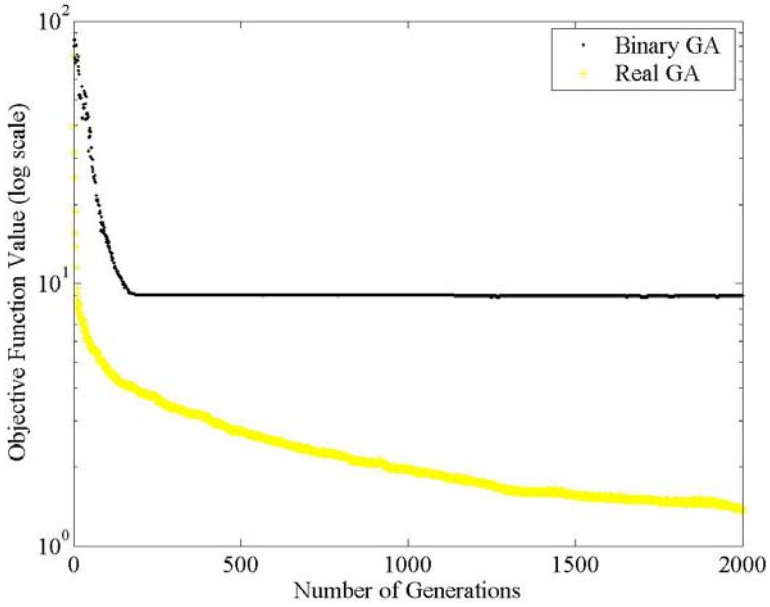

Figure 4.10 Comparison of BGA and RGA objective function values.

4.2.5 Simulated Annealing Description

In contrast to a GA, which uses a population of different decision variable combinations, an SA uses only one population member made up of decision variables (velocities or elevations). These decision variables are perturbated. The general method of an SA is shown in Figure 4.11.

Firstly, an initial 'temperature' is set based on the fitness of the initial decision variable values. These values may either be random, or the values obtained from the previously run GA. The temperature is analogous to the temperature used in the annealing of a metal. First it is increased to melt the metal, or in this case increased to allow a certain number of solutions that have a worse fitness to be accepted. The temperature is then decreased slowly, in the case of a metal, allowing a crystalline structure to form. In the computer optimisation, this allows a near optimal solution to be found.

The calculation of the initial temperature is based on research by Cunha and Sousa (1999; 2001). It is calculated using Eq. (4.16), where an arbitrary factor called the elasticity of acceptance (a) is assigned a value between 0.1 and 0.3.

$$t_1 = -\frac{0 \cdot 1 c_0}{\ln a} \qquad (4.16)$$

where t_1 is the starting temperature; and

c_o is the initial fitness.

This starting temperature is applied for the number of perturbations required at each temperature, after which the number of perturbations accepted can be assessed. If the number of perturbations accepted is 40% or more, as suggested by Johnson et al. (1989), then the temperature may be decreased. Otherwise, the temperature must be increased until this criterion is satisfied.

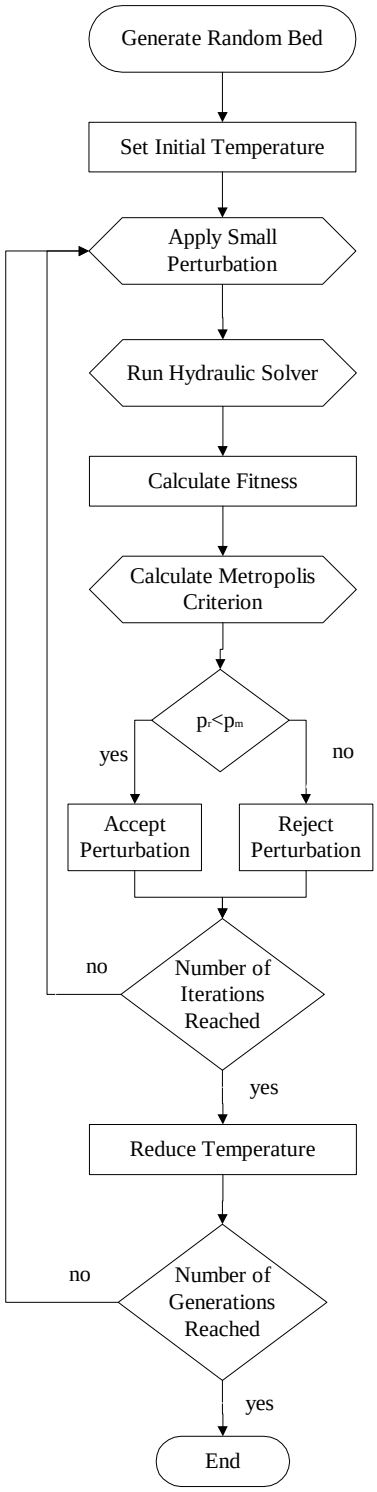

Figure 4.11 Simulated annealing methodology.

where p_r is a random number between 0 and 1; and

p_m is the metropolis criterion.

At each temperature a number of perturbations are made using random mutation, where a single decision variable is chosen on the string and mutated to a different value within specified limits. The difference in fitness between the new and old string is calculated and the Metropolis criterion is applied, as shown in Eq. (4.17) (Metropolis et al., 1953). The Metropolis criterion is dependent on temperature. At a higher temperature it will allow a greater number of solutions that increase the objective function value to be accepted, allowing the search method to jump out of local minima. As the search progresses, the method should be nearing the global minimum, so it becomes increasingly less likely that a worse solution will be accepted.

$$A_{ij}(c) = \begin{cases} \exp\left(\dfrac{-\Delta f}{c}\right) & if \quad \Delta f > 0 \\ 1 & if \quad \Delta f \le 0 \end{cases} \qquad (4.17)$$

where $A_{ij}(c)$ is the Metropolis criterion;

c is the current temperature; and

Δf is the difference in objective function value between the decision variables with and without the perturbation.

The Metropolis criterion is compared to a randomly generated probability of acceptance value. If this random number is less than the metropolis criterion, the perturbation is accepted and another is tested. This continues for a number of iterations. The temperature is then decreased, by multiplying it by a cooling factor, and the process begins again. Temperature decreases continue until a specified temperature is reached. Further information on the method involved in SA can be found in Kirkpatrick et al (1983), Černý (1985) and Cunha and Sousa (2001).

4.2.6 Real Genetic Algorithm and Simulated Annealing Comparison

A similar test case was utilised to compare the performance of the RGA and SA, but this time a longer channel was used in order to increase the number of decision variables the optimisation modules were solving for. In this test case the RGA decreased the value of the objective function significantly at the start of the run, but stagnated after a while. Stagnation occurred as the GA became insensitive in the vicinity of the global minimum, when the solution space has

multiple dimensions and a very shallow gradient near the global minimum. The SA was able to significantly reduce the objective function value as can be seen in Figure 4.12.

A similar trend was observed by Ruessink (2005b), as the GA he employed to optimise cross-shore model parameters was quick to approach a global minimum, but then slowed down, taking a large amount of computational time to reach a final convergence value as it did not exploit local information. Because of this, Ruessink used a local search method, namely a downhill simplex algorithm, to obtain final minimum parameter values. Similarly Franconi and Jennison (1997) observed the stagnation of the GA after 100 generations, when the probability of improving single decision variables in the population member, while still keeping other decision variables that had already been improved, reduced. They observed that in this situation SA was advantageous as it was able to alter only one decision variable at a time. To accommodate both these features they developed a hybrid model where crossover was utilised as part of a simulated annealing module. This was discussed previously in Chapter Two, Section 2.3.

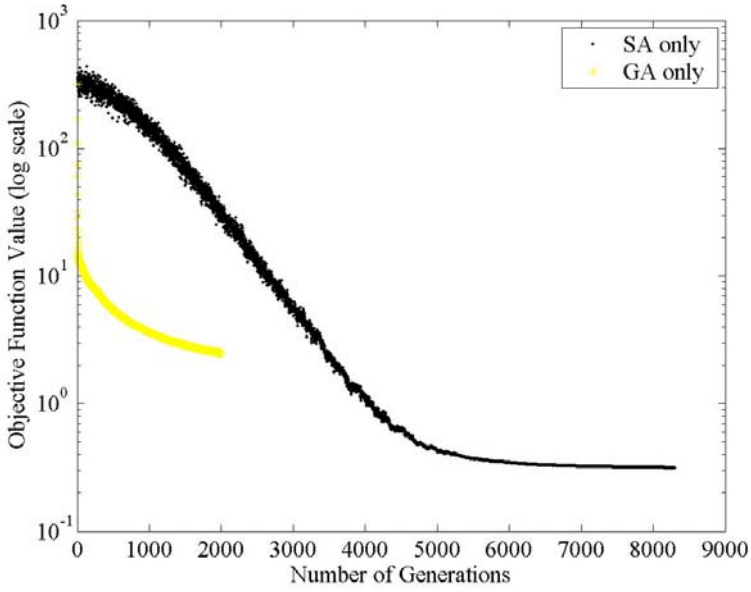

Figure 4.12 Comparison of RGA and SA optimisation models. The objective function value is given in log scale.

The parameters used in the SA were not as robust as in the GA and required some tuning. The best parameters found were a final temperature of $1x10^{-6}$, a cooling factor of 0.998, an acceptance elasticity of 0.2 and 800 iterations before a temperature decrease. A comparison of the RGA and SA can be seen in Figure 4.12. The SA did not drop initially as fast as the RGA and was more dependent

on starting configurations. It was also more computationally intensive, taking longer to reach a similar fitness level. Athias et al. (2000) also observed this problem of the SA being overly computationally intensive in their comparison of a BGA and SA.

The example shown in Figure 4.13(b) uses a longer channel than Figure 4.9. This increases the number of decision variables, which reduces the ability of the GA to find a near optimal solution.

The SA achieved a lower objective function and so predicted a more realistic velocity pattern than the GA. The GA was able to correctly predict the velocity near the boundary, but has trouble near the centre of the channel. It has made these velocities small, so that the error between velocities near the plate is small in proportion. The error between the larger velocities near the boundaries was larger however, which contributed to the overall error. It only reduced the error of the velocities near the boundary by correcting the velocities interacting with the boundary in column 8 and –7.5, followed by 7.5 and –6, and so on. The different merits of each optimisation method led to the idea that a combined approach may be more beneficial. The combined approach used is discussed in the following sections.

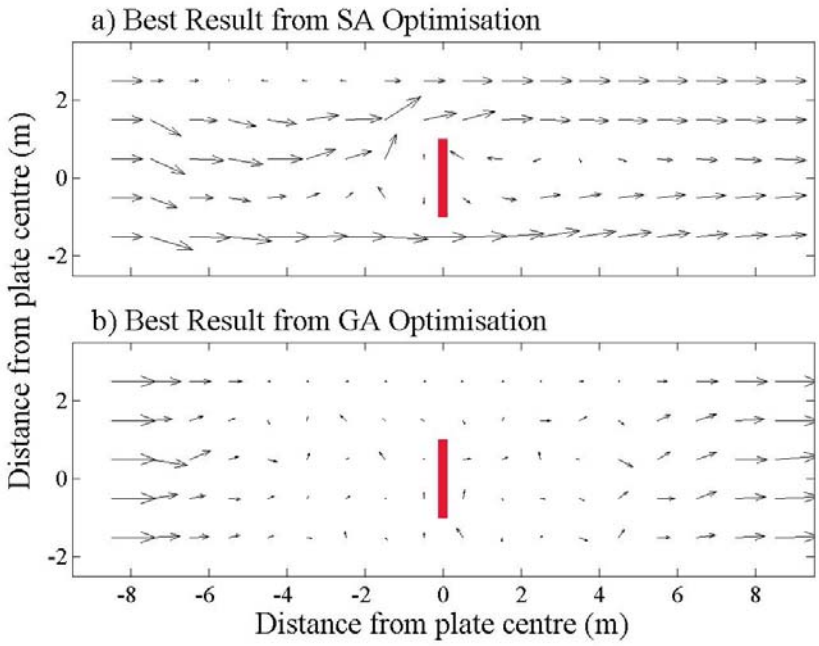

Figure 4.13 Velocity pattern around a plate – best results from (a) SA and (b) RGA in a long channel with a plate asymmetrically positioned.

4.2.7 Combination Comparison

Based on this result and comparisons of the two separate methods by previous researchers such Wang and Zheng (1998) and Ingber and Rosen (1992), a hybrid RGA-SA optimisation method was developed. The RGA was run initially until stagnation occurred, at which point the SA algorithm took over, using the best configuration of the RGA as its starting configuration. This was able to reduce the objective function significantly in a shorter time than a SA model alone, as can be seen in Figure 4.12 and in Table 4.1.

As can be seen in Figure 4.14 and Table 4.1, the combined optimisation method was able to find a better solution in less computation time, as the number of objective function calls was reduced. This reduced the run time. The run time could be further reduced if the SA took over after fewer GA generations as the SA requires fewer function calls as it is applied to a single population member rather than an entire population. It would be possible to apply the SA earlier in this example as the GA was stagnant for numerous generations before the SA began.

Table 4.1 Values obtained using GA or SA only and RGA-SA optimisation combination.

	Number of Function Calls ($x10^6$)	Minimum Objective Function Value ($x10^{-1}$)
GA Only	0.8	63
GA-SA Combination	4.38	2.29
SA Only	6.64	3.17

An indication of the original randomly chosen velocity pattern is shown in Figure 4.15, with the final pattern in Figure 4.16 (b). This suggests that the optimisation method is quite powerful, as it can find a realistic answer from a starting point that is very unrealistic. The end flow pattern may not be identical to that obtained using a traditional model, as can be seen in Figure 4.16, but the directions are correct, even if some of the magnitudes are different. The important knowledge gained in this investigation is that the optimisation method was able to predict a general pattern when it started from something completely wrong. This is advantageous, as when the optimisation methods are transported into the morphological prediction model, they should be able to predict a general morphology that conforms to a minimum objective function.

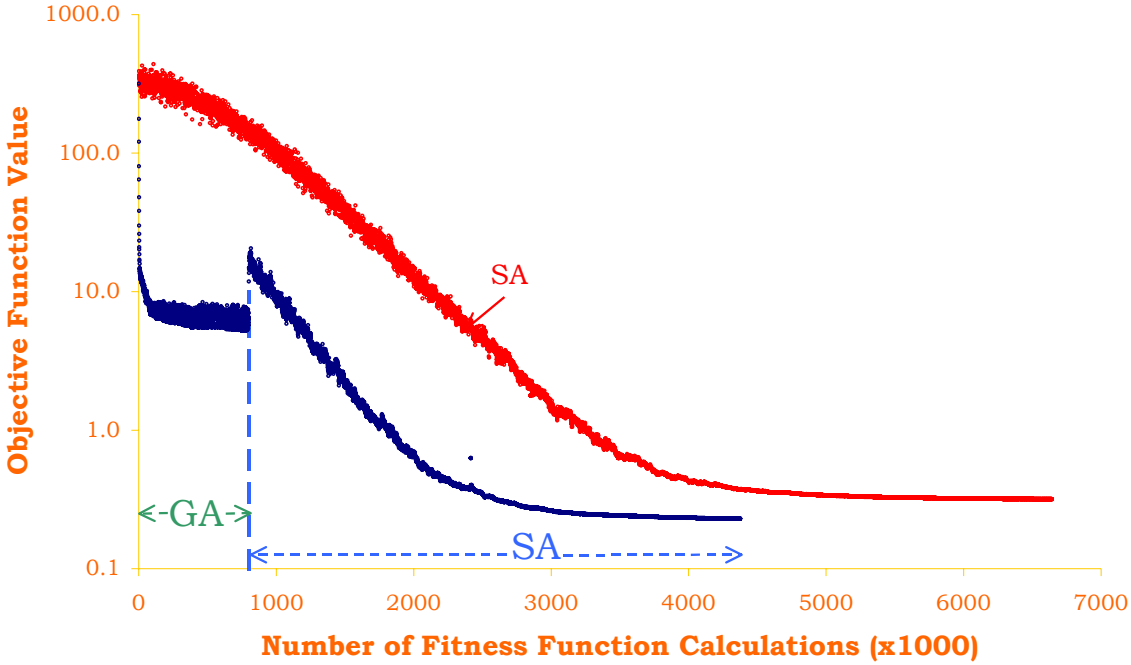

Figure 4.14 Comparison of combined RGA-SA and SA optimisation models.

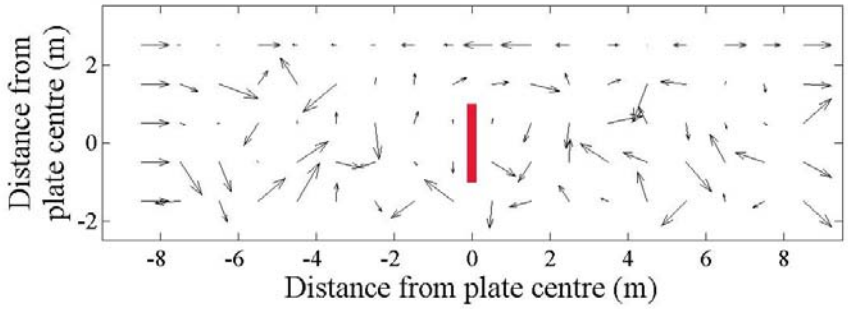

Figure 4.15 Initial random velocity pattern in a long channel around an asymmetrically positioned plate.

A comparison was made between the optimisation deduced velocity pattern and the pattern found using a standard finite difference flow solver, the same program used in Chapter Three and discussed in Section 3.3. The pattern obtained by the hydraulic solver is shown in Figure 4.16. A quantitative comparison of the two velocity patterns was made using the RMAE (see Eq. 4.18), the preferred method of velocity comparisons suggested by van Rijn et al. (2003). The RMAE is preferred as it less susceptible to the presence of outliers (van Rijn et al., 2003).

$$RMAE = \frac{\left(\left|V_c - V_m\right| - \Delta V_m\right)}{\left|V_m\right|} \qquad (4.18)$$

where V_c is the computed velocity;

V_m is the measured velocity; and

ΔV_c is the error in the measured velocity.

A value of 0.28 was obtained when the velocity vectors were compared, giving a good result (see Table 4.2). If the longitudinal and transverse components are examined separately, values of 0.41 and 0.83 respectively result. The longitudinal value sits within the reasonable bracket, but the transverse velocity predictions are rated as bad. This is due to their small size, with values being close to zero, an area where the RMAE is extremely sensitive (van Rijn et al., 2003).

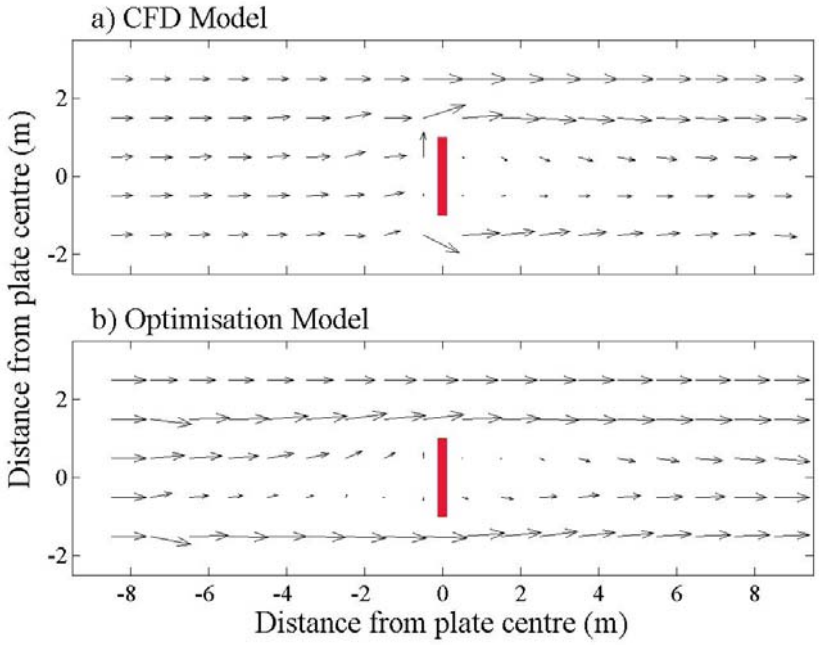

Figure 4.16 Velocity pattern around an asymmetrically positioned plate – obtained using (a) a standard hydrodynamic solver and (b) the RGA-SA optimisation model.

Table 4.2 Comparison table for values of RMAE velocity, as suggested by van Rijn et al. (2003)

Relative Performance of Model	RMAE Value Obtained
excellent	<0.1
good	0.1 - 0.3
reasonable	0.3 - 0.5
poor	0.5 - 0.7
bad	>0.7

4.2.8 Optimisation Conclusions

The results of this exercise suggested that the most efficient optimisation method to tackle problems with a large number of decision variables was a combined RGA and SA. This optimisation program was developed to be very flexible and adaptable, with the basic structure able to be transported to various problems, for which different objective functions could be added. It also allowed for the trial and comparison of numerous different types of crossover and mutation, as well as the comparison of three different optimisation techniques, namely BGAs, RGAs and SA algorithms.

The ability of the optimisation method to predict correct flow patterns starting from an initial randomly determined population, with a large number of decision variables, means that the method should be able to be transferred to the prediction of equilibrium morphologies, the situations of interest in this thesis.

4.3 Unidirectional Lagoon Modelling

A unidirectional lagoon was modelled utilising a similar set up as in Chapter Three, Section 3.9, and as discussed in the Method Section earlier in this Chapter (Section 4.1) (see Figure 4.17). A number of different objective function descriptions were trialled. Firstly an objective function using both local and global energy dissipation minimisation is discussed (Eq. 4.6). This is followed by a description of sensitivity analyses undertaken, where only the local (the sum of Eqs. 4.4 and 4.5) or the global energy dissipation (Eq. 4.3) was considered, as well as the application of multiple SAs to further reduce the objective function value. A copy of the code used to optimise this example can be found in Appendix C.

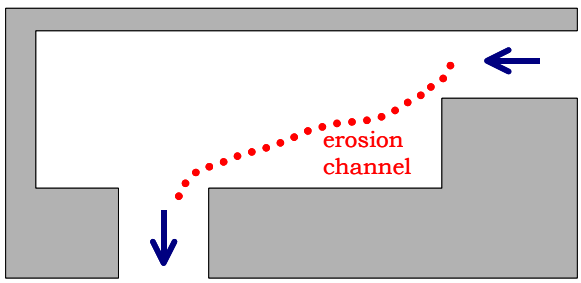

Figure 4.17 Example of field-sized unidirectional lagoon.

4.3.1 Optimisation Parameters Utilised

Based on experience from the flow field optimisation experiment, combined RGA and SA were used to find near equilibrium morphologies. In general, a GA with a

population of 50 was initially used to narrow the search space to an area where the pinpointing of a near optimal solution was possible and the search had stagnated. An example of one result from such an optimisation can be seen in Figure 4.18.

The GA used tournament selection, average crossover, random mutation and elitist selection. The crossover and mutation probability values used were 0.8 and 0.01 respectively.

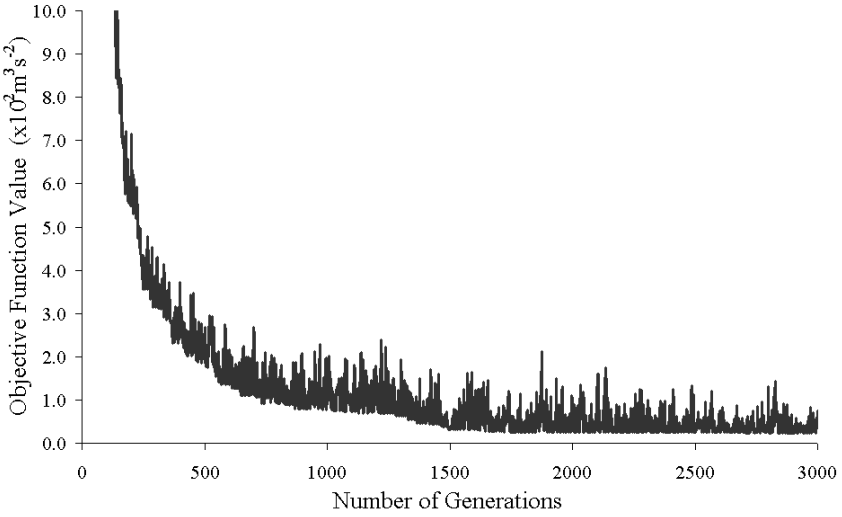

Figure 4.18 Path of objective function during GA optimisation.

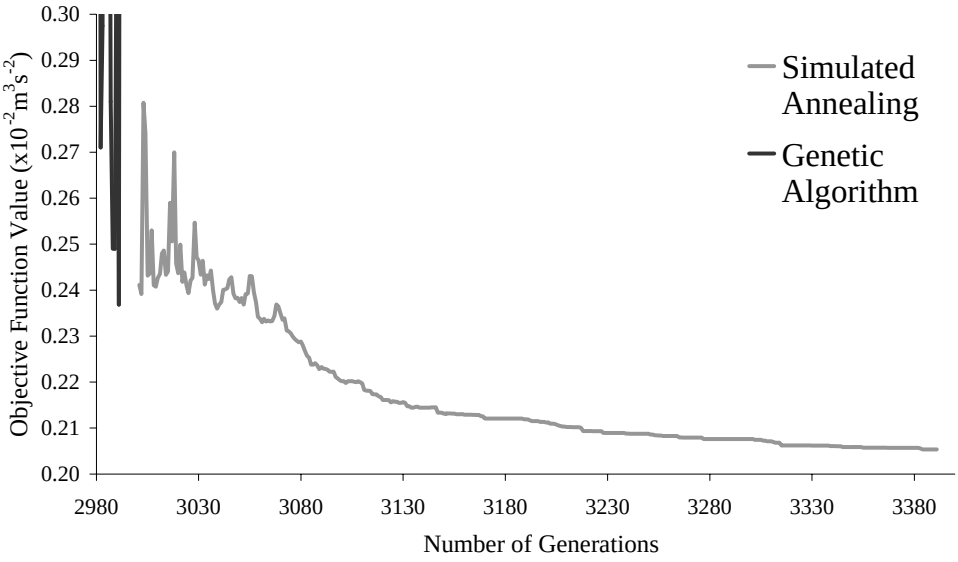

Figure 4.19 Path of objective function found when an SA was employed after the GA shown in Figure 4.18 stagnated.

After the GA had stagnated in its attempt to locate a near global minimum, an SA was employed using the best result from the GA optimisation. The path of the

subsequent SA employed following the GA pictured in Figure 4.18 can be seen in Figure 4.19.

4.3.2 Results of General Application

The fittest population member from the initial population of randomly chosen morphology is shown in Figure 4.20 (a). As can be seen here, the initial solution started as a highly improbable morphology, not at all in equilibrium configuration. Very quickly the optimisation modules were able to refine the morphologies, pinpointing a path between the entrance and exit of the flow. This can be demonstrated by a comparison of the morphology in Figure 4.20 (b), where after only 200 generations, a channel is clearly visible.

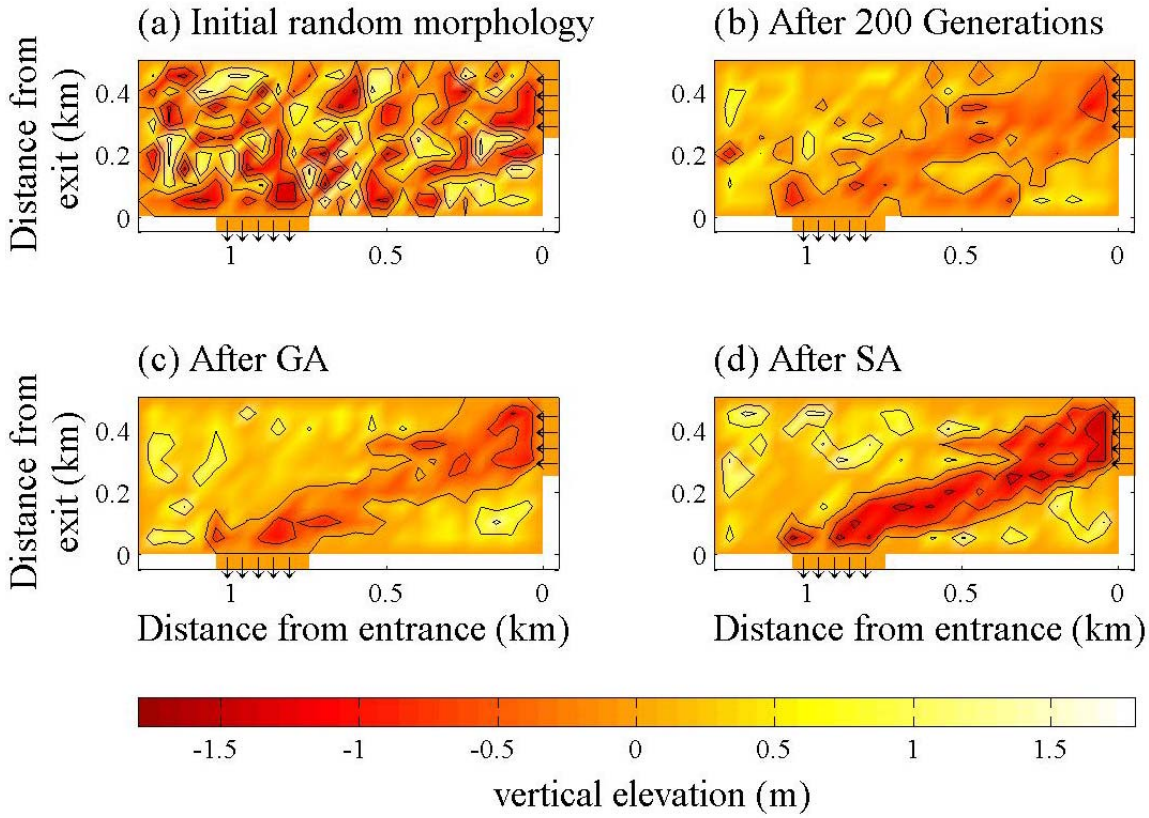

Figure 4.20 (a) Initial random morphology utilised in the GA optimisation model, (b) morphology after 200 GA generations, (c) morphology after the GA optimisation and (d) morphology after the SA optimisation.

The results of morphologies found after the GA and SA stages for a simplified lagoon with unidirectional flow can be seen in Figure 4.20 (c) and (d) respectively. As can be seen in these figures, the SA is able to refine the initial GA morphology, further defining the channel formed between the flow entrance and

exit of the lagoon. Both solutions did not contain any excess erosion or deposition, but the resultant morphology found using the SA involved a smaller energy dissipation rate associated with the water travelling through the system. After 200 generations, the GA had refined the initially random morphology to a noticeable channel connecting the entrance and exit of the lagoon. It then continued to refine this morphology for the remainder of the generations evaluated.

This entropy-based method is advantageous as the optimisation module is able to find a plausible morphology without the steps required to move from an initial to a final morphology, purely on a comparison and refinement of random bed morphologies favouring those that minimise energy dissipation. As discussed in Section 4.3.4, if only the global energy dissipation of the water is minimised, that is the difference in energy between the entrance and exit of the lagoon, the system tends towards a single deep lagoon. On the other hand, if only the local sediment component is considered, that is the velocity penalty, the system forms a basic channel but is not at liberty to refine it. This shows that to minimise energy dissipation and satisfy minimum entropy production, the global energy and local penalties must both be considered.

4.3.3 Sequential SAs

To test the sensitivity of the model to finding a near optimal solution after the application of only one GA and SA, further SAs were run using the best morphology found in the previous SA and a different set of random numbers. The paths of these subsequent SAs can be seen in Figure 4.21.

As can be seen in Figure 4.21, the first few SAs do improve the fitness of the solution marginally. On the lagoon morphology itself they act to deepen the channel between the flow entrance and exit. This can be seen by the lagoon morphology obtained after the third SA application as shown in Figure 4.22.

After the application of the third sequential SA, no further changes in the objective function value are observed. This suggests that the near optimal solution has been found. The SA is better able to refine the solution as it changes only one depth at a time before recalculating the objective function. The GA on the other hand can change multiple elevations at once, so it is harder for it to refine finer details. The SA is not able to find the near optimal solution as efficiently if it is used from the start in predicting the general area in the solution space where the optimum objective function value occurs. Similar points to this

were noted by Wang and Zheng (1998) in their comparison of GA and SA suitability for groundwater management optimisation.

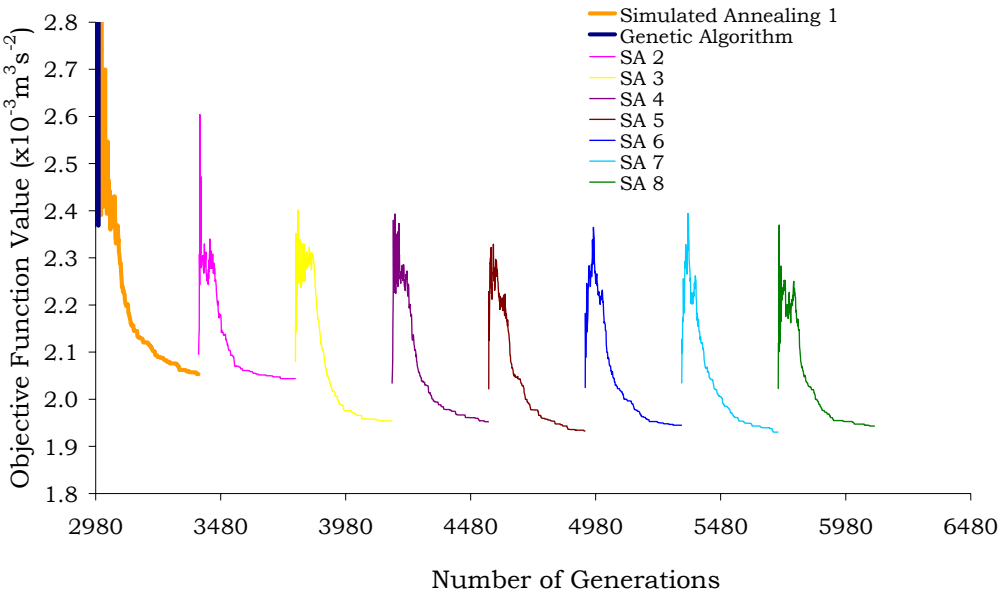

Figure 4.21 Path of Objective Function Values for Sequential SAs.

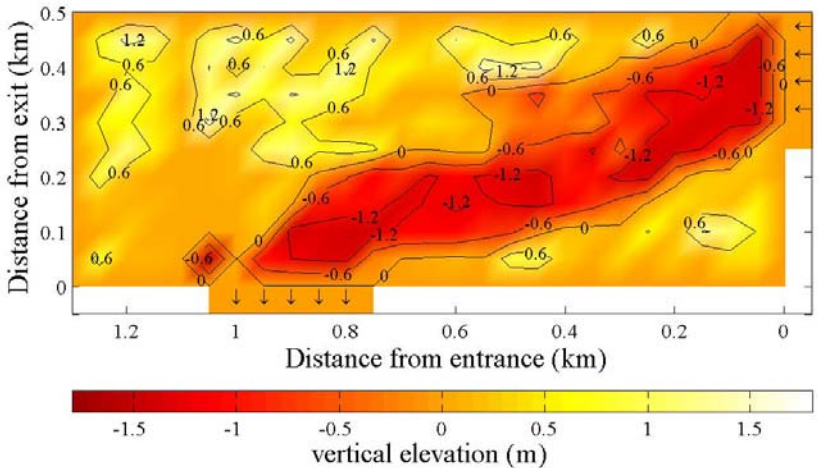

Figure 4.22 Resultant morphology from the application of a GA and three SAs sequentially.

4.3.4 Sensitivity Analyses

To test the requirement of both global energy dissipation and local penalty minimisation, a number of tests were considered where only one of the components was optimised. The results of these tests are outlined in this Section.

Initially an objective function (Eq. 4.3) using only minimum global energy difference between the entrance and exit channels of the lagoon was considered.

The optimum morphology developed using a GA based on this objective function can be seen in Figure 4.23.

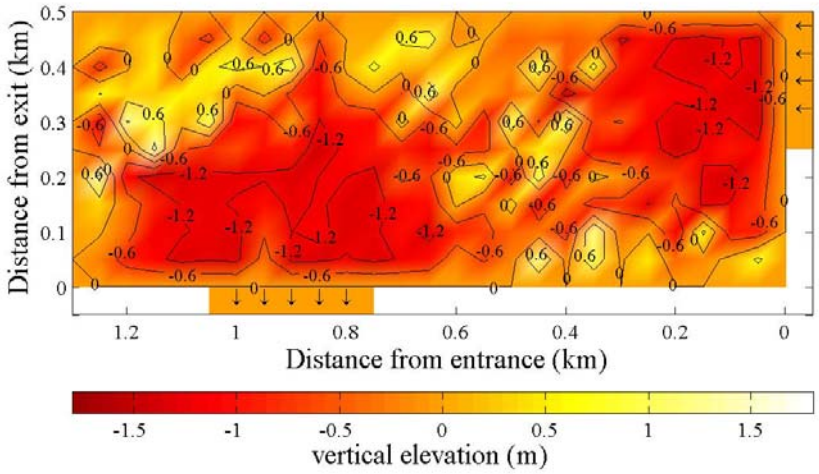

Figure 4.23 Morphology of lagoon with only global energy dissipation component in objective function value calculations.

This morphology appeared to create a channel between the entrance and exit, but also acted to erode the whole area as much as possible, thereby deepening the flow and reducing the loss of energy from the system. A volume of 120m^3 of sediment was lost from the system.

When only global energy difference is considered, and the system tries to deepen itself as much as possible, no account of the energy requirement to erode the whole system is taken into consideration. The energy used to mould the bed is not taken into consideration, only the energy lost from the water. Including a critical velocity penalty component for excess erosion in the objective function helps to account for areas where the velocity is not high enough to erode and transport.

Penalties were also necessary in areas where deposition became so great that velocities over the deposition mounds were strong enough to erode sediment, meaning that physical constraints were no longer adhered to.

A number of other variations in the calculation of the objective function values were considered, using the same critical velocity penalties but different global energy derivations. Firstly the points used to obtain an average entrance and exit energy were moved one point closer to the model boundaries. The results of this variation can be seen in Figure 4.24 after the SA application. This variation obtained a similar morphology to the use of the points further from the model boundaries as was expected as the same erodible bed area was considered.

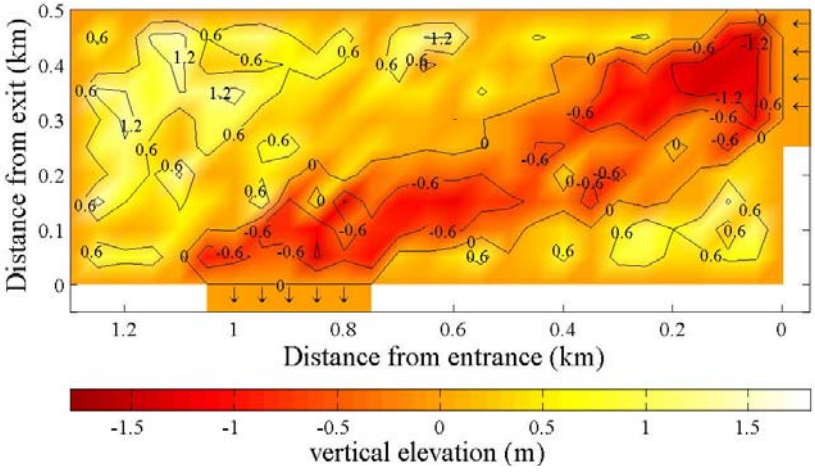

Figure 4.24 Morphology of lagoon after SA using points closer to the boundaries for energy calculations.

A second variation was trialled, where the global energy was determined by an average velocity and water level rather than averaging the individual point energies. The average velocity was based on the sum of velocities and areas divided by the sum of the areas. The resultant morphology of this variation after the SA optimisation can be seen in Figure 4.25. Again, a similar morphology was found. This is shown when a comparison of the objective function value paths is made, as in Figure 4.26, where both methods reach a similar objective function value after SA optimisation.

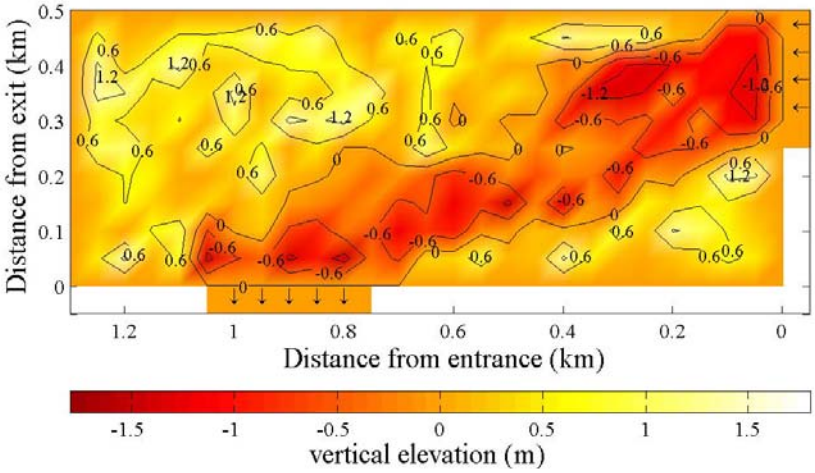

Figure 4.25 Morphology of lagoon using average water elevations and velocities after SA optimisation.

A test was undertaken using critical velocity penalties only in the objective function (Eqs. 4.4 and 4.5), thereby disregarding the energy lost globally over the system via the water. The morphology that resulted after a GA run of 3000 generations with a population of 50 can be seen in Figure 4.27.

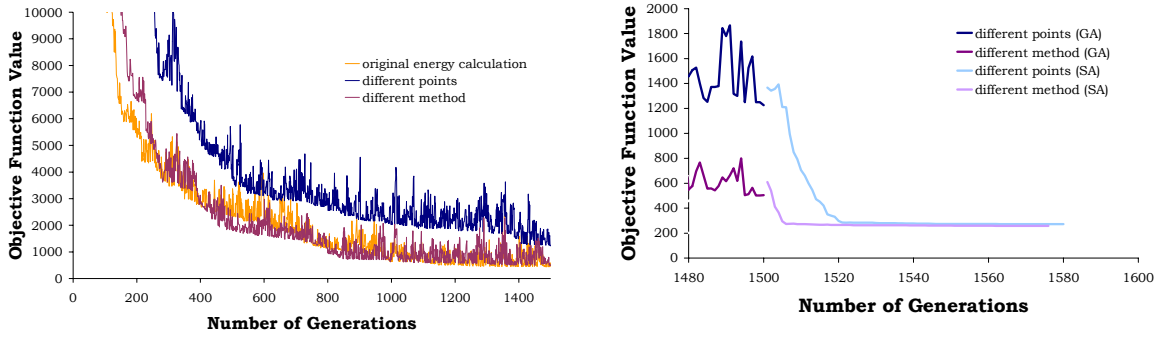

Figure 4.26 Comparison of objective value paths using different energy dissipation descriptions.

Upon first glance, it appears that the critical velocity only objective function obtains a similar value. However, once it reaches a state where all the erosion areas have a value of critical velocity and all the deposition areas a velocity below critical, there is nothing more that it can do to improve. There is nothing to guide the solution as to whether it should erode in a different area, erode a deeper channel, deposit a higher mound, or multiple mounds. The deposition of an additional mound may cause a greater loss in energy over the system, but have no effect on the velocity rule, as the velocity is below critical. This suggests that it is vital to include both components in the calculation of the objective function. Figure 4.28 illustrates the point that when the SA is run, the morphology of the lagoon does not improve, as the threshold of the objective function has already been reached, and further changes that do not worsen the critical velocity penalty will have no effect.

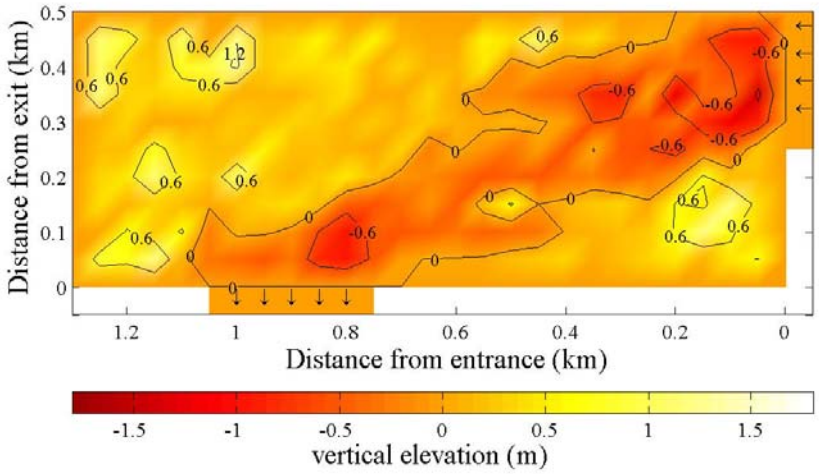

Figure 4.27 Morphology of lagoon using critical velocity penalties only as objective function after GA.

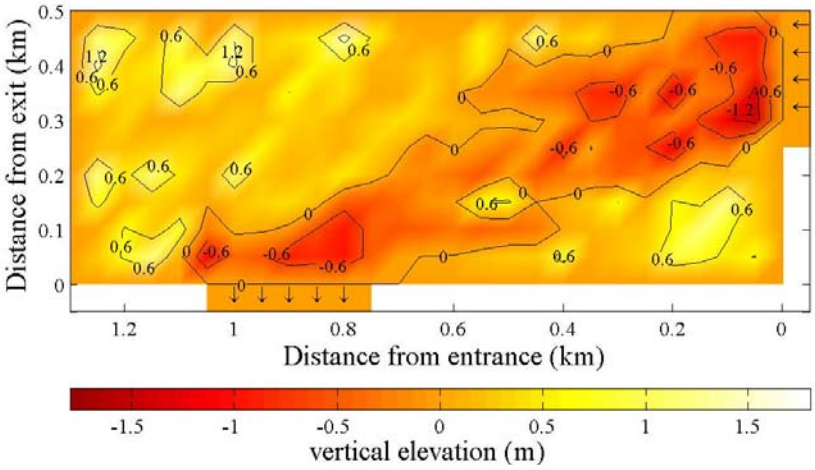

Figure 4.28 Morphology of lagoon using critical velocity penalties only as objective function after SA.

Figure 4.29 and Figure 4.30 show the different paths taken by the optimisation routines with and without the inclusion of global energy in the objective function. As can be seen in Figure 4.29, the objective function reaches a minimum value of 200 after approximately 1700 generations. After this it oscillates around this minimum value. When the SA is run, there is no way that it can improve on the objective function value as the minimum has already been reached, so it just finds a different combination of bed elevations that obey the critical velocity rules. This can be seen by the flat objective function path with each temperature drop in Figure 4.30. It has no way of knowing if one combination is better than the others as the energy of the water flowing through the lagoon is not included, just the minimisation of excess sediment erosion and deposition.

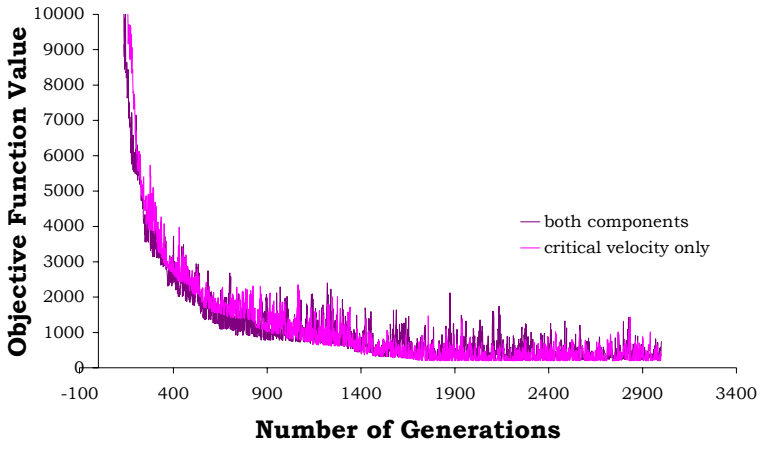

Figure 4.29 Path of objective function values for optimisation with and without global energy using GA only.

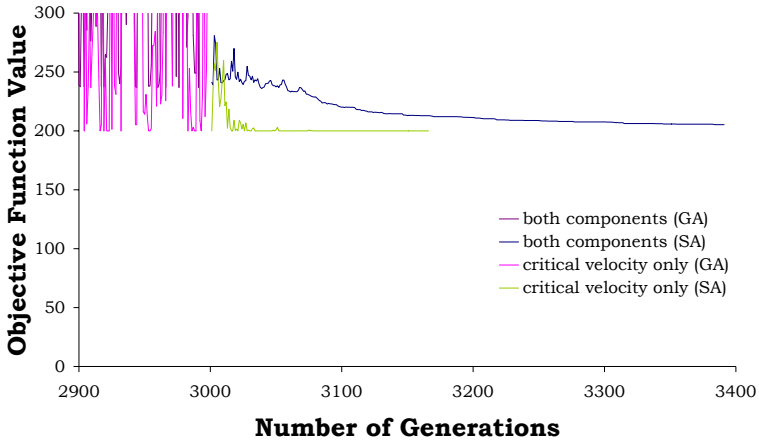

Figure 4.30 Enlargement of SA results for path of objective function values for optimisation with and without global energy.

This research suggests that the use of an optimisation algorithm together with an objective function in which the global energy of the water and a penalty for excess erosion and deposition of sediment, is able to predict a minimum energy equilibrium morphology and the path a river might take in a lagoon on its way between the river entrance and the exit to the coastal zone. This finding is important as it represents a new result into the prediction of two-dimensional moving bed morphologies between two defined boundaries. It builds upon work done to predict river systems but exceeds these by considering complete morphologies, and is at a much finer scale than previous research.

4.3.5 Discussion of Unidirectional Lagoon Results

The model detailed in the previous sections represents a method to predict equilibrium morphologies that differs from widely used traditional process-based models. In fact, it is the focus on determining an equilibrium position that differentiates the new approach from traditional models. The entropy-based method still incorporates essential processes, but is not implemented by time stepping, instead predicting the equilibrium morphology directly, based on the processes considered to have a major influence. It is based on universal principles, applicable over a wide range of disciplines. The entropy-based method is designed to jump straight to stable morphologies, so it is not restricted to predicting short-term phenomena based on computational time and error amplification constraints. Conversely, the lack of time-step does limit the model to some degree, as it is harder to estimate the time taken for a system to recover from a perturbation such as channel dredging which significantly alters the system. In general, it is not possible to estimate the exact time period taken for equilibrium to occur. However, if long-term estimates are the goal, the exact time

taken to form an equilibrium morphology may not be as important as the actual morphological prediction.

The self-organisation-based method discussed in Chapter Three, Section 3.1 and the entropy-based method differ in their steps to obtaining a stable morphology but give similar predictions. The cellular automata model rearranges sediment gradually using rules and sediment increments, while the entropy-based method finds the equilibrium morphology using optimisation based refinements that minimise energy dissipation independent of sediment transport. If a different magnitude of forcing is applied, i.e. the inflow velocity is reduced or increased, the amount of erosion in both models is reduced or increased respectively. As an optimisation tool the SA works better if it uses an optimised morphology as its starting place. The further application of the SA improved the energy minimisation of the system marginally until no further changes were observed after three successive SA applications. The GA could be stopped after a fewer number of generations and the SA applied earlier, to give similar results, as the GA was able to refine the morphologies very quickly to pinpoint a channel after a hundred generations, but the further refinement of the morphologies was slow using the GA.

The entropy-based method is more advantageous than the self-organisation-based method, as it considers the energy requirements of the sediment and water, both locally and globally. The cellular automata model is based on the self-organisation of the sediment. The organisation of the sediment enables the energy dissipation of the system to be minimised to some extent through the feedback associated with the rule application and hydraulics updating. The self-organisation-based method is limited by the amount of sediment in the system initially. The structuring of its rules allows for clear water scour only, as no sediment may enter the system; sediment may leave by the exit only. The entropy-based method on the other hand is limited only by the conditions imposed by the maximum and minimum values assigned to the random values chosen as the sediment elevations. This means that it has some leeway as to the addition or depletion of sediment in the system, so some sediment may be introduced or lost to fill or erode holes respectively. This means that the fixed amount of sediment that the self-organised system starts with is not a limiting factor in the entropy system, as this initial mass is not conserved.

Grid size of the models has some effect on the morphologies obtained. The use of a finer grid will create a larger solution space that the optimisation routines must search through. This means that much longer run times are required to refine the solution and the routines are more likely to become stuck in local minima.

The disadvantage to using a coarse grid size is that the solutions do not give fine detail of the morphological formations. Any formations with a characteristic size less than, or approximately the same size as the grid size will not be observed in the model, so only large scale formations can be studied with the large grid size currently utilised in the model.

The research undertaken in this PhD develops the application of entropy principles in a two-dimensional, medium scale, coastal situation. This approach is general and different components can be enhanced if needed by the use of more advanced threshold velocity formulas and hydraulic models.

The results portrayed in this Section show that the use of simplified lagoon geometries, energy dissipation descriptions and critical velocity formulae, together with unidirectional flow give a morphology resembling bed shapes studied in the field. The attack of the problem from a different angle, predicting equilibrium morphologies without directly stepping through time, shows great promise. In the following Section, the model is expanded to be applicable to simplified lagoons where flow reversal occurs.

4.4 Expansion to Reversing Flow Lagoon Modelling

The next aim of the research was to build upon knowledge gained in the unidirectional flow entropy-based lagoon model, discussed in the previous Section (Section 4.3), and develop a model which could apply similar principles to a lagoon with reversing flow, moving through the lagoon from both the river and ocean entrances. This was an additional step to bridge the knowledge gap and develop an equilibrium morphology prediction model capable of analysing situations where flow oscillated through a lagoon without the need to time-step and apply traditional methodologies. The following sections outline the additional methodology required to make the model applicable to flow reversal situations. Results are then given using similar parameters to that of the unidirectional lagoon, followed by a discussion including some sensitivity analyses.

4.4.1 Method

The inclusion of flow reversal was done on a simplified basis. The objective function was modified to include the addition of two different components. The first component was the calculation of the objective function with flow moving through the system from the river to the ocean. The second component used the same calculations as the first, but with the flow in the reverse direction, i.e. flow

now entered the basin area from the previously defined exit and exited from the previously defined entrance (see Figure 4.31 for a graphical description). This is shown in Eq. 4.19.

$$OF = \left(E_s \times P_d \times P_e\right)_{flood} \times \left(E_s \times P_d \times P_e\right)_{ebb} \qquad (4.19)$$

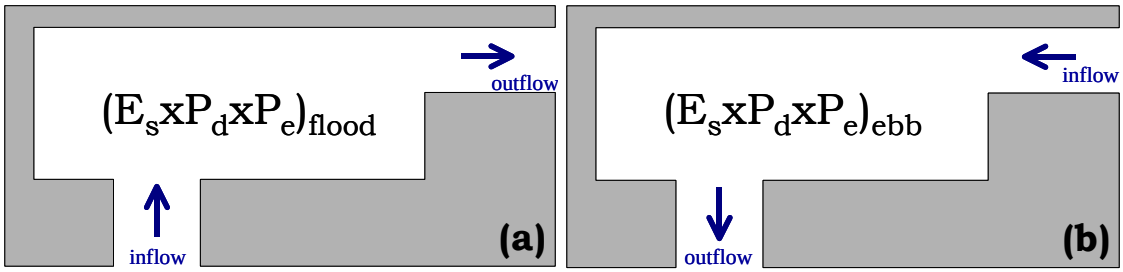

Figure 4.31 Description of flow reversal lagoon objective function definitions.

This method was designed to mimic the maximum extremes of flood and ebb oscillating flow. If the system is at equilibrium when the maximum flows in each direction occur then the flow in between these times should not have the energy to influence the morphology as greatly and can be ignored for this simplified initial study. This approach is necessary as there is no use of time in the optimisation method, a snapshot of morphology is analysed with the inclusion of excess erosion and deposition penalties, without the need to see what will happen over a tidal period.

4.4.2 Results

The initial random morphologies with the smallest objective function values from the unidirectional lagoon study in Section 4.3 and the flow reversal situation were identical. This suggests that the other random morphologies do not include an erosion hole at the entrance and exit channels, hence do not minimise energy dissipation or the penalties as much. The objective function values are slightly different however, as the flow reversal objective function (Eq. 4.19) includes the addition of the energy dissipation of the flow travelling in the opposite direction. In this Section, the reversing flow example used the same flow ($240 m^3 s^{-1}$) and water elevation (1.5m) in both directions. These conditions were the same used in the lagoon example discussed previously.

After 600 generations (see Figure 4.32 (a) and (b)), a difference between the unidirectional and flow reversal situation optimisations was noticeable. With the inclusion of flow reversal, the original exit, now also an entrance was deeper.

The system also began forming a wider channel with two main arms instead of one.

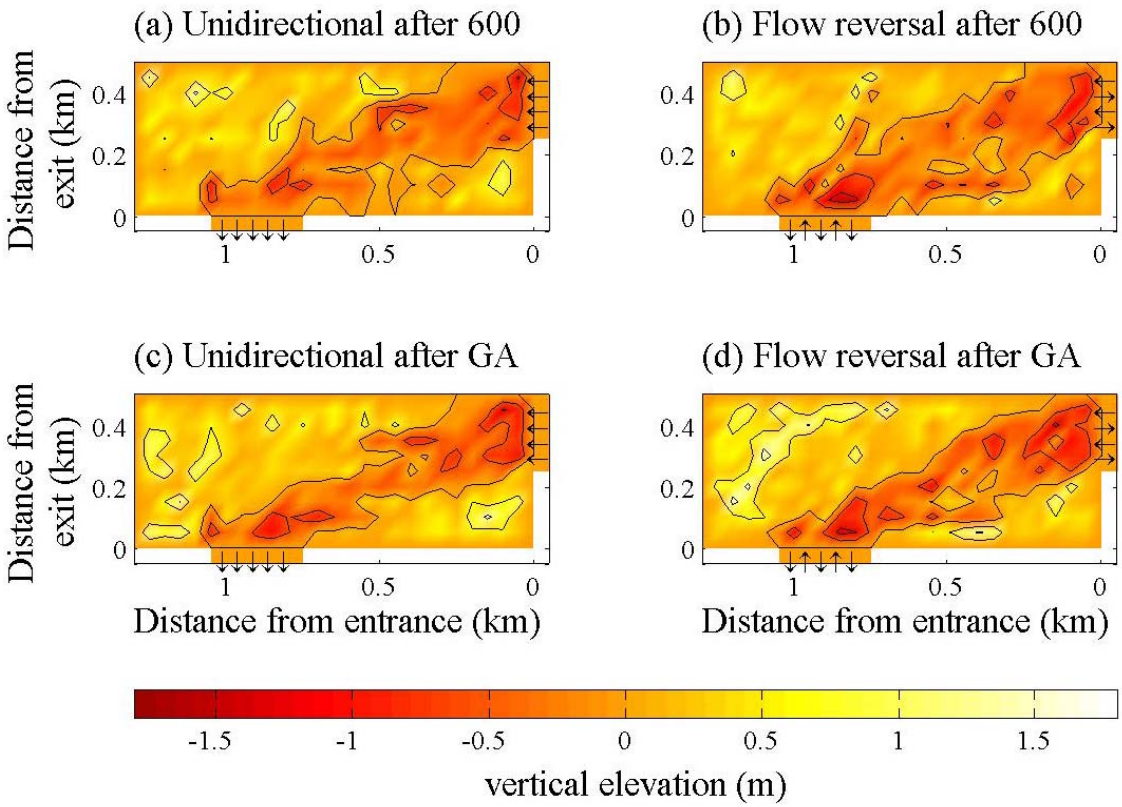

Figure 4.32 Best morphologies after 600 and 3000 generations for unidirectional and flow reversal GA optimisation.

After 3000 generations, the formation of a split channel containing islands with the inclusion of flow reversal was more pronounced. The channels at either end of the lagoon had deepened and banks of sand had formed in areas where flow was minimal. The morphologies of the best objective function value after the GA optimisation can be seen in Figure 4.32 (c) and (d).

The inclusion of flow reversal allowed a wider channel to form between the exits and entrances. It also formed a system with greater depositional mounds in the areas unaffected by the flow, thereby channelling the flow through the centre of the lagoon.

The results of further optimisation using an SA can be seen in Figure 4.33. The SA refined the solution, deepening the channel found using the GA. In the example where flow reversal was considered, deposition mounds helped to channel the water through the erosion pathway.

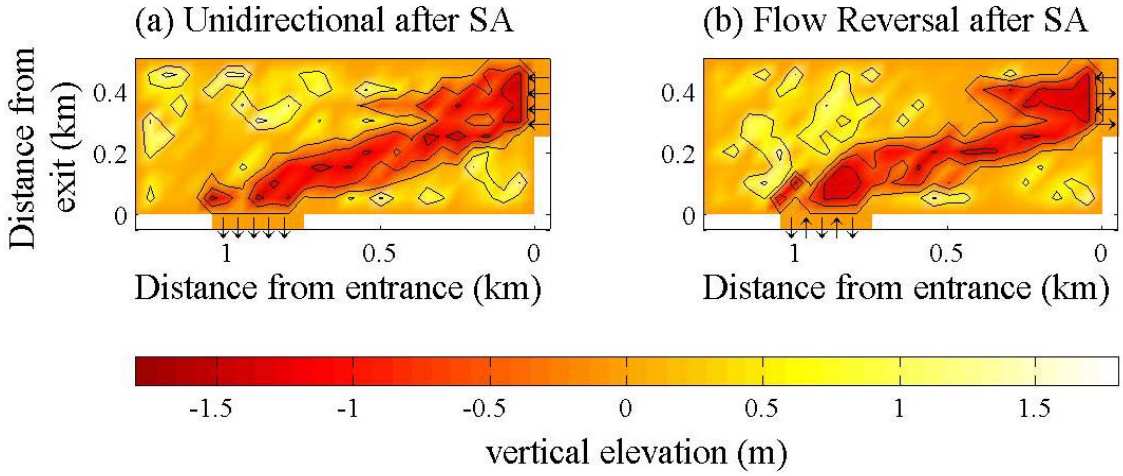

Figure 4.33 Best morphologies for unidirectional and flow reversal SA optimisation at final temperature.

Good qualitative agreement was found between the computer modelled morphology and the morphology found in the River Murray tidal inlet, the morphology which is shown in the aerial photograph in Figure 4.34.

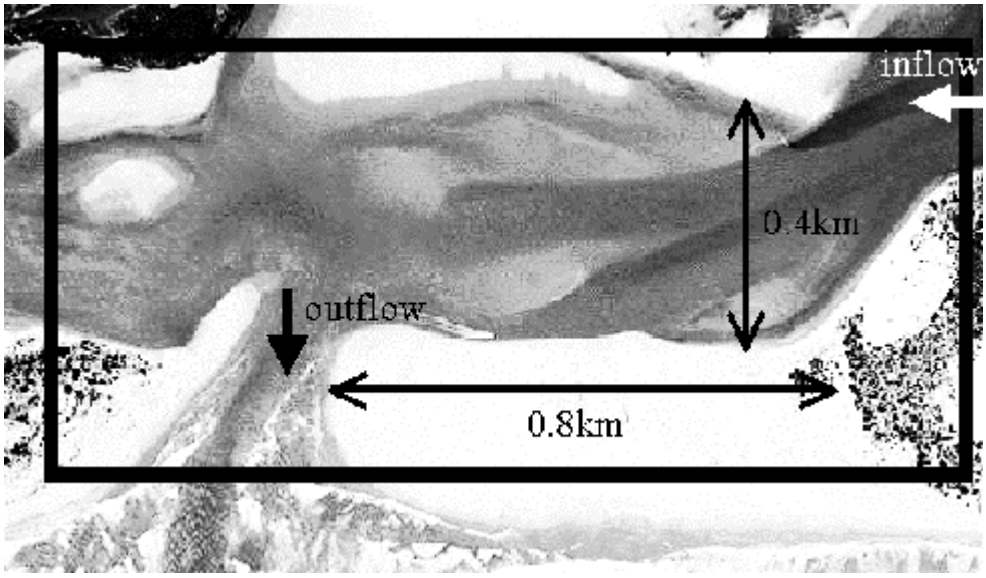

Figure 4.34 Aerial photograph of a small section of the River Murray mouth for comparison with the modelled morphology.

Figure 4.35 shows the descent of the objective function value for GA optimisation of a flow reversal lagoon system. As can be seen in this Figure, the objective function when combined with reversing flow was greater than that of the unidirectional flow. It was greater than double the unidirectional flow value, so the penalty for reverse flow was greater than flow in the initial direction. As in the unidirectional case, the path of the objective function value in the GA reached a plateau in the multidirectional example, where the GA optimisation

was unable to reduce greatly the objective function value. At this point the solution stagnated and a SA was applied to further refine the minimum energy expenditure morphology.

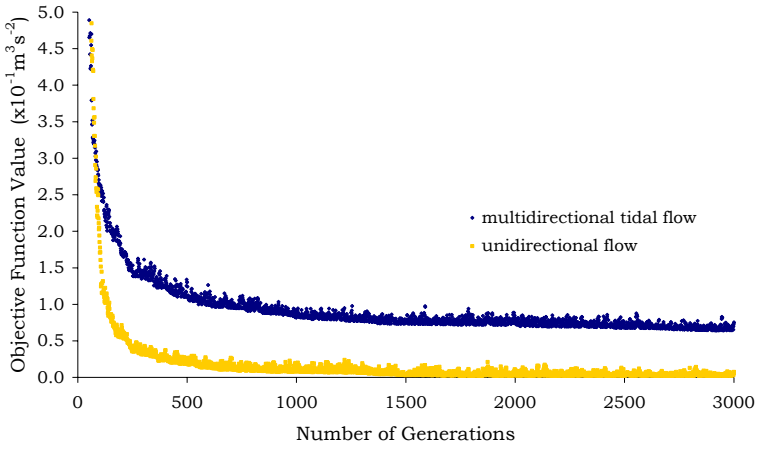

Figure 4.35 Path of Objective Function Value for GA optimisation.

The SA was able to reduce the objective function obtained by the GA optimisation. Initially, there was a small peak in objective function value, where the SA accepted worse fitness solutions. The SA was then able to refine the morphology, as can be seen in the plots of the SA optimised morphology, where the SA enabled the formation of a deeper and wider channel than the initial GA optimised result. The path of the SA objective function value results are shown in Figure 4.36.

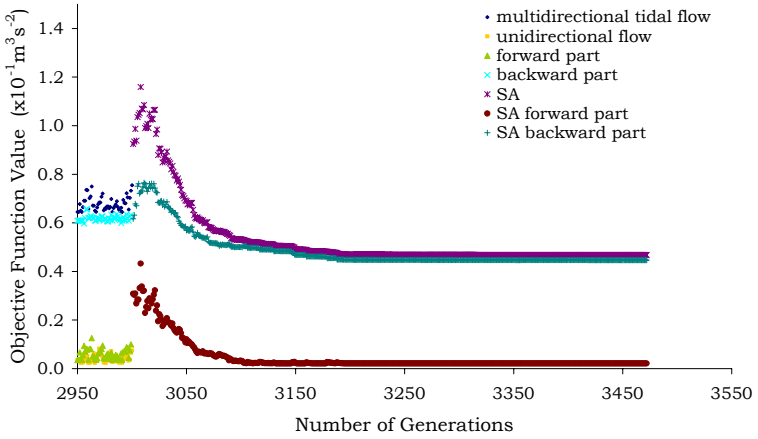

Figure 4.36 Path of Objective Function Value for SA optimisation.

During the second SA optimisation (see Figure 4.37), the channel deepened, particularly at the entrances/exits and a greater amount of deposition occurred in the areas of low flow. This is what would be expected in nature, as low energy flow would deposit sediment it no longer has the energy to carry. This would act

as positive feedback, reinforcing the already developing channel and concentrating the main flow away from the deposition area. Particularly in the area where flow enters from the sea side, there is a large deposit of sand, where the initial energy of the water racing through the lagoon mouth is lost and large amounts of sand carried into the system are deposited. In this further optimisation, the overshoot of the channel has been diminished, with the flow being directed more strongly on a path to the exit.

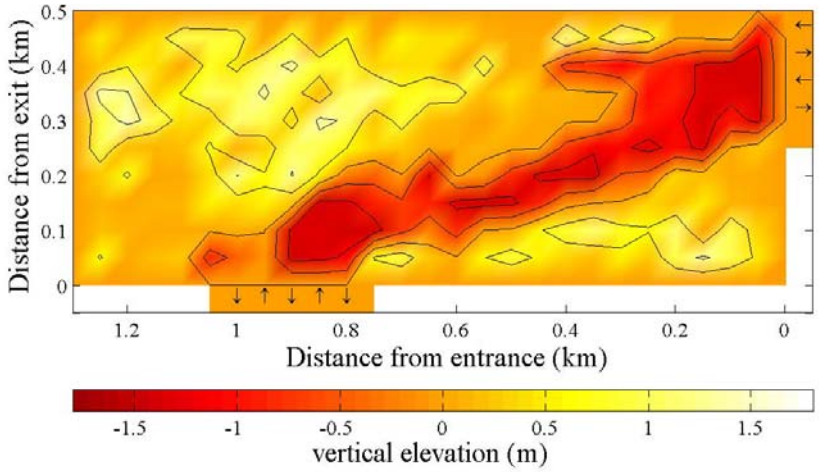

Figure 4.37 Best morphology for flow reversal second SA optimisation.

The way the optimisation operates, from many grid points all starting in random positions, means that some points may take longer to form the correct morphology with regard to those around them, and some kind of smoothing effect may be necessary to reduce the pitting or rippling effect present over some areas of the morphology.

The next step in this branch of research, not covered in the current PhD study scope, would be to include some kind of sediment budget function, limiting the system to the correct amount of sediment loss/gain.

4.5 Application to a Lagoon Laboratory Experiment

A different approach was also studied to test the limits and applicability of the entropy-based method to two-dimensional morphological modelling. Instead of the large scale field situations discussed previously in this Chapter, a small scale laboratory sized application was analysed.

The situation analysed included fixed narrow entrance and exit channels diagonally opposite each other, with a wide moveable sand bed between them.

This is show in Figure 4.38. Flow was unidirectional as in the initial lagoon situation investigated in Section 4.3.

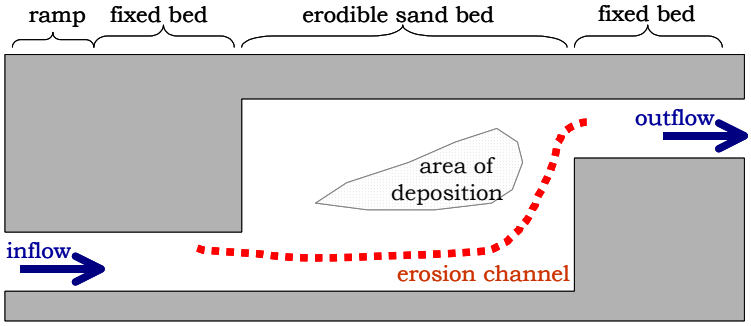

Figure 4.38 Laboratory lagoon setup description.

The laboratory study description and results with which to compare the model are detailed in the Section below. This is followed by a description of the results found when a traditional process-based model was utilised. The entropy-based modelling method modification requirements are then detailed followed by the results obtained using this method and a comparison between the abilities of the entropy-based and traditional model to predict the results acquired in the laboratory.

4.5.1 Laboratory Experiment

A laboratory study was undertaken in a unidirectional flow flume, with a lagoon type situation installed in its centre, to assess the applicability of an entropy-based model to a small-scale morphological equilibrium prediction.

A 1.8m long movable sand bed was placed in the centre of a flume 10m long and 0.914m wide. Narrow fixed beds were constructed at the entrance and exit, each one metre long by 265mm and 260mm wide respectively, placed diagonally opposite each other at either end of the sand bed. Fine sand with a d_{50} of 240µm filled the boxed area to a depth of 150mm. Before each experiment, the bed was flattened and levelled, using a levelling block and spirit level, to create a negligible bed slope in both the longitudinal and transverse directions. A sketch of the laboratory setup can be seen in Figure 4.39.

The flume was gradually filled with water, with care taken to ensure the sand on the bed was not disturbed. Once water in the flume reached a pre-determined level, the flow was increased to the desired rate. Flow entered the boxed area at approximately 4.8Ls^{-1}. The exit depth of the box was 52mm.

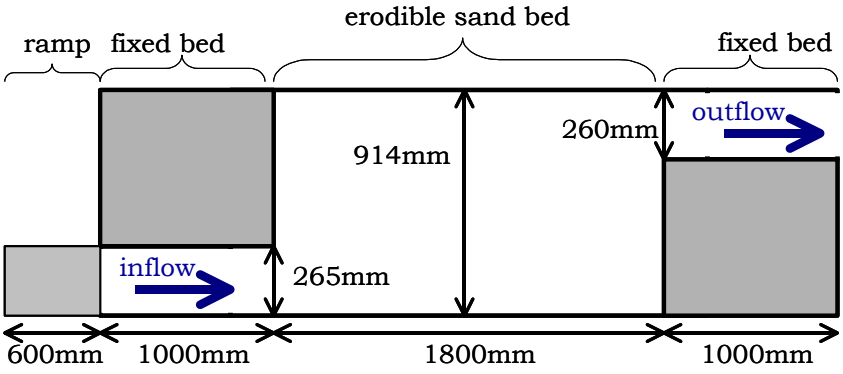

Figure 4.39 Laboratory setup for a lagoon experiment.

The experiment was left to evolve for 75 hours, until an equilibrium morphology had developed. A number of shorter experiments were investigated, but it was found that the time required for equilibrium to be reached was greater than the 6 to 8 hours that these experiments were run for. Some results of these short experiments are given in Appendix D. Once the system had been left for sufficient time for an equilibrium state to develop, the flow was stopped and the morphology was measured using a laser scanner. The scanner used was the same one used in the constriction and obstruction experiments as described in Chapter Three, Sections 3.2 and 3.5. The scanner measured points every 2mm longwise and 20mm across the flume, with a vertical accuracy of ± 0.2mm. A scan of the equilibrium morphology can be seen in Figure 4.40.

A channel developed along the wall adjacent to the flow entrance, bending around to the flow exit when it reached the back wall. To one side of the channel, a rippled bed was formed. The ripples were small and appeared to be stationary. Enough sediment had been eroded and left the system, so that the remaining morphology was at elevations where velocities were below the critical velocities required to erode sediment. There were no steep sloped areas in this configuration, apart from near the exit channel, where the critical angle of repose may have had an effect in shaping the steepness of the erosion hole, and allowing avalanches to transport sediment towards the exit channel. The area on the opposite side to the entrance channel changed little from its original morphology. This is because the flow here was very weak and unable to reach a critical velocity to move sediment, even at the beginning of the experiment.

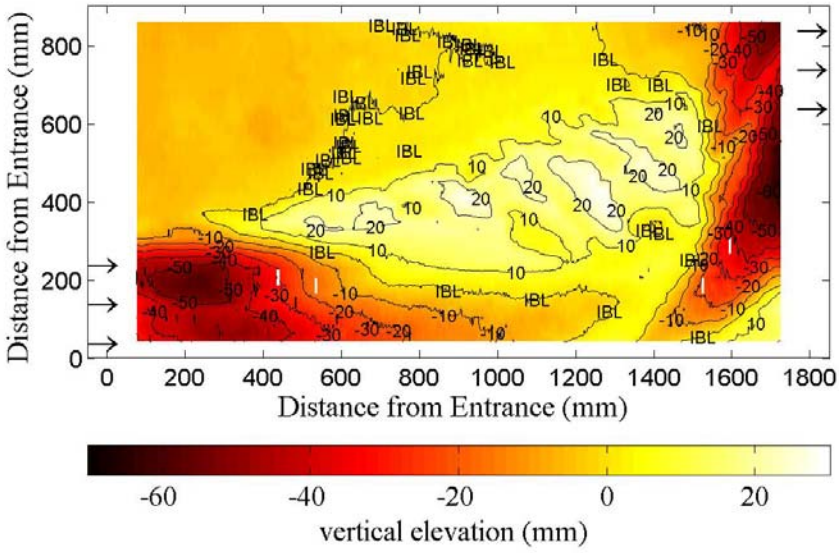

Figure 4.40 Equilibrium morphology after 75 hours of flow down the flume. IBL represents the initial average bed level.

4.5.2 Traditional Modelling Results

The set-up observed in the laboratory was analysed using a traditional processed-based medium-term morphological model. The model consisted of a Navier-Stokes flow model and a sediment transport routine. The flow model was the same one that was described in Section 3.3. The equilibrium morphology was determined when the model had been applied over a sufficient time that the rate of bed change reached a nominal cut-off value. The model used a relaxed critical velocity criterion, where the velocity value required to initiate motion was halved. This enabled the channel to fully develop.

The morphology predicted by the traditional model is shown in Figure 4.41. The area of erosion near the lagoon entrance and exit was predicted well, but the depositional mound is not far enough away from the entrance and the erosion channel is moving towards the centre of the lagoon, rather than being along one side as it was observed in the laboratory. One possible reason for this might be the inability of the traditional model to correct its flow pattern after initial movement to the left due to a poor representation of the jet leaving the entrance.

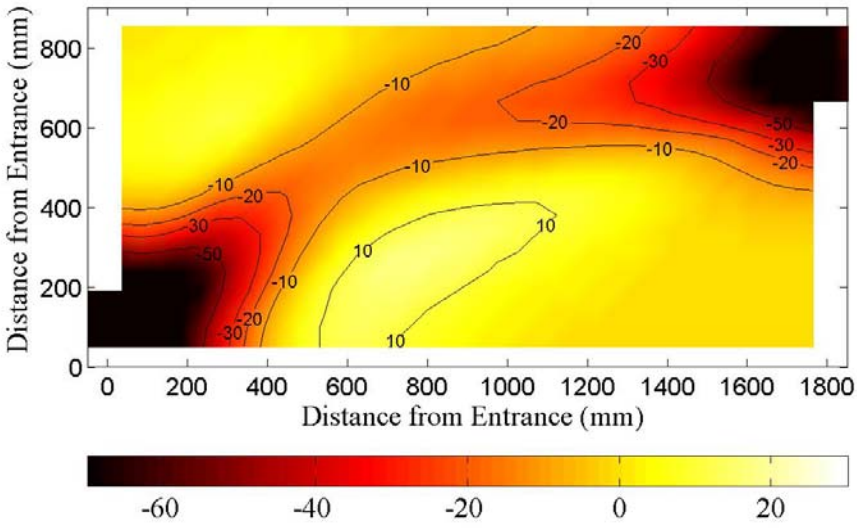

Figure 4.41 Equilibrium morphology prediction using traditional process-based model. The initial average bed level was zero.

4.5.3 Modelling Modifications

As an alternative to traditional process-based models, a method based on entropy principles was developed, using methodology similar to that of the large scale unidirectional lagoon model, discussed in Section 4.1 earlier, but with a modification to the objective function to include the effect of angle of repose avalanching. Avalanching is a phenomenon discussed in Chapter Three, Section 3.1, where it is included in the cellular automata models using a rule. Here it is included via a penalty as discussed in the remainder of this Section.

The objective function used to derive the equilibrium morphologies can be broken down into three components. Similarly to the field-sized lagoon example, global energy dissipation and local excess erosion and deposition penalties are considered. Section 4.1 has details of these objective function components (Eq. 4.6). The third component of the objective function is a physical constraint, with a penalty applied if the angle between two predicted elevations is greater than the critical angle of repose, as this represents a physically unrealistic situation.

The third, additional component was not required in the large-scale lagoon modelling, as the elevation difference was small relative to the grid size. However, in the small scale laboratory case study the possible elevation difference, determined by the limits, was greater than the angle of repose, so had to be controlled by a penalty. To evaluate the penalty, the angles between neighbouring elevation points were calculated and compared to the critical angle of repose. In areas where the angle between two elevations does exceed the angle

of repose, a penalty is applied. The penalty is calculated as the sum of the excess elevation differences above the value that equals that of the critical angle of repose (see Eq. 4.20). The overall objective function is then calculated (Eq. 4.21). Again, a scaling factor is used, in order to reduce computation time and allow more efficient optimisation, as discussed in Section 4.1.

$$P_a = \theta \sum \Delta z - \Delta z_{crit} \qquad \text{for } \Delta z > \Delta z_{crit} \tag{4.20}$$

$$OF = E_s \times P_d \times P_e \times P_a \tag{4.21}$$

where P_a is the excess steepness penalty (m);

Δz_{crit} is the elevation difference that corresponds to a critical angle of repose (m); and

θ is a scaling factor.

4.5.4 Modelling Results and Laboratory Comparison

Global search methods were used, namely a combined GA and SA, to minimise the overall objective function value (Eq. 4.21). An RGA was initially used to narrow the search space to an area where the pinpointing of a near optimal solution was possible and the search had stagnated, as can be seen in Figure 4.42 and Figure 4.43. The GA used tournament selection, average crossover, random mutation and elitist selection. The crossover and mutation probability values used were 0.8 and 0.01 respectively.

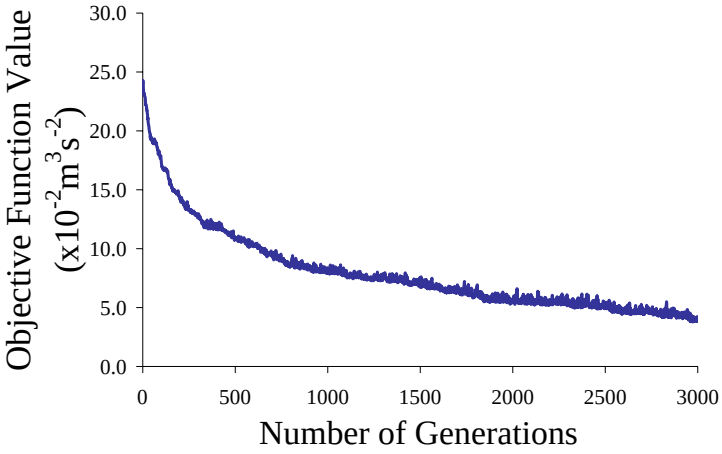

Figure 4.42 Path of objective function using a GA optimisation routine.

Next, an SA was applied to the best GA bed profile, until a near optimal morphology with minimal energy dissipation was found. Figure 4.42 portrays the path of the objective function during the GA optimisation. Initially, the value was reduced dramatically, with the GA quickly discarding energy inefficient solutions.

It can be seen that after 1000 generations, the rate of improvement of the solution fitness was minimal, with the objective function value reducing only slightly after 2000 generations. Figure 4.43 shows the ability of the SA to further refine the solution, enabling an energy efficient morphology to be deduced, from a point where the GA was no longer actively reducing the objective function value.

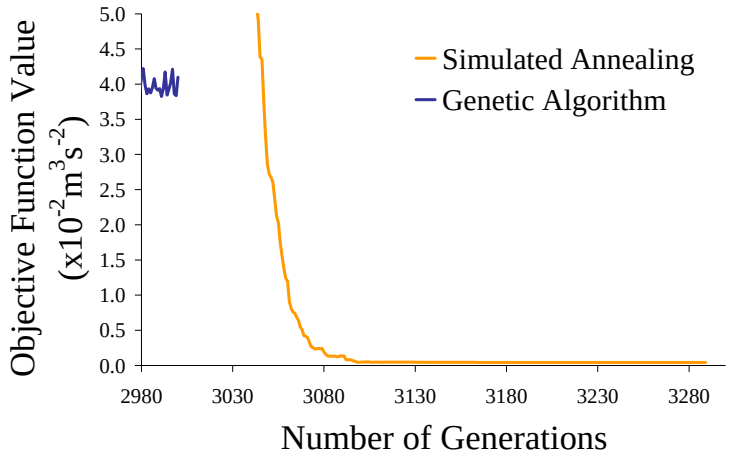

Figure 4.43 Further enhancement of the minimum objective function by the SA, after the stagnation of the GA routine.

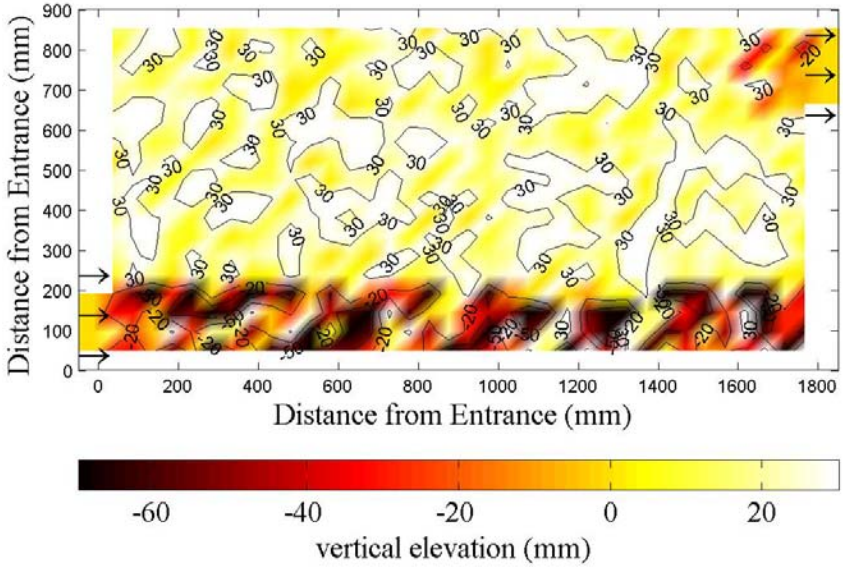

Figure 4.44 Equilibrium morphology prediction after SA optimisation.

The morphology found after the GA and SA stages for a simplified lagoon with unidirectional flow can be seen in Figure 4.44. The SA was able to refine the initial GA morphology, further defining the channel formed along the side of the lagoon. It was also able to rid the solution of areas of erosion where there was little flow, on the opposite side of the lagoon to the entrance. The erosion

channel was predicted by the entropy-based model to occur in the correct position with respect to the laboratory results, but excess deposition was predicted due to the formulation of the objective function as mentioned earlier in Section 4.3. The discontinuity in the erosion channel, where some slight deposition occurs between erosion pockets, is due to the coarseness of the grid size used. This discontinuity became more prominent if the coarseness of the grid was increased. The computation time required for a finer grid was uneconomical.

4.5.5 Discussion and Conclusions

The application of entropy-based modelling method has the ability to predict equilibrium morphologies of a laboratory unidirectional lagoon situation, without the need for time-stepping and the errors that may be associated with this process.

With the inclusion of the physical constraint on the maximum acceptable slope, the morphology formed was more realistic and smoother. The availability of excess sediment due to the bounds of the randomly derived elevations accounts for the excess sediment deposition in areas of little flow, after the SA optimisation. The addition of a sediment balance penalty to the objective function may help alleviate this, but would also increase the computational search time involved and is beyond the scope of the current research. The inclusion of other physical constraints would also be beneficial in further studies of this kind.

The model is better able to predict erosional channels than it is the rippled deposition mound. This is because the energy description and penalty inclusions are aimed at modelling erosional environment, where sediment is removed from the system, but is not readily available to enter the system. This is where they have been used in the past for longitudinal river profiles (Leopold and Langbein, 1962) or river delta profiles (Wright et al., 1973). The van Rijn (1984) critical velocity equation is used to determine where sediment will be eroded, but no similar equivalent is used to determine if sediment will be deposited. As the model is not time based, there is no exact relationship as to where eroded sediment will be deposited. The prediction of depositional mounds is an area where further study is required and is beyond the scope of this PhD research. A number of methods were trialled to investigate ways to include deposition in the objective function, including the use of the Exner equation (Eq. 4.22). Due to the Exner equation's reliance on time, or the servere limiting of the solution space by the inclusion of a sediment balance penalty, these trials were unsuccessful.

$$\left(1-\lambda_p\right)\frac{\partial \eta}{\partial t} = -\frac{\partial q_t}{\partial x} \qquad (4.22)$$

where λ_p is the bed porosity;

η is the bed elevation (m);

t is time (s);

q_t is volumetric total sediment transport rate per unit streamwidth (m^2s^{-1}); and

x is the streamwise distance (m).

The entropy method is better able to predict the erosion channel that forms along one side of the lagoon, than the conventional model that it was compared to. Both use the same hydrodynamic model, but the entropy-based model is able to introduce sediment deposition mounds and erosion holes randomly throughout the bed until an energy efficient solution is found. The conventional model starts off with flow moving in slightly the wrong direction, sideways rather than in a forward jet, and is not able to correct itself, falsely making a channel move along the shortest route possible to the exit of the lagoon. This shows that the entropy-based method is superior in its prediction of erosion, as it is able to use the same tools as a traditional model, but predicts a more realistic erosion channel.

This study is the first step to make a comparison between entropy-based equilibrium modelling and laboratory simulations. The narrowness of the flume used in the laboratory tests may have contributed to an enhancement of three-dimensional processes, not predicted by the two-dimensional hydraulic solver. The inclusion of a three-dimensional model to improve on the energy calculations in the objective function could improve the quantitative agreement between the two models but is beyond the scope of this initial investigative study.

4.6 Summary

Inspired by river morphological researchers, that have in the past been able to predict large-scale river profiles or planform river patterns, a method has been developed utilising entropy-based principles in conjunction with physical penalties in an optimisation framework, to predict small-scale equilibrium morphologies associated with lagoons both in the field and laboratory. The work discussed in this Chapter has made use of entropy-based principles to define models that do not time-step, but make use of direct comparisons between morphologies based on objective functions to determine which morphologies are closer to equilibrium. Three difference examples have been studied, a laboratory

sized unidirectional lagoon, a unidirectional field-sized lagoon and a field-sized lagoon with flow reversal.

Firstly, an objective function (Eq. 4.6) was developed and outlined in detail, utilising entropy-based principles and physically based penalties. The objective function minimised energy dissipation globally throughout the system by considering the energy lost between water entering and exiting the lagoon. Physical penalties were applied if velocities were too high over deposition mounds or too low in erosion holes, based on the van Rijn (1984) critical velocity equation. Thirdly, in laboratory scale formulations, a third physical penalty was introduced, to prevent the optimisation routine from choosing solutions where the angle between two grid cells violated a maximum angle of repose.

The velocities of random morphologies were then generated over the case study being analysed, and objective function values calculated. These random morphologies were then optimised using global search techniques, namely GAs and SA, to pinpoint near optimal solutions corresponding to equilibrium morphologies.

The method was found to predict the erosion channels well, as this is similar to the erosion environment the method was originally designed for in river morphology. The method differed from river morphological attempts as it examined morphology in a lot finer detail, added extra complexity in its optimisation methods and looked at the complete morphology, rather than the one-dimensional profile or planform of rivers. It also included physical constraints in its optimisation method, a further improvement on past river research.

In the field of coastal modelling, this research adds a new way of predicting equilibrium morphologies that differs from the traditional process-based time-step model approach. The new method utilises optimisation and objective functions that describe the environments most likely configuration at equilibrium. The method predicts morphologies directly, rather than moving towards them gradually. It eliminates the possibility of errors being introduced with each sediment transport calculation over a single time-step, as no sediment transport occurs, snapshots of morphologies are compared directly to decipher which is closer to equilibrium.

In the next Chapter the use of optimisation to compare potential equilibrium morphologies directly is applied to a coastal breakwater example.