Using Genetic Algorithms to Optimise Current Drive in STEP

S.P. Marsden, F.J. Casson, S. Freethy, T. Wilson, B. Patel, E. Tholerus
CCFE, Culham Science Centre, Abingdon, OX14 3DB, UK

The Spherical Tokamak for Energy Production (STEP) programme aims to deliver a commercially viable fusion energy plant. The reactor will be fully non-inductive, utilising microwave heating in the form of Electron Cyclotron Radio Heating (ECRH) and Electron Bernstein Wave (EBW) heating. Current will be driven by the microwave heating, with 80% to 90% of the total current being driven by the plasma pressure.

The safety factor, q, is a function of the plasma current and an important profile to control. The microwave heating profile must thus be designed to generate a favourable q profile. The process of optimising the microwave heating profile is a highly non-linear problem as changing the power deposition affects the local temperature and density, which in turn affect the current drive efficiency and bootstrap current.

Simulations were performed using JETTO to calculate the current driven and the fully diffused self consistent current profile for a given ECRH power deposition. The time required for each of these simulations inhibits the use of traditional optimisation techniques that rely on taking many sequential steps in parameter space.

A Genetic Algorithm (GA) is an optimisation method inspired by natural selection. A population of points in parameter space are considered in parallel. Those which are judged by the algorithm to have performed the best are combined to create the next generation. The process is iterated until a sufficiently optimal solution is found.

Performing a simulation in JETTO requires of order hours. As GAs test a number of points in parameter space simultaneously, they are able to make use of parallel processing to quickly converge on a solution.

In order to optimise the microwave heating profiles, they first needed to be broken down into a set of parameters which could be treated analogous to genes. This was done by defining a number of points and linearly interpolating between each pair of neighbours. The profile was then normalised to a predefined heating power. Overall this parameterisation gave an unrealistically angular distribution, but was sufficient at this stage with a more refinement performed at a later stage.

The initial population provided to the GA were randomly generated within predefined limits. Figure 1 shows the initially generated ECRH profiles and their resultant q, Electron Cyclotron

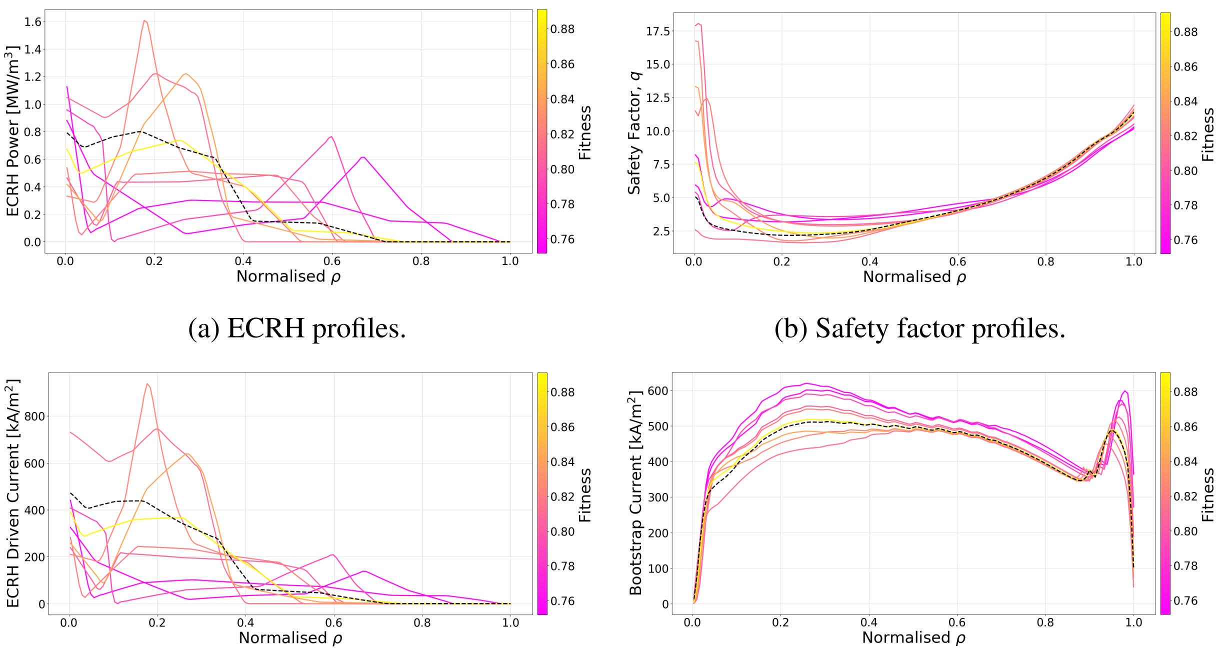

(a) ECRH profiles. (b) Safety factor profiles.

(c) ECCD profiles. (d) Bootstrap current profiles.

Figure 1: Initial randomly generated ECRH profiles and the resulting q, ECCD, and bootstrap current profiles.

Current Drive (ECCD), and bootstrap current profiles.

After each generation it was necessary for the GA to be able to rank the members of the population. The best performing individuals being bred to create the next generation. There were many features of the q profile which were desired, for instance the location and value of the minimum. A fitness was assigned separately to each of these features between 0 (the worst case scenario) and 1 (the ideal case). The final fitness function was then defined as a weighted average of all of the individual fitness elements. The fitness function and a breakdown the highest weighted constituent elements is shown in fig. 2.

The best performing 20% of runs from a generation were bred to create the next generation. For each member of the next generation between one and three parents were selected. For each parameter of the ECRH profile, the value was taken from one of the parents at random. Additionally each parameter had a probability of mutating, whereby the value would change by a small amount. This mutation was used as a means of further exploring parameter space. Were this mutation too small the final result would not reflect a global maximum fitness, were it too large the GA would fail to converge at all.

Using 30 individuals per generation the GA was able to converge on a sufficiently optimal solution in fewer than 10 generations. Figure 3 shows the fitness value of the best performing individual from each generation. The optimal ECRH profiles and resulting q profiles after 10

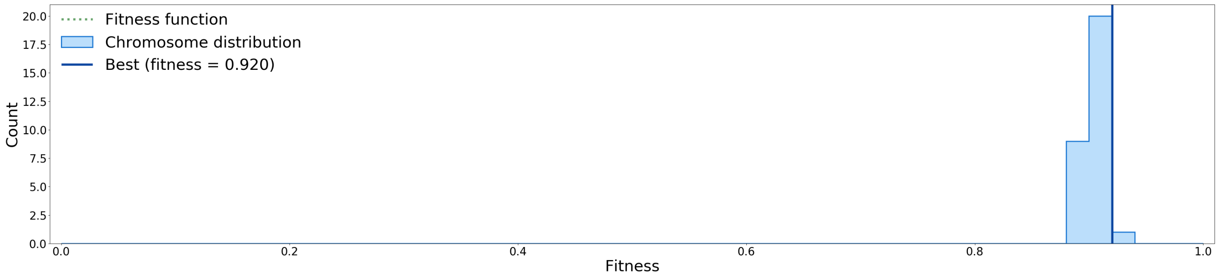

(a) The combined fitness.

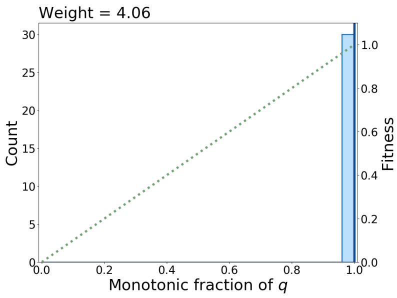

(b) The fraction of ρ over which q is increasing.

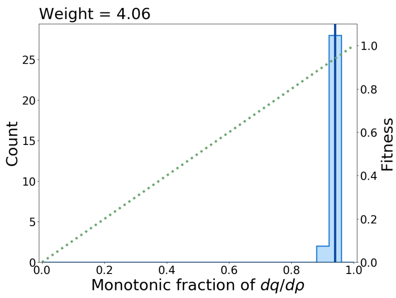

(c) The fraction of ρ over with $\frac{dq}{d\rho}$ is increaing.

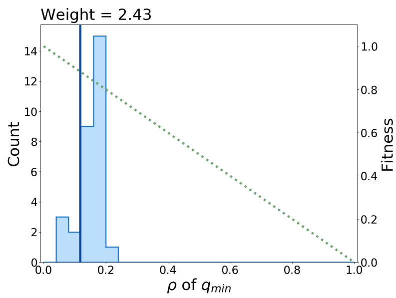

(d) The value of ρ at which q_{min} is acheived.

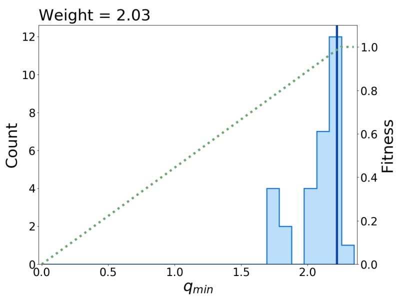

(e) The minimum value of q.

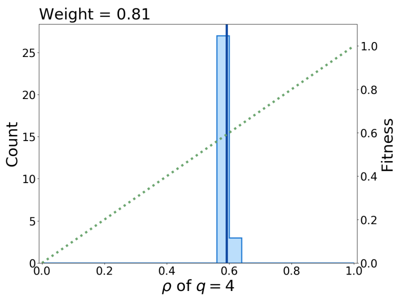

(f) The value of ρ at which the tail of the q profile exceeds 4.

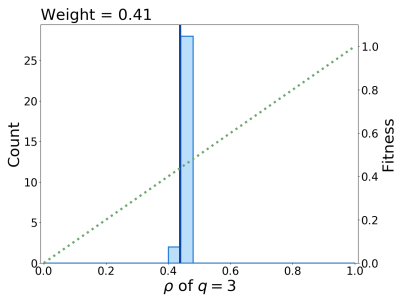

(g) The value of ρ at which the tail of the q profile exceeds 3.

Figure 2: The combined fitness function and the highest weighted fitness elements. The histograms show the values achieved by each member in the population after five generations. The member with the highest combined fitness is indicated in each of the fitness elements with a solid line. The dotted line shows the fitness awarded for each element.

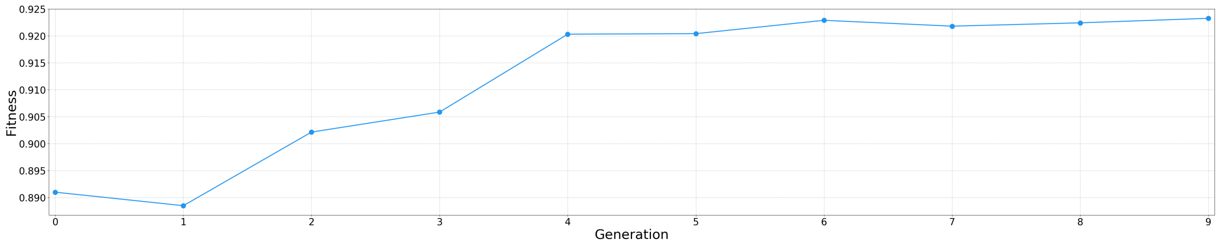

Figure 3: The fitness of the best performing individual from each generation.

generations are shown in fig. 4. Though this required hundreds of simulations to be run, depending on JETTO settings the optimisation to be completed within a day thanks to the parallel nature of the GA.

In order to ensure that the solution reflected a global optimum, the random seed responsible for generating the initial generation was varied. In each instance very similar optimal solutions were reached. This demonstrated that the mutation rate was sufficiently large to allow for the parameter space to be adequately explored.

Shown here is the purely ECRH optimisation. This GA implementation is currently being applied in the case of a combined ECRH and EBW scenario, optimising both heating methods simultaneously. The GA was designed in a general manner and allows any JETTO inputs to be optimised and any JETTO outputs to be used in the fitness function. Currently pellet injection parameters and total fusion power are under consideration as parameters which may benefit from this tool's application.

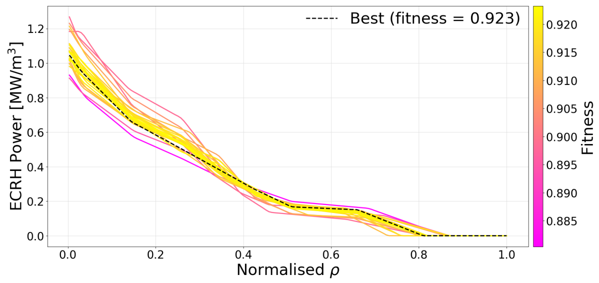

(a) ECRH profiles.

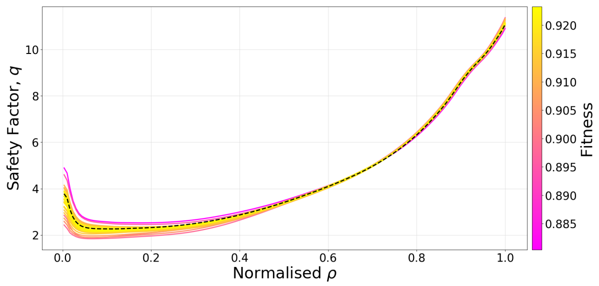

(b) Safety factor profiles.

Figure 4: The ECRH and q profiles after 10 generations. The individual with the highest fitness score is indicated by the dotted line.