This is a peer-reviewed, author's accepted mansucript of the following conference output: Onslow, M., & Stock, A. (2024).
Genetic algorithms for design of optimal velocity tracking controllers including PTO efficiencies. In 2024 UKACC 14th International Conference on Contro
(UKACC International Conference on Control (CONTROL); Vol. 2024). IEEE.
https://doi.org/10.1109/control60310.2024.10532033

# Genetic Algorithms for Design of Optimal Velocity Tracking Controllers Including PTO Efficiencies

Matthew Onslow
Centre for Doctoral Training in Wind and Marine Energy Systems and Structures,
Electrical and Electronic Engineering Department,
University of Strathclyde, Glasgow,
G1 1XQ, United Kingdom

Adam Stock
Institute of Mechanical, Process and Energy Engineering,
School of Engineering and Physical Sciences,
Heriot-Watt University, Edinburgh,
EH14 4AS, United Kingdom

***Abstract*— Genetic algorithms use the ideas of Darwinian evolutionary theory to find the optimal solution to a design problem. Here they are utilised in two scenarios. Firstly, finding the optimal power take-off (PTO) force for maximising the electrical power output of a device by accounting for PTO efficiencies. The genetic algorithm finds a solution marginally faster than a brute forcing method with the added benefit of not being constrained to a discrete grid of test points, hypothetically leading to a more accurate result.**

**Secondly, these optimal power take-off forces are used with another genetic algorithm to fit a transfer function for use as part of a previously designed adapted optimal velocity tracking controller that accounts for PTO efficiencies. Along with the reduced requirement for control engineering expertise, the resultant transfer function is found to have a smaller average phase error, when compared to a manually fitted transfer function. Simulations are undertaken that find that using a genetic algorithm derived transfer function results in approximately the same, or better energy capture when compared to the manually fitted transfer function, depending on the sea state, with the largest improvement being an increase of 5.93%. These methods form the basis of a potential control co-design methodology.**

## I. Introduction

It has been shown that efficient control of a wave energy converter (WEC) can increase its mechanical power output by 14 - 50% [1]; it is therefore important to design an effective controller for use in any WEC. The theoretical optimal control strategy is the so-called 'complex-conjugate control' [2]; however, implementation is difficult so various attempts have been made to approximate it [3], [4]. One approach is optimal velocity tracking (OVT) [5], in which the optimal velocity of the WEC to maximise energy capture is calculated and tracked via a disturbance rejection controller.

Figure 1 shows the structure of an OVT controller; here the WEC's position, velocity, and acceleration, shown collectively as $\underline{x}$, are inputted into a model of the WEC's intrinsic impedance, Z, that estimates the total force on the WEC, $F_Z = F_{\text{ex}} + F_{\text{PTO}}$. Subtracting the power take-off (PTO) force gives an estimate of the wave excitation force. This is then fed into an extended Kalman filter (EKF) to find estimates of both the excitation force amplitude and frequency, $\tilde{A}_{\text{ex}}$ and $\tilde{\omega}_{\text{ex}}$; these are used to set the gain of

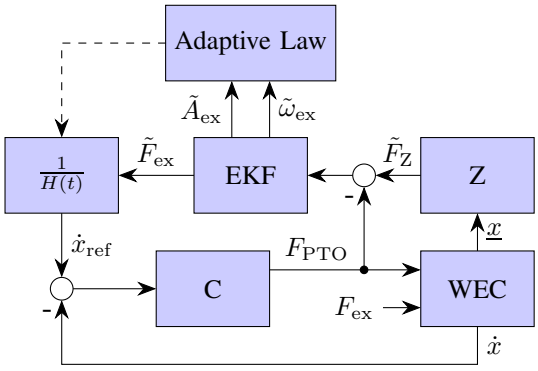


Fig. 1. Standard Optimal Velocity Tracking structure. Tildes represent estimated values. Adapted from [6].

$1/H(t)$ through an adaptive law, such that:

$$\dot{x}_{\text{ref}} = \frac{\tilde{F}_{\text{ex}}}{H(t)} \approx \frac{F_{\text{ex}}}{2B(\omega)} = \dot{x}_{\text{opt}}, \tag{1}$$

where $B(\omega)$ is the frequency dependant radiation resistance of the WEC, and $\dot{x}_{\text{opt}}$ is the theoretical optimum velocity. The controller C is then used to minimise the error between $\dot{x}_{\text{ref}}$ and $\dot{x}$. For a more detailed explanation of OVT, the reader is directed to [5], [6].

However, complex-conjugate control (and therefore OVT) works to optimise the mechanical energy capture of the device, but this does not necessarily result in the maximum possible electrical energy output once the efficiency of the PTO is taken into account. Prior work [7], developed a modified OVT controller that can account for these efficiencies and work towards optimal electrical power output. In this approach, the model of the intrinsic impedance, Z, is replaced with a model that instead gives an estimate of $H(t)v_{\text{ref}}$ where $v_{\text{ref}}$ is the optimal velocity for maximising electrical power output, this alternative model will be referred to as the 'whole system dynamics model' (WSDM), $M_v$.

Considering the regular wave case, a WEC's mechanical intrinsic admittance (the inverse of the mechanical impedance) is represented by $Y = 1/Z$, which, for a regular wave with frequency $\omega$, can be represented by a complex number $A_Y e^{i\phi_Y}$, where the amplitude $A_Y$ and phase $\phi_Y$

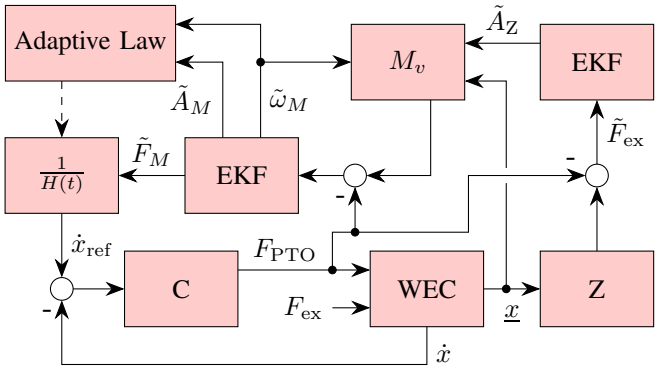


Fig. 2. Adapted Optimal Velocity Tracking structure. Tildes represent estimated values. Adapted from [7].

are functions of $\omega$. Defining the input excitation force in the time domain as $F_{\text{ex}} = A_{\text{ex}}e^{i\omega t}$, and the PTO force as $F_{\text{PTO}} = A_{\text{PTO}}e^{i(\omega t+\phi_{\text{PTO}})}$, the resultant WEC velocity is given by:

$$A_{\dot{x}}e^{i(\omega t+\phi_{\dot{x}})} = \left(A_{\text{ex}}e^{i\omega t} + A_{\text{PTO}}e^{i(\omega t+\phi_{\text{PTO}})}\right) A_Y e^{i\phi_Y}, \tag{2}$$

and the electrical energy output of the WEC is:

$$E_{\text{elec}} = -F_{\text{PTO}}\dot{x} \cdot \eta(\dot{x}, F_{\text{PTO}}), \tag{3}$$

where $\eta$ is a lookup table of the efficiency of the PTO, here it is a function of $\dot{x}$ and $F_{\text{PTO}}$ since they are the main drivers of efficiency for electrical drives. For any given frequency, there is a combination of $A_{\text{PTO}}$ and $\phi_{\text{PTO}}$ that maximises, $E_{\text{elec}}$. The value of the WDSM at that frequency can then be found from:

$$\frac{1}{M_v} = \frac{1}{2B(\omega)} + \frac{A_{\text{PTO}}}{A_{\dot{x}}}e^{i(\phi_{\text{PTO}}-\phi_{\dot{x}})}. \tag{4}$$

Repeating this for a range of frequencies can provide a series of reference points that a transfer function for $M_v$ can be fitted to. The adapted OVT controller both increased electrical power output, and reduced loads on the WEC, which could allow for reductions in capital and operational expenditure [7]. The adapted structure is shown in Figure 2.

In [7] the model $M_v$ is derived using a 'brute force' approach, in which a wide range of $A_{\text{PTO}}$ and $\phi_{\text{PTO}}$ values are tested to find the best pairing. The model $M_v$ is then approximated in transfer function form through manual fitting. The goal of this work was to automate this process with an optimisation algorithm; one example of these is genetic algorithms (GAs), which have previously been applied to optimise other components of a WEC [8]. This work seeks to investigate the application of GAs to another component of a WEC.

Analytical methods exist for fitting transfer functions in literature [9], and MATLAB has functions available [10]. However these methods tend to favour the fit of the gain over the phase, and it was found in [7] that the fitting of the phase was more important for the overall performance of the controller.

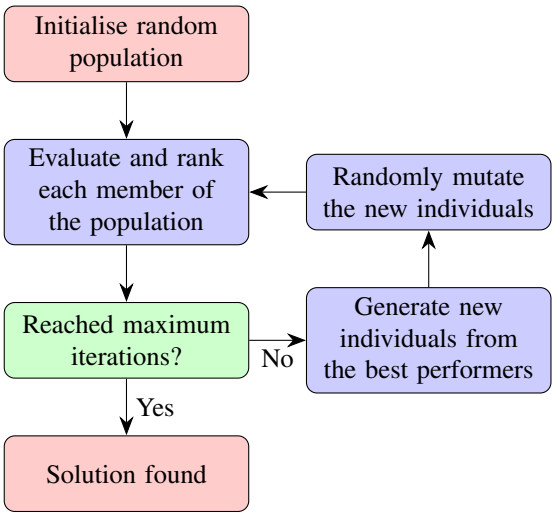


Fig. 3. Simplified routine for computing a genetic algorithm.

The rest of this paper is structured as follows: Section II describes the methodology of creating a GA, and specific considerations for both uses; Section III analyses the performance of the GAs; finally, Section IV gives some conclusions and discussion on the application of GAs and co-design to wave energy devices.

## II. Methodology

### A. The simulated WEC

The WEC used in this work is from the IMPACT toolbox [11], 'WEC2' is a large (6.5 MW), single-body, submerged heaving point absorber, it is the same WEC used in [7], and was chosen for ease of comparison of methodological differences.

Alongside the WEC, a model of the Trident PowerPod [12] PTO is used; it is not designed to be used with this WEC nor designed for use with an OVT controller. Instead it is selected to highlight the challenges of PTO efficiency, control for electrical power output and co-design [7]. To provide the required power to match the WEC, 18 sets of six Trident LGF30 PTOs were used, allowing for a peak force of 3.24 MN, with a maximum stroke of 10 m and a maximum allowable velocity of $\pm$ 2 m/s.

### B. Genetic Algorithms

A GA uses the ideas of Darwinian evolutionary theory to iterate towards an optimal solution. Each member of the population of possible solutions has a genome, in which each gene corresponds to a design parameter to be optimised (for example in the case of optimising the PTO force, the amplitude and phase of that force would each be genes), and a 'cost' which is used to determine the performance of that individual's genome. A simplified structure is shown in Figure 3, with a more detailed description of this process to follow.

First, the initial population is randomised. Then on each iteration, two parents are chosen to have their genomes merged to create new individuals. The selection criteria used

here is called 'roulette-wheel selection', in this any individual has the chance of being selected but the better performing individuals have better chances of doing so. The weighting for individual '$a$' is:

$$W_a = e^{-\beta \bar{C}_a}, \tag{5}$$

where $\beta$ is a selection pressure factor that can be chosen to further encourage the selection of better performing individuals, and

$$\bar{C}_a = \frac{C_a}{\sum_{i=1}^{N_{\text{pop}}} C_i}, \tag{6}$$

where $C_a$ is the cost of individual $a$, and $N_{\text{pop}}$ is the number of individuals in the population.

There is a chance that a GA fails to find the optimal solution; one common cause for this is the population coalescing around a local minima. From testing, it was found that this problem was a common occurrence for the use case; to rectify this, an 'incest prevention algorithm' (IPA) is added [13]. The IPA is used after the selection of the first parent in the crossover stage. Every individual's genome is compared to that of the first selected parent to create a similarity score, for individual '$a$':

$$S_a = \frac{\left(\sum_{i=1}^{N_{\text{gene}}} |g_{a,i} - p_i|\right)^3}{\sum_{j=1}^{N_{\text{pop}}} \sum_{k=1}^{N_{\text{gene}}} |g_{j,k} - p_k|}, \tag{7}$$

where $g_{a,i}$ is the i-th gene of individual $a$, $p_i$ is the i-th gene of the first selected parent, and $N_{\text{gene}}$ is the number of genes in a genome. The final weightings for the selection of the second parent are then $S_a W_a$. The inclusion of the IPA works to ensure that the population is not overran with lots of very similar individuals from the repeated breeding of two similar, fairly well performing individuals.

Once the two parents are selected, they are used in a process known as crossover, to create a new individual; the method used in this work is known as uniform crossover, in which two children are created, their genomes are:

$$\begin{aligned} G_{\text{child1}} &= G_{\text{parent1}} \cdot \mathrm{X} + G_{\text{parent2}} \cdot (1 - \mathrm{X}) \\ G_{\text{child2}} &= G_{\text{parent1}} \cdot (1 - \mathrm{X}) + G_{\text{parent2}} \cdot \mathrm{X} \end{aligned} \tag{8}$$

where X is randomly chosen from a uniform distribution of range $[-\gamma, 1 + \gamma]$. The extension $\gamma$ is used to encourage further exploration of the design space to help speed up the algorithm. To further help mitigate against stagnation, a child's genome is compared to its parents, before being added to the general population, and should it not have a cost T% better than both of them, the child is discarded. The cost improvement requirement, T, is linearly reduced every iteration such that $\mathrm{T} = 0$ on the final iteration, to allow for finer improvements towards the end of the optimisation.

Following crossover, the new individuals are then 'mutated' to further encourage exploration of the design space. Each gene within the genome has the chance to be mutated, according to the mutation rate; if selected, they are randomly mutated on a normal distribution with a mean of the gene's original value and a given standard deviation. A summary of all the parameters used in both GAs are shown in Table I.

TABLE I
SUMMARY OF THE GENETIC ALGORITHM PARAMETERS USED

| **Parameter** | **Value** | |
|---|---|---|
| | Optimal $F_{\text{PTO}}$ | WSDM fitting |
| Population size | 60 | 150 |
| Number of genes | 2 | 15 |
| Gene bounds | $A_{\text{PTO}} : [0, 3.24 \cdot 10^6]$<br>$\phi_{\text{PTO}} : [0, 2\pi]$ | $[-10, 0]$ |
| Maximum iterations | 50 | 20,000 |
| Mutation rate | 1/2 | 1/15 |
| Mutation standard deviation | 15 | 1 |
| Parent selection pressure, $\beta$ | 5 | 3 |
| Crossover extension, $\gamma$ | 0.25 | 0.25 |
| Starting child cost improvement requirement, T | 0% | 20% |

*1) Finding the optimal PTO force:* The GA for finding the optimal PTO force is the simpler of the two. Each individual requires only two genes in their genome, one for each of $A_{\text{PTO}}$ and $\phi_{\text{PTO}}$. The optimisation is defined as:

$$\begin{aligned} \underset{A_{\text{PTO}}, \phi_{\text{PTO}}}{\text{argmin}} \quad & -E_{\text{elec}}(A_{\text{PTO}}, \phi_{\text{PTO}}) \\ \text{subject to} \quad & F_{\text{PTO}} < 3.24 \cdot 10^6 \\ & |v| < 2. \end{aligned} \tag{9}$$

The cost for each individual is the negative of electrical energy (Equation 3); normal convention in optimisation is to minimise the cost, so the algorithm is working to drive the cost further into the negative. However, care needs to be taken in accounting for the PTO's constraints. Accounting for the force limit is trivial, an upper bound is placed on the value of $A_{\text{PTO}}$. The velocity limit is more complicated, as there is no way to directly limit the speed of the WEC. The velocity of a possible solution is calculated using Equation 2, and if it is larger than the velocity limit, the individual's cost is set to zero; since the cost will always be negative, this will be worse than any valid solution. With these checks, the controller should not demand any force that would cause the device to operate beyond the PTO's limits; however any implementation using this tuning should also include a supervisory controller to ensure there is no operation beyond the limits, see [14] for more details.

*2) Fitting transfer functions:* For the second problem, the aim is to automatically fit a transfer function to the optimal PTO force magnitude and phases reference points over a range of frequencies. The optimisation is defined as:

$$\begin{aligned} \min \quad & \frac{1}{N_\omega} \sum_{i=1}^{N_\omega} |\phi(\omega_i) - \phi_{\text{ref}}(\omega_i)| \\ \text{subject to} \quad & \text{real}(Z) < 0 \\ & \text{real}(P) < 0, \end{aligned} \tag{10}$$

where $\phi(\omega_i)$ is the phase of the transfer function at frequency $\omega_i$, $\phi_{\text{ref}}(\omega_i)$ is the reference phase at that frequency (found from the previous section), $N_\omega$ is the number of reference

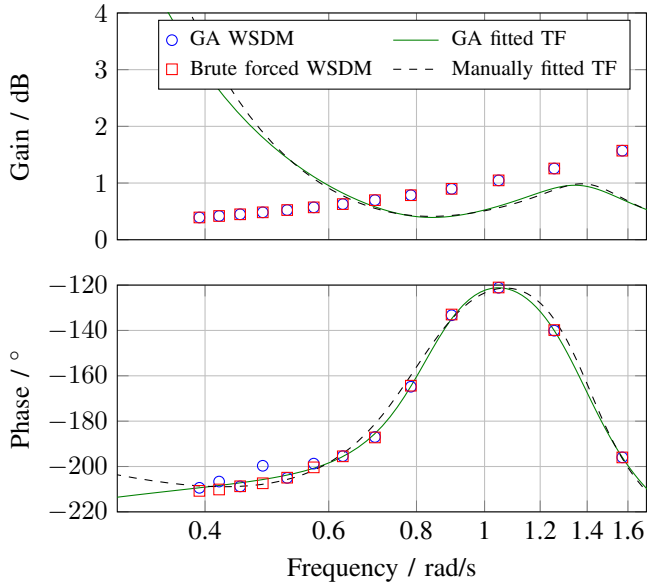


Fig. 4. Bode plot of transfer functions fitted to the WSDMs.

TABLE II
COMPARISON BETWEEN THE WSDM METHODOLOGIES

| Comparison | Value |
|---|---|
| Average phase disagreement | 1.22° |
| Maximum phase disagreement | 7.71° |
| Average percentage difference in gain | $1.8 \cdot 10^{-7}\%$ |
| Maximum percentage difference in gain | $8.6 \cdot 10^{-7}\%$ |

frequency points, and $Z$ and $P$ are the zeros and poles of the transfer function respectively. In [7], the quality of the fit of the phase was found to be more important than that of the gain in creating a well performing controller as errors in the gain can be corrected later in the controller through gain scheduling. The average absolute phase error was therefore chosen to be the cost function of the GA.

Creating a transfer function requires a much larger GA than that of the optimal PTO force; the genome needs to contain a gene for the location of each pole and zero of the transfer function. The size of the transfer function is defined by the user, as well as the number of complex conjugate pairs of poles and zeros in the transfer function. For fitting to WEC2, a transfer function with seven zeros, including two complex conjugate pairs, and eight poles, including one complex conjugate pair, is used.

To create a complex conjugate pair of zeros, a pair of genes, $g_a$ and $g_b$, are used. The system then places two zeros at locations: $Z_1 = g_a + ig_b$, and $Z_2 = g_a - ig_b$; the same method is applied to complex conjugate poles.

## III. RESULTS

### A. Finding the optimal PTO force

Shown in Figure 4, are two WSDMs, one calculated via the GA, and the other via a method of brute forcing. 13 equally spaced periods in the range 4 to 16 seconds (the operating range of the WEC) were calculated. The time to run using GAs was marginally faster, with the methods taking approximately 55 and 60 seconds for the GA and brute forcing methods respectively. The final results of the two agree fairly well, the level of agreement is shown in Table II. The main advantage of using a GA for this task is that it treats the design space as a continuous plane, as opposed to the brute forcing that uses a series of discrete points; this means that should the optimal point lie away from any of the discrete points, the brute forcing will fail to return a good result, whereas the GA can find the optimal point anywhere on the grid. The two WSDMs broadly agree at most frequency points; the largest disagreement between the two appears at approximately 0.5 rad/s, it appears likely that the GA failed to find the optimal value for this frequency.

### B. Fitting transfer functions

The fitted transfer functions for the WSDMs are also plotted in Figure 4. Although the two WSDMs are very similar, it should be noted that the manually fitted transfer function was fitted to the brute forced WSDM, and the GA fitted transfer function was fitted to the GA derived WSDM.

Table III shows a comparison of the fitting of both the GA derived transfer function to the GA derived WSDM, and the fitting of the manually fitted transfer function to the brute forced WSDM. The manually derived transfer function here is not being presented as the best possible but instead an example of what an experienced control engineer could create within a reasonable timeframe. From approximately 40 minutes of runtime on a standard laptop, the GA was able to produce a transfer function with an average phase error about 60% of the manual fit, and a similar maximum phase error. Gain errors were also higher in the manually fitted transfer function; however, it should be noted that in the implementation of the controller, the gains are scheduled to that of their respective WSDMs, so achieving a good gain fit was not a priority in either the GA or manual fitting cases. Each frequency point is weighted equally for this comparison; in reality some frequencies will contribute more to overall power production than others. Calculating the importance of each frequency would depend on a multitude of factors, for example: WEC body geometry; site conditions; and, PTO efficiency making it highly specific to each use case. However, in general, a better overall fit would likely result in more power.

The GA requires a user to define the number of poles, zeros, and complex conjugate pairs of each. For the example in Figure 4, the transfer function had seven zeros, including two complex conjugate pairs, and eight poles, including one complex conjugate pair. The output plotted in the figure is:

$$\frac{0.0324(s+4.602)(s+10.52)(s+0.144)}{s(s+1.912)(s+2.226)(s+21.04)(s+7.61\cdot10^{-4})} \times \frac{(s^2+0.322s+0.689)(s^2+0.009s+85.07)}{(s+1.26\cdot10^{-6})(s^2+0.549s+1.9)}.$$

TABLE III

ANALYSING THE FIT OF THE TRANSFER FUNCTIONS

| | Comparison | Value |
|---|---|---|
| GA | Average absolute phase error | $1.24°$ |
| | Maximum absolute phase error | $6.07°$ |
| | Average percentage difference in gain | $878\%$ |
| | Maximum percentage difference in gain | $2.41 \cdot 10^3\%$ |
| Manual | Average absolute phase error | $1.96°$ |
| | Maximum absolute phase error | $5.96°$ |
| | Average percentage difference in gain | $954\%$ |
| | Maximum percentage difference in gain | $2.77 \cdot 10^3\%$ |

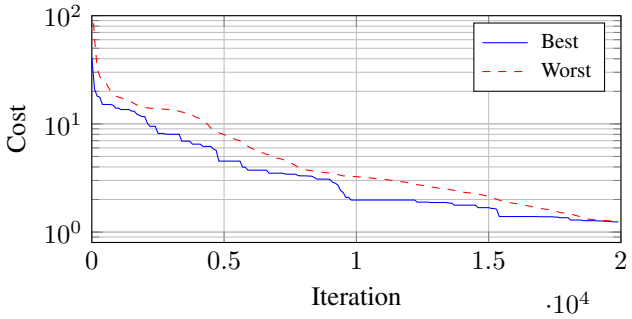


Fig. 5. Costs of the best and worst performing individuals of each iteration of the genetic algorithm to fit the transfer function.

This transfer function is not presented as the perfect solution, and a user could make improvements, such as changing any poles and zeros located very close to zero to pure integrators and differentiators, and potentially cancelling out very close pole-zero pairs; this should lead to computational runtime improvements with only a small effect on the shape of the transfer function.

Figure 5 shows the cost of the best and worst performing individuals of each iteration, as can be seen, as the algorithm progresses, the worst performing individual tends towards the best performing showing the algorithm is converging on a solution; running the algorithm multiple times outputs a very similar resultant transfer functions, showing the algorithm is consistently converging.

### C. Power Simulations

Simulations were conducted using WEC-Sim [15]. Each simulation was run for 1000 s with a time step of 0.01 s; the controller was started after 300 s. Data shown is from seconds 500 - 1000 to remove transient behaviours. The wave spectra followed a Pierson-Moskowitz spectrum, multiple simulations were run with varying peak periods in the range 5 - 13 s.

Shown in Figure 6 is selection of time series plots of the captured energy over time by the new GA controller and the old manually fitted controller, as well as a standard OVT implementation. The GA controller generally performs marginally better, its captured energy slowly pulls ahead of the manually fitted controller over the length of the simulation; this is despite it also consistently being the one of the worst performing controllers in terms of mechanical power output, showing it is better accounting for the PTO's efficiencies. This is as expected since the average phase error was smaller in the GA transfer function. As previously mentioned, for both the GA and manually fitted controllers, their gain was scheduled to the WSDM and were therefore identical, meaning any differences lie solely in the phase of the transfer function. Both implementations of the adapted OVT controller consistently outperformed a standard OVT controller.

Figure 7 shows the total energy captured by each controller in a range of sea states. The GA controller is consistently better performing either the same or better than the manual controller, with the best improvement being nearly 6%.

## IV. CONCLUSIONS & DISCUSSION

Ongoing work shows that the GA method is an effective way of estimating the WSDM as results of simulations (using the OVT control method that is still under development) using both manually fitted and GA fitted functions give similar outputs. Although the improvements are small in terms of energy improvement and runtime, there are methodological improvements that improve the ease of use of this controller. It is possible for engineers without an intimate knowledge of control theory to utilise not only advanced control methods, but advanced control methods optimised for their device.

The gain of the GA fitted transfer function was worse than that of the manually fitted one. As the GA currently works, there is no costing associated with the fit of the gain, so including this could help to improve the fit of the gain. However, it is likely that this will also result in an increase in phase error; therefore, gain scheduling is likely the better choice to control the fit of the gain.

Currently, simulations have only been conducted for a singular significant wave height of 2 m, for a more thorough analysis, simulations should be conducted for a range of excitation forces. In the full adapted OVT controller a selection of different transfer functions are used, scheduled on the estimated excitation force amplitude. The variability across wave amplitudes is due to PTO force limits and PTO efficiencies. For each significant wave height a different transfer function approximation is used.

Aims for future work include expanding the transfer function fitting GA to have the ability to also optimise the number of poles and zeros in the transfer function to further improve usability for a non-control engineer user. Currently, a user would have two options, firstly to try and brute force the solution by attempting all combinations to attempt to find the best fitting one, but with four parameters available to be changed (number of: poles, zeros, complex conjugate poles, complex conjugate zeros). This is quite a large design space that would be very time consuming to work through. The other option is to create a very high order transfer function to over engineer the solution, which would be a much more complex design problem; making it harder for the GA to find a good solution. There is also the possibility that this high

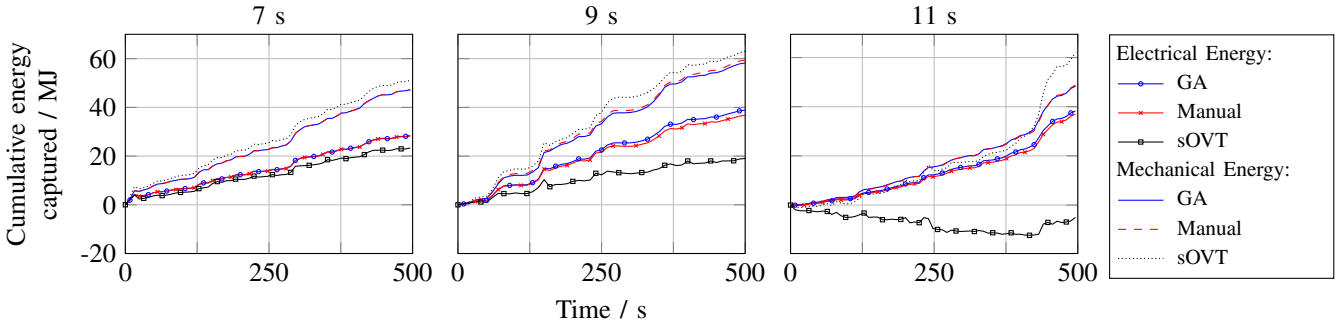


Fig. 6. Time series of the energy captured by three controllers in a WEC-Sim simulation: 'GA' uses the GA derived transfer function; 'Manual' uses the manual derived; and, 'sOVT' uses a standard OVT controller (shown in Figure 1); both the mechanical and electrical energy captured for each controller is shown. Each plot uses a different sea spectrum, all are Pierson-Moskowitz with a significant wave height of 2 m and plot titles show the peak period.

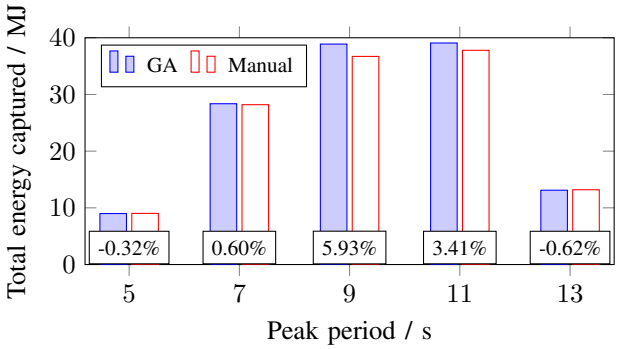


Fig. 7. Total energy captured for a variety of Pierson-Moskowitz spectra with varying peak periods by two controllers in a WEC-Sim simulation; one using the GA derived transfer function, and the other using a manually fitted transfer function. All wave spectra have the same significant wave height of 2 m. Percentages show the improvement from manual to GA.

order solution can introduce complex, oscillatory behaviours between the reference points, something that would most probably not be realistic but arises solely from the high order of the solution.

Another point on usability is the values of the GA parameters (listed in Table I) and what the 'optimal' values would be; the parameters could be fine tuned to the point that the algorithm reaches the solution just at the end of the simulation, this would result in improved runtimes of the simulation. However, should the algorithm be reapplied to another, more complex problem then it is possible that the algorithm fails to find a good solution; a user would then have to rerun the algorithm, possibly multiple times after varying the various parameters, a process that could end up taking longer, and would require more active work, than running a singular more conservatively tuned GA.

This work has shown the capability of GAs in the tuning of controllers, previous literature has proven the capability of GAs to optimise the geometry of a WEC [8]. Expanding the use of GAs to automate the design of the PTO would allow for the co-design of all of these components, hypothetically leading to even greater improvements than what can be achieved by accounting for component efficiencies at the controller design stage.

## Acknowledgements

This work has been supported by the EPSRC funded CDT in Wind and Marine Energy Systems and Structures (EP/S023801/1), and alongside the EPSRC funded Holistic Advanced Prototyping and Interfacing for Wave Energy Control (HAPiWEC) project (EP/V040987/1).

## References

[1] Y. Hong, R. Waters, C. Boström, M. Eriksson, J. Engström, and M. Leijon, "Review on electrical control strategies for wave energy converting systems," *Renewable and Sustainable Energy Reviews*, December 2014.

[2] J. Falnes and A. Kurniawan, *Ocean Waves and Oscillating Systems: Linear Interactions Including Wave-Energy Extraction*, 2nd ed. Cambridge University Press, 2020.

[3] J. Hals, J. Falnes, and T. Moan, "A comparison of selected strategies for adaptive control of wave energy converters," *Journal of Offshore Mechanics and Arctic Engineering*, August 2011.

[4] R. G. Coe, G. Bacelli, and D. Forbush, "A practical approach to wave energy modeling and control," *Renewable and Sustainable Energy Reviews*, March 2021.

[5] F. Fusco and J. V. Ringwood, "A simple and effective real-time controller for wave energy converters," *IEEE Transactions on Sustainable Energy*, January 2013.

[6] A. Stock and C. Gonzalez, "Design of optimal velocity tracking controllers for one and two-body point absorber wave energy converters," *Renewable Energy*, September 2020.

[7] A. Stock, "A practical metbod to account for power take-off efficiency in optimal velocity tracking control of wave energy converters," *Renewable Energy - Under Review*.

[8] A. P. McCabe, "Constrained optimization of the shape of a wave energy collector by genetic algorithm," *Renewable Energy*, October 2013.

[9] D. Valério and M. D. Ortigueira, "Identifying a transfer function from a frequency response," *Journal of Computational and Nonlinear Dynamics*, January 2008.

[10] MathWorks, "invfreqs." [Online]. Available: https://uk.mathworks.com/help/signal/ref/invfreqs.html

[11] A. Stock, C. Gonzalez, and D. Robb, "Integrated marine point absorber control tool final stage 2 project report," SgurrControl Ltd., Tech. Rep., 2018.

[12] "PowerPod technical datasheet," Trident Energy Ltd., Tech. Rep., 2014.

[13] L. J. Eshelman and J. D. Schaffer, "Preventing premature convergence in genetic algorithms by preventing incest," in *Fourth International Conference on Genetic Algorithms*, July 1991.

[14] A. Stock, N. Tom, and C. Gonzalez, "Adapting optimal velocity tracking control to account for wec constraints and power-take-off efficiencies," in *21st IFAC World Congress*, July 2020.

[15] A. Keester, J. Leon, N. Tom, K. Ruehl, D. Ogden, D. Forbush, J. Grasberger, S. Husain, and M. Topper, "Recent developments in the wec-sim open-source design tool," National Renewable Energy Laboratory, Tech. Rep., 2022.