# Optimisation of the Absolute Sine for Dynamic Induction Control

## Power Maximisation of a Wind Farm Through Infinite and Finite Horizon Control

J.A.G. Kuppers

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

# Optimisation of the Absolute Sine for Dynamic Induction Control

## Power Maximisation of a Wind Farm Through Infinite and Finite Horizon Control

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

J.A.G. Kuppers

January 10, 2024

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of Technology


# Abstract

The clustering of wind turbines in a wind farm results in overall efficiency losses as downstream wind turbines operate in the wake of their upstream neighbours. Wind farm flow control (WFFC) strategies have emerged to reduce these wake effects with the goal of maximising overall performance. Dynamic induction control (DIC) aims to enhance the wake breakdown and restore the wake's energy content through dynamic thrust variations. The control signals are often found through the economic model predictive control (EMPC) method, which relies on an internal model to incorporate future system behaviour in the determination of the next optimal control input. These models are designed to capture the most dominant wake characteristics while remaining computationally efficient. We employ the two-dimensional free-vortex wake (FVW) model presented from [1], which models the wake through vortex element pairs released from the edges of the actuator disc. The power of the two-turbine wind farm is maximised through EMPC, improving performance by 9.64% over greedy control simulations. However, the EMPC method inherits finite horizon effects, resulting in large control horizons to optimise. In this study, we address these limitations by employing an absolute sine parameterisation in the FVW model to limit the finite horizon effects and reduce the dimension of the optimisation problem. The significant dimension reduction allows for a grid search to find the optimal infinite horizon steady-state solution, improving the mean steady-state performance by 2.43% over the baseline results from [1]. Additionally, we focus on converging towards this optimum through finite horizon EMPC optimising over the amplitude and offset. Grid search analyses reveal sensitivity towards initialisation due to the appearance of local minima around the infinite horizon optimum. A maximum success rate is realised for very large control horizons, maximising the probability of converging towards the infinite horizon optimum. Accounting for the inherited system delay in the objective function also realises a maximum success rate but for shorter control horizons, which significantly decreases the simulation time. The final controller design terminates simulations ten times faster through the absolute sine parameterisation compared to the baseline simulation from [1] while maximising the probability of convergence towards the infinite horizon optimum.

# Contents

# Acknowledgements

I would like to thank my supervisors prof.dr.ir. J.W. van Wingerden and ir. M.J. van den Broek for their guidance during the research and writing of this thesis. The definition of the project was difficult at the beginning, and it helped me find the gap in the literature. By asking critical questions and providing good suggestions, you really helped me shape this project. Particularly, I would like to thank ir. M.J. van den Broek for our weekly meetings and quick responses every time I struggled. I wish you all the best of luck in your new adventure and future career.

Furthermore, I would like to thank my thesis committee, consisting of prof.dr.ir. J.W. van Wingerden, dr.ir. S.P. Mulders, ir. A.A.W. van Vondelen, and ir. M.J. van den Broek for their critical questions and honest judgement during the discussion and defence of this thesis.

Delft, University of Technology
January 10, 2024

J.A.G. Kuppers

# Chapter 1

# Introduction to wind farm flow control

One of the critical challenges the world is facing is the threat of global warming. Due to their exothermic nature, fossil fuels have been the primary energy source for the past century. However, the emissions that are released into the atmosphere as a result of the combustion have contributed to average global temperature rises. Sustainable alternatives must replace fossil fuels to reduce these emissions to achieve the goals of the Paris Agreement, which aims at keeping the global average temperature rise below 2 degrees Celsius [2]. Together with solar, hydrodynamic, and many other power sources, wind energy is considered to account for the larger part of the energy demand in Europe in the near future [3].

In recent years, global investments in the wind energy sector have grown, resulting in a vast increase in installed power capacity. In 2022, the global installed capacity was 823 GW, which is 13% more compared to 2021 [4]. Onshore wind turbines account for the majority of the power generation, but the market is slowly shifting towards offshore. Offshore wind is stronger and less turbulent than onshore wind because no uneven terrains such as mountains and cities influence the flow [5]. Furthermore, due to the increasing wind turbine dimensions, logistical challenges arise during onshore construction. Transporting blades or tower pieces larger than one hundred meters is not always possible due to bridges or corners in the road [6]. Loading them onto ships eliminates these transportation limitations but increases the operation and maintenance costs because they are costly to use [6, 7]. The wind turbines also deteriorate faster due to the salty environment and harsh operational conditions, increasing these costs. Currently, offshore levelised costs of energy are higher than onshore, and to make offshore wind farms more appealing, these costs must be reduced [8].

## 1-1 Individual wind turbine control

Horizontal axis wind turbines (HAWT), which will be further referred to as wind turbines, are systems initially designed to operate in their optimal setpoints. This control method is referred to as greedy control [9]. With the greedy control approach, the wind turbine setpoints change according to the varying wind conditions. Figure 1-1 shows the nacelle, including the

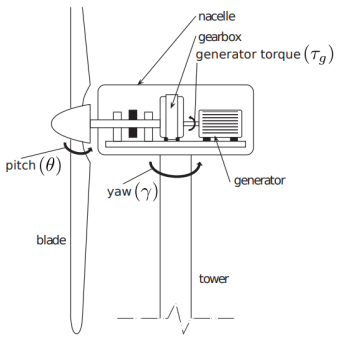


**Figure 1-1:** Schematic representation of a wind turbine including the control variables $\tau_g$, $\theta$, and $\gamma$. Image is retrieved from [11].

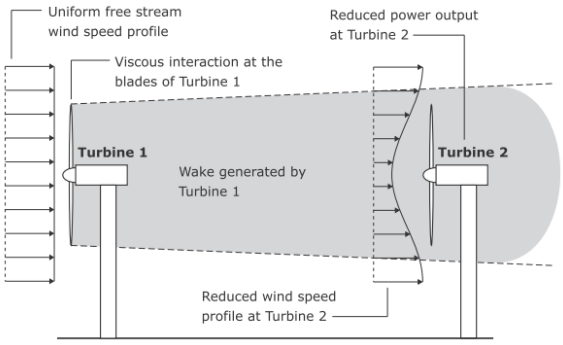


**Figure 1-2:** Schematic representation of the wake reaching the downstream turbine. The velocity deficit in the wake is indicated with the arrows. The image is retrieved from [6].

control variables. The first variable is the generator torque $\tau_g$, employed in below-rated wind conditions to keep the tip-speed ratio at the optimal design value [10]. The rated wind speed is where the wind turbine becomes overly loaded as the blades capture too much wind. The second control variable, the blade pitch angle $\theta$, pivots the blades out of the wind in above-rated wind conditions to limit the loads and maintain a constant power production. Lastly, the rotor's yaw angle $\gamma$ is utilised to rotate the whole rotor perpendicular to the wind to maximise the surface area if the wind direction changes.

The setpoints of the control variables determine the amount of energy extracted from the wind. The extraction causes a velocity deficit in the wind that flows through the wind turbine. Under the assumption that this mass of air is contained in a control volume, a new variable is introduced [12]. The (axial) induction factor $a$ is the normalised change in wind speed between the freestream wind speed $u_\infty$ and the wind speed at the rotor $u_r$ defined as

$$a = \frac{u_\infty - u_r}{u_\infty}. \tag{1-1}$$

The greedy control induction factor equals the optimal wind deficit ratio. It is derived through the power coefficient $C_P$, which is the wind turbine's power production normalised by the freestream power [12]. It is defined as

$$C_P = \frac{P}{P_\infty} = \frac{P}{\frac{1}{2}\rho A_r u_\infty^3} = \frac{2\rho A_r u_\infty^3 a(1-a)^2}{\frac{1}{2}\rho A_r u_\infty^3} = 4a(1-a)^2, \tag{1-2}$$

where $\rho$ is the air density and $A_r$ the rotor swept area. The optimal induction factor is found through its derivative

$$\frac{\partial C_P}{\partial a} = 4\,(a-1)\,(3a-1)\,, \tag{1-3}$$

which results in $a_{opt} = (1, \frac{1}{3})$. Since it is not possible to have a power coefficient of one, meaning that $u_r = u_\infty$, the optimal induction factor is $a_g = 0.33$. It has been used to derive the theoretical limit of the steady-state energy extraction known as the Betz limit [12]. For

$a_g$, the power coefficient is equal to $C_P = 0.593$, and no wind turbine has been designed yet that has exceeded this limit. The limitation is the expansion of the air tube right before it reaches the wind turbine, which causes a reduction in velocity.

Another measure for the wind turbine's energy extraction is the thrust coefficient ($C_T$). It is the normalised thrust force similar to the definition of the power coefficient and is given as

$$C_T = \frac{T}{T_\infty} = \frac{T}{\frac{1}{2}\rho A_r u_\infty^2} = \frac{2\rho A_r u_\infty^2 a(1-a)}{\frac{1}{2}\rho A_r u_\infty^2} = 4a(1-a). \tag{1-4}$$

### 1-1-1 Wake losses behind a wind turbine

Inherent to operating a wind farm with greedy control is the absence of communication between the wind turbines. The adverse effects experienced by the downstream wind turbines caused by their upstream neighbours are disregarded. As the upstream wind turbine extracts energy from the wind, it leaves a trail of low-energy, turbulent wind [10]. This region is referred to as the wake and is marked in grey in the schematic representation in Figure 1-2. Wind turbines operating within the wake of their upstream neighbour face limitations in performance due to these wind conditions, which results in reduced energy extraction, thereby decreasing the total wind farm's energy capture [8, 13]. Furthermore, the turbulence increases the loads, limiting endurance and reliability.

In response, wind farm flow control (WFFC) approaches have emerged to increase the wind farm's efficiency [14, 8]. This can be achieved by minimising the structural loading on the wind turbines to reduce the operation and maintenance costs [15, 16, 17], or maximise the wind farm's power generation [18, 19, 20]. In this study, the latter will be the control objective since it is considered the most beneficial [21].

## 1-2 Wind farm flow control

The wind farm's power production can be increased by conducting control strategies that mitigate the wake effects on downstream wind turbines. Roughly three strategies can be distinguished: power derating, wake steering, and dynamic induction control, which are pursued by (quasi-)statically or dynamically changing the wind turbine's setpoints [8]. The former implies that the setpoints slowly change over time. The latter approach aims to include faster flow dynamics and account for wake changes caused by the dynamic setpoint variations.

### 1-2-1 Power derating

Power derating is a static control strategy where an upstream wind turbine operates with a reduced thrust coefficient and, therefore, energy extraction. This causes the wake to have higher energy content, allowing the downstream neighbour to operate in better conditions and increasing its power production [14]. The goal of power derating is to compensate for the losses of the upstream wind turbine and realise an overall power production increase. The approach is illustrated in Figure 1-3 and is also referred to as axial induction control because of the direct relationship between the induction factor and thrust coefficient [6]. Initially,

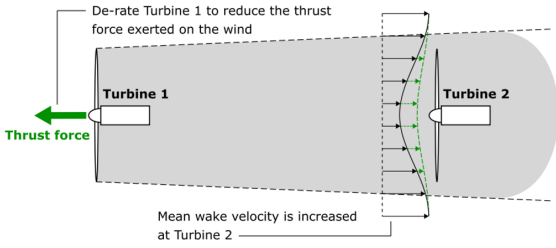


**Figure 1-3:** Schematic representation of power derating where a reduced thrust of the upstream wind turbine allows the downstream turbine to improve performance. The image is retrieved from [6].

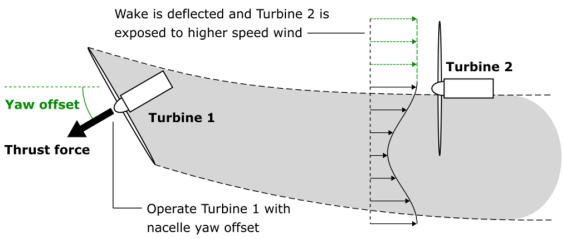


**Figure 1-4:** Schematic representation of wake steering where a yaw offset results in wake redirection. The image is retrieved from [6].

power derating seemed to be a promising strategy for power maximisation [22, 23]. However, these gains are attributed to their use of wake engineering models [24], which tend to be less accurate due to simplifications and parameterisations. Recent studies in field tests [25], wind tunnel experiments [26], and studies using highly accurate large-eddy simulations (LES) [27] have questioned the applicability of power derating.

### 1-2-2 Wake steering

With the second control strategy, the wake is redirected around the downstream wind turbine by misaligning the upstream one from the incoming wind direction [8]. Wake steering relies on a yaw offset, which imposes the redirection through asymmetric loading along the blades [12]. The result is a reduced overlap between the second wind turbine and the wake as shown in Figure 1-4, allowing for better wind conditions to operate in. Again, the aim is for the downstream turbine's power to increase and compensate for the losses of the upstream turbine. The control strategy has been applied in many studies [28, 29, 30, 31] in which power gains were found. The potential has been validated in field tests [32, 33, 34], wind tunnel experiments [35, 36], and LES studies [37].

### 1-2-3 Dynamic induction control

The third control strategy focuses on dynamic variations of the wind turbine's control variables to further improve the previously mentioned static strategies. Dynamic induction control (DIC) aims to enhance the wake breakdown, enabling downstream wind turbines to operate in higher-energy wind conditions without significantly affecting the power production of the upstream turbine. The wake breakdown is enhanced by dynamically varying the thrust, which is characterised by the pulse. The pulse results in velocity deficit variations in the wake shown in Figure 1-5, which causes the higher-energy surrounding air to get mixed into the wake, thereby restoring the wake's energy content [8].

A special form of wake mixing is performed in [38], where the dynamic variations of the wind turbine's control variables lead to a faster wake recovery due to increased turbulence. Through individual pitch control (IPC), a helix shape appears in the wake visualised in Figure 1-5. IPC has been used in previous studies [39], but mainly for load reduction strategies [38].

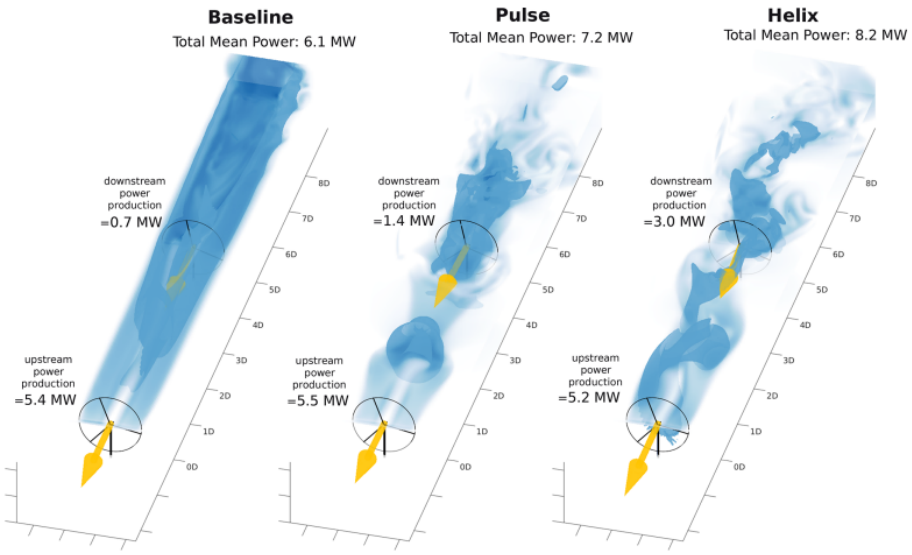


**Figure 1-5:** The wake propagations from dynamic induction control resulting in the pulse [20] and individual pitch control resulting in the helix [38] compared to greedy control conditions. The image is retrieved from [8].

Periodically exciting the upstream wind turbines can enhance the wake's energy recovery [20]. A sinusoidal perturbation is imposed on the greedy optimal local thrust coefficient $C'_{\mathrm{T}} = 2$ such that the controlled wind turbines are either over- or under-inducting. In the case of over-induction, the wind turbine's thrust coefficient (or induction factor) exceeds the greedy optimal value, resulting in increased energy extraction and higher wind velocity deficits in the wake. Conversely, for a turbine operating in under-induction, the wind turbine is derated, and the velocity deficits are smaller. A periodic signal is found to increase the power production of a two-turbine wind farm in [1]. Instead of imposing a sinusoidal thrust variation around the greedy optimal value, the controller converges to a periodic steady-state through online optimisation.

The potential of DIC is validated in wind tunnel experiments [9]. The results strengthen the conclusions drawn in [20] that the first downstream wind turbine is responsible for the largest power gain, which implies that it benefits from the wake recovery.

## 1-3 Control designs for dynamic induction control

DIC will be the control strategy in this study, and its potential will be further investigated. The power increases can be realised through a proper design of the controller. Various methods exist to design controllers, and they can be categorised into model-based and model-free depending on whether they rely on an internal model to represent the plant.

### 1-3-1 Model-free control

Model-free controllers disregard output or state measurements to analyse the system's performance and determine new control inputs. The control sequences are often pre-defined and

implemented through feed-forward, open-loop control [6]. Studies employing these methods have shown that periodic control signals enhance the mixing of the wake through the wake velocity deficits. Similar to the earlier mentioned sinusoidal perturbation on the thrust coefficient in [20], the collective blade pitch angles are dynamically varied in [40]. Both studies focus on vortex rings that are shed from the wind turbines and enhance the wake breakdown.

Periodic, pre-defined signals for the blade pitch and torque repetitively excite the upstream wind turbine in the computationally expensive LES framework in [24]. Through closed-loop simulation, earlier in the study, the control variables converge to a periodic steady-state where one period is identified and utilised for open-loop simulation. These pre-defined sequences reproduce the closed-loop features and reduce the computational expense as no online optimisation is required.

### 1-3-2 Model-based control

The last study obtains the periodic steady-state control signals through model-based control. Model-based control determines the optimal control variable setpoints through dynamic optimisation using a wind farm model [6]. These controllers are typically implemented in a closed-loop feedback control scheme, making them more robust towards disturbances than model-free, open-loop control. This robustness stems from the ability of model-based control to adapt to evolving system conditions, especially in dynamic and unpredictable operating environments like wind farms. This is a flexibility that pre-defined signals in model-free control methods cannot account for.

#### Model predictive control

The most common model-based control approach for DIC is model predictive control (MPC). It is widely used for its compatibility with constraints during optimisation [6]. The aim is to find the optimal control sequence by considering the future behaviour of the system obtained through model simulation. Regular model predictive control can only consider linear systems. For a discrete-time, one-dimensional, single-input-single-output system, the system's state space is defined as

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k. \end{aligned} \tag{1-5}$$

The primary applications of MPC are to stabilise the system or to minimise the error between the system's output $y_k$ and some reference signal $r_k$, marked orange and red in Figure 1-6, respectively [41]. The model is used to approximate the system's behaviour in the future over a horizon of $N_\mathrm{h}$ steps from the measured current states $x_k$ as shown in the block diagram in Figure 1-7. Ideally, $N_\mathrm{h}$ is infinitely long, so the controller has complete knowledge of the future behaviour. However, this would highly increase the computational expense. Since the controller is bound to the time step of the discretised system, the problem is solved in a receding manner using a finite horizon [41].

At every time step, the objective function $J$ in Equation 1-6 is defined as a linear combination of weighted squares of $(x_k, u_k)$ with $k \in \{k_0, k_1, \ldots, k_{N_\mathrm{h}}\}$. The future states are retrieved from the measured state through the linear state equation. The optimal control sequence of length

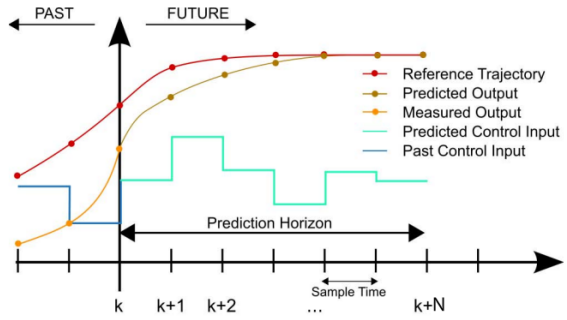


**Figure 1-6:** Visualisations of all MPC-related signals. The error between the reference and output signal must be minimised through control actions determined through optimisation. The image is retrieved from [42].

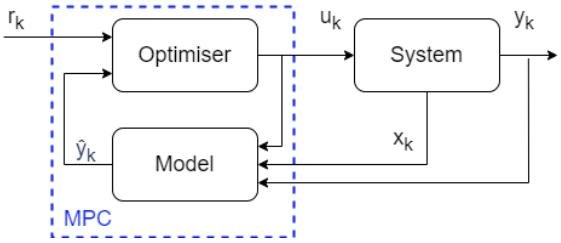


**Figure 1-7:** Block diagram including the true system and the MPC controller. The optimisation and internal model block result in the model predictive controller.

$N_\mathrm{h}$ is then found through online optimisation of $J$. The first control input of the sequence is fed to the system, and the cycle is repeated.

$$\begin{aligned} \min_{u_k} J(x_k, u_k) &= \min_{u_k} \frac{1}{2} \sum_{k=k_0}^{k_0+N_\mathrm{h}} x_k Q x_k + u_k R u_k \\ s.t. \quad x_{k+1} &= A x_k + B u_k \end{aligned} \tag{1-6}$$

The weights $Q$ and $R$ are squared matrices that penalise the state or control input. They are assumed to be positive semi-definite and positive definite, respectively. For large $Q$ with respect to $R$, the controller is determined for fast convergence at the expense of large control input values. Conversely, relatively high values for $R$ limit the control action and the time to drive towards the reference value will increase. This standard objective form allows for quadratic optimisation, which is guaranteed to find the optimal solution, given the constraints [43].

MPC is a control method extensively employed in wind farm flow control [13]. However, the objective function becomes nonlinear when the wake dynamics are incorporated. For regular MPC to still work, the model can be simplified to a linear representation [44], although this can result in higher inaccuracies. Furthermore, the nonlinear model can be linearised around the current operating points [45]. A drawback is that the linearised model is only accurate in a small region around the operating points, making it sensitive to inaccuracies. This can be addressed by incorporating the nonlinear behaviour into the objective function and performing Nonlinear MPC (NMPC), which is similar to MPC in its objective, but nonlinear objective functions and constraints can be included.

### 1-3-3 Economic model predictive control

For the control objective in this study, a different form of (N)MPC is employed. It better coincides with the control objectives rather than controlling towards some reference value or

stabilising the system. With economic MPC (EMPC), the goal is to maximise the economic performance of the system by optimising some objective function that incorporates the nonlinear wake dynamics [46]. The optimisation step and the wake dynamics model are two important blocks of the EMPC controller as illustrated in Figure 1-7 and will therefore be elaborated on.

### Optimisation in EMPC

During optimisation, the algorithm seeks to identify the control sequence that maximises the economic performance within the control horizon $N_\mathrm{h}$ at every time step. As mentioned, $N_\mathrm{h}$ determines the length of the control sequence and, therefore, also heavily impacts the search space dimension of the optimisation. In WFFC, this often translates into very high-dimensional optimisation problems due to the large control horizon required. The control horizon must be sufficiently long to allow the model to capture the wake propagation far downstream, gathering essential information on its effects on other wind turbines.

The optimisation goal is to find the global optimum. However, in high-dimensional search spaces with embedded nonlinear dynamics, nonconvex and multimodal properties are very likely to be present [47], which implies the existence of multiple local minima within the search space [48]. These are solutions where optimisation algorithms terminate in their search for the global optimal solution [41]. The search space can also contain flat regions where the algorithm's convergence and propagation can be rather slow [41]. Various optimisation methods exist, and the choice depends on the problem. Gradient-free, nature-inspired (NI) algorithms have been utilised in WFFC [48]. However, most studies rely on gradient-based (GB) optimisation with numerically computed gradients [1, 49, 50].

NI optimisation methods iteratively explore large parts of the search space and are inspired by evolutionary processes in nature [51]. These algorithms randomly select new search directions, comparing and potentially combining them with the best solution to improve further. The approach facilitates exploration and mitigates the risk of remaining near a local minimum. NI algorithms require a substantial number of objective function evaluations, making them suitable for objective functions that are computationally inexpensive. However, due to their random nature, finding the global optimum cannot be guaranteed. Different solutions will be found for the same initialisation, and trajectories will never be the same. Furthermore, their reliance on randomness can be inefficient. In such cases, the optimisation method must excel at exploiting rather than exploring.

Contrarily, the trajectories of GB algorithms are always the same due to the additional information included. They are more effective in local exploitation and converge towards a solution in fewer iterations [52]. While generally computationally expensive, the gradient provides valuable information about the local curvature of the search space [47]. This second-level information is beneficial in determining new search directions during the optimisation process. GB optimisation methods are prone to get stuck in local minima as the gradient becomes zero in those regions. They are sensitive to initialisation, and in flat regions, GB methods may struggle to converge rapidly as the gradient approximates zero, resulting in minimal value updates.

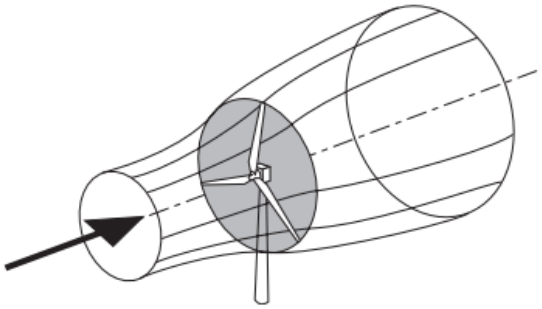

**Figure 1-8:** Illustration of a wind turbine considered as an actuator disc. The wake expansion right before the disc indicates the wind speed reduction. The image is retrieved from [12].

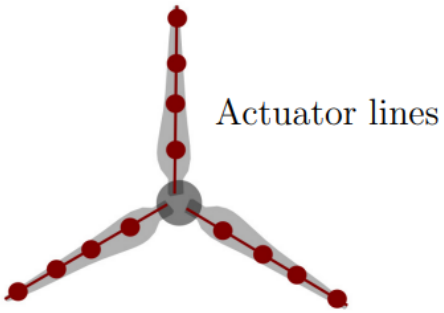


**Figure 1-9:** Visualisation of the actuator line model where a discrete number of points characterise the blades of a wind turbine. The image is retrieved from [54].

## 1-4 Control-oriented wind farm models

The second block in Figure 1-7 is the internal model, which represents the true system to gather knowledge on its future behaviour. In WFFC, the wind farm can be represented in two or three dimensions. Furthermore, two main model components can be distinguished: the wind turbine and the wake [53]. The relevant models used for DIC with the aim of power maximisation will be discussed.

### 1-4-1 Wind turbine modelling

Wind farm models depend on the interaction between wind turbines and the fluid. This interaction significantly impacts the wind fields and their propagation [6]. Contrarily, to establish a relation between the wind flow and the power generation, two wind turbine models are commonly employed [53]. Here, the flow field is the input related to the loading, forces and, consequently, power generation.

#### The actuator disc

In the actuator disc model (ADM), the entire rotor swept area is considered to be a thin disc as shown in Figure 1-8 and is based on mass conservation theories of a controlled volume tube of wind passing through the disc [12]. The model assumes a constant wind speed and uniform pressure for the wind approaching the actuator disc. As the wind interacts with the disc, energy is extracted, resulting in a pressure drop, a decrease in wind speed, and an expansion of the wake area—the wake expansion results from a constant mass flow rate [12]. Due to the pressure drop, the power production can be derived from the velocity deficit over the disc. The induction factor and (local) thrust coefficient, defined earlier to relate the deficit to power, are thus utilised as control variables when wind turbines are modelled using the ADM [1, 20].

#### The actuator line

The second wind turbine model introduces a higher level of complexity because the forces are radially distributed along the individual blades [53, 55]. This approach accounts for the rotating dynamics of the blades, representing the wake as lines [55]. The number of lines corresponds to the number of blades, typically three for modern wind turbines [12]. While this model still simplifies the actual shape of the blades, it allows for the torque and pitch angles to serve as control variables [24].

Each line is characterised by a set of points, as depicted in Figure 1-9. The computational expense of the model increases exponentially with the number of points, as dynamic properties in the wake model define each point. For a large number of points, this implies a high computational demand. Due to its more accurate representation of the blades and the circulation of the wake, especially compared to the ADM, the actuator line model (ALM) is often preferred as the wind turbine model in LES studies [24, 56, 57].

### 1-4-2 Wake modelling

The wind is a continuous flow which includes local pressure and velocity gradients [53]. Attempting to encompass all details in a wake model is impractical and would be too computationally expensive. The goal, therefore, is to identify a wake model that captures the essential characteristics, making it suitable for online, model-based control.

#### Leapfrogging

In DIC, the objective is to enhance the wake breakdown by creating velocity deficit variations. Through dynamic setpoint changes, the wake recovery is increased, allowing downstream wind turbines to benefit from higher-energy content wind. Along the wake boundary layer, vortex structures are created which surround the wake, resulting in the entrainment of the surrounding air [8]. Figure 1-10 is a snapshot taken from a single-wake model in LES and shows these vortex structures right behind the wind turbine. They are circular regions of rotating fluids that take a helical shape due to the blade's rotational motion and the velocity field induced by the vortex structures themselves [58]. They rotate around one imaginary axis and the axis which points in the downstream direction as shown in the figure. Under the assumption that the rotor is uniformly loaded, as in the ADM, the latter rotation is not captured by the model changing the helical structures to vortex rings.

These vortex rings are shed from the tip of the blades with varying wind speeds and intensities due to the dynamic thrust variations. Therefore, the wake's boundary contains alternating vorticity intensities, which results in vortex roll-ups of subsequent rings [40]. These roll-ups arise from interactions between rings where higher-speed rings roll over the lower-speed ones. Figure 1-10 also shows these interactions further downstream, where the higher-speed and lower-speed sections are marked in red and blue, respectively. For the remainder of the report, this phenomenon is referred to as leapfrogging [59].

### Wind farm flow models

Low-fidelity, steady-state, parametric models cannot describe these flow dynamics and interactions because they rely on properly tuning a limited number of design parameters. These simplifications increase the simulation efficiency but cause the parametric models to fail in capturing the dynamic wake flows required for DIC.

Instead, dynamic flow models can be utilised to describe the shedding vortex rings to enhance the wake breakdown. The most accurate are high-fidelity models whose governing equations rely on computational fluid dynamics (CFD). They are characterised by a set of partial differential equations referred to as the Navier-Stokes equations [60]. These three-dimensional, nonlinear equations of motion describe the flow of the fluid and include phenomena such as turbulence and wake expansion. An analytic solution for the equations has yet to be discovered when considering the boundary or inflow conditions without relying on substantial assumptions [53]. A common assumption is that the flow is assumed to be incompressible, meaning constant density.

High-fidelity models capture the large wake structures by solving the Navier-Stokes equations on a large number of cells on a grid, while the finer structures and stresses are approximated using subgrid models [60], which are referred to as eddies. The sophisticated, high-fidelity models are hence called large-eddy simulations (LES). Studies have utilised LES as a control-oriented model for DIC through dynamic variations of $C'_{\mathrm{T}}$ [61, 62] and IPC [38].

In [20], an LES study is conducted in which the wind turbines are subjected to periodic variations of the local thrust coefficient around the greedy control value of $C'_{\mathrm{T}} = 2$. The concept behind the periodic variation is to imitate the quasi-periodic shedding of vortices along the edges of the wake, which relies on modelling the wind turbines with the ADM. Equation 1-7 introduces a sinusoidal perturbation to the optimal thrust coefficient, causing the upstream wind turbines to either over- or under-induct.

$$C'_{\mathrm{T}}(t) = 2 + A \sin\left(2\pi \frac{St \cdot u_\infty}{D} t\right). \tag{1-7}$$

The parameterisation of $C'_{\mathrm{T}}$ allows for a significant reduction in the optimisation dimension. Instead of solving the control problem in a receding horizon manner with EMPC, an offline parameter sweep is performed to identify the optimal amplitude $A$ and the Strouhal number $St$, which is a non-dimensional number describing oscillations in a flow [63]. Multiple test cases are conducted with a mean power increase between two to five percent for amplitude and Strouhal number $(A, St) = (1.5, 0.25)$.

An extension is proposed in [24] by analysing a two-turbine wind farm in an LES with the ALM as opposed to the ADM in [20]. The reasoning is that the ADM results in simplified turbulence representations in the near wake, providing a simplified representation of the wake dynamics. A smaller wind farm is analysed to allow for sufficiently refined resolutions in the LES. The first wind turbine is controlled with EMPC to maximise the wind farm's power output in a closed-loop manner. The solution converges to a periodic steady-state control sequence with a Strouhal number of $St = 0.38$. This signal is then implemented as the open-loop control sequence to reproduce the closed-loop optimal control features. Modelling the wake with sufficiently refined resolutions comes at the expense of high computational demand.

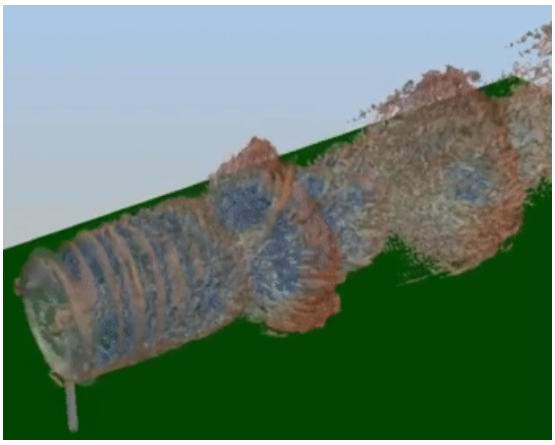

**Figure 1-10:** A snapshot of a single-wake model in LES performing DIC. We observe the vortex rings behind the wind turbine. Their alternating wind speeds result in the leapfrogging as the higher-speed sections jump over the lower-speed ones causing the entrainment of the surrounding air. The image is retrieved from [64].

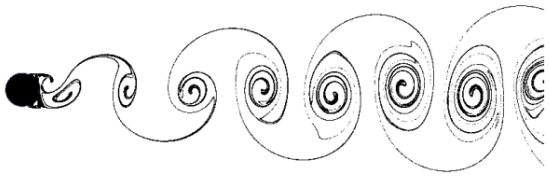

**Figure 1-11:** An illustration of the release of vortex elements behind an object. The speed and pressure difference behind the object result in the rotational motion. The image is retrieved from [65].

This is due to the large number of states that represent the flow. It can take up to weeks to perform a single simulation [53]. The fidelity of LES is, therefore, too high for model-based control.

### 1-4-3 Free-vortex wake method

Wake instabilities can also be characterised by modelling the vortices along the wake boundary [66]. These vortices are swirling structures formed at the tip and root of the blades due to the difference in wind speed and pressure that arises during blade rotation. Figure 1-11 illustrates these structures as they are released from an object. Apart from the incompressible flow assumption, the free-vortex wake (FVW) model also relies on the assumption of inviscid flow, meaning it has no viscosity. This assumption holds in low-speed aerodynamics and low Reynolds numbers typical in wind farm flows [67]. The vortices' circulation strength is determined by the vorticity formulation of the Navier-Stokes equations [60], which is related to the thrust. The velocity of the vortices is induced by others, as described by the Biot-Savart law [12], which originally characterises the flow and circulation of a magnetic field but is applied in aerodynamics to describe the induced velocity of vortex elements [67]. Vortex models prove more accurate in near-wake regions due to the distortion of vortices in the far-wake caused by their interactions [68].

In [69], a three-dimensional FVW model is proposed, where the wind turbines are modelled with the ADM, enabling the description of vortex rings through vortex filaments. An adapted version is presented in [1], both in two and three dimensions. A discrete number of vortex rings represents the wake shed from the actuator discs. In a two-turbine wind farm, the downstream wind turbine operates in greedy optimal conditions, while the upstream one is dynamically controlled. The optimal control sequence is determined through EMPC. The

three-dimensional model is utilised for wake steering. The vortex modelling results in wake curling, known as the kidney shape. It is also found in other studies [29, 69, 70]. A three-dimensional vortex ring becomes a vortex pair in 2D, significantly reducing the number of data points required to describe the wake. Through online optimisation of the induction factor, the 2D model serves as the internal model to perform DIC. The results in [1] show that the two-dimensional model can potentially increase the power production of the wind farm with DIC while significantly increasing the simulation efficiency compared to LES.

## 1-5 Research scope

Hence, this research project will employ the two-dimensional FVW model [1] to explore the potential of dynamic induction control further to maximise the power production of the two-turbine wind farm. The remainder of the report is structured as follows.

In Chapter 2, the model and results proposed in [1] are discussed. It should be noted that both the model and results are presented by the writers of [1], and they will serve as the baseline case from which this research will start. The simulation of 300 time steps in the paper ran for over 10 hours on a cluster with two cores. The study will determine whether any improvements can be made. Additionally, the control method has some limitations, which will be elaborated on after discussing the baseline results.

Furthermore, sinusoidal functions are introduced in Chapter 3 to approximate the periodic steady-state control signals to reduce the optimisation's dimensionality and allow for gradient-free optimisation. The optimal steady-state values of the parameterised control variables are determined through infinite horizon control. Then, the parameterisation is implemented in the economic model predictive control framework to obtain initial results in Chapter 4. They are further explored in Chapter 5 where the control horizon is increased to mitigate the finite horizon effects and converge the wind farm towards the infinite horizon optimal periodic steady-state. In Chapter 6, the delay present in the system is approximated to eliminate irrelevant segments in the objective function evaluation. It results in the optimal control design, which maximises the probability of convergence towards the infinite horizon optimum.

# Chapter 2

# Free vortex wake model

This study will utilise the two-dimensional, control-oriented wake model proposed in [1] to extend the analysis of the wake effects on a wind farm. The wind farm consists of two wind turbines standing in line. The objective is to maximise the wind farm's power production, which is pursued by performing DIC. The main model features are described in the next section, after which the results mentioned in [1] will be discussed. These will serve as the baseline for this study, and the results will be compared.

## 2-1 Description of the main model features

The aerodynamic wake model in [1] describes the propagation of a single wake through a wind farm consisting of two inline wind turbines. The upstream wind turbine is modelled with the ADM. It is assumed to be uniformly loaded, resulting in rings of swirling vortices being released along the discs' edges at every time step. This corresponds to pairs of these vortices forming the rings in a two-dimensional representation. To enhance the wake mixing, the vortex rings should interact so that leapfrogging occurs. The second turbine's performance is evaluated without considering its effects on the wake. The wind is assumed to be uniform, and both turbines are aligned such that their rotor surface is perpendicular to the flow direction. The non-dimensional time step is set to $h = 0.2$.

The FVW model is defined as a discrete-time nonlinear state space as

$$\begin{aligned} \boldsymbol{q}_{k+1} &= f(\boldsymbol{q}_k, \boldsymbol{m}_k) \\ \boldsymbol{y}_k &= g(\boldsymbol{q}_k, \boldsymbol{m}_k), \end{aligned} \tag{2-1}$$

where $k$ is the current time step. The states $\boldsymbol{q}_k$ and control inputs $\boldsymbol{m}_k$ are used for the state update and to calculate the output vector $\boldsymbol{y}_k$.

The general input definition includes the two wind turbines' induction factor $a$ and yaw angle $\psi$. The lack of power gain for varying yaw angles in the two-dimensional FVW model is

illustrated in [1] through a yaw angle parameter sweep. Therefore, the $\psi$ for both wind turbines will remain constant at zero in this study. The wake is excited through dynamic setpoint variations of the upstream wind turbine, and the second one operates in greedy optimal conditions $a_1 = a_g = 0.33$. The control input vector $\boldsymbol{m}$ therefore becomes

$$\boldsymbol{m} = \begin{bmatrix} a_0 \end{bmatrix}. \tag{2-2}$$

The state vector $\boldsymbol{q}_k$ contains the discretised wake states $\boldsymbol{X}$ and $\boldsymbol{\Gamma}$, the stored free-stream velocities $\boldsymbol{U}$, and control inputs from the previous time step $\boldsymbol{M}$ as

$$\boldsymbol{q} = \begin{bmatrix} \boldsymbol{X} \\ \boldsymbol{\Gamma} \\ \boldsymbol{U} \\ \boldsymbol{M} \end{bmatrix}. \tag{2-3}$$

In [1], 60 vortex pairs encapture the entire wake. At time $k$, each element has a position $\boldsymbol{x}_k$ and strength $\Gamma_k$. The vortex positions are updated through a discrete-time addition of the time step $h$ multiplied by the sum of the induced velocity vector $\boldsymbol{u}_{ind}(\boldsymbol{x}_k, \boldsymbol{q}_k)$ and the uniform inflow vector $\boldsymbol{u}_\infty$. At any point, $\boldsymbol{u}_{ind}$ is defined as the sum of induced velocities $\boldsymbol{u}_j$ of all contributing elements in the system. Each vortex element's velocity at position $\boldsymbol{x}_0$ is induced by another at position $\boldsymbol{x}_1$, calculated by the two-dimensional Biot Savart law as

$$\boldsymbol{u}_j(\boldsymbol{x}_0, \boldsymbol{x}_1) = \begin{bmatrix} -r_\mathrm{y} \\ r_\mathrm{x} \end{bmatrix} \left( \frac{\Gamma}{2\pi} \frac{1}{\|\boldsymbol{r}\|^2} \right) \left( 1 - \exp\left( -\frac{\|\boldsymbol{r}\|^2}{\sigma^2} \right) \right), \tag{2-4}$$

with $\boldsymbol{r} = \boldsymbol{x}_1 - \boldsymbol{x}_0$ the relative position and $\Gamma$ the vortices' strength. The final term in the equation is a Gaussian core, which prevents any numerical errors when two elements are very close to each other and $\boldsymbol{r}$ approximates zero.

The vortices are released from the upstream actuator disc containing a circulation strength. The initial strength of both elements is directly related to $a_0$ through the following expression

$$\Gamma = \frac{1}{2} C'_\mathrm{T}(a_0) \left( \boldsymbol{u}_\mathrm{r} \cdot \boldsymbol{e}_\mathrm{x} \right)^2 h, \tag{2-5}$$

where $C'_\mathrm{T}$ is the local thrust coefficient and $\boldsymbol{u}_\mathrm{r}$ is the average wind speed vector. The direction of the strength is orthogonal to the rotor disc through the dot product with the unit vector $\boldsymbol{e}_\mathrm{x} \in \mathbb{R}^2$ which points in the downstream direction. In the FVW model, the strength is assumed to be inherited and remains constant as the vortices propagate downstream.

Finally, the output vector contains the power of the two wind turbines as

$$\boldsymbol{y} = \begin{bmatrix} P_0 \\ P_1 \end{bmatrix}, \tag{2-6}$$

where the power of wind turbine $P_i$ is calculated as

$$P_i = \frac{1}{2} C'_\mathrm{P}(a_i) \rho A_\mathrm{r} \left( \boldsymbol{u}_\mathrm{r} \cdot \boldsymbol{e}_\mathrm{x} \right)^3. \tag{2-7}$$

$C'_{\mathrm{P}}$ is the local power coefficient which is derived from $C_{\mathrm{P}}$ in Equation 1-2 by considering the power calculation of the production in terms of $\boldsymbol{u}_{\mathrm{r}}$ instead of $\boldsymbol{u}_{\infty}$

$$P = \frac{1}{2} C_{\mathrm{P}}(a) \rho A_{\mathrm{r}} \left(\boldsymbol{u}_{\infty} \cdot \boldsymbol{e}_{\mathrm{x}}\right)^3 = \frac{1}{2} C_{\mathrm{P}}(a) \rho A_{\mathrm{r}} \frac{\left(\boldsymbol{u}_{\mathrm{r}} \cdot \boldsymbol{e}_{\mathrm{x}}\right)^3}{(1-a)^3}. \tag{2-8}$$

$C'_{\mathrm{P}}(a)$ then becomes

$$C'_{\mathrm{P}}(a) = \frac{P_i}{\frac{1}{2}\rho A_{\mathrm{r}} \left(\boldsymbol{u}_{\mathrm{r}} \cdot \boldsymbol{e}_{\mathrm{x}}\right)^3} \frac{C_{\mathrm{P}}}{(1-a)^3} = \frac{C_{\mathrm{P}}}{(1-a)^3} = \frac{4a(1-a)^2}{(1-a)^3} = \frac{4a}{1-a} \tag{2-9}$$

### 2-1-1 Economic model predictive control in the free vortex wake model

As mentioned, the FVW model is proposed as a control-oriented model that can be utilised for model-based control methods. In [1], EMPC is deployed to maximise the power production of the wind farm. The objective function in Equation 2-10 consists of two terms. The first is the sum of all powers produced by the two turbines from the current time step $k_0$ over the control horizon $N_{\mathrm{h}} = 100$. The second term is the difference in control values between subsequent time steps to penalise the control inputs. It is a physical constraint added to the control problem because the setpoint changes in wind turbines cannot be too large.

The output weights are set negative to $\boldsymbol{Q} = \begin{bmatrix} -1 & -1 \end{bmatrix}$ to maximise the power output and the input weights positive to $\boldsymbol{R} = \begin{bmatrix} 10 \end{bmatrix}$ to minimise the input changes. This configuration has been found through a parameter sweep and resulted in the best balance between power outputs and the control input penalties. In this study, the same objective function configuration will therefore be used.

$$\min_{\boldsymbol{m}_k} \sum_{k=k_0}^{k_0+N_{\mathrm{h}}} J\left(\boldsymbol{q}_k, \boldsymbol{m}_k\right) = \min_{\boldsymbol{m}_k} \sum_{k=k_0}^{k_0+N_{\mathrm{h}}} \boldsymbol{Q}\boldsymbol{y}_k\left(\boldsymbol{q}_k, \boldsymbol{m}_k\right) + \boldsymbol{\Delta}\boldsymbol{m}_k^{\mathrm{T}} \boldsymbol{R} \boldsymbol{\Delta}\boldsymbol{m}_k \tag{2-10}$$

The optimal control input sequence $\boldsymbol{m} = (\boldsymbol{m}_{k_0}, \boldsymbol{m}_{k_0+1}, ..., \boldsymbol{m}_{N_{\mathrm{h}}})$ is found through gradient-based optimisation. In [1], the gradient is determined with the adjoint method where, at every time step, the gradient is constructed through backward propagation of adjoint states [71]. It is a method that computes the exact gradient of a discrete objective function.

## 2-2 Baseline results

With the FVW model and EMPC controller stated, the results from [1] are discussed. They will serve as the baseline results for this thesis project. The research done in this project will be compared to the baseline results to see whether any improvements are made. The results shown in this section are based on different processing of the same data obtained from [1].

### 2-2-1 IEA 15 MW reference wind turbine

For convenience, all units in the model have been non-dimensionalised by the normalised free-stream velocity $u_\infty$ and the rotor diameter $D$. For a better sense of the results, the IEA 15 MW wind turbine is chosen as the reference wind turbine because compared to the state-of-the-art wind turbines produced today, it is similar in size [72]. The results in this report will be displayed by adding its dimensions back to the model. The parameters are provided in Table 2-1. The power coefficient of the reference turbine is not equal to $C_{\mathrm{P,max}}$, which is accounted for by multiplying the power production with the ratio of the two power coefficients. Furthermore, some minor losses are present in the drive train of the wind turbine. Therefore, the power production slightly reduces through the multiplication with the efficiency $\eta$.

**Table 2-1:** Design parameters to dimensionalise the results from the simulations.

| Name | Property | Value | Unit |
|---|---|---|---|
| Diameter | $D$ | 240 | m |
| Wind speed | $u_\infty$ | 10.59 | $\frac{\mathrm{m}}{\mathrm{s}}$ |
| Density | $\rho$ | 1.225 | $\frac{\mathrm{kg}}{\mathrm{m}^3}$ |
| True power coefficient | $C_{\mathrm{P}}$ | 0.489 | – |
| Maximum power coefficient | $C_{\mathrm{P,max}}$ | 0.593 | – |
| Drive train efficiency | $\eta$ | 0.965 | – |

### 2-2-2 Baseline greedy simulation

The initial simulation shows both wind turbines operating in greedy conditions at $a_0 = a_1 = a_{\mathrm{g}}$, and no dynamic control on the first turbine is active. Figure 2-1 shows the two power productions. The wind turbines have high power production for the first twenty seconds. However, this results from transient effects that must leave the system. This transient behaviour is a characteristic in which the model searches for its steady state and is related to a delay in the system. The vortices being released at the beginning of the simulation need time to propagate all the way downstream. Before the wake has reached $\mathrm{T}_1$, it is considered operating in the surrounding air conditions and does not experience any wake effects. The power production of $\mathrm{T}_1$ is therefore high and similar to the steady state value of $\mathrm{T}_0$ at the beginning of the simulation. This modelling effect can be removed by initialising many time steps before starting the simulation.

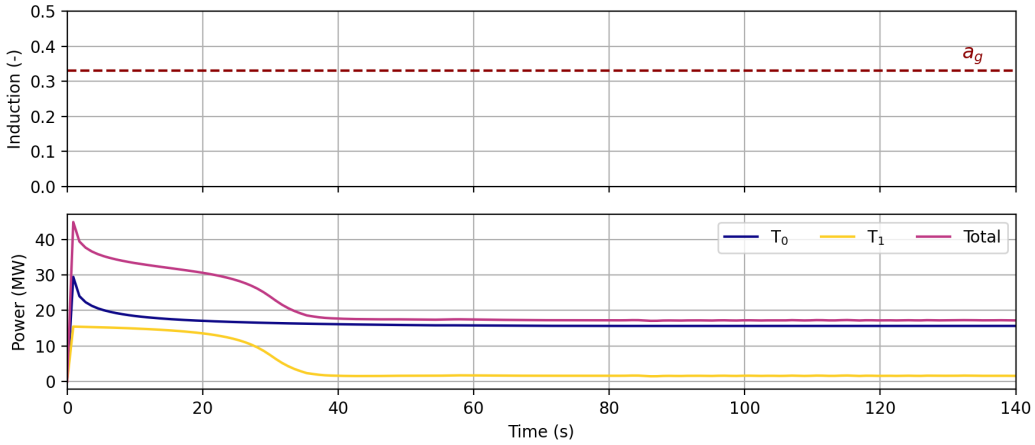


**Figure 2-1:** The power production of both wind turbines over time. The initial high performance is due to transient behaviour where the wake has not propagated through the entire wind farm. Data retrieved from [1].

After the transient simulation, the system converges to the expected steady state. The wind turbines produce a constant total power output of $P = 17.2\,\text{MW}$, shown in Figure 2-2. The figure clarifies the significant difference in power production between the two wind turbines attributed to the different wind conditions in which they operate. Being responsible for over 90% of the total output, the overall power production is heavily emphasised on $T_0$.

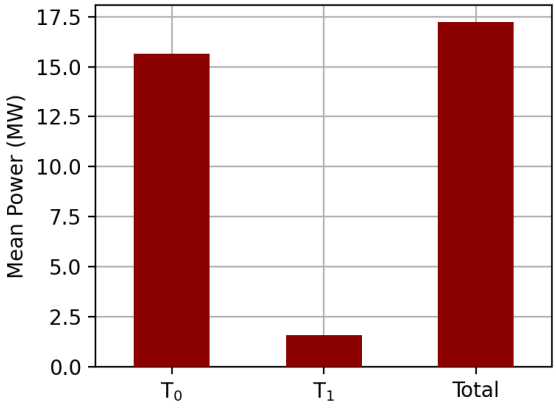


**Figure 2-2:** Greedy, steady-state power production of the two wind turbines and their total. The power production decreases due to operation in the wake stands out. $T_0$ produces a mean power of $P = 15.5\,\text{MW}$ and accounts for over 90% of the total power production.

The velocity field in Figure 2-3 illustrates the wind speed deficits in the wake, and it is clear that $T_1$ operates in reduced wind speed conditions. The assumption of a uniformly loaded actuator disc is reflected in the wake through the even lateral wind speed distributions. Additionally, the velocity deficit contributes to the expected expansion of the wake by mass conservation laws.

Furthermore, the wind speed reduction right before the upstream wind turbine can be noticed.

As mentioned in Chapter 1, this is due to the tube of air that flows through wind turbines expanding right before it reaches the turbine. One could say that the wind is blocked or held up in a way by the wind turbine before it passes through.

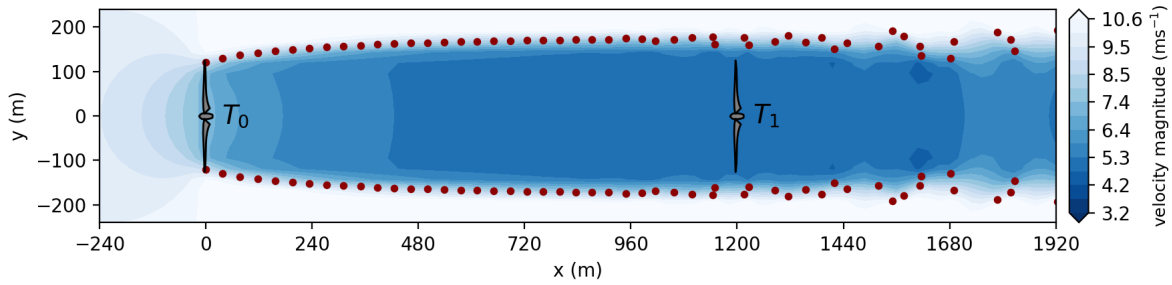


**Figure 2-3:** Snapshot of the two-dimensional wake flow velocity field, including the vortices along the boundaries marked in red. The wake breakdown is initiated around $T_1$, where the vortices start deviating from their previous path.

A final observation is the wake breakdown occurs around the second wind turbine and further downstream. The wake remains intact at around $x = 1200\,\text{m}$, and no interaction with the surrounding air is observed. Further downstream, the wake begins to break down, illustrated by the vortices deviating from their previous path. This marks the wake recovery process as the mixing of the wake with the surrounding air is initiated.

### 2-2-3 Baseline results for dynamic induction control

Figure 2-3 illustrates the considerable distance the wake propagates before recovery is initiated. As mentioned, DIC aims to enhance this wake recovery through dynamic thrust variations, which is achieved in the following simulations by deploying the economic model predictive controller discussed in Section 2-1-1. The FVW model serves as the internal model, capturing future system behaviour within the control horizon $N_\text{h}$. Through online optimisation, the control input is determined and implemented in the model as the true wind farm plant.

The optimised control signal derived by the controller is shown in Figure 2-4. The initial phase of the simulation ($t \leq 450\,\text{s}$) includes transient behaviour as the controller searches for a (periodic) steady state. Besides minor alterations, the optimised control signal follows a periodic signal for the remainder of the simulation, indicating convergence to this periodic steady state. The periodicity is reflected in the power productions of $\text{T}_0$ and $\text{T}_1$. As expected, higher induction factor values increase power production for $\text{T}_0$ and decrease it for $\text{T}_1$, which occurs later in the simulation due to the time delay.

The periodic induction factor signal is compared to the results in [20]. The mean induction factor is $\bar{a}_\text{base} = 0.3$ corresponds to a local thrust coefficient value of $C'_\text{T} = 1.75$ which is slightly smaller than the mean value in [20]. Furthermore, the periodic signal related to Equation 1-7 and characterised by the amplitude and Strouhal number $(A, St) = (0.87, 0.2)$. The value of $St$ is similar to the findings of $St = 0.25$, but the amplitude $A$ is reduced to almost half. This could be attributed to the penalty on the control input changes in Equation 2-10,

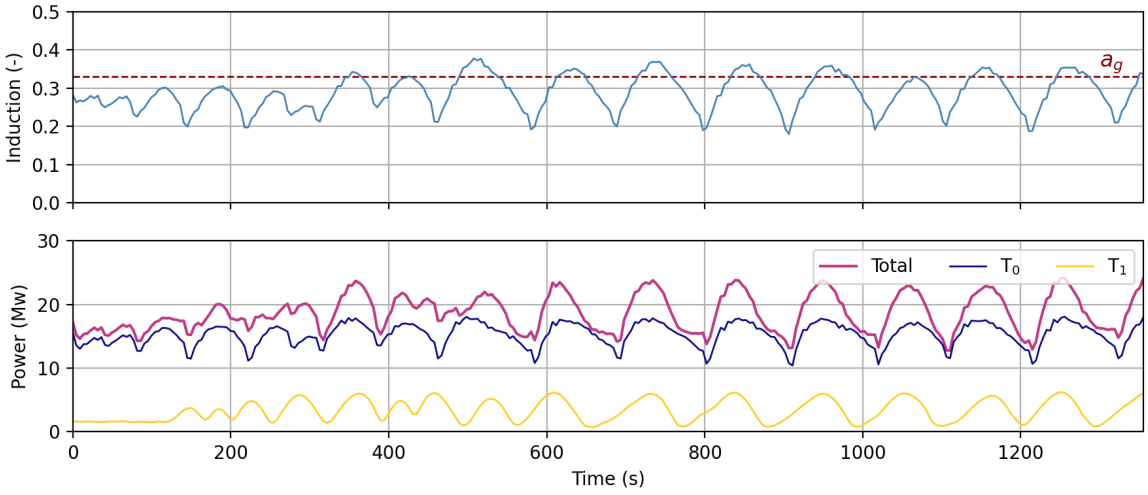


**Figure 2-4:** The wind farm's power production over time for the baseline DIC simulation. The system operates in a periodic steady-state with a non-dimensional frequency of $f \approx 0.2$ for the final two-thirds of the simulation.

which prevents steep slopes in the input signal. Higher amplitudes would correspond to larger input changes per time step.

A lower mean induction factor value implies a reduced mean wind velocity deficit in the wake. $T_0$ extracts less energy from the wind on average, allowing for a greater emphasis on $T_1$. Figure 2-5 illustrates the mean power productions of both simulations. The mean power of the DIC baseline simulation is calculated from the final two-thirds of the simulation. The figure illustrates the slight reduction in performance of $T_0$ compared to the greedy simulation. However, it is compensated for by the increased emphasis on $T_1$, which is now responsible for 19% of the mean power production, whereas this was under 10% in greedy control. The overall mean power production has increased to $P = 18.9\,\text{MW}$, which is a 9.64% increase compared to the greedy control simulation. The DIC baseline results are compared to the maximum achievable power production under power derating in [1], corresponding to an induction factor of $a_0 = 0.257$. The DIC baseline control simulation realises a 6% power gain in this context.

Snapshots of the periodic steady-state velocity field are depicted in Figure 2-6. The wake breakdown initiates around $x = 720\,\text{m}$ downstream (corresponding to three rotor diameters), presenting a significant improvement compared to the greedy velocity field in Figure 2-3. The periodicity of the control signal is reflected in the wake flow. As mentioned, a Strouhal number of about $St = 0.20$ is concluded from Figure 2-4. Consequently, the snapshots are captured every five time steps from the last full period in the simulation. The initial and final snapshots represent approximately the same stage in the velocity field.

Near $T_1$, leapfrogging occurs as a clustered group of low-speed vortices is rolled over by a group of higher-speed vortices, expanding the wake. The interaction leaves a region of contracted wake where the surrounding air is drawn in, indicated by the brighter colours of the flow field.

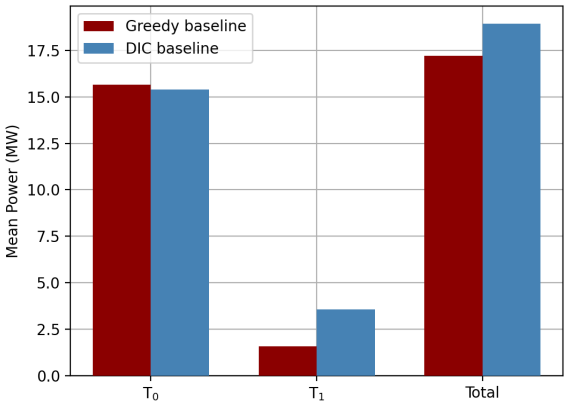


**Figure 2-5:** Mean power production comparison of the DIC baseline simulation with the greedy control simulation. The minor reduction in the mean performance of $T_0$ is compensated by the mean power increase of $T_1$, realising a 10% gain in mean power production.

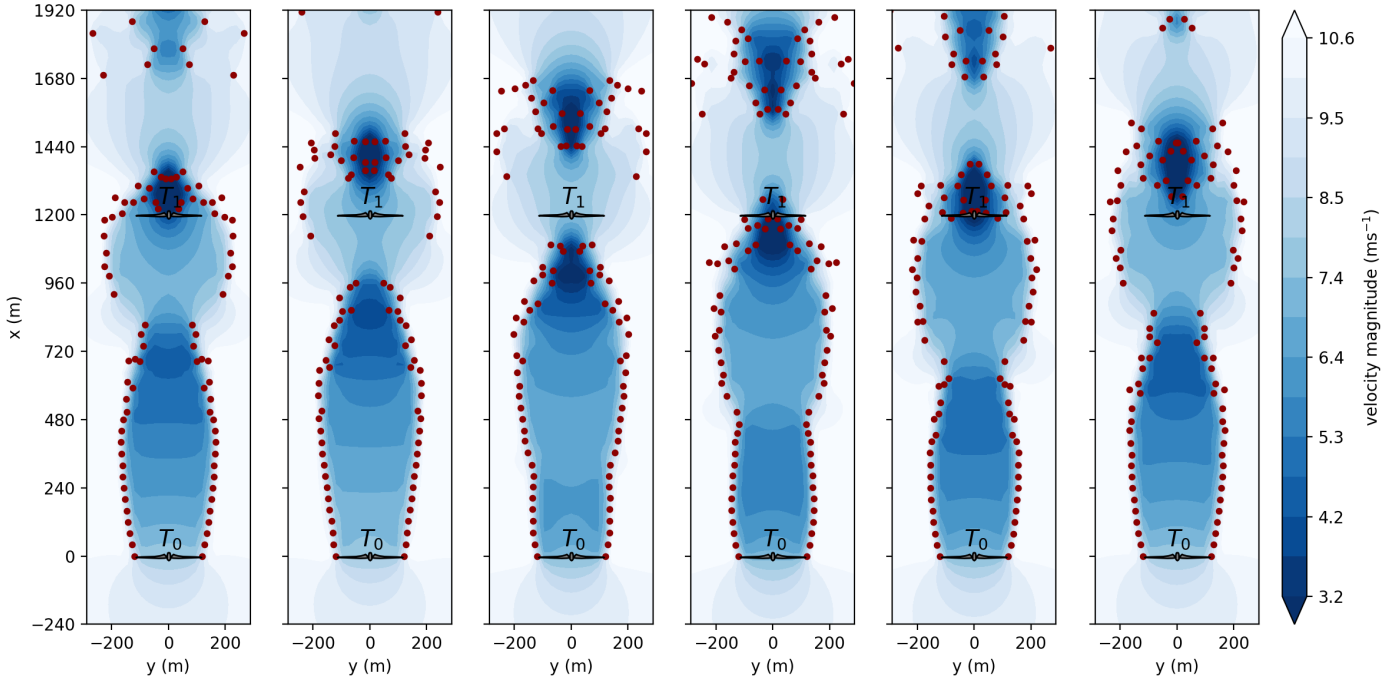


**Figure 2-6:** Wake flow velocity field snapshots of the DIC baseline simulation. The snapshots are taken five time steps apart, starting from $t = 1210\,\mathrm{s}$ until $t = 1320\,\mathrm{s}$. The snapshots show the propagation of the wake. The roll-up of the high-speed wake over the low-speed section marks the leapfrogging.

## 2-3 Baseline model features

The tradeoff between model accuracy and computational demand limits the real-time implementation of model-based control methods. This gives rise to the need for control-oriented models that capture the most dominant wake characteristics. The results from the previous chapter show that the FVW model defined in [1] can serve as the internal model in the EMPC framework. Compared to LES, the computational expense is more efficient. However, a simulation of 300 time steps ran over ten hours on a cluster computer with two cores. This leads to the following research question:

What features in the dynamic induction control simulations employing the free-vortex wake model can be improved to increase the computational efficiency while guaranteeing similar or even better results?

### 2-3-1 Finite horizon effect

In the EMPC optimal control problem, a (periodic) steady-state solution known as the turnpike exists [73]. The objective is to remain in the vicinity of the turnpike for the majority of the control horizon [74]. Due to finite horizon effects, the final segment of the control sequence may diverge from the turnpike, influenced by inherent properties within the objective. Increasing the control horizon mitigates the finite horizon effects and increases the segment of the control horizon in the vicinity of the turnpike.

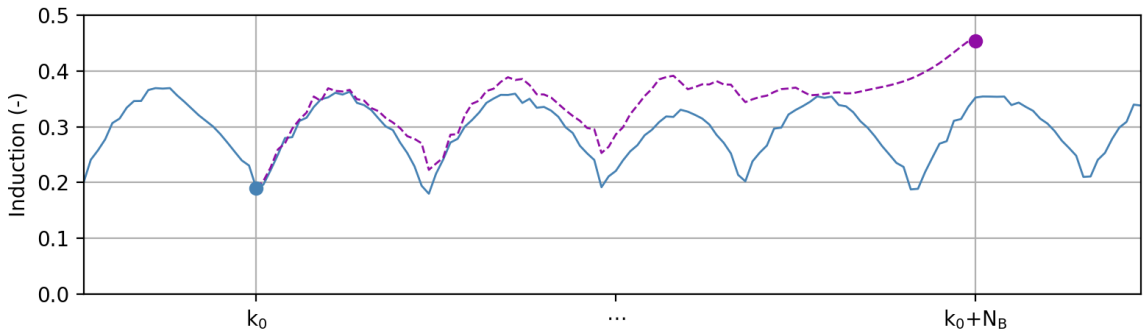


**Figure 2-7:** The resulting control horizon sequence of the optimisation at the current time step for $N_\mathrm{h} = N_\mathrm{B}$ marked in purple. It deviates from the actual control sequence because of the finite horizon effects, which cause the rise of $a_0$.

Figure 2-7 shows an example of a control horizon sequence at time step $k_0$ alongside the optimal control signal that embodies the turnpike. The sequence closely aligns with the turnpike signal for the majority of its duration but deviates towards the end of the horizon. The definition Equation 2-10 includes the power production sum of the two wind turbines for the entire control horizon. The horizon is chosen to be sufficiently long to incorporate the effects of the excited wake on $\mathrm{T}_1$ in the objective function. This is the time delay observed earlier in Section 2-2-2. However, the wind that flows through $\mathrm{T}_0$ in the final segment of the control horizon will never reach $\mathrm{T}_1$. Considering the end of the horizon to be the end of time, the controller maximises the power production of $\mathrm{T}_0$ from that wake segment. This

effect results in the turnpike tail at the end of the control sequence. A drawback of this finite horizon effect is a significant increase in the control horizon. The finite horizon effect length adds to the already large control horizon to include accurate information on the wake effects on downstream wind turbines.

### 2-3-2 Optimisation

The length of the control horizon determines the optimisation dimension in [1] given the presence of one control variable $a_0$. To ensure that control sequences predominantly remain near the turnpike, a baseline control horizon of $N_\mathrm{h} = 100$ is selected, resulting in a similar search space dimension. This large-scale search space is analysed in [1] and contains multiple local minima and flat regions due to the incorporated nonlinear wake dynamics in the objective function. As discussed in Section 1-3-3, the optimisation algorithm faces challenges related to local minima, considering it has found the optimal solution. Therefore, the high dimensionality and multimodality of the search space introduce uncertainty as to whether the found solution is the global optimum.

In [1], a gradient-based optimisation method is employed to incorporate search space curvature information into the optimisation scheme. The Adam optimiser achieves this through momentum terms containing gradient information from the previous iterations [75], reducing sensitivity towards local minima and flat regions. Moreover, the gradient information increases the convergence speed towards optima due to the included curvature information. However, a drawback is the computation time of the gradient. In [1], the gradient is computed with the adjoint method, which is an efficient gradient computation method but increases the computational expense of single model evaluations. Therefore, reformulating the objective to allow for gradient-free optimisation might improve the simulation efficiency.

In this study, the control signal is parameterised to assess the features discussed in this section. The reasoning is threefold. First, the finite horizon effects can be limited by constraining the control signal to a predefined function. Second, the optimisation dimension significantly reduces the degrees of freedom of the parameterisation function, potentially enhancing certainty about the optimality of the solutions. Finally, the parameterised function allows for a gradient-free optimisation algorithm, potentially improving the computational performance of the model evaluation.

# Chapter 3

# Infinite horizon control with parameterisation

The studies on DIC conclude that wake mixing can be enhanced through periodic excitation of upstream wind turbines. Most studies found this through simulations with model-based controllers, whereas some have focused on obtaining similar results by simplifying and approximating the optimal control signal. In [24], the closed-loop simulation converged to a periodic steady-state. An open-loop control signal was derived by extracting one period and removing higher-frequency components. This simplified signal resulted in comparable power gains to the closed-loop scenario.

Similarly, in [20], a sinusoidal perturbation around the optimal local thrust coefficient induced vortex instabilities in the wake, leading to power gains ranging from two to five percent. The open-loop control signal in this study was a parameterised function defined by the amplitude and Strouhal number. The approach in this report draws inspiration from this predefined sinusoidal function. Such functions are a logical choice due to the inherent periodicity associated with DIC.

## 3-1 Parameterisation of the control sequence

Two potential sinusoidal functions can serve as base functions: a regular sine and an absolute sine function. The discretised functions are depicted in Equation 3-1 and are characterised by three parameters: the amplitude $A$, offset $O$, and frequency $f$.

$$a_0(k) = A \sin(2\pi f k h) + O \tag{3-1a}$$

$$a_0(k) = A \left|\sin(\pi f k h)\right| + O \tag{3-1b}$$

Figure 3-1 shows an example control signal for both functions. They are compared to the control horizon sequence from Figure 2-7, which shows the effect of the turnpike tail. By constraining the control horizon sequences to the predefined shape of a sinusoid, no increasing trend will be present in the control horizon sequence, which is caused by the finite horizon effects.

The optimisation variables will be $A$, $O$, and $N_\mathrm{h}$. Therefore, the sinusoidal functions are reformulated to express $f$ in terms of $N_\mathrm{h}$. The frequency is equal to one over the time period $T_\mathrm{f}$ which is equal to the number of samples per period $p_\mathrm{L}$ multiplied with the time step as

$$f = \frac{1}{T_\mathrm{f}} = \frac{1}{p_\mathrm{L} \cdot h}. \tag{3-2}$$

Because the period length is the control horizon divided by the number of periods $n$ as

$$p_\mathrm{L} = \frac{N_\mathrm{h}}{n}, \tag{3-3}$$

the sinusoidal functions in terms of $N_\mathrm{h}$ become

$$a_0(k) = A \sin\left(2\pi \frac{n}{N_\mathrm{h} \cdot h} k\right) + O, \tag{3-4a}$$

$$a_0(k) = A \left|\sin\left(\pi \frac{n}{N_\mathrm{h} \cdot h} k\right)\right| + O. \tag{3-4b}$$

$N_\mathrm{h}$ must be an integer because it represents the number of control inputs in the control horizon. It is chosen as the optimisation variable instead of $f$ because rounding errors arise by relating the frequency to horizon length. An integer number of periods is chosen to avoid any tails at the end of the control horizon that possibly cause any finite horizon effects.

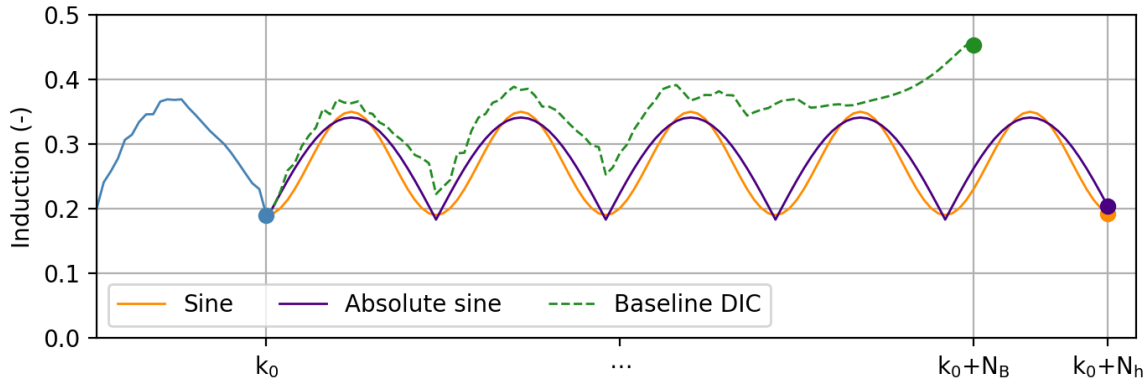


**Figure 3-1:** Examples of the sinusoidal control signals of length $N_\mathrm{h}$. The signals are compared to a control sequence found in the DIC baseline simulation with control horizon length $N_\mathrm{B}$. The predefined shape of the sinusoids limits the mean increase caused by the finite horizon effects.

The Strouhal number $St$ is a dimensionless number used in fluid dynamics to characterise the behaviour of oscillating flows, such as the occurrence of vortex shedding. It is defined as

$$St = \frac{f \cdot L}{U}, \tag{3-5}$$

where the characteristic length $L$ is equal to the rotor diameter $D$, and the flow velocity $U$ is equal to the free-stream velocity $u_\infty$. Since the model is non-dimensional, these model parameters are equal to one, which leads to the special case of $St = f$.

## 3-2 Infinite horizon control through grid search analysis

The performance of the parameterised functions is analysed to determine which sinusoidal function will be deployed within the EMPC framework. The final two-thirds of the optimal control signal in Figure 2-4 seem to fit the absolute sine better. However, applying the regular sine has resulted in power gains [20].

In Section 1-3-2, it was noted that infinite horizon, model-based control can be computationally demanding, as it involves iterating over an infinitely long horizon at each time step to determine a new control input. However, infinite horizon control can be performed without high computational demands in a system converging to a periodic steady state. The parameterisation allows the periodicity to be described by its three parameters. The system will operate in a periodic steady state by defining a parameter configuration and running the simulation sufficiently long to eliminate transient behaviour. A configuration is denoted as $\boldsymbol{v} = (A, O, N_\mathrm{h})$ and the optimal combination of optimisation variables is determined by evaluating the objective over an integer number of periods. The three-dimensional optimisation problem allows a grid search to find the optimal values. Because the control horizon is one of the optimisation variables, convergence towards an infinitely long control horizon is prevented by normalising the objective function in Equation 2-10 as

$$\min_{\boldsymbol{v}} J_\mathrm{new}\left(\boldsymbol{q}_k, \boldsymbol{v}\right) = \min_{\boldsymbol{v}} \frac{1}{N_\mathrm{h}} \sum_{k=k_0}^{k_0+N_\mathrm{h}} J\left(\boldsymbol{q}_k, \boldsymbol{v}\right). \tag{3-6}$$

### 3-2-1 Grid search analysis

In [1], an exploration of the objective function has shown that the high-dimensional search space contains multiple local minima and flat regions. The large degrees of freedom result in various configurations that give (local) optimal cost values. Through the parameterisation, these degrees of freedom are significantly reduced from one hundred to three, limiting the number of optimal solution configurations. This is attributed to the predefined shape of the control signals. Varying one optimisation variable impacts the entire shape of the control sequence. Therefore, both landscapes have become smoother through the parameterisation.

With the newly defined objective function, grid searches are conducted, including both parameterisation functions. A range of values is defined for each optimisation variable, and the FVW model is simulated for each configuration. Before evaluating the cost over the control horizon, the transient behaviour is eliminated by controlling the upstream wind turbine sufficiently long with the same sinusoid configuration. A range of integer numbers is defined for the control horizon, and the optimal horizon length is determined for $n = 5$ periods. A refined

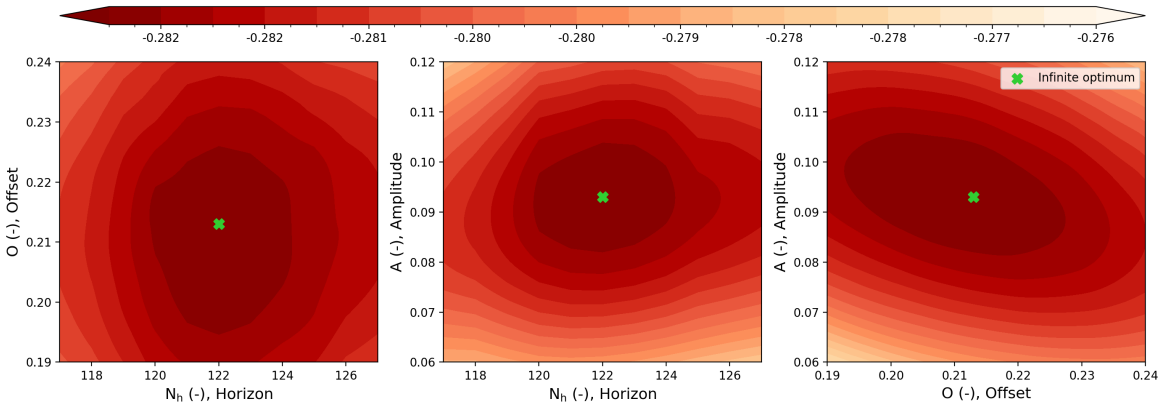


**Figure 3-2:** Grid search to find the optimal parameter configuration for the regular sine function for infinite horizon control. The optimum is marked with the green cross.

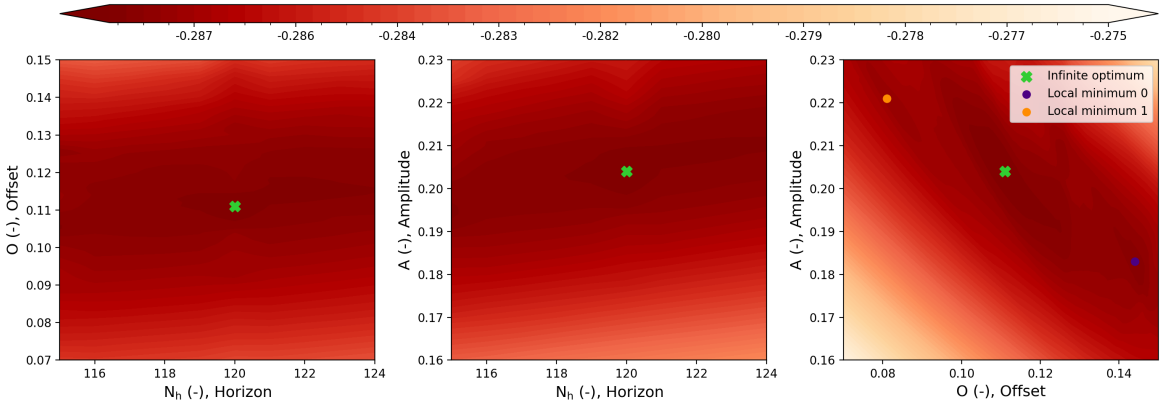


**Figure 3-3:** Grid search to find the optimal parameter configuration for the absolute sine function for infinite horizon control. The optimum is marked with the green cross. Subsequently, two local minima are marked as well.

grid with steps of 0.001 is defined for the amplitude and offset to allow detailed visualisation of the search spaces.

Figure 3-2 shows the search space of the regular sine function, which has become a large convex bowl in all directions containing one global minimum. The absolute sine search space in Figure 3-3 is nonconvex with one global optimum. The grid search shows two local minima marked in orange and purple for a constant control horizon. The number of levels in this figure is significantly higher than for the regular sine to visualise these local minima. No other minima were found by increasing the number of levels in Figure 3-2. Therefore, to distinguish between the levels, a smaller number is chosen. The optimal values for both grid searches are

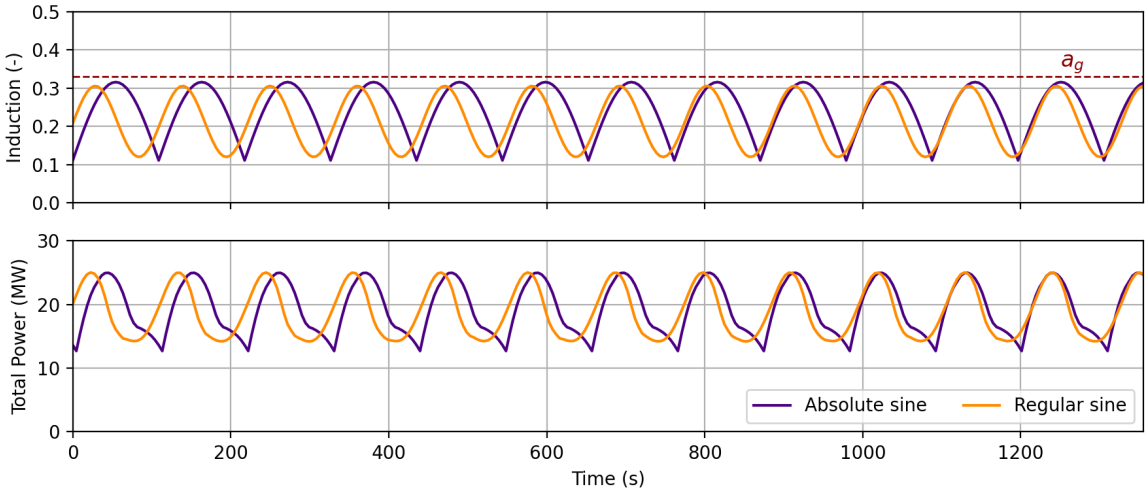


**Figure 3-4:** The infinite horizon simulations for both parameterisations. The half-phase shift over the entire simulation marks the difference in optimal frequency. Both signals do not exceed $a_{\mathrm{g}}$, which suggests a reduced value for $\bar{a}_{\mathrm{S}}$ and $\bar{a}_{\mathrm{AS}}$, and a steady derate component in the infinite horizon optimal control signals.

$$\boldsymbol{v}_{\mathrm{S},\infty} = (A, O, N_{\mathrm{h}}) = (0.093, 0.213, 122), \tag{3-7a}$$

$$\boldsymbol{v}_{\mathrm{AS},\infty} = (A, O, N_{\mathrm{h}}) = (0.204, 0.111, 120), \tag{3-7b}$$

where the subscript S denotes the regular sine and AS the absolute sine. The frequencies associated with the control horizons are

$$\begin{aligned} f_{\mathrm{S}} &= 0.205, \\ f_{\mathrm{AS}} &= 0.208, \end{aligned} \tag{3-8}$$

which are both very similar values to the Strouhal number from [1]. This alignment suggests that the optimal oscillation number for the shedding of vortices is determined for these wind conditions and the wind farm layout configuration.

Figure 3-4 shows the sinusoidal optimal control signals for the parameterisations, and we observe the minor difference in frequency because of the half-phase shift over the entire simulation. As anticipated, the periodic control input signals result in periodic power productions. The optimal sine configuration $\boldsymbol{v}_{\mathrm{S},\infty}$ results in a sinusoid that consistently remains below the greedy value $a_{\mathrm{g}}$, with its maximum reaching $a_{\mathrm{S}} = 0.306$. This contrasts with the behaviour observed from the DIC baseline simulation in Figure 2-4, where the peaks exceed $a_{\mathrm{g}}$. As a result, the mean induction factor drops significantly to $\bar{a}_{\mathrm{S}} = 0.213$ whereas the mean induction factor value in the DIC baseline is equal to $\bar{a}_{\mathrm{base}} = 0.3$. The drop in mean induction factor to $\bar{a}_{\mathrm{AS}} = 0.241$ is less significant for the absolute sine. Consequently, its maximum is $a_{\mathrm{AS}} = 0.315$. The significant reductions in the mean induction factor suggest that the infinite horizon optimal control signals also include a steady derate component combining the two control strategies.

This is reflected in the mean power productions for $\mathrm{T}_0$, shown in Figure 3-5. The figure compares the mean power productions of the two infinite horizon simulations with the periodic

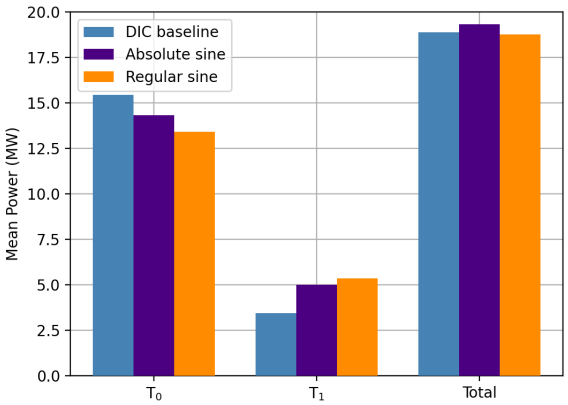


**Figure 3-5:** Power production comparison between the DIC baseline simulation and the infinite horizon simulations for the regular and absolute sine. The reduced values for $\bar{a}_{\mathrm{S}}$ and $\bar{a}_{\mathrm{AS}}$ reflect in the mean power decrease for $\mathrm{T}_0$ and the increase for $\mathrm{T}_1$.

steady-state of the DIC baseline simulation. The combination of the steady derate component and the sinusoidal excitation increases power production for $\mathrm{T}_1$ because the wake has a reduced mean velocity deficit. For the regular and absolute sine, mean power increases of 57% and 48% are realised, respectively. $\mathrm{T}_1$ accounts for 28% and 26% of their respective total power production, which is a vast increase compared to the 18% of the DIC baseline simulation.

Figure 3-6 and 3-7 show the wake propagations through the wind farm for the regular and absolute sine. A snapshot is taken every five time steps, showing the propagation through one entire period of the sinusoids. As a result, the initial and final snapshots are very similar, indicating an evident periodicity in the wake. Both figures show very symmetric wakes. The vortices propagate in a very structured manner and remain aligned for the larger part downstream. This indicates reduced turbulence, which could result from the smooth excitation functions employed. Furthermore, the colour variations behind $\mathrm{T}_0$ reflect the periodic thrust variations. The brighter colours in both figures compared to Figure 2-6 are the result of the reduced value for $\bar{a}_{\mathrm{S}}$ and $\bar{a}_{\mathrm{AS}}$.

The wake's breakdown is initiated further downstream than in the DIC baseline simulation, where the vortices start deviating from the wake's boundary at $x = 720\,\mathrm{m}$ from $\mathrm{T}_0$. However, the mixing with the ambient air through leapfrogging is visible in both figures. Between $x = 720$ and $x = 960\,\mathrm{m}$, a high-speed vortex ring is located and approaches a slower-evolving vortex ring near $\mathrm{T}_1$. As time proceeds, they start to interact, and the high-speed vortex ring rolls over the other. It leaves a tight part of the wake the surrounding air is drawn into. A subsequent low-speed vortex ring approaches again to get rolled over by another high-speed ring.

Both parameterisations result in a shift in emphasis on the wind turbines' power productions caused by their reduced mean induction factor values. The increase in mean production of $\mathrm{T}_1$ for the regular sine does not compensate for the losses made by $\mathrm{T}_0$. It results in a mean power production of $P = 18.77\,\mathrm{MW}$, which is a minor decrease of 0.06% compared to the DIC baseline simulation. The optimal infinite horizon configuration for the absolute sine better balances the emphasis and results in a mean total power production of $P = 19.34\,\mathrm{MW}$, an

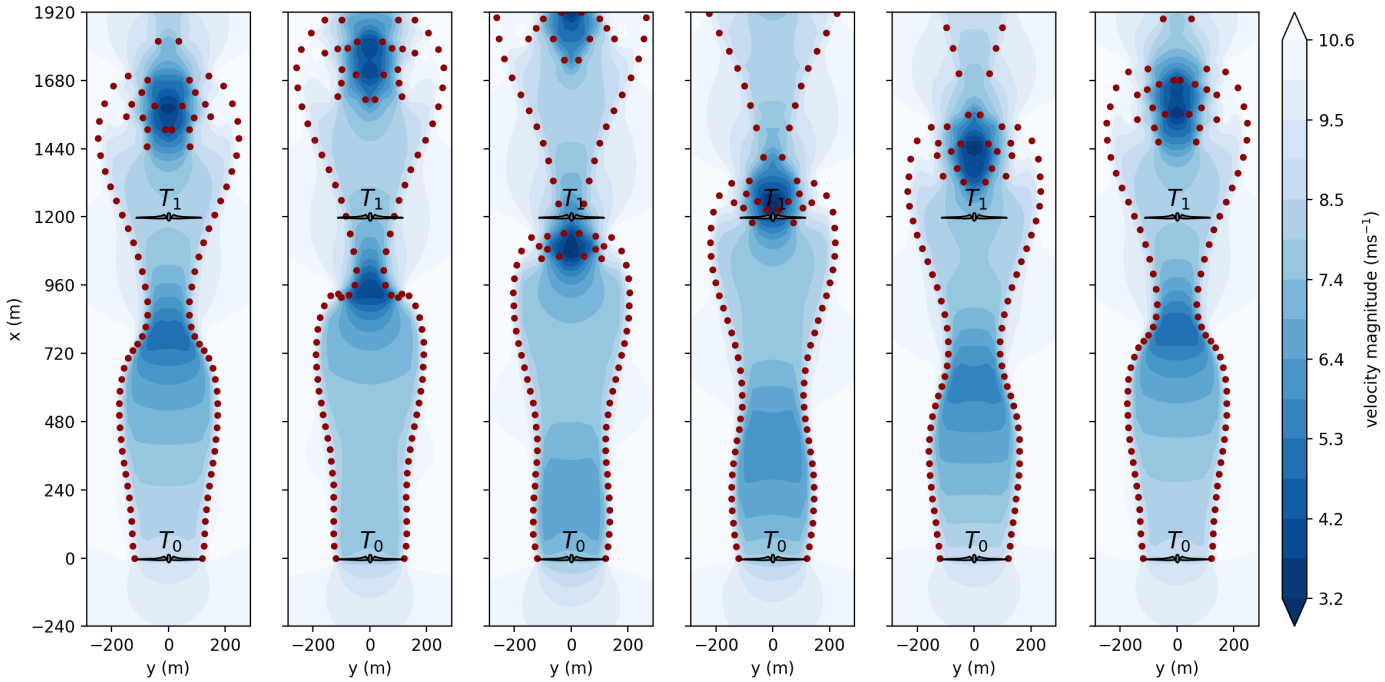


**Figure 3-6:** Snapshots of the wake flow through the wind farm for the regular sine infinite horizon simulation. The bright colours indicate a reduced mean velocity deficit in the wake caused by the decrease in $\bar{a}_{\mathrm{S}}$. It results in the significant performance improvement of $\mathrm{T}_1$.

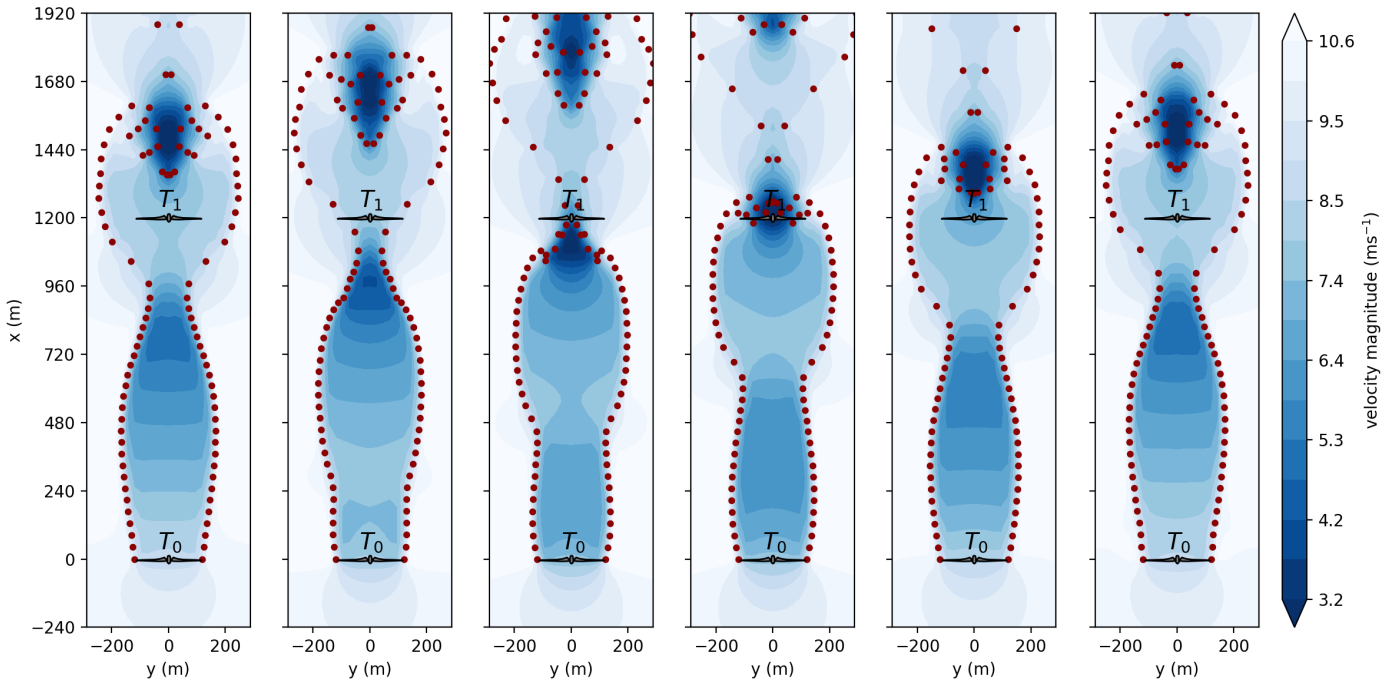


**Figure 3-7:** Snapshots of the wake flow through the wind farm for the absolute sine infinite horizon simulation. The darker colours compared to the regular sine snapshots indicate a relative increase in $\bar{a}_{\mathrm{AS}}$ compared to the regular sine simulation.

increase of 2.43% compared to the DIC baseline simulation and a 12.3% gain over to greedy control.

The improved performance of the absolute sine function can be attributed to its shape. The DIC baseline control signal converged towards a periodic signal similar to the absolute sine. Moreover, the absolute sine function exhibits steeper slopes, resulting in a higher mean induction factor value. This is due to the minor time that $T_0$ is excited by very low induction factor values, thereby preventing excessive power production reductions. The steeper slopes incur a larger penalty on control input changes, which seem to be outbalanced by the improved performance. This is supported by the colour bars in Figure 3-2 and 3-3, where the objective function cost of the absolute sine infinite optimum is better compared to the regular sine.

## 3-3 Performance comparison minima for absolute sine

The absolute sine will serve as the parameterisation function for the remainder of this study. The goal is to design receding horizon controllers that can maximise the probability of convergence towards the infinite optimum value $\boldsymbol{v}_{\mathrm{AS},\infty}$. $\boldsymbol{v}_{\mathrm{AS},\infty}$ will be referred to as $\boldsymbol{v}_\infty$ for the remainder of this report. Figure 3-3 shows the nonconvexity of the absolute sine search space, where two local minima were identified. The steady-state performances for these configurations are compared to the infinite horizon optimum in this section because, in the subsequent chapter, the finite horizon controllers might converge towards these steady-state values. Optimisation algorithms will manoeuvre through similar search spaces, searching for the infinite optimum, and they might terminate in these local minima. It is, therefore, interesting to analyse the performance difference when performing infinite horizon control with these configurations.

Similar simulations are conducted for these configurations, as in the previous section. The transient and optimal control signal are defined by the configurations in Equation 3-9. The transient simulation is sufficiently long to eliminate all transient behaviour in the system and ensure the system is already in its periodic steady state when the simulation starts.

$$\boldsymbol{v}_{\mathrm{local},0} = (A, O, N_{\mathrm{h}}) = (0.183, 0.144, 120) \tag{3-9a}$$

$$\boldsymbol{v}_{\mathrm{local},1} = (A, O, N_{\mathrm{h}}) = (0.221, 0.081, 120) \tag{3-9b}$$

The first local minimum $\boldsymbol{v}_{\mathrm{local},0}$ has a smaller amplitude and higher offset compared to the infinite optimum, which increases the mean induction factor value of $\bar{a} = 0.261$. This increase improves $T_0$'s performance as more energy is extracted from the wind. Contrarily, the higher velocity deficits in the wake cause the mean performance reduction for $T_1$ in Figure 3-8. The figure shows the mean power differences between the infinite horizon local minima $\boldsymbol{v}_{\mathrm{local},0}$ and $\boldsymbol{v}_{\mathrm{local},1}$ with the mean performance of the infinite horizon optimum $\boldsymbol{v}_\infty$. The gains of $T_0$ cannot compensate for the reduction of $T_1$, which accounts for 22% of the total power production. A total mean power of $P = 19.26\,\mathrm{MW}$ is produced, which is a 0.39% reduction compared to the infinite horizon optimum. However, compared to the DIC baseline simulation, a gain of 2.04% is realised.

The second local minimum $\boldsymbol{v}_{\mathrm{local},1}$ is located on the other side of $\boldsymbol{v}_\infty$ in the search space, which results in the opposite relative performance in Figure 3-8. A higher amplitude and

lower offset result in a mean induction factor value of $\bar{a} = 0.221$. This causes a significant reduction in $T_0$'s power production, allowing $T_1$ to almost gain the same. $T_1$ now accounts for 30% of the total power production, realising a total mean power of $P = 19.33\,\text{MW}$, which is an almost negligible loss of 0.04% compared to the infinite horizon optimum and a 2.40% gain compared to the DIC baseline simulation. From this, we conclude that gaining overall performance by reducing $T_0$'s power production also has limitations through the absolute sine parameterisation, although the differences are minimal. Because lowering $\bar{a}$ eventually results in overall power losses again.

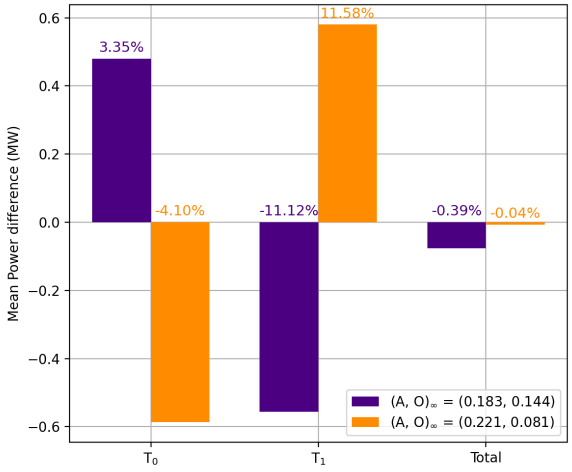


**Figure 3-8:** The mean power difference of the infinite horizon local minima $\boldsymbol{v}_{\text{local},0}$ and $\boldsymbol{v}_{\text{local},1}$ with the mean performance of the infinite horizon optimum $\boldsymbol{v}_\infty$. The relative gains are also shown. The local minima lie on opposite sides of $\boldsymbol{v}_\infty$ in the search space, resulting in the opposite performance differences.

The increased penalty on the control input changes is another reason this configuration is a local minimum compared to the infinite horizon global minimum. The higher amplitude causes steeper slopes in the absolute sine and larger control input changes. Since this is still a limitation, the infinite horizon optimum outperforms this infinite horizon configuration.

However, both local minima result in improved steady-state performance compared to the DIC baseline simulation, while the simulations were completed in the order of seconds. Furthermore, it is interesting that the two local minima result in opposite balances between the wind turbines' power productions. The increase in $T_0$'s power production for $\boldsymbol{v}_{\text{local},0}$ causes the overall wind farm to reduce its overall performance because of $T_1$'s decreased production. Although the smaller amplitude results in a smaller penalty on the control input difference, power production is still the most important component in the objective function. Notably, both configurations improve the wind farm's steady-state performance compared to the DIC baseline simulation.

# Chapter 4

# Finite horizon control with parameterisation

In the previous chapter, the wind farm's steady-state performance improved by employing the absolute sine parameterisation function, yielding a 2.43% power gain over the DIC baseline simulation. However, the current wind conditions in the FVW model are a simplified version of reality. There are no disturbances in the system and no wind speed or direction changes because the wind inflow is assumed to be uniform. The simplification allows for the determination of the optimal infinite horizon control configuration.

## 4-1 Economic model predictive control with the absolute sine

While the predefined control signal proves optimal for these particular wind conditions, its limitation lies in the inability to adapt to varying wind speeds or directions. Real-world wind farm controllers face continuous exposure to disturbances, necessitating a control solution capable of robustness and adaptability. EMPC is a commonly employed strategy to address this because its iterative design enables the controller to adapt to unexpected state measurement changes, which are the starting point for the internal model evaluation every time step [46].

In the following chapters, no disturbances are added, but a further step is taken to control the two-turbine wind farm in the FVW model through EMPC. The controller defined in [1] is redesigned to achieve convergence towards the infinite horizon steady-state values $\boldsymbol{v}_\infty$ which we identified to result in the best periodic steady-state performance.

### 4-1-1 Initialisations for all simulations

Because the simulations are conducted for the same wind conditions and wind farm layout, the frequency in the absolute sine parameterisation is set to $f = 0.208$ because it corresponds to the optimal Strouhal number for the vortex ring shedding. By allowing two degrees of freedom in the optimisation $\boldsymbol{v} = (A, O)$, the parameterised control horizon sequence becomes

$$\begin{aligned} a_0(k) &= |A \sin(\pi f h k)| + O \\ &= |A \sin(\pi \cdot 0.208 \cdot hk)| + O. \end{aligned} \tag{4-1}$$

For the remainder of the simulations in this report, two initialisation points are defined to approach the infinite optimum from two opposite sides in the search space, the bottom left and the top right, given as

$$\boldsymbol{v}_{0,k=0} = (A, O)_{0,k=0} = (0.0, 0.33), \tag{4-2a}$$

$$\boldsymbol{v}_{0,k=1} = (A, O)_{1,k=1} = (0.33, 0.0). \tag{4-2b}$$

Their interpretation is that for $\boldsymbol{v}_{0,k=0}$ the absolute sine will be a constant signal without an amplitude, whereas $\boldsymbol{v}_{1,k=0}$ initialises a control horizon signal without any offset and only an amplitude. The parameterisation introduced constraints on the degrees of freedom of the control signal, resulting in a smooth search space containing two local minima near the infinite optimum. Therefore, optimisation is expected to be concentrated in the region around the infinite optimum. Approaching this region from two sides of the search space will give better insights into the sensitivity towards initialisation.

The initialisations are chosen such that they align with the transient control input signal $a_{\mathrm{t}} = a_{\mathrm{g}}$. The model transients are removed by letting $\mathrm{T}_0$ operate in greedy conditions. Before starting dynamic control at $k = 0$, the wake must have propagated through the wind to obtain proper results. The last time step of the initialisation transient removal simulation marks the start of the simulation at $k = 0$. This time step must be chosen appropriately to align the first absolute sine control horizon sequence with the transient control input signal. Otherwise, the induction factor changes become too large, and the physical limits of the wind turbine could be reached or surpassed.

### 4-1-2 Beetle antennae search optimisation algorithm

The optimisation scheme in EMPC aims towards rapid convergence to the (global) solution. Although the adjoint states are not computed, evaluating the objective function in Equation 3-6 remains the limitation in the optimisation scheme. Therefore, the following optimisation algorithm contains a minimal number of function evaluations while achieving convergence in the search space because it conducts only three function evaluations per iteration. The beetle antennae search (BAS) algorithm is inspired by the searching behaviour of the longhorn beetle [76]. BAS mimics the beetle's use of long antennae for detecting prey, partners, or warnings. This is accomplished by pointing each antenna to a random side of the body and determining which side experiences the strongest sent.

At time step $t$, the beetle is at state $\boldsymbol{x}_t$. The algorithm determines a new search direction by selecting the best fitness function $J_{\mathrm{new}}$ at two randomly found states. A random, normalised vector $\boldsymbol{b}$ is defined as

$$\boldsymbol{b}_t = \frac{\mathrm{rnd}(n_{\mathrm{v}}, 1)}{\| \mathrm{rnd}(n_{\mathrm{v}}, 1) \|}, \tag{4-3}$$

where rnd $\sim \mathcal{N}(0,1)$ is a random sample of dimension $n_\mathrm{v}$ from a normal distribution. The two random search states are then formulated as follows

$$\boldsymbol{x}_r = \boldsymbol{x}_t + d_t \boldsymbol{b}_t, \tag{4-4a}$$

$$\boldsymbol{x}_l = \boldsymbol{x}_t - d_t \boldsymbol{b}_t, \tag{4-4b}$$

where $d_t$ is the antennae's length. The fitness function is evaluated in both directions to find the strongest sent. The state update is then formulated as

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t - \delta_t \boldsymbol{b}_t \cdot \operatorname{sign}\left(J_\mathrm{new}\left(\boldsymbol{x}_r\right) - J_\mathrm{new}\left(\boldsymbol{x}_l\right)\right), \tag{4-5}$$

with $\delta_t$ the step size, which accounts for the convergence speed. The antennae's length and step size are updated as

$$d_{t+1} = 0.93 d_t + 0.005, \tag{4-6a}$$

$$\delta_{t+1} = 0.95 \delta_t. \tag{4-6b}$$

The determination of constant values involves an iterative parameter-tuning process. The antennae's length and step size are initialised at $(d_0, \delta_0) = (0.2, 0.01)$. The termination condition for the algorithm is met when the maximum number of iterations equal to ten is reached. The algorithm returns the state corresponding to the best fitness function value $J_\mathrm{new}\left(\boldsymbol{x}_\mathrm{best}\right)$.

As discussed in Section 1-3-3, NI algorithms, including the BAS algorithm, cannot guarantee global convergence due to their reliance on random search. This inherent characteristic stems from the need for the algorithm to discover additional information to determine its convergence path. Therefore, a controller that can maximise the probability of convergence towards the infinite horizon optimum will be designed, minimising the probability of converging to other local minima.

### 4-1-3 Performance with initial control horizon

The BAS algorithm serves as the optimisation block in Figure 1-7, which optimises the following objective function

$$\min_{\boldsymbol{v}_k} J_\mathrm{new}\left(\boldsymbol{q}_k, \boldsymbol{v}_k\right) = \min_{\boldsymbol{v}_k} \frac{1}{N_\mathrm{h}} \sum_{k=k_0}^{k_0+N_\mathrm{h}} J\left(\boldsymbol{q}_k, \boldsymbol{v}_k\right), \tag{4-7}$$

for a control horizon of $N_\mathrm{h} = 120$. The objective function is similar to the one defined in Equation 3-6. However, $\boldsymbol{v}_k$ is updated every time step. The value for $N_\mathrm{h}$ closely aligns with the control horizon length from [1], balancing between computational demand and ensuring sufficient wake propagation information from the FVW model. Moreover, the parameterisation allows for the gradient-free BAS algorithm, neglecting the expensive computation of the adjoint states to compute the gradient. Ten simulations are conducted for both initialisations,

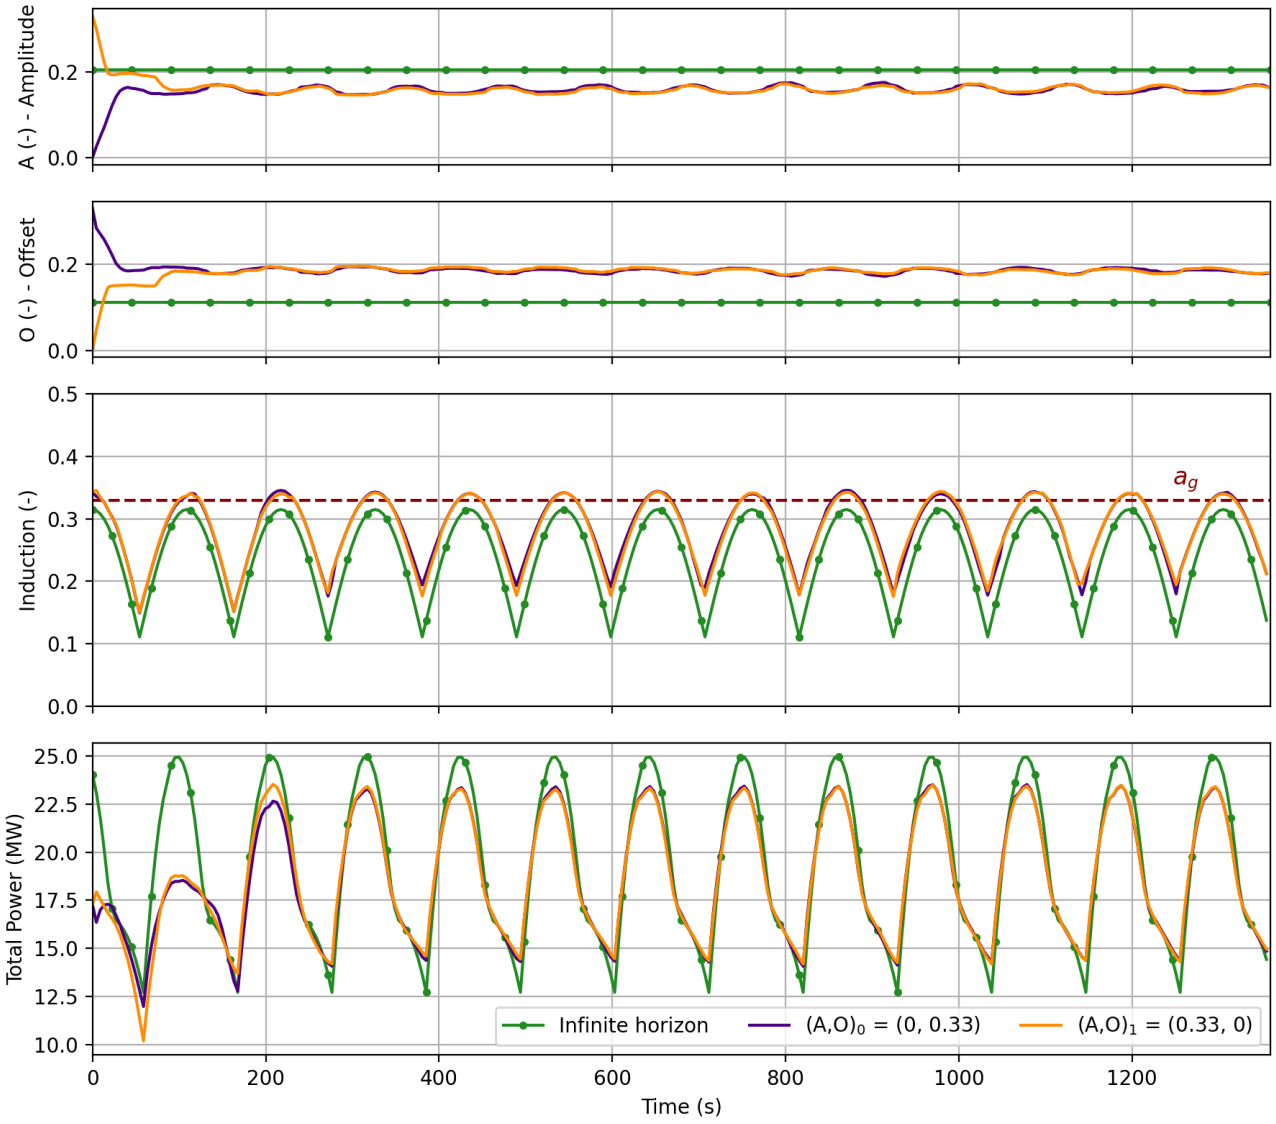


**Figure 4-1:** All mean signals concerning the ten receding horizon control simulations. Their infinite horizon equivalents are also illustrated for comparison. The quick convergence results in a periodic steady-state operation for the largest part of the simulation. The benefit of a reduced value for $\bar{a}$ is not noticed over the short control horizon, resulting in the induction factor exceeding $a_g$.

providing a more certain evaluation of the solution. The simulations run for the same number of time steps to compare the duration to the DIC baseline simulation.

The first two subplots in Figure 4-1 show the mean optimal control signals. Although they initiate from opposite directions in the search space, they quickly converge towards the same periodic steady state. The observed periodicity is attributed to a programming error leading to the omission of one control input in the final period of the absolute sine control sequences. Further clarification is provided in Section 7-2. Despite the periodicity, the optimiser efficiently identifies the paths towards the same optimal values and remains in that neighbourhood for the largest part of the simulation. The corresponding mean values and standard deviations for both initialisations are

$$\boldsymbol{v}_{0,\mathrm{ss}} = (A, O)_{0,\mathrm{ss}} = (0.159 \pm 0.0077, 0.184 \pm 0.0076), \tag{4-8a}$$

$$\boldsymbol{v}_{1,\mathrm{ss}} = (A, O)_{1,\mathrm{ss}} = (0.159 \pm 0.0078, 0.184 \pm 0.0079), \tag{4-8b}$$

The minor standard deviations and similar mean values indicate that the optimiser found the same periodic steady state in all twenty simulations. This suggests the existence of one minimum and no sensitivity towards initialisation. The optimisations algorithm converged towards different solutions from the infinite horizon optimum in Equation 3-7b. The reduced mean amplitude and increased mean offset result in a lifted absolute sine control signal compared to the infinite horizon simulation. The mean induction factor value rises to $\bar{a} = 0.283$, and the mean maximum value of $a_0 = 0.343$ exceeds the greedy value. This is not desired because the power coefficient of $\mathrm{T}_0$ reduces for induction factor values higher than $a_\mathrm{g}$. The wake deficits in those segments are larger, which decreases the overall performance of the wind farm. Figure 4-1 shows some difference in the initial power signals, which also converge to the same periodic values for the larger part of the simulation. The initial differences illustrate the system delay because the power production's periodic steady state occurs around $t = 300\,\mathrm{s}$, whereas the control variables converge in under $t = 200\,\mathrm{s}$.

Figure 4-2 illustrates the shift in emphasis compared to the infinite horizon optimal simulation. The power production of $\mathrm{T}_0$ and the velocity deficit in the wake increase for larger mean induction factor values but decrease the mean performance for $\mathrm{T}_1$. Therefore, $\mathrm{T}_1$ accounts for 20% of the total power production, whereas 26% was achieved in the infinite horizon simulation. The power increases of $\mathrm{T}_0$ do not compensate for the losses of $\mathrm{T}_1$ and result in mean power productions of $P = 19.0\,\mathrm{MW}$, a performance decrease of 1.4%.

While the mean induction factor value is smaller than $\bar{a}_\mathrm{base} = 0.3$, the absolute sine approaches the steady-state signal from the DIC baseline simulation, indicating a similar power production balance. Figure 4-3 compares the infinite and finite absolute sine simulations with the DIC baseline simulation. The mean power differences and the relative power gains are shown. The improved performance of $\mathrm{T}_0$ with respect to the infinite horizon simulation results in relative power losses of little over 1%. This increase for $\mathrm{T}_0$ causes the relative power gain for $\mathrm{T}_1$ to reduce. However, the smaller value for $\bar{a}$ allows better operating conditions for $\mathrm{T}_1$. It compensates for the losses of $\mathrm{T}_0$, thereby resulting in a minor overall performance increase of 0.78% and 0.69% for the respective initialisations.

The wake flow through the wind farm from one of the simulations is displayed in Figure 4-4. Snapshots are taken at the same instances as in previous figures between $t = 1210$ and $t =$

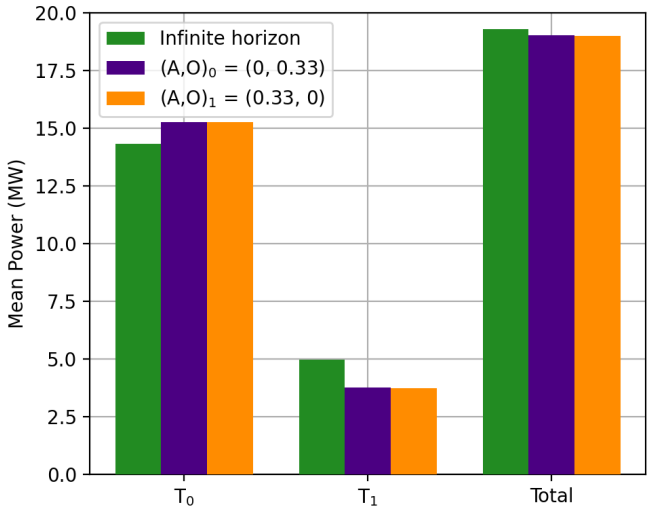


**Figure 4-2:** The mean power productions for both initialisations in steady-state compared to the infinite horizon simulation. The mean powers result from ten simulations to account for the uncertainty caused by the optimisation algorithm.

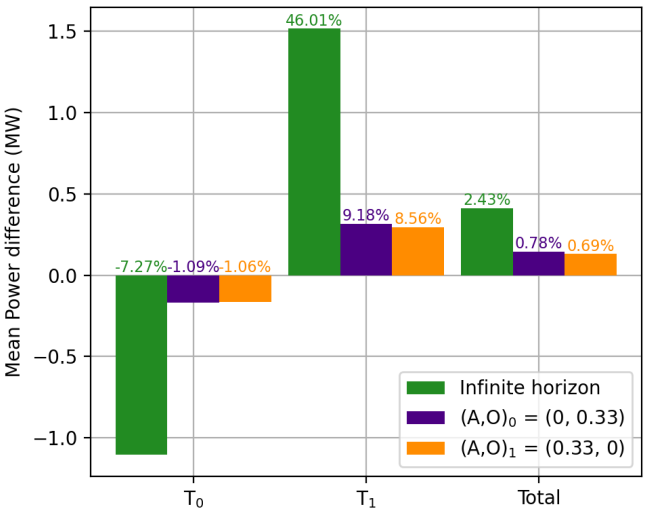


**Figure 4-3:** The performance of the infinite horizon and the two finite horizon simulations are compared to the DIC baseline simulation mean power production. We observe the slightly better performance of the finite horizon simulations.

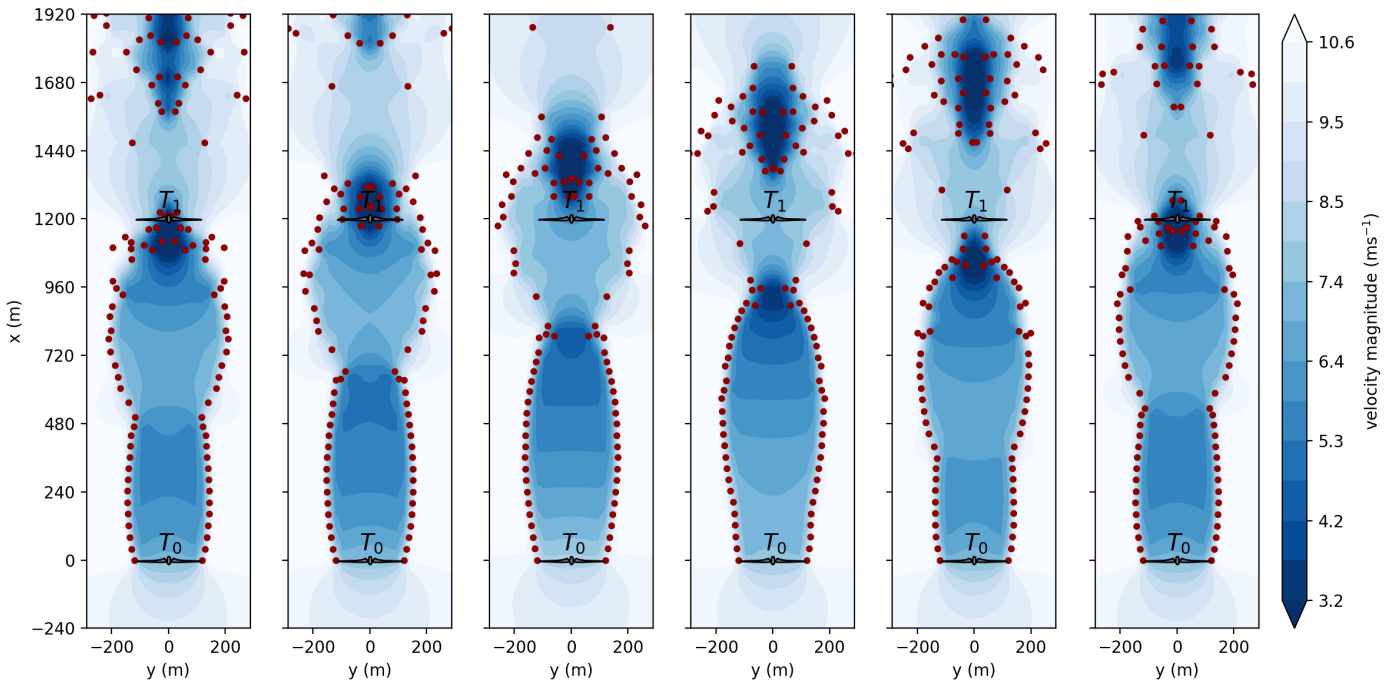


**Figure 4-4:** Snapshots of the wake flow through the wind farm for the initial finite horizon simulation with a control horizon $N_{\mathrm{h}} = 120$. The vortices start deviating from the wake boundary once the wake breakdown is initiated, which indicates increased turbulence in the wake compared to the smooth, infinite horizon control signal.

1320 s into the simulation. The system's periodicity is again evident through the resemblance of the initial and final subfigures. Furthermore, the colour variations behind $T_0$ reflect the periodic thrust variation. The leapfrogging occurs similarly to the infinite horizon control simulations, but this optimal control signal causes the vortices to deviate more from the wake boundary. A darker colour trend through the wake is visible, indicating the mean induction factor value increase.

In conclusion, the current control design, employing a receding horizon control approach, achieves performance similar to the DIC baseline simulation under the same wind conditions and wind farm configuration. On average, the simulations were completed in around 25 minutes, which is a significant reduction in simulation time by a factor of 30. While the solutions may not converge to the infinite optimum with the current control design, the observed performance is promising.

## 4-2 Initial time step search space analysis for convergence

A search space analysis is conducted to further analyse why this controller with a control horizon of $N_\mathrm{h} = 120$ is not converging towards the infinite optimum. Although the problem is time-varying, causing the search space to change at every instant, this analysis might offer a better understanding of the landscape the optimisation algorithm navigates through.

The landscape of the first time step is the same for both initialisations because of the consistent use of the same transient removal control signal $a_\mathrm{t} = a_\mathrm{g}$. Therefore, a single grid search is performed. The left contour plot in Figure 4-5 supports the preliminary conclusions from the previous section that no minimum is located near the infinite optimum point. The initial optimal cost is located at $\boldsymbol{v} = (0.163, 0.174)$, which is already in the neighbourhood of the mean steady-state values in Equation 4-8.

As discussed in Section 4-1-3, the optimal control signals converge quickly towards their periodic steady states. While the system converges, the search space changes, and the optimal cost shifts towards the mean steady-state value. When the system operates in a steady state long enough, the landscapes will stop changing. Due to the minor periodicity that remains in the control signals, it might imply minor changes over time without any significance. The middle and right contour plot in Figure 4-5 show similar final landscapes for both initialisations. Both search spaces have become large convex bowls with their optimum at $\boldsymbol{v} = (0.144, 0.194)$. Their resemblance supports the conclusion that the solution converged towards the same periodic steady-state values in all twenty simulations. As a result, control variable changes will similarly influence both systems.

From these figures, we conclude that the cost improves over time. The darker shades in the subfigures corresponding to the final simulation steps indicate better performance. This improvement is reflected in the power production signals in Figure 4-1, where the initial power production is worse than the mean in periodic steady-state attributed to transient behaviour. However, this controller design will never reach the infinite horizon configuration. For both initialisations, the optimal cost shifts away over time as the benefit of low offset and high amplitude cannot be noticed over the horizon. Instead, the solution converges towards a smaller amplitude and higher offset to increase the mean induction factor value. Although

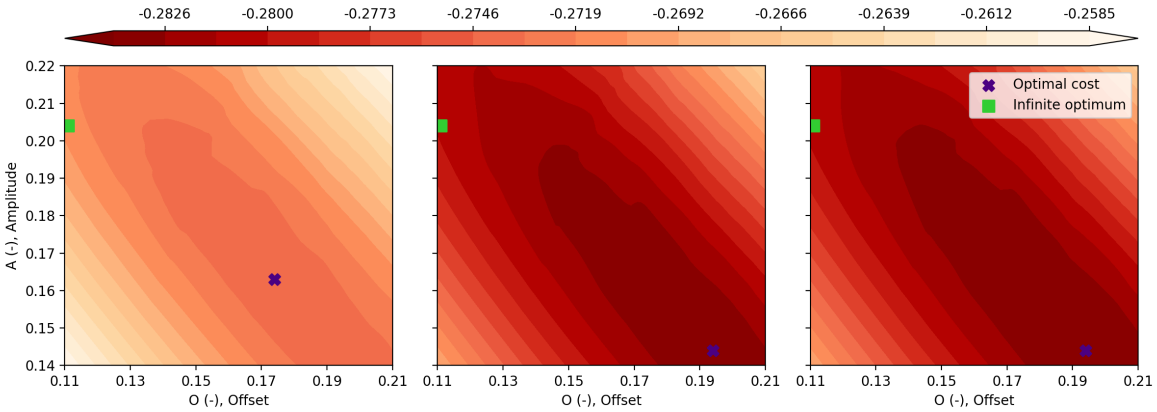


**Figure 4-5:** Contour plots of the initial and final time steps for the finite horizon simulations with $N_\mathrm{h} = 120$. The similarity between the latter two supports the conclusion on the convergence for both initialisations. Compared to the left contour plot, their darker shades indicate performance improvements over the simulation.

the infinite horizon optimum achieved proper power production balance between the two wind turbines, the finite horizon controller does not optimise towards this benefit.

Using a similar control horizon length, as defined in [1], has resulted in comparable periodic steady-state values but very different from the infinite horizon optimum. This could stem from some finite horizon effects that remain present in the objective definition. In the final segment of the control horizon sequence, the excitation of a section of the wake that never reaches $\mathrm{T}_1$ is included in the objective. For $N_\mathrm{h} = 120$, this section could have a significant influence because the optimisation algorithm maximises the final induction factor values corresponding to this section. As a result, the entire absolute sine is lifted, reflected by the increased mean induction factor value $\bar{a}$.

# Chapter 5

# Increase control horizon for convergence guarantee

The performance with the controller design in the previous chapter is promising, realising minor power production improvements while the simulations are completed 30 times faster. However, we consider the control horizon length too short for the controller to encounter the long-term benefits of a reduced mean induction factor value resulting in the lifted absolute sine signal. The goal remains to design a controller that maximises the power production of the two-turbine wind farm in the FVW model. The infinite horizon optimum $\boldsymbol{v}_\infty$ improves the steady-state performance by 2.43%. Hence, the influence of increasing the control horizon is explored to analyse whether the probability of converging towards the infinite horizon optimum can be maximised for increasing control horizons. Increasing the control horizon will increase the computational expense of the model. As discussed in Section 2-2-3, the horizon cannot be increased to infinity as the tradeoff between computational expense and performance remains. Therefore, the relevant simulations are compared for their increase in simulation time.

## 5-1 Grid search analysis for increasing control horizons

If the global optimum at the first time step of the simulation is at the desired location, the probability of convergence increases. This is attributed to the quick convergence in every simulation so far. Therefore, grid searches are conducted for increasing control horizons to analyse the shift towards the infinite optimum and seek the horizon for which the optimal cost is located at the infinite optimum. The control horizons increase with 48 control inputs per search space, corresponding to two periods per horizon.

Figure A-1 to A-4 show that increasing the number of periods in the control horizon instantly shifts the optimum from the previous optimum located at $\boldsymbol{v} = (0.163, 0.174)$ to $\boldsymbol{v} = (0.184, 0.146)$ very near $\boldsymbol{v}_{\text{local},0}$. For $N_\text{h} = 168$ (or $n = 7$), the two points are located inside one large valley, but further increasing the control horizon results in one minimum

around $\boldsymbol{v} = (0.184, 0.146)$. In Section 3-3, it was found that steady-state performance for this configuration results in 0.4% loss over the infinite horizon setpoints but a 2.04% gain over the DIC baseline simulation.

$\boldsymbol{v}_{\text{local},0}$ remains the optimum of the search spaces for very long control horizons. Figure A-5 to A-21 illustrate the search spaces for control horizons $312 \leq N_{\text{h}} \leq 1032$. As the control horizon increases, the search space becomes nonconvex with two local minima. For $264 \leq N_{\text{h}} \leq 648$, a local minimum located at $\boldsymbol{v} = (0.195, 0.125)$ arises which later vanishes into the same valley as $\boldsymbol{v}_{\infty}$. For $N_{\text{h}} = 840$, the optimal cost shifts into the valley around $\boldsymbol{v}_{\infty}$, where it remains at $\boldsymbol{v}_{\text{opt}} = (0.205, 0.111)$ for $N_{\text{h}} \longrightarrow \infty$. This new global optimum is approximately the same as the infinite horizon optimum values.

Figure 5-1 illustrates the shift in optimal cost value by plotting the cost at location $\boldsymbol{v}_{\text{local},0} = (0.183, 0.144)$ and $\boldsymbol{v}_{\text{opt}} = (0.205, 0.111)$ over the control horizon lengths. Furthermore, for better comparison, the absolute value of the normalised difference $J_{\text{norm}}$ defined as

$$J_{\text{norm}} = \frac{J_{\text{new}}\left(\boldsymbol{v}_{\text{local},0}\right) - J_{\text{new}}(\boldsymbol{v}_{\text{opt}})}{J_{\text{new}}\left(\boldsymbol{v}_{\text{opt}}\right)} \tag{5-1}$$

is displayed as well. The shift in slope on the logarithmic scale indicates the shift in optimal cost for $N_{\text{h}} = 840$.

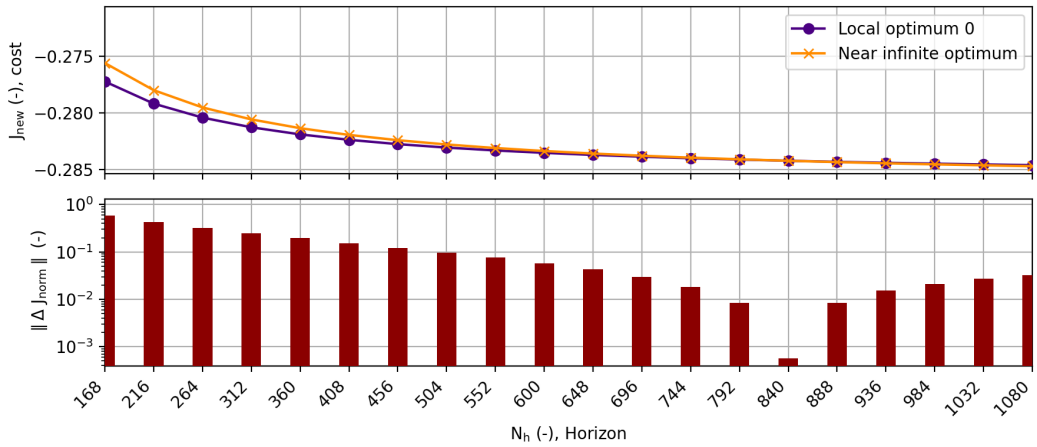


**Figure 5-1:** Cost comparison between $\boldsymbol{v}_{\text{opt}}$ and $\boldsymbol{v}_{\text{local},0}$ for increasing control horizons. The optimal cost shift towards $\boldsymbol{v}_{\infty}$ for a control horizon of $N_{\text{h}} = 840$.

Figure A-22 shows the search space for the control horizon $N_{\text{h}} = 1512$. It shows that the valley around $\boldsymbol{v}_{\text{local},0}$ does not vanish. The similarity to the infinite horizon landscape in Figure 3-3 seems logical as the horizon is increased towards infinity. However, this means that convergence towards the infinite horizon optimum cannot be guaranteed because there is always a probability the optimisation algorithm terminates in the local minimum.

## 5-2 Simulations for increased control horizons

By increasing the control horizon, the search space optimum of the first simulation time step shifts towards optima observed in the infinite horizon simulations. The long-term benefit of

reduced mean induction factor values by mitigating the finite horizon effects is realised for a control horizon larger than or equal to $N_\mathrm{h} = 840$. For shorter control horizons, the search space optimum is located at the local minimum $\boldsymbol{v}_{\mathrm{local},0}$, which also results in power gains compared to the DIC baseline simulation.

In this section, simulations are conducted with the same EMPC controller from Section 4-1-3 but for longer control horizons to analyse the performance and determine whether the probability of convergence towards steady-state solutions that improve performance can be maximised. The objective function from Equation 2-10 is used with the same configuration of the BAS algorithm with ten iterations per time step. First, the control horizon is increased with a few periods to analyse the effects of the optimal cost shift towards $\boldsymbol{v}_{\mathrm{local},0}$. Then simulations are conducted for control horizons around $N_\mathrm{h} = 840$ to analyse the performance of controllers that have their initial optimal cost located at $\boldsymbol{v}_\infty$.

### 5-2-1 Local minimum convergence guarantee for increased control horizon

Similar to the analysis in Section 4-1-3, ten simulations are conducted per initialisation. For $N_\mathrm{h} = 168$ (or $n = 7$), Figure A-2 shows the search space at the beginning of the simulation, containing a valley in which two minima lie. This makes the problem sensitive towards the initialisations as indicated by the mean steady values

$$\boldsymbol{v}_{0,\mathrm{ss}} = (A, O)_{0,\mathrm{ss}} = (0.165 \pm 0.0014, 0.171 \pm 0.0012), \tag{5-2a}$$

$$\boldsymbol{v}_{1,\mathrm{ss}} = (A, O)_{1,\mathrm{ss}} = (0.188 \pm 0.0012, 0.147 \pm 0.0007). \tag{5-2b}$$

Figure B-1 shows the mean power performance difference with the DIC baseline simulation for both initialisations. Although both result in a minor performance increase, the first initialisation outperforms the other. However, the main conclusion is that for this control design, convergence towards $\boldsymbol{v}_{\mathrm{local},0}$ cannot be guaranteed.

Therefore, two periods are added to the control horizon such that $N_\mathrm{h} = 216$ (or $n = 9$) and again ten simulations are conducted. Figure 5-2 shows that the amplitude and offset converge quickly towards a steady state, similar to the initial receding horizon simulations. A minor periodicity is still present but has reduced, indicating a mitigated influence of the omission of the final control input in the control horizon. For this control design, both mean steady-state values are very close to the local optimum at

$$\boldsymbol{v}_{0,\mathrm{ss}} = (A, O)_{0,\mathrm{ss}} = (0.186 \pm 0.0011, 0.145 \pm 0.0008), \tag{5-3a}$$

$$\boldsymbol{v}_{1,\mathrm{ss}} = (A, O)_{1,\mathrm{ss}} = (0.186 \pm 0.0012, 0.145 \pm 0.0007). \tag{5-3b}$$

Their similar mean steady-state values result in the same mean induction factor value $\bar{a}_\mathrm{g} = 0.264$, which is approximately the same as $\bar{a} = 0.261$ for $\boldsymbol{v}_{\mathrm{local},0}$. Due to the nearly constant steady-state values for $A$ and $O$, both absolute sines do not show any strange peaks or drops. With a maximum of $a_0 = 0.331$, they both slightly exceed the greedy control value, which is not favourable for the mean performance because it reduces the system's efficiency. However, the smaller mean induction factor value compared to the simulations with $N_\mathrm{h} = 120$ results in a reduced mean power production of $\mathrm{T}_0$ in Figure 5-3. The figure compares the

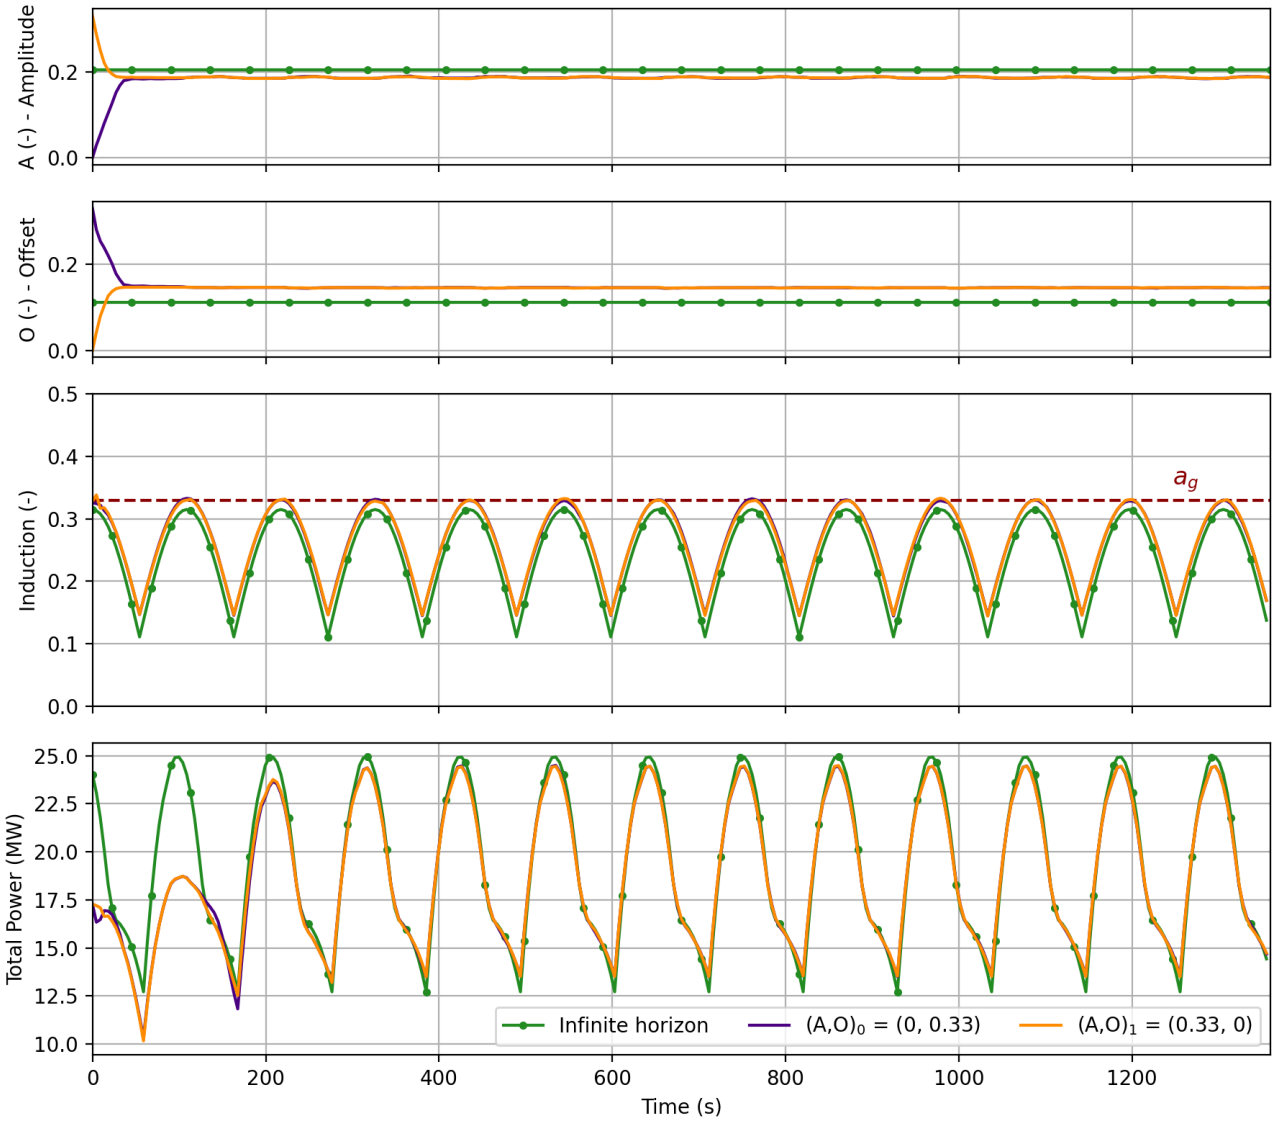


**Figure 5-2:** All mean signals concerning the ten finite-horizon control simulations. The system converges quickly for both initialisations, remaining in a periodic steady-state for the largest part of the simulation. The steady-state values are very near $\boldsymbol{v}_{\text{local},0}$.

power production differences between the DIC baseline simulation and the simulations with the absolute sine parameterisation, expressed in mean power differences and percentages. As anticipated, the 3.7% mean power decrease of $T_0$ is compensated by the mean power production of $T_1$ improving by over 27.5%, causing $T_1$ to be responsible for 23% of the total power production for both initialisations.

For this controller design, the total mean power production is $P = 19.25\,\text{MW}$, which is almost equal to the mean infinite horizon performance for $\boldsymbol{v}_{\text{local},0}$ at $P = 19.26\,\text{MW}$. It corresponds to an improvement of 2.0% over the DIC baseline simulation and a 1.4% increase compared to the initial control design in Section 4-1-3. The difference between the infinite horizon optimum is minimal at 0.4%.

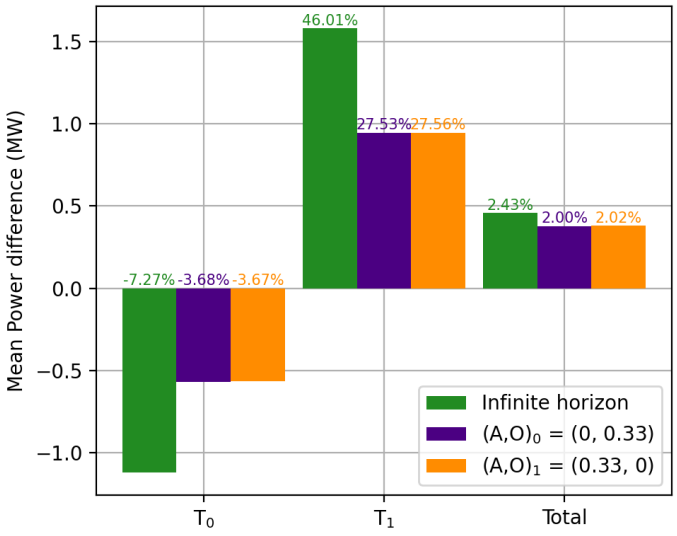


**Figure 5-3:** The mean power production differences where the DIC baseline simulation results are subtracted from the infinite and finite horizon simulations for comparison. The differences are also expressed in the relative power gain percentages.

We conclude that the performance of this control design with a control horizon of $N_\text{h} = 216$ achieves similar steady-state performance as the infinite horizon local optimum $\boldsymbol{v}_{\text{local},0}$. The improvement of 1.4% to the initial control design comes at the expense of an increased simulation time. Whereas a simulation of 300 steps with a control horizon of $N_\text{h} = 120$ took approximately 25 minutes to complete, these simulations with $N_\text{h} = 216$ completed in around 45 minutes. This is still significantly faster than the ten hours for the DIC baseline simulation.

### 5-2-2 Infinite optimum convergence guarantee for increased control horizon

The second set of simulations is conducted for a range of control horizons around $N_\text{h} = 840$ to analyse whether a maximum probability of convergence towards the infinite horizon optimum is realised for the shortest control horizon possible. The search space analysis showed that there are two valleys the optimisation algorithm can terminate, which are around the minima $\boldsymbol{v}_\infty$ and $\boldsymbol{v}_{\text{local},0}$. The previous section determined that convergence towards the latter had a

100% success rate for the control horizon $N_h = 216$. However, the goal remains to design a finite horizon controller to slightly improve the wind farm's mean power production and overcome the 0.4% gap in performance.

Because the BAS optimisation algorithm relies on a random search, the conclusions must be supported by running the same simulations multiple times. For this study, time and simulation limitations allowed us to run one hundred simulations per initialisation for the same control design. This will give some statistical support to whether the convergence was successful, which is when the amplitude and offset steady-state values are within a three percent confidence interval of the infinite optimum values $\boldsymbol{v}_\infty = (0.204, 0.111)$.

### Results for the first initialisation

Figure 5-4 shows the rate of success for the simulations initialised at $\boldsymbol{v}_0 = (0.0, 0.33)$. The simulations were conducted for a range of control horizons way longer than was intended as they were supposed to be conducted around $N_h = 840$. However, the conclusions remain the same, whereas none control designs with these control horizons can guarantee convergence towards the infinite optimum. There will always be a probability the solution is near the local optimum $\boldsymbol{v}_{\text{local},0}$.

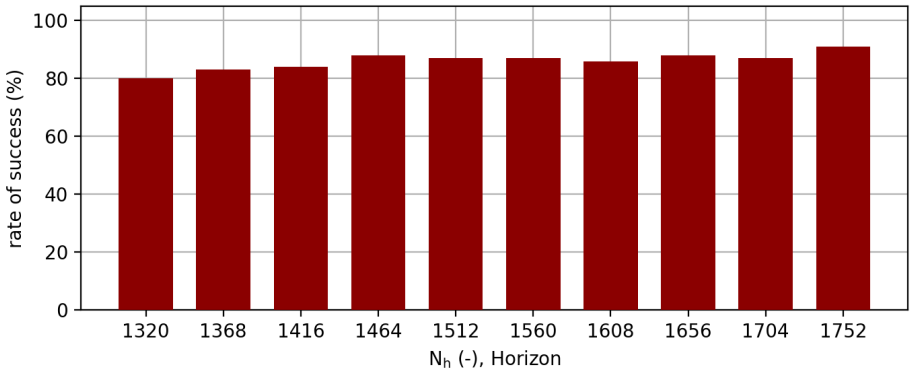


**Figure 5-4:** The rate of success for very long control horizons for the first initialisation $\boldsymbol{v}_0 = (0.0, 0.33)$. Even for these large control horizons, the rate of success does not exceed 90% as a result of the local minimum $\boldsymbol{v}_{\text{local},0}$ that blocks the path towards $\boldsymbol{v}_\infty$.

We discussed this at the beginning of the chapter due to the local minima that would not vanish for these extremely long control horizons. The slope of the success slowly increases, possibly suggesting that it will reach a probability of 100%. However, the landscapes show similarities with the infinite horizon landscape for control horizons that approach infinity. As this still contains local minima, a 100% success rate cannot be guaranteed for simulations initialised at $\boldsymbol{v}_{0,k=0} = (0.0, 0.33)$. We conclude there is a large probability of converging towards the infinite horizon optimum. However, for these periods, the probability does not exceed 90%. Furthermore, these control horizon lengths cause high increases in computational expense, leading to similar durations for the simulations to complete compared to the DIC baseline simulation.

### Results for the second initialisation

The other one hundred simulations start at the second initialisation $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$. Very different results are obtained by approaching the infinite horizon optimum from this side of the search space. Figure 5-5 illustrates the rate of success for control horizons shorter than $N_\mathrm{h} = 840$, which shows the maximum probability for convergence towards the infinite horizon optimum for $N_\mathrm{h} \geq 552$ (or $n \geq 23$). This is attributed to the fact that for these control horizons, the local minimum at $\boldsymbol{v}_\infty$ arises in Figure A-10. Increasing the control horizon to $N_\mathrm{h} = 840$ causes the shift where $\boldsymbol{v}_\infty$ becomes the global optimum. However, for the initialisation at $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$, this shift is not necessary to achieve a maximum success rate because the optimisation algorithm conveniently terminates locally at $\boldsymbol{v}_\infty$ rather than converging towards the global optimum at $\boldsymbol{v}_{\mathrm{local},0}$.

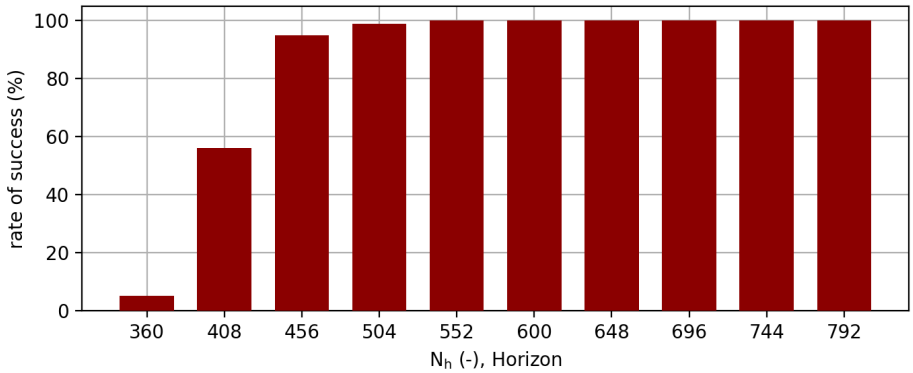


**Figure 5-5:** The rate of success for very long control horizons for the second initialisation $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$. The optimal cost shift occurs for $N_\mathrm{h} = 840$, but a 100% success rate is realised for $N_\mathrm{h} \geq 552$ because of the appearance of a local valley around $\boldsymbol{v}_\infty$.

However, the control horizon $N_\mathrm{h} = 552$ remains very large, causing high computational demands. The simulations finished in under three hours on average, leaving a minor advantage gained by the increased efficiency of the gradient-free model evaluations. Although this control configuration has a maximum success rate of convergence towards the infinite horizon optimum, the minor improved simulation time compared to the DIC baseline simulation suggests that further research on controller designs is necessary. The objective function is redefined in the next section to see whether the finite horizon effects can be eliminated.

The results for both initialisations suggest that the problem is sensitive towards initialisation because of the local minimum at $\boldsymbol{v}_{\mathrm{local},0}$. By defining the two initialisation points on the opposite side of the search space, the initial induction factor value aligns with the transient control signal at $a_\mathrm{g} = 0.33$, and two very different initial absolute sine functions are constructed. The second initialisation $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$, corresponding to large amplitude without any offset, gives the best results as a maximum success rate is achieved. This leads to the conclusion that the best results are obtained by initialising with high amplitude and low offset.

# Chapter 6

# Delay elimination for convergence guarantee

The delay in the system is attributed to the time the wake propagates from one wind turbine to the other. The control horizon must be chosen sufficiently long to capture the effects of $T_0$'s setpoint changes on $T_1$. So far, the objective function contained all power productions and control input changes in the control horizon, including terms corresponding to wake that never reaches $T_1$. At the end of Section 4-2, it was discussed that the finite horizon effects were still present as the final segment of the control horizon sequence caused a lift of the entire absolute sine sequence to maximise the energy extraction from the final wake section. Moreover, the horizon contains power production evaluations of $T_1$, which have not been extracted from wake sections excited by the control inputs in the control horizon.

## 6-1 Eliminate delay from the objective

In the previous chapter, the effects of these sections were mitigated by increasing the control horizon such that they became insignificant. In this chapter, they are approximated such that the irrelevant segments in the objective function can be eliminated, and the optimisation algorithm solely considers the values corresponding to the relevant wake. Figure 6-1 illustrates the objective function per iterate and the separate terms at time step $k_0$ over the control horizon $N_\mathrm{h}$. The irrelevant segments in the control input penalty and power production sequences are marked.

The final segment of the control horizon, which has resulted in the finite horizon effects, is neglected, eliminating the control input penalties and the corresponding power productions for $T_0$. Furthermore, the initial segment's cost is reduced because the delayed power productions of $T_1$ from previous time steps are eliminated. The middle section, where all three objective terms are active, is the most significant and allows for proper evaluation of the wind farm's behaviour within the timeframe of the control horizon.

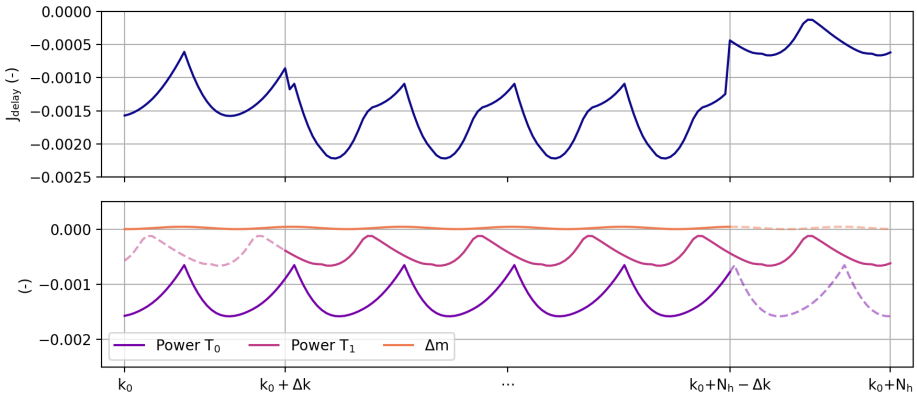


**Figure 6-1:** A cost sample at time step $k_0$ over the control horizon $N_\mathrm{h}$ and its components. The irrelevant segments are eliminated by approximating the average wake speed to account for the system delay in the objective function.

### 6-1-1 Approximate the time delay for objective redefinition

A method is proposed in [77] where a time delay evaluation between two subsequent wind turbines is conducted. The duration of the wake propagation downstream is computed by approximating the average wake's velocity. A two-point time correlation coefficient calculates the delays between multiple points located at integer rotor diameters downstream. The average wake's velocity is then found by plotting these distances over their corresponding delays. The result is an approximately linear line corresponding to a constant wind speed of $u_* = 0.7u_\infty$. The wake's time delay $\Delta t$ is therefore characterised as

$$\Delta t = \frac{5D}{u_*} = \frac{5D}{0.7u_\infty} = \frac{5}{0.7}. \tag{6-1}$$

The number of time steps in the delay is determined by adding the discretisation time step $h = 0.2$ as

$$\Delta k = \frac{\Delta t}{h} = 36. \tag{6-2}$$

With the derived time delay, the normalised objective function in Equation 3-6 is redefined by including $\Delta k$ into the in- and output weights as

$$\begin{aligned} &\min_{\boldsymbol{v}_k} J_\mathrm{delay}\left(\boldsymbol{q}_k, \boldsymbol{v}_k, \Delta k\right) = \min_{\boldsymbol{v}_k} J_\mathrm{new}\left(\boldsymbol{q}_k, \boldsymbol{v}_k, \Delta k\right) = \\ &\min_{\boldsymbol{v}_k} \frac{1}{N_\mathrm{h}} \sum_{k=k_0}^{k_0+N_\mathrm{h}} \mathbf{Q}\left(\Delta k\right) \boldsymbol{y}_k\left(\boldsymbol{q}_k, \boldsymbol{v}_k\right) + \boldsymbol{\Delta m}_k^\mathrm{T}\left(\boldsymbol{v}_k\right) \mathbf{R}\left(\Delta k\right) \boldsymbol{\Delta m}_k\left(\boldsymbol{v}_k\right), \end{aligned} \tag{6-3}$$

such that the controller can optimise towards the relevant control horizon components.

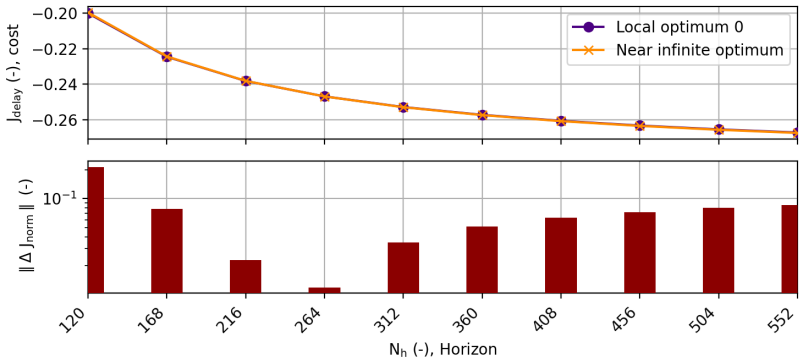


**Figure 6-2:** Cost comparison between $\boldsymbol{v}_{\text{opt}}$ and $\boldsymbol{v}_{\text{local},0}$ for increasing control horizons. The optimal cost shift towards $\boldsymbol{v}_{\infty}$ for a control horizon of $N_{\text{h}} = 264$.

## 6-2 Grid search analysis for eliminated delay objective

A similar strategy for maximising the probability of convergence towards the infinite optimum is conducted for this controller design. Finding the control horizon for which the optimal cost of the initial time step search space is located at the infinite horizon point will give optimal direction to the optimisation algorithm. We, therefore, conduct a grid search analysis for search spaces at $k = 0$ for a range of control horizons. Similar to the analysis in Chapter 5, the horizons increase by 48 control inputs per search space corresponding to two periods added per step. Figures A-23 to A-36 show the optimal cost shift from $\boldsymbol{v} = (0.186, 0.144)$ for $N_{\text{h}} = 120$ to $\boldsymbol{v}_{\text{opt}} = (0.205, 0.111)$ for $N_{\text{h}} = 744$. For the normalised objective function in the previous chapter, the local minimum shifted towards the local minimum $\boldsymbol{v}_{\text{local},0}$ for $N_{\text{h}} = 168$ but due to the elimination of the delay, this occurs for $N_{\text{h}} = 120$.

The optimal cost shifts from the valley around $\boldsymbol{v}_{\text{local},0}$ to a point near $\boldsymbol{v}_{\infty}$ for the control horizon $N_{\text{h}} = 264$ (or $n = 11$). It is located at $\boldsymbol{v} = (0.210, 0.109)$, within the three percent range defined earlier. This is a significant improvement compared to the previous case where the optimal cost shifted for the control horizon $N_{\text{h}} = 840$ (or $n = 35$). This shows the significant reduction of the finite horizon effects and the shift towards a global optimum corresponding to an increased emphasis on $\text{T}_1$ for shorter control horizons. Figures A-26 to A-36 show that the optimal cost remains near the infinite optimum value for $N_{\text{h}} \longrightarrow \infty$.

The analysis of the search spaces is supported in Figure 6-2 by comparing the measured cost differences at $\boldsymbol{v}_{\infty}$ and $\boldsymbol{v}_{\text{local},0}$ and plotting them over the control horizon lengths. Subsequently, the absolute value of the normalised cost difference $J_{\text{norm}}$ from Equation 5-1 is shown. We notice the slope change, indicating the optimum shift for $N_{\text{h}} = 264$.

Similar to Chapter 5, further increasing the control horizon does not result in the disappearance of the local valley around $\boldsymbol{v}_{\text{local},0}$. Instead, it results in an additional local valley around $\boldsymbol{v}_{\text{local},1} = (0.221, 0.081)$ which arises for control horizons $N_{\text{h}} \geq 504$ (or $n \geq 21$). An interesting note is that we observed the local minimum while conducting the grid search to find the infinite horizon optimum in Figure 3-3, which did not arise in the search spaces in Section 5-1. However, the existence of this local minimum decreases the probability of converging towards $\boldsymbol{v}_{\infty}$.

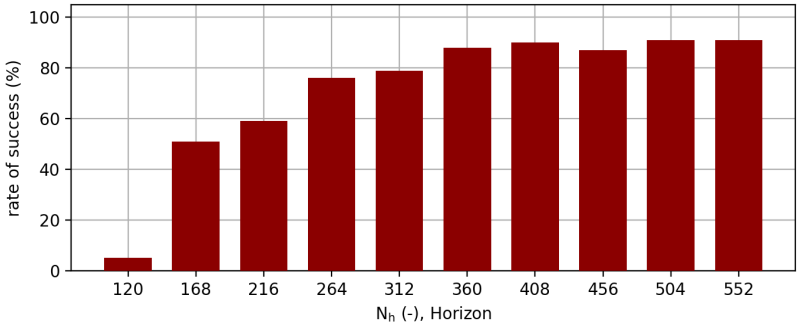


**Figure 6-3:** The rate of success for very long control horizons for the first initialisation $\boldsymbol{v}_{0,k=0} = (0.0, 0.33)$. Due to the local minimum $\boldsymbol{v}_{\text{local},0}$ that does not disappear, the probability of terminating here remains, and the maximum success rate saturates at 91%.

## 6-3 Simulations for increasing control horizons

To support the considerations on the convergence for the range of control horizons, each is analysed by conducting one hundred simulations. We consider the convergence a success if the offset and amplitude steady-state values lie within the three percent range of the infinite horizon optimum value. Further control characteristics include the optimisation algorithm BAS, which remains in default configuration from Section 4-1-2.

### 6-3-1 Results for the first initialisation

Figure 6-3 shows the rate of success of all simulations initialising in $\boldsymbol{v}_{0,k=0} = (0.0, 0.33)$. Notice that the rate of success saturates at 91% for control horizons $N_\text{h} \geq 408$ apart from the minor decrease at $N_\text{h} \geq 456$. This supports the conclusion from the search space analysis that the local minimum does not disappear by increasing the control horizon. It is again found to block the convergence towards the infinite optimum.

We, therefore, conclude that the probability of converging towards the infinite horizon optimum from $\boldsymbol{v}_{0,k=0} = (0.0, 0.33)$ cannot be maximised for this controller design. Although the resulting performance does improve compared to the DIC baseline simulation when the solution converges towards $\boldsymbol{v}_{\text{local},0}$, we will not further consider this initialisation.

### 6-3-2 Results for the second initialisation

Figure 6-4 shows the limited success for the initial two control horizons and the rise to a 100% success rate for $N_\text{h} = 264$ (or $n = 11$). This was anticipated because the optimal shift towards the infinite horizon optimum occurred for this control horizon. Since there is no local minimum to block the convergence towards $\boldsymbol{v}_\infty$, the perfect rate of success is logical. Therefore, we conclude that a maximum convergence probability is realised for this initialisation for control horizons $264 \leq N_\text{h} \leq 408$. For $N_\text{h} \geq 456$, we observe a decrease in the rate of success. This is attributed to the appearance of an additional local valley around $\boldsymbol{v}_{\text{local},1} = (0.221, 0.081)$ noticed in the search space analysis. We determined the valley to

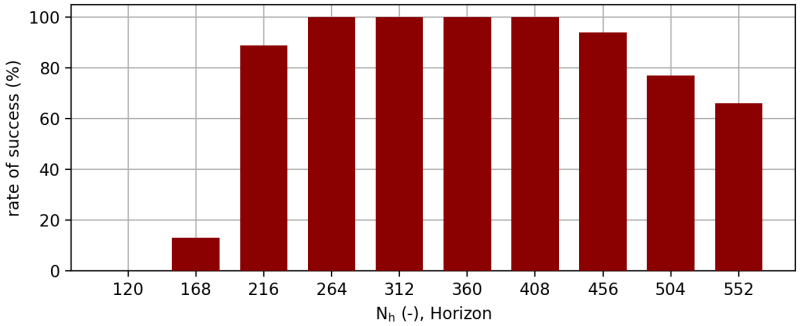


**Figure 6-4:** The rate of success for very long control horizons for the first initialisation $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$. The optimal cost shifts for $N_\mathrm{h} = 264$, which is why a 100% success rate is realised.

arise for $N_\mathrm{h} \geq 504$, but it affects the convergence towards the infinite horizon optimum for shorter control horizons. However, the valley should not block the convergence towards the infinite horizon optimum for the final controller design because it does not arise for shorter control horizons.

Compared to Figure 5-5, the control horizon realising a maximum success rate has significantly reduced from $N_\mathrm{h} = 552$, which results in an extensive simulation time reduction. The simulations ran for 300 time steps, corresponding to 23 minutes, by adding the dimensions of the reference turbine. For the controller design with $N_\mathrm{h} = 264$, completing this simulation in a little over an hour means we have come closer to real-time implementation. The completion time is ten times faster than the DIC baseline simulation and three times faster than the previous normalised objective without considering the delay.

For the normalised objective function, the maximum success rate at $N_\mathrm{h} = 552$ is shorter than the horizon that contained the optimal cost shift at $N_\mathrm{h} = 840$. This was attributed to the appearance of the local valley around the infinite horizon optimum value. Since the optimal cost is located within the valley of the infinite horizon, the convergence guarantee for the current simulations is higher. We, therefore, conclude that the optimal controller design that maximises the probability of convergence towards the infinite horizon optimum evaluates the objective function that eliminates the irrelevant segments over a control horizon of $N_\mathrm{h} = 264$.

## 6-4 Final results with optimal control design

Finally, the performance of this controller is analysed. Figure 6-5 shows the mean performance of the controller from all one hundred simulations. We consider the minor periodicity in the amplitude and offset signal to be the result of the omission of the final control input. The fast convergence of the control variables supports the conclusions on the search space that it is a large, steep convex bowl apart from the region around the infinite horizon optimum. Furthermore, the fast convergence is attributed to the number of iterations per time step that the BAS algorithm searches for better objective function values. If this number is reduced, the simulation efficiency increases and the next step towards real-time implementation is achieved. However, the resulting induction factor and power production signals match the infinite horizon well for the most significant part of the simulation, resulting in the same mean

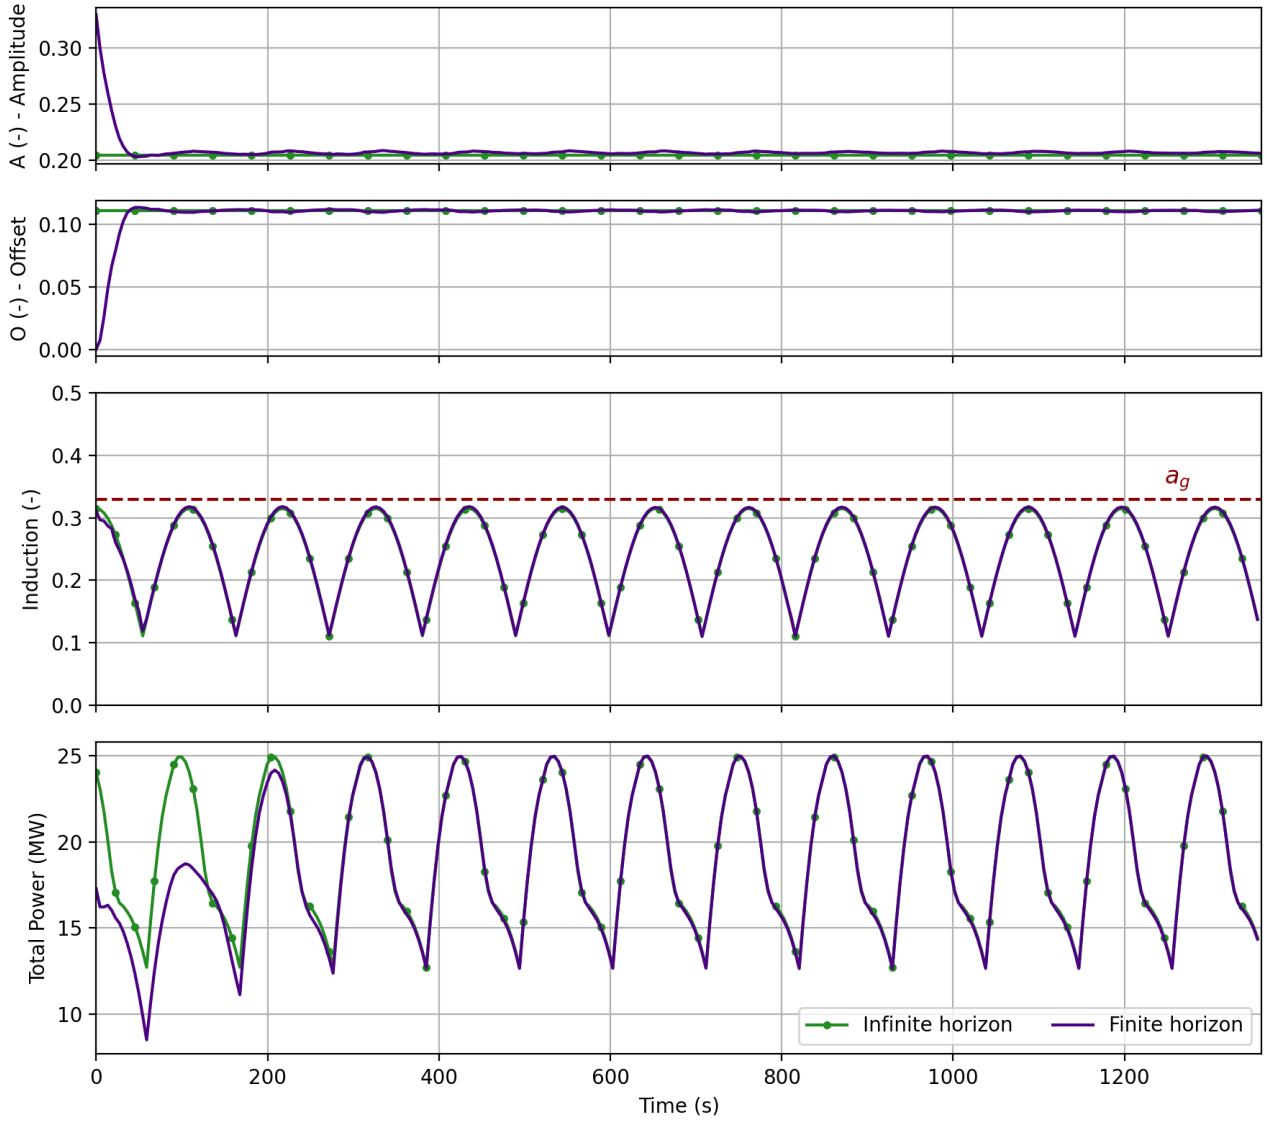


**Figure 6-5:** Mean signals concerning the optimal finite-horizon control simulations. The system quickly converges towards the infinite horizon optimum values, resulting in similar periodic steady-state performance.

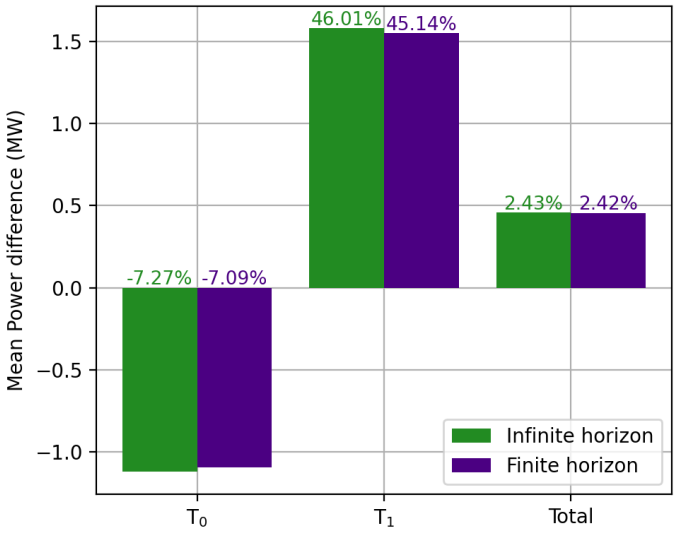


**Figure 6-6:** The mean power production differences where the DIC baseline simulation results are subtracted from the infinite and optimal finite-horizon simulations for comparison. The relative power gains are also shown. The negligible difference in total power indicates that the goal is achieved by designing a finite horizon controller that realises the infinite horizon steady-state performance.

induction factor value of $\bar{a} = 0.241$. We notice the resemblance in the periodic steady-state induction factor and power production signals between the infinite and finite horizon signals, which indicate similar performance.

Figure 6-6 shows the mean power differences and relative power gains compared to the DIC baseline simulation. The figure confirms the resemblance in performance because the mean power production realised by the finite horizon controller is $P = 19.33\,\text{MW}$. We consider it to be very similar to the infinite horizon simulation, resulting in a mean total power production of $P = 19.34\,\text{MW}$. Both controller designs improve the steady-state performance of the two-turbine wind farm in the FVW model by over 2.4% with respect to the DIC baseline simulation in [1].

Finally, Figure 6-7 shows the flow through the wind farm over one period of excitation. The similarity in the control signal is reflected in the flow of the wake. We notice the leapfrogging, where a higher-speed wake section is located at $x = 720\,\text{m}$ downstream in the first snapshot, which rolls over the lower-speed section in the subsequent snapshots.

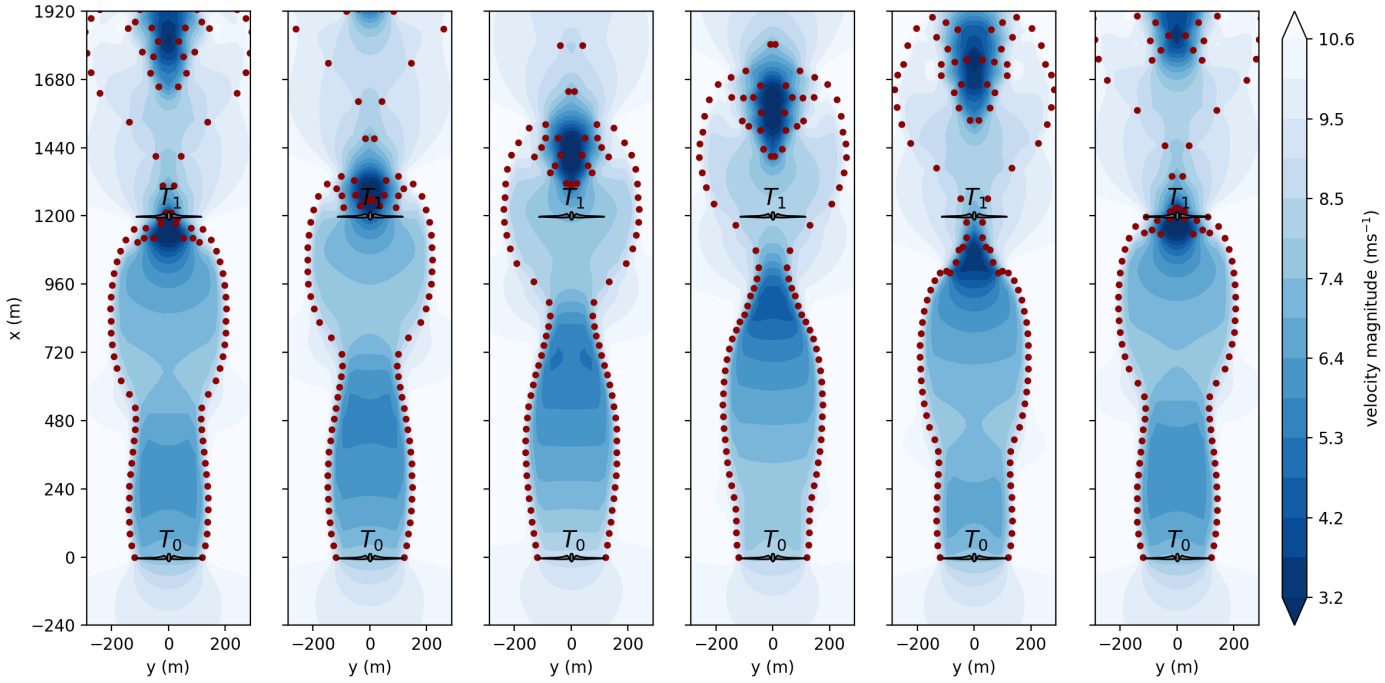


**Figure 6-7:** Snapshots of the wake flow through the wind farm for the optimal finite-horizon controller design. The snapshots show the roll-up of the high-speed vortex ring over the low-speed ring, allowing surrounding air to be drawn in the wake.

# Chapter 7

# Conclusion, discussion, and recommendations

## 7-1 Conclusion

The objective for this master thesis was to explore the possibilities of sinusoidal excitation of wind turbines in a two-dimensional, single-wake model presented from [1]. The control-oriented, free-vortex wake model simulates the release of swirling vortices from the edges of the actuator disc. By adding the dimensions of the 15 MW reference wind turbine, the results from the simulations were more straightforward to interpret. In a two-turbine wind farm, the upstream wind turbine is excited through an absolute sine function, which reflected the optimal control signal from [1] best compared to a regular sine function. In the baseline simulation, optimisation of the control sequences resulted in a turnpike tail due to finite horizon effects in the internal model evaluation. The parameterisation allowed for better constraints on the control horizon sequences through the predefined shape. It also resulted in a significant dimension reduction from one hundred control inputs in the control horizon to three design parameters: the amplitude, offset and control horizon.

The optimal infinite horizon configuration $\boldsymbol{v}_\infty = (A, O, N_\mathrm{h})_\infty = (0.204, 0.111, 120)$, where the optimal control horizon length was determined for $n = 5$ periods. $\boldsymbol{v}_\infty$ was determined by conducting a grid search and resulted in a total mean power production of $P = 19.34\,\mathrm{MW}$, which is a 2.43% performance increase compared to the periodic steady-state in the dynamic induction control baseline simulation. The emphasis on power production shifted more towards the downstream wind turbine by reducing the mean induction factor value from $\bar{a}_\mathrm{base} = 0.3$ to $\bar{a} = 0.241$. This suggests that the optimal control signal combines the control strategies of dynamic induction control and power derating. The absolute sine parameterisation also shed light on the shape of the search space, which revealed two local minima in the neighbourhood of the infinite optimum and a large convex bowl outside of that region. The value for the control horizon corresponded to the optimal frequency for vortex shedding and was thereby set for the remainder of the report, reducing the optimisation dimension to two. The infinite

horizon simulation proved the potential of the parameterisation due to the mean power increase and the visible leapfrogging in the figures displaying the wake flow through the wind farm.

We then implemented the absolute sine into the model-based, finite horizon controller to conduct economic model predictive control. Two initialisations were defined on opposite sides of $\boldsymbol{v}_\infty$ in the search space at $\boldsymbol{v}_{0,k=0} = (0.0, 033)$ and $\boldsymbol{v}_{1,k=0} = (0.33, 0.0)$. They were defined to align the initial induction factor with the control signal of the transient removal simulation where both wind turbines operated in greedy conditions. The parameterisation of the control signal allowed for the gradient-free beetle antennae search algorithm as the optimisation algorithm, which did not require expensive gradient computations but was based on a random search method to optimise the amplitude and offset every time step.

For a control horizon of $N_\mathrm{h} = 120$, the resulting performance was similar to the DIC baseline simulation while reducing the simulation time by a factor of 30. Although the parameterisation allowed for a constrained shape of the control sequence, the final section of the power productions in the objective function lifted the total absolute sine. This resulted in a mean induction factor increase to $\bar{a} = 0.283$. From there, we concluded that the finite horizon effects remained present for the current definition of the objective function. Increasing the control horizon reduced the significance of this objective section because starting in both initialisations and for a control horizon of $N_\mathrm{h} = 216$, all two hundred simulations quickly converged towards the local minimum $\boldsymbol{v}_{\mathrm{local},0} = (0.183, 0.144)$. These simulations resulted in a total mean, steady-state power production of $P = 19.25\,\mathrm{MW}$, which is a 2.04% performance increase compared to the DIC baseline simulation, while completion occurred 15 times faster.

Especially the local minimum at $\boldsymbol{v}_{\mathrm{local},0}$ prevented the finite horizon controllers from converging towards the infinite horizon optimum, making the problem sensitive towards initialisation. Increasing the control horizon could not guarantee the convergence because this local minimum would not vanish. From the one hundred simulations for the large control horizons, a probability of over 10% remained that the steady-state value would converge towards this local minimum when initialising in $\boldsymbol{v}_{0,k=0}$. Contrarily, initialising in $\boldsymbol{v}_{1,k=0}$ resulted in a 100% success rate of convergence towards the infinite optimum for control horizons $N_\mathrm{h} \geq 552$. These large control horizons increased the computational demand, resulting in an average completion of almost 3 hours.

A redesign of the objective function also resulted in a maximum success rate of convergence but for shorter control horizons. The final method accounted for the delay in the system and mitigated the finite horizon effects by eliminating the irrelevant components in the objective function. This was done by approximating the average speed of the wake to derive the sections of the objective function that could be eliminated. Through this redesign, convergence towards the infinite horizon optimum could saturated at 91% for increasing control horizons when initialising in $\boldsymbol{v}_{0,k=0}$ due to the same valley around the local minimum $\boldsymbol{v}_{\mathrm{local},0}$ the convergence also terminated in. Contrarily, initialising in $\boldsymbol{v}_{1,k=0}$ resulted in a 100% success rate for a control horizon of $N_\mathrm{h} = 264$. All one hundred simulations converged accordingly while completing the simulations in little over an hour on average. Through the dimensionalisation with the 15 MW reference turbine, 23 minutes of wind farm simulation was conducted, improving towards real-time implementation. We, therefore, conclude that this final economic model predictive control design resulted in the best performance.

## 7-2 Discussion

The parameterisation has shown promising results in infinite and finite horizon simulations. However, the methodology and analyses have resulted in some research limitations. The regular and absolute sine were destined to be the parameterisation functions because literature has shown the potential of the sine and the shape of the optimal control signal in [1] resembled the absolute sine shape. However, the degrees of freedom reduction might have been too significant. During the study, other options rather than the regular and absolute sine were not considered. For example, adding higher-order terms or combinations of sinusoids with varying frequencies might have better captured the optimal control signal. Moreover, other base functions could have captured the control sequences while significantly reducing the dimensionality.

The proposed beetle antennae search algorithm was chosen because it required minimal function evaluations per iteration while remaining effective. Throughout the research, we analysed many search spaces, which had smoothened due to the parameterisation. The simple shape allowed the algorithm to easily manoeuvre towards the region of minima, which is the main reason it was effective. However, this also resulted in little attention to the optimal algorithm tuning, mainly to the termination condition, because the maximum number of iterations could have been smaller. The shape allowed for a guaranteed convergence towards the region of minima, and by increasing the maximum number of iterations per time step, convergence occurred faster. However, once the steady-state values were found, these iterations slowed down the simulation because they no longer had any effect. The guaranteed convergence could have allowed other termination conditions to be added to the algorithm. For example, a minimal step size condition that would terminate when a minimum was reached.

The inherent randomness remained a complex topic throughout the report. The objective of this thesis was to maximise the wind farm's power. Through the infinite horizon simulation, we knew where the algorithm had to converge and tried shaping the controller to maximise the probability of convergence. Many simulations supported strong considerations that suggest initialising on the side of zero offset and high amplitude, but hard conclusions cannot be drawn. Unfortunately, the sensitivity towards initialisation remains.

Towards the end of the thesis, we noticed a minor error in the coding of the control sequence, which was attributed to the omission of the final control input in the control horizon sequence. This resulted in a non-periodic signal that also caused the cost function to incorporate minor periodicities. The minor periodicities in the optimal control signals for shorter control horizons are attributed to this error. Luckily, the error had little influence due to its insignificance when the control horizon was increased.

A final note is that too much focus has been on ensuring convergence towards the infinite horizon optimum by conducting numerous simulations for very high control horizons without success. This is why Figure 5-4 shows these high control horizons.

## 7-3 Recommendations

Although promising results have been achieved, this section contains recommendations for directions in future work.

1. The number of initialisation points could be increased to explore the sensitivity towards initialisation further. From the search space analyses and the extreme initialisation points we utilised in the study, we consider a wide range of initialisations to converge to the infinite horizon optimum for the final control design where the time delay is included. This can be further analysed by initialising from a larger number of points. This range cannot include initialisations in the direction of $(A, O) = (0.0, 0.0)$ because the physical bounds would be exceeded through the large corresponding induction factor changes.

2. The optimal tuning of the optimisation algorithm can also be addressed. As mentioned in the discussion, limited attention was paid to the optimal tuning of this algorithm. A more thorough parameter sweep over the hyperparameters and an extra termination condition might improve the algorithm's performance, leading to faster simulations. Consequently, the convergence for a minimum number of iterations can be analysed. The search space shape guarantees convergence towards the region of minima. Computing one or two steps per iteration would significantly reduce the simulation time.

3. The non-dimensionality of the FVW model resulted in a particular case in which the Strouhal number was equal to the sine frequency. The optimal Strouhal number for vortex shedding was found by finding the optimal frequency for the particular wind conditions and wind turbine spacing. Preliminary conclusions were drawn on the parameter relations because changing the conditions or wind farm layout will influence the frequency. It remains interesting to analyse whether this Strouhal number will remain constant.

4. This study focused on one particular wind condition without any variations. It resulted in quick convergence towards the periodic steady state the system remained for the largest part of the simulation. Future work could address this by including wind variations in the simulation to determine the controller's response. Sudden wind speed changes might necessitate the online determination of the optimal frequency, which increases the optimisation dimension to three. It would also require the implementation of a mixed-integer optimisation algorithm because the control sequence remains constrained to an integer number of control inputs in the horizon. Instead of optimising over the control horizon every time step, the number of samples per period could be defined as the optimisation variable.

5. Currently, no background turbulence is present. Therefore, a turbulence term could be added to the wind, for example, random noise, to analyse the performance in disturbed flow.

# Appendix A

# Grid searches for the grid search analyses

The various search spaces that are utilised for the grid search analyses in Chapter 5 and 6 are displayed. The search spaces are constructed for varying control horizons over varying ranges to visualise the different local minima around the infinite horizon optimum $v_\infty$. It was not possible to run all simulations for constant ranges because time was limited. Since the optimal cost shifted as the control horizon increased, the particular ranges of interest were determined and analysed for each grid search simulation. The grid is refined as much as possible by defining steps of 0.001 for the amplitude and offset, which also caused the grid searches to become time-consuming.

The grid search analyses are conducted by analysing the different minima, which arose by increasing the number of levels in the figure. Therefore, this number varied to properly distinguish between the minima, ranging from 40 to 60.

## A-1 Grid searches to determine optimal shift towards infinite optimum for normalised objective function

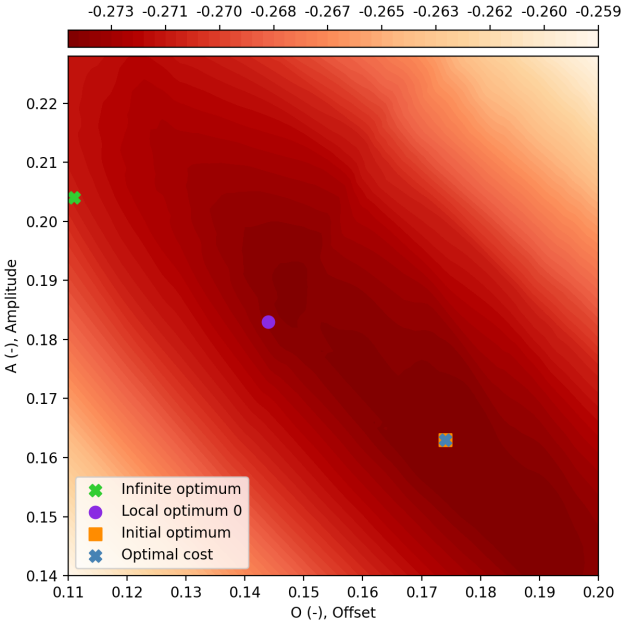


**Figure A-1:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 120$ (or $n = 5$). The optimal cost is located at $v = (0.163, 0.174)$. It results in similar shape behaviour compared to the DIC baseline simulation, not accounting for the long-term benefit of a decreased value for $\bar{a}$.

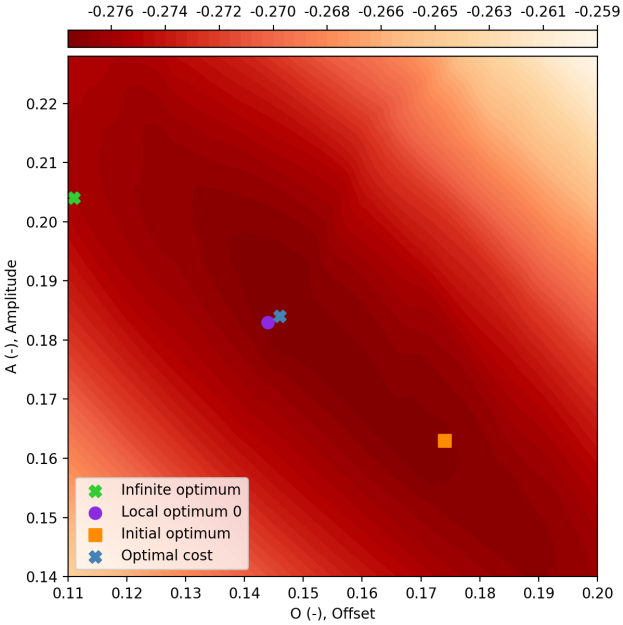


**Figure A-2:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 168$ (or $n = 7$). The optimal cost shifts and is located very near the local minimum $v_{\text{local},0}$.

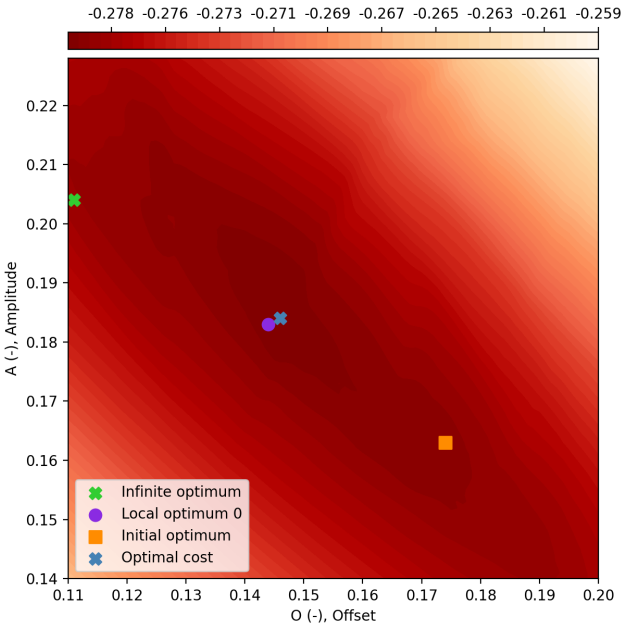


**Figure A-3:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 216$ (or $n = 9$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

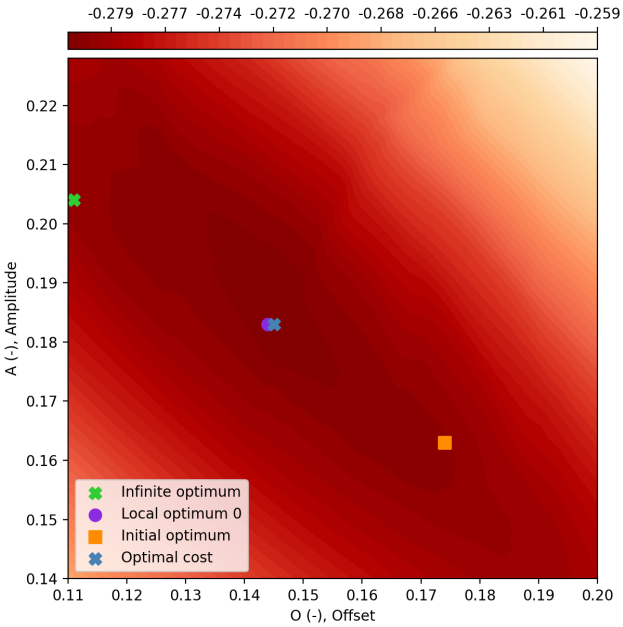


**Figure A-4:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 264$ (or $n = 11$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

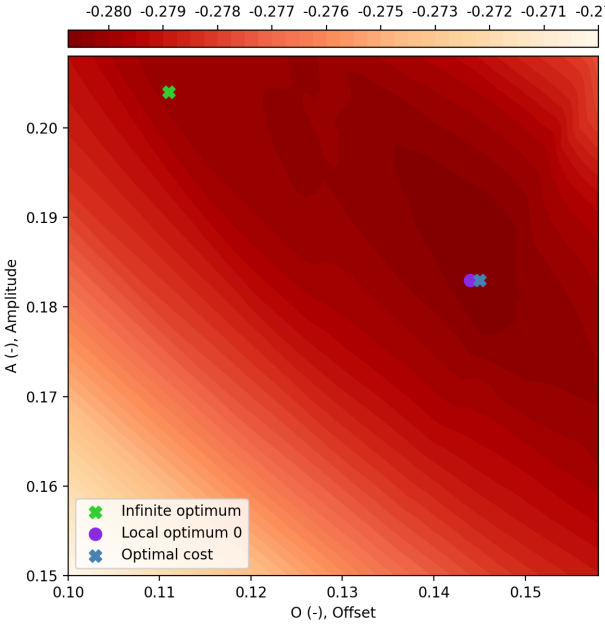


**Figure A-5:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 312$ (or $n = 13$). The optimal cost is located very near the local minimum $v_{\text{local},0}$ and $v_{\infty}$ is displayed as well.

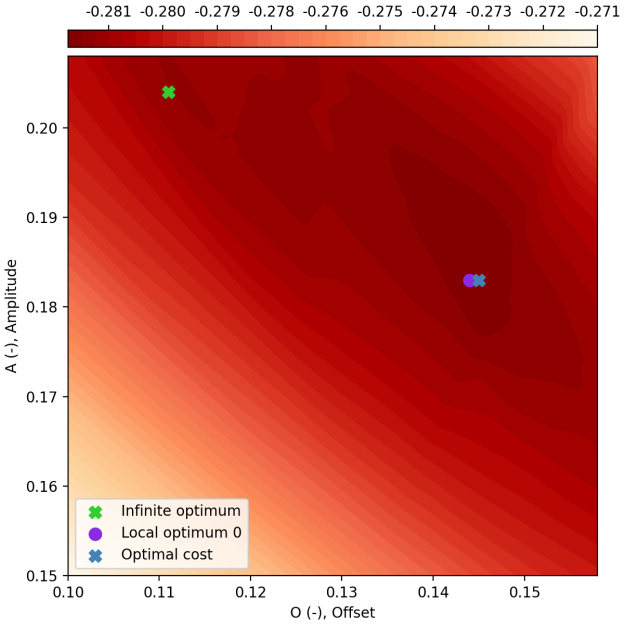


**Figure A-6:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 360$ (or $n = 15$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

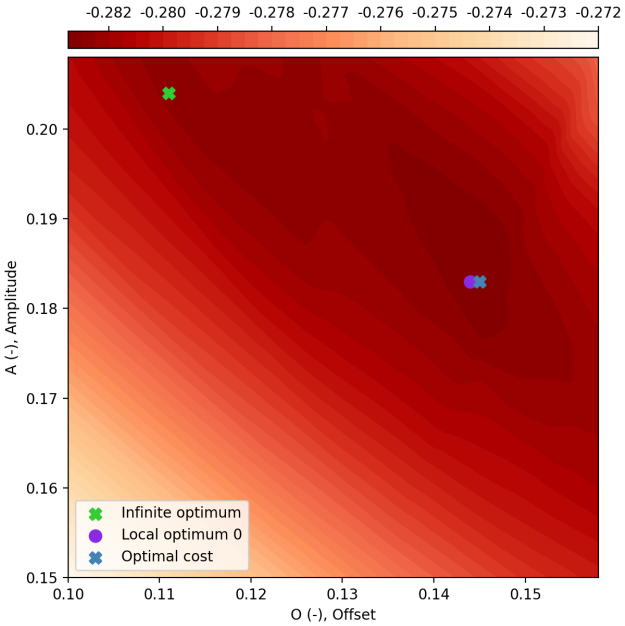


**Figure A-7:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 408$ (or $n = 17$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

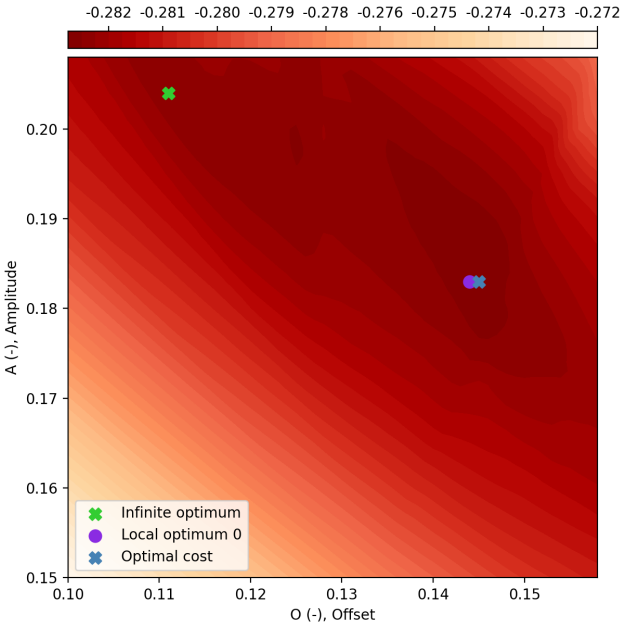


**Figure A-8:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 456$ (or $n = 19$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

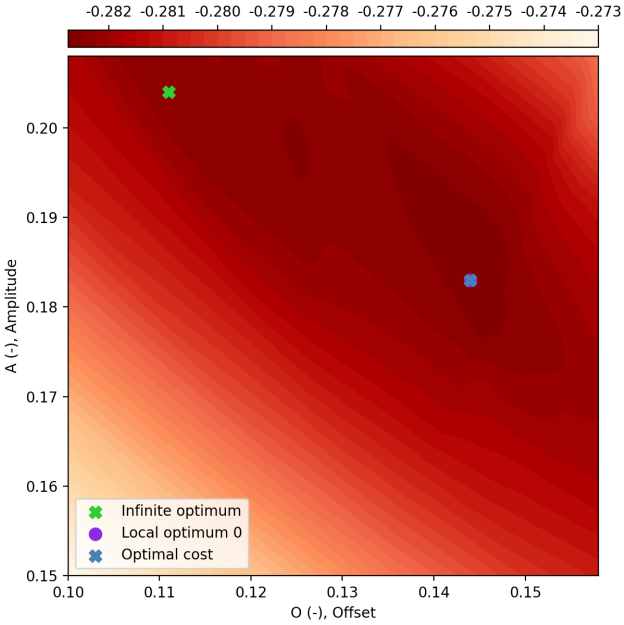


**Figure A-9:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 504$ (or $n = 21$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

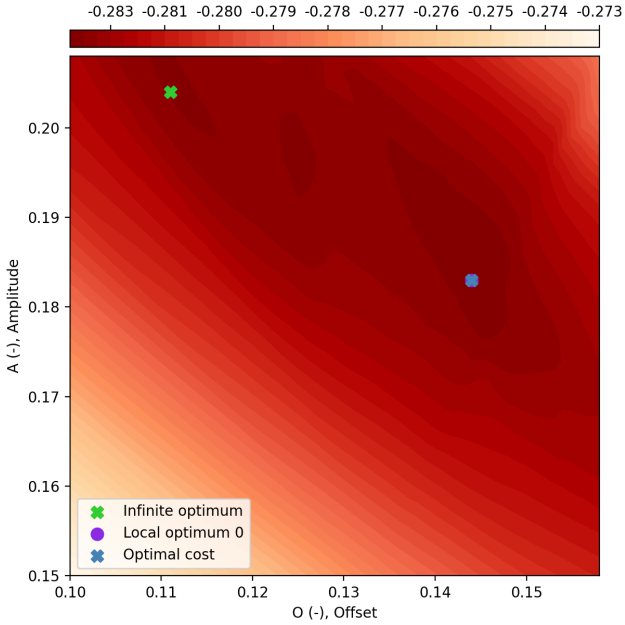


**Figure A-10:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 552$ (or $n = 23$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

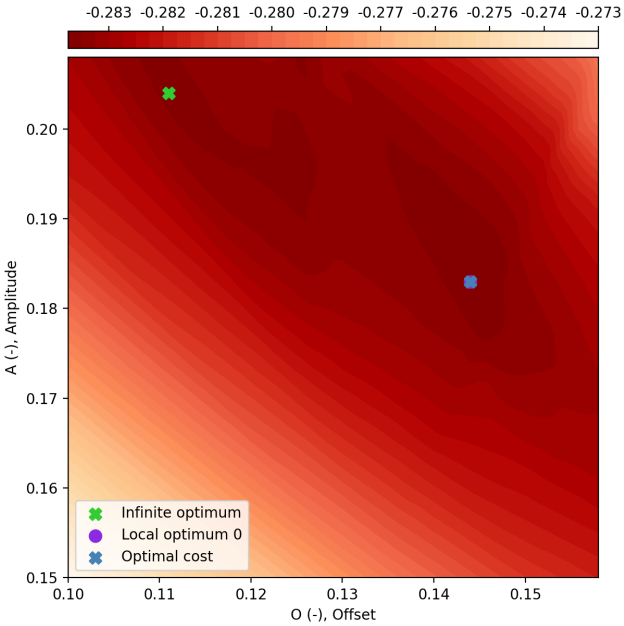


**Figure A-11:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 600$ (or $n = 25$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

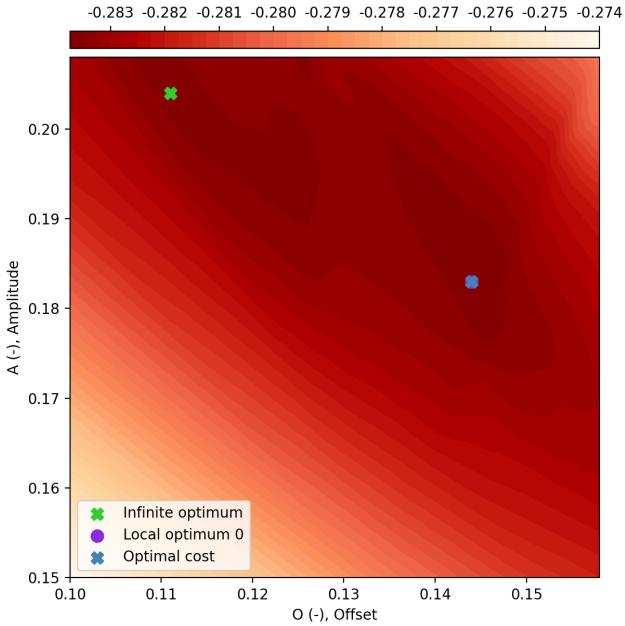


**Figure A-12:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 648$ (or $n = 27$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

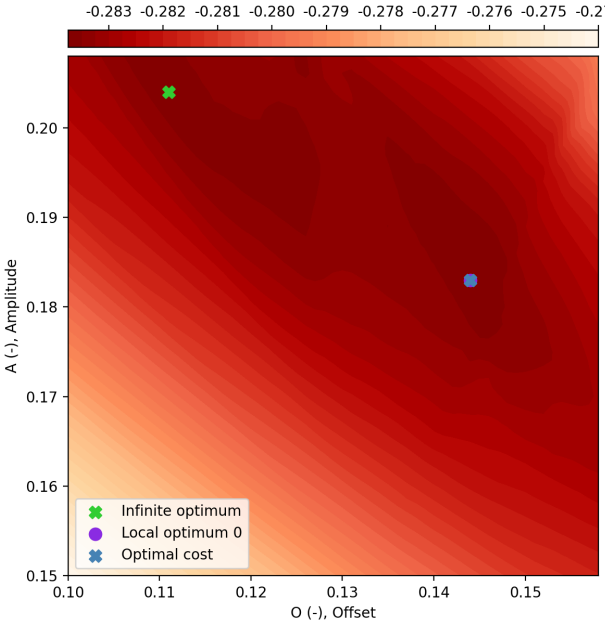


**Figure A-13:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 696$ (or $n = 29$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

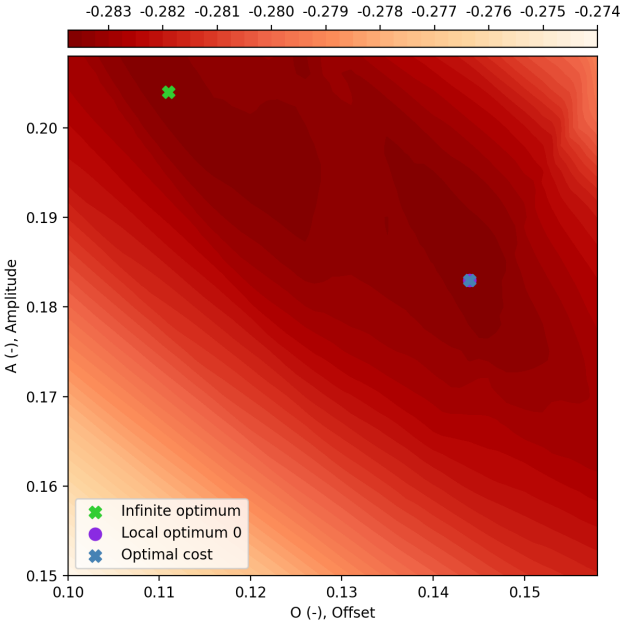


**Figure A-14:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 744$ (or $n = 31$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

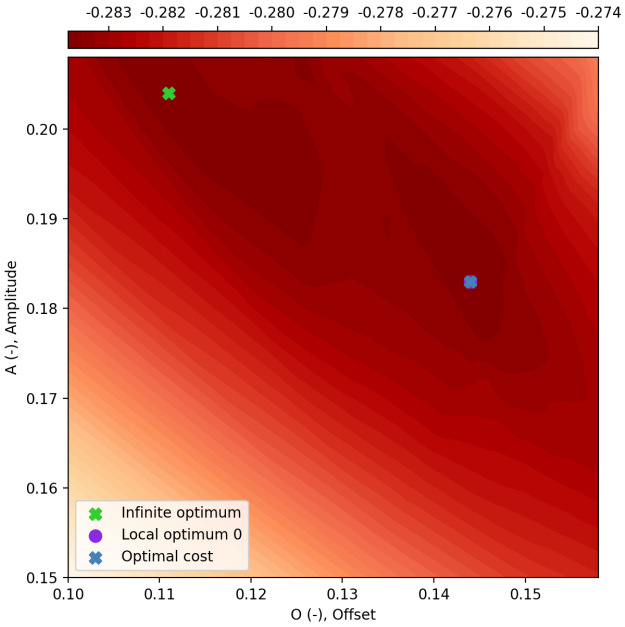


**Figure A-15:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 792$ (or $n = 33$). The optimal cost is located at the local minimum $v_{\text{local},0}$.

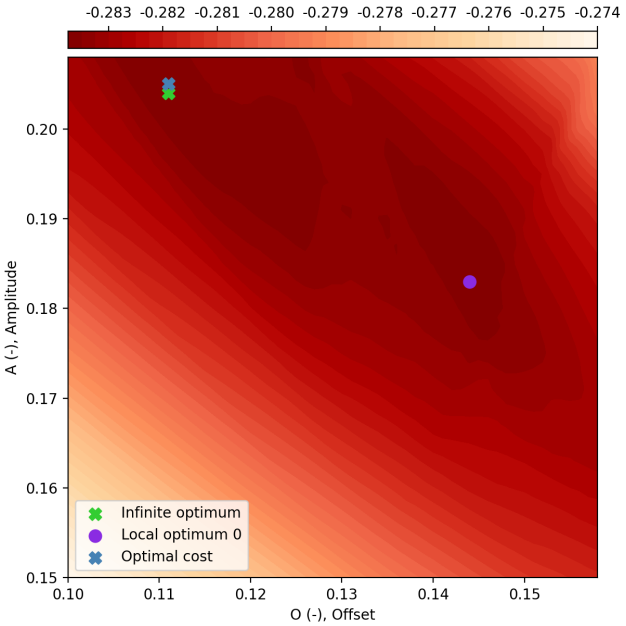


**Figure A-16:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 840$ (or $n = 35$). The optimal cost shifts and is located very near $v_{\infty}$.

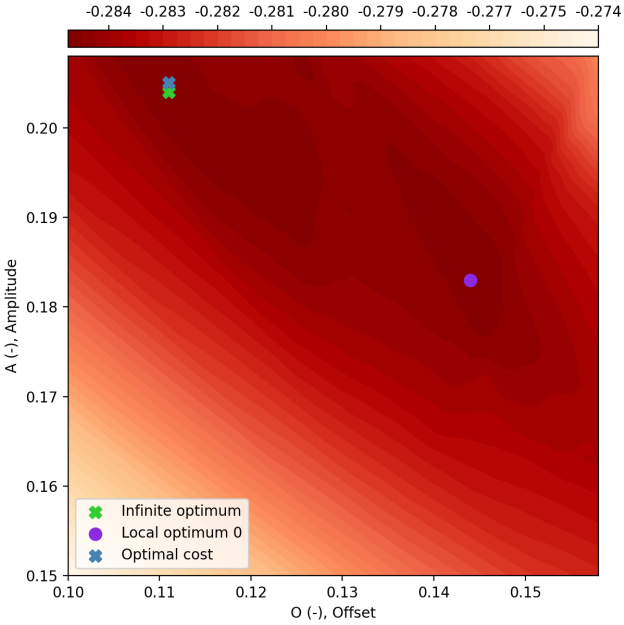


**Figure A-17:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 888$ (or $n = 37$). The optimal cost shifts and is located very near $v_{\infty}$.

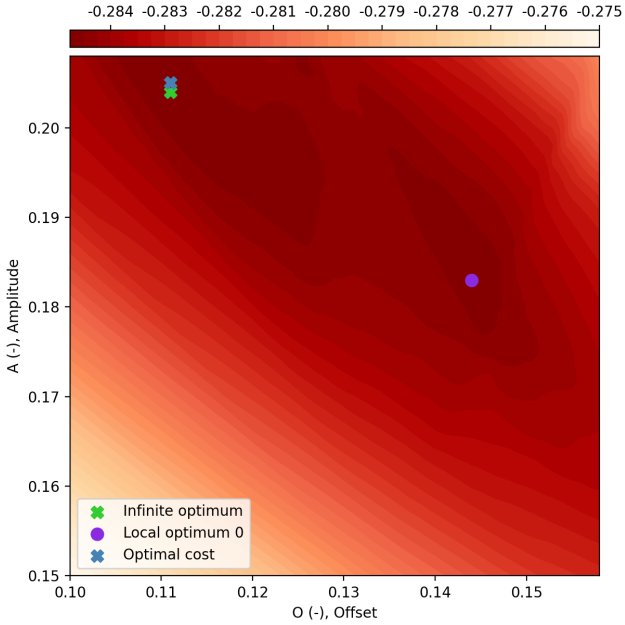


**Figure A-18:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 936$ (or $n = 39$). The optimal cost shifts and is located very near $v_{\infty}$.

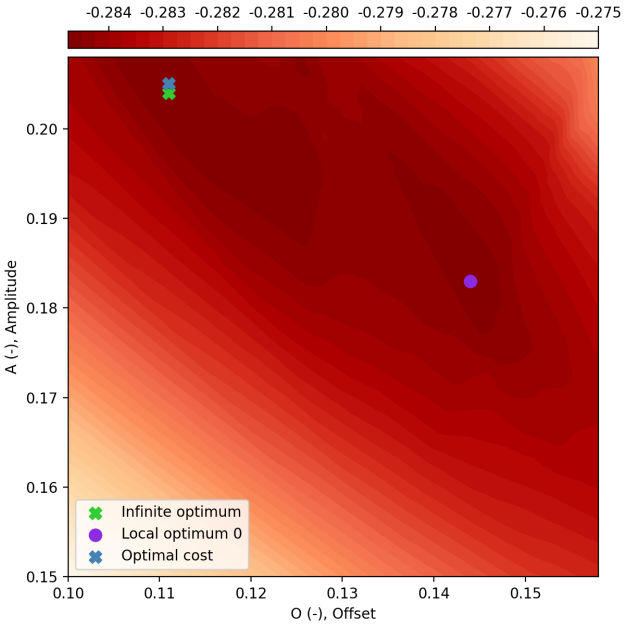


**Figure A-19:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 984$ (or $n = 41$). The optimal cost shifts and is located very near $v_{\infty}$.

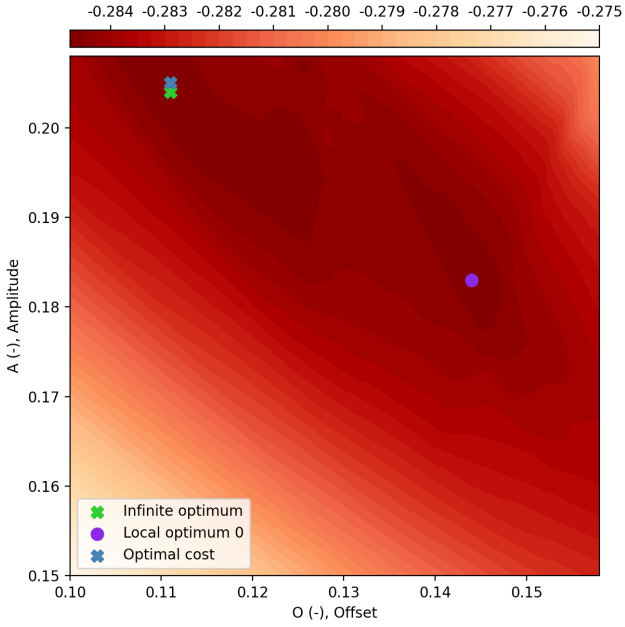


**Figure A-20:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 1032$ (or $n = 43$). The optimal cost shifts and is located very near $v_{\infty}$.

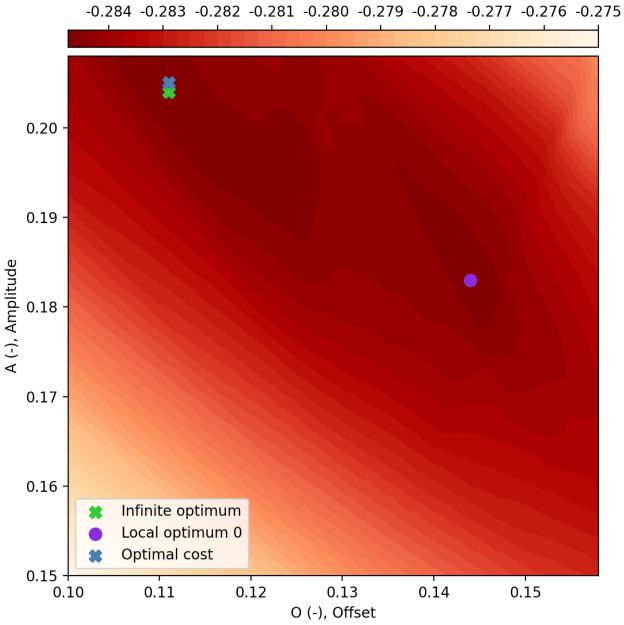


**Figure A-21:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 1080$ (or $n = 45$). The optimal cost shifts and is located very near $v_{\infty}$. Note the very limited change with respect to Figure A-16.

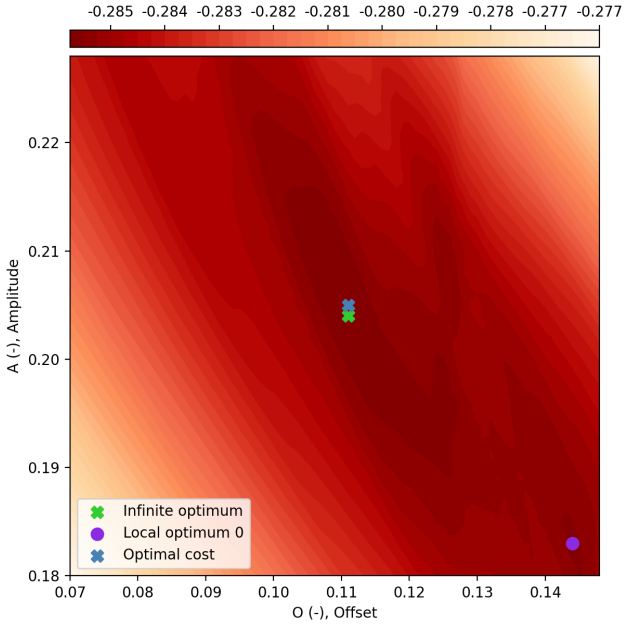


**Figure A-22:** Initial search space for $J_{\text{new}}$ and the control horizon $N_{\text{h}} = 1512$ (or $n = 63$). The optimal cost shifts and is located very near $v_{\infty}$. Note the different ranges of the y-axis with respect to all previous search spaces.

## A-2 Grid searches to determine optimal shift towards infinite optimum for objective function eliminating delays

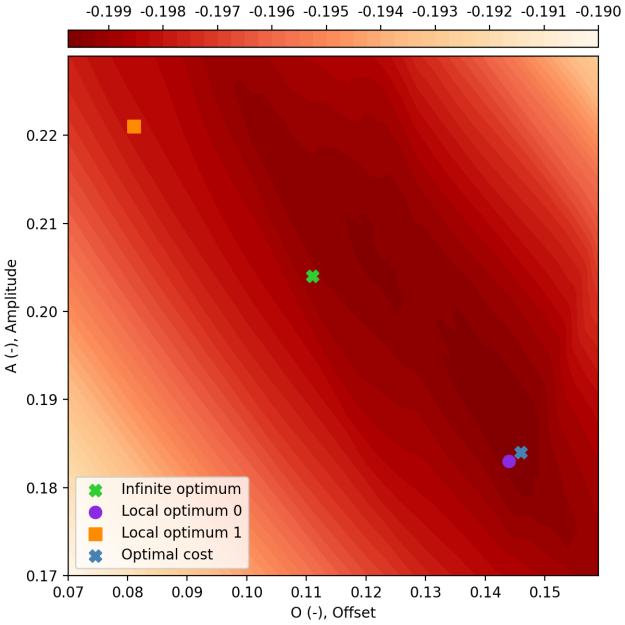


**Figure A-23:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 120$ (or $n = 5$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

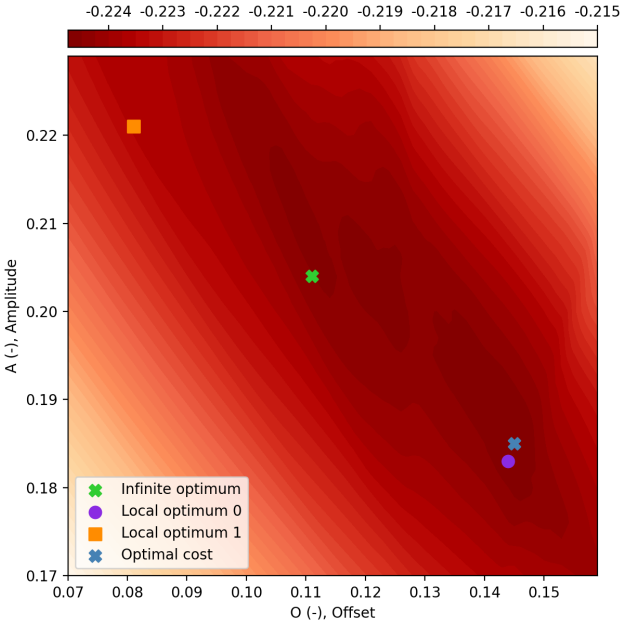


**Figure A-24:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 168$ (or $n = 7$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

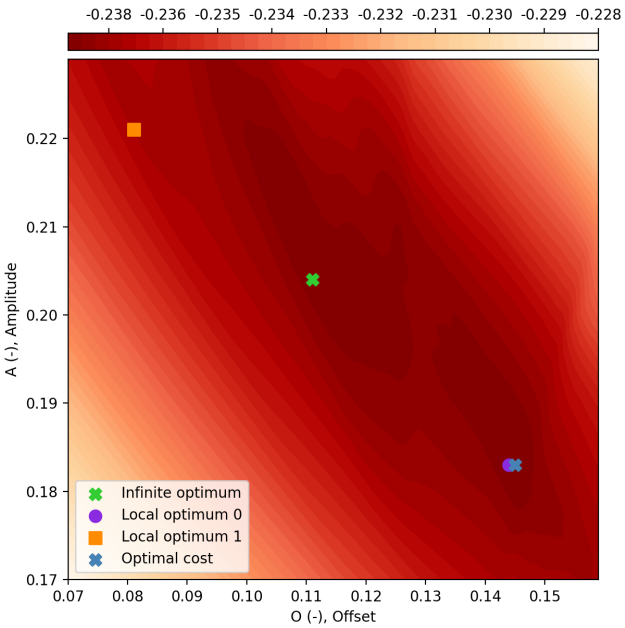


**Figure A-25:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 216$ (or $n = 9$). The optimal cost is located very near the local minimum $v_{\text{local},0}$.

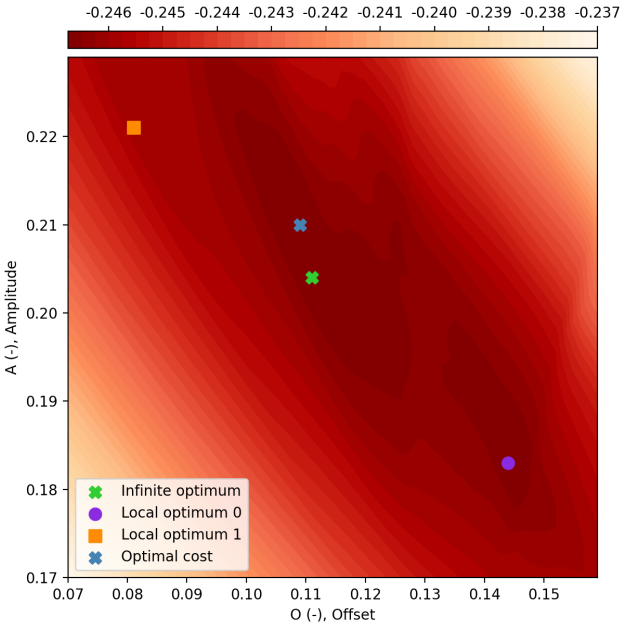


**Figure A-26:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 264$ (or $n = 11$). The optimal cost shifts towards $v_{\infty}$ and is now located within its three percent range, where it remains for the remainder of the control horizons.

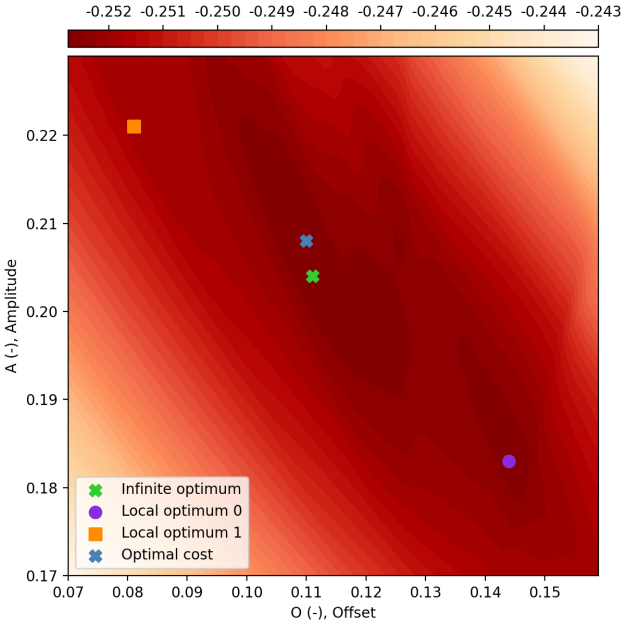


**Figure A-27:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 312$ (or $n = 13$). The optimal cost further shifts towards the infinite horizon optimum $v_{\infty}$.

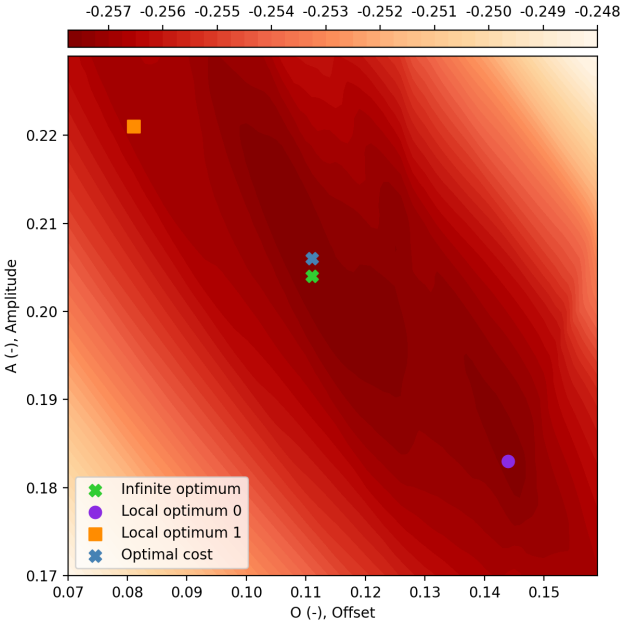


**Figure A-28:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 360$ (or $n = 15$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

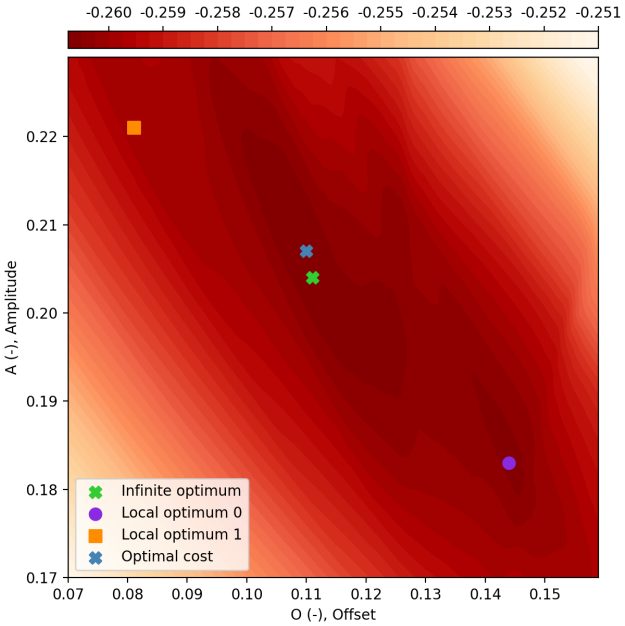


**Figure A-29:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 408$ (or $n = 17$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

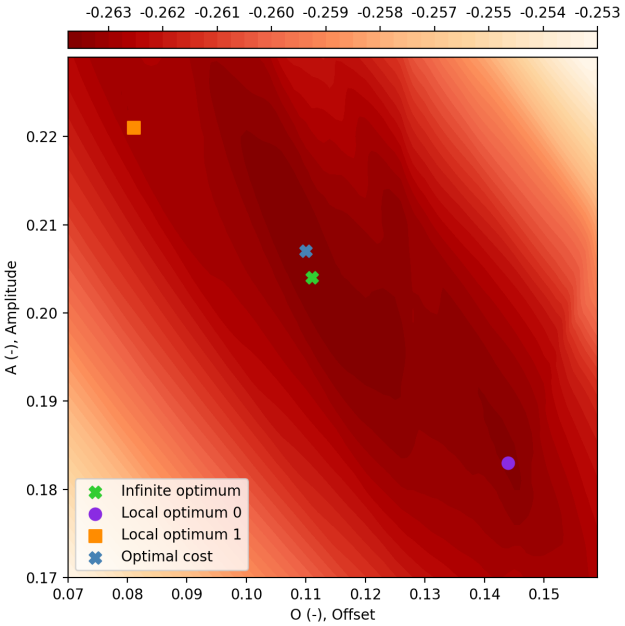


**Figure A-30:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 456$ (or $n = 19$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

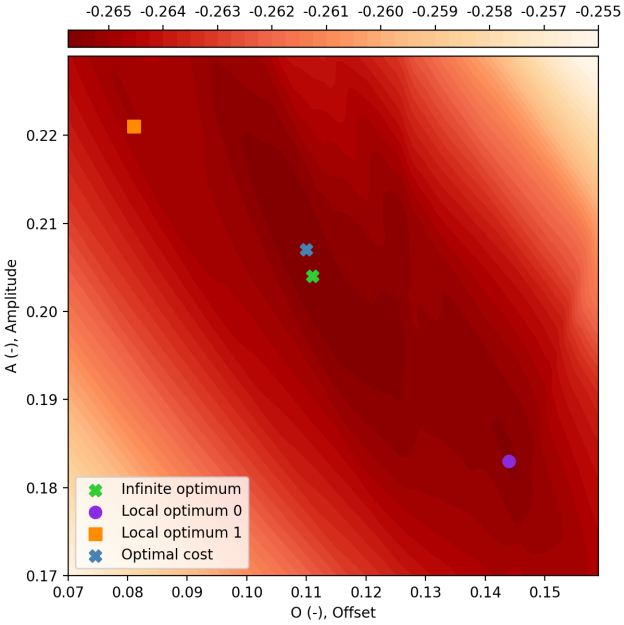


**Figure A-31:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 504$ (or $n = 21$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

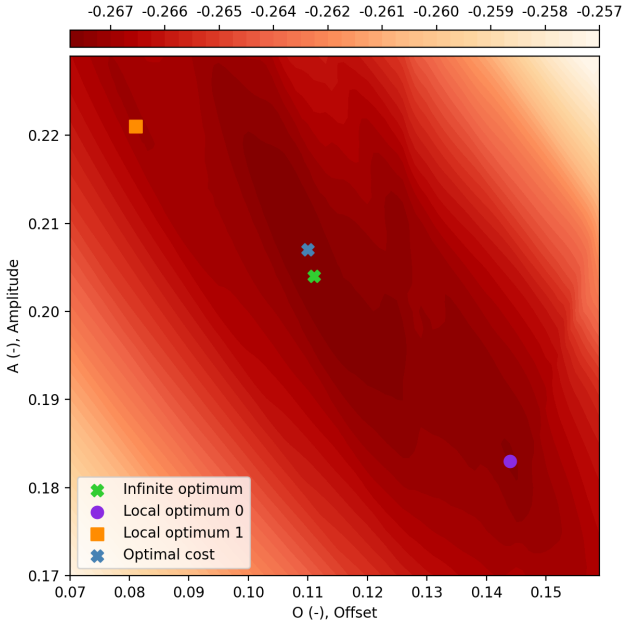


**Figure A-32:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 552$ (or $n = 23$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

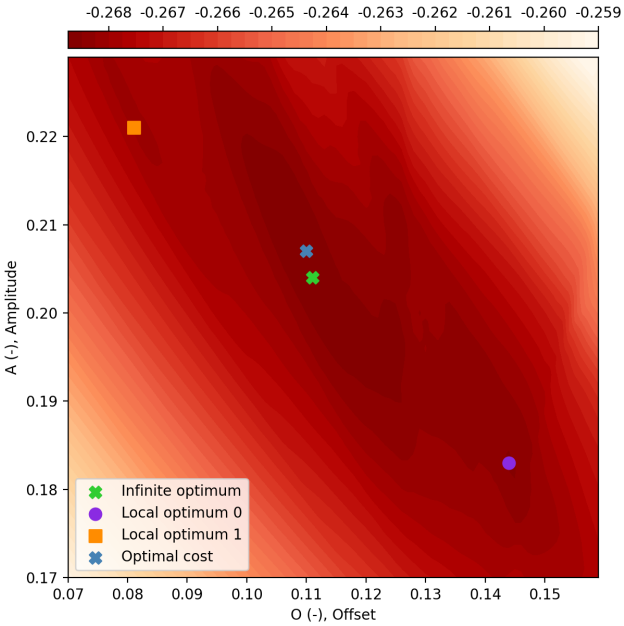


**Figure A-33:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 600$ (or $n = 25$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

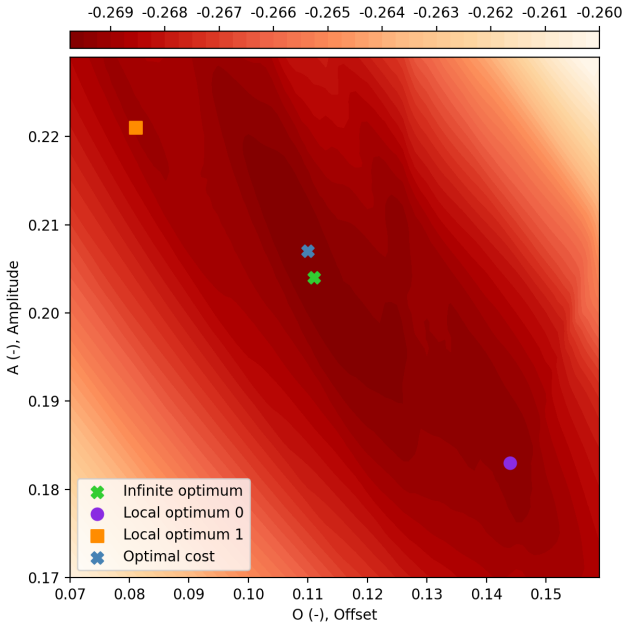


**Figure A-34:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 648$ (or $n = 27$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

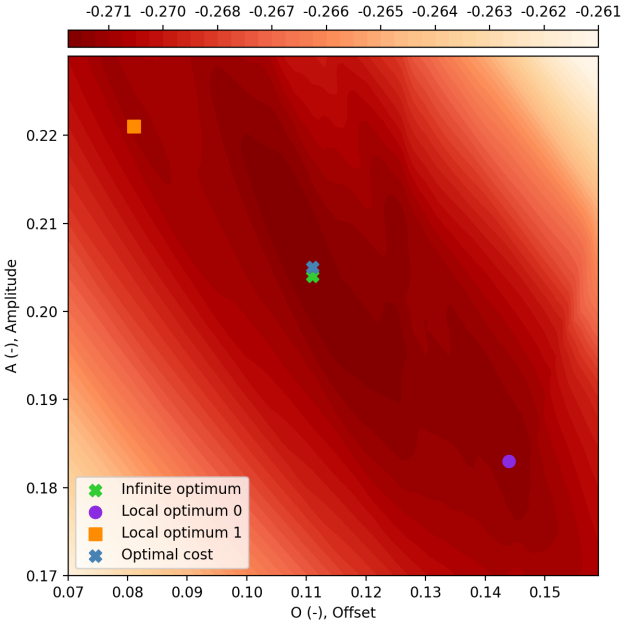


**Figure A-35:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 696$ (or $n = 29$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

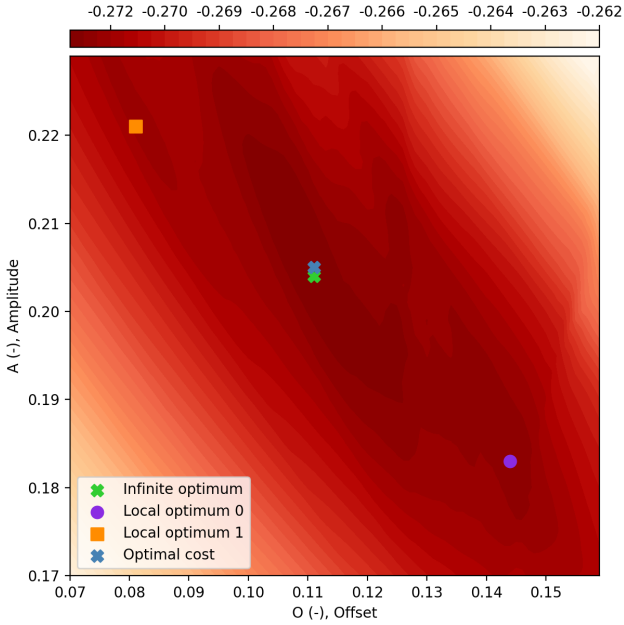


**Figure A-36:** Initial search space for $J_{\text{delay}}$ and the control horizon $N_{\text{h}} = 744$ (or $n = 31$). The optimal cost is located within the three percent range of the infinite horizon optimum $v_{\infty}$.

# Appendix B

# Additional power production results

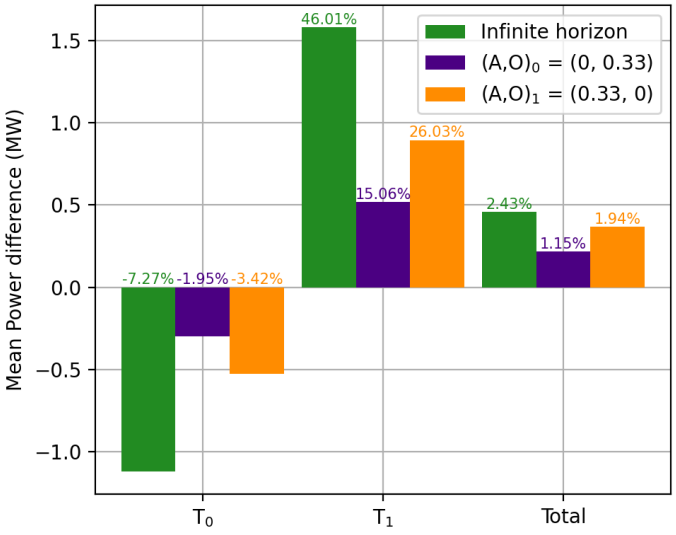


**Figure B-1:** The mean power production differences where the DIC baseline simulation results are subtracted from the infinite and finite horizon simulations for comparison. The relative power gains are also shown for the normalised objective function and a control horizon of $N_{\mathrm{h}} = 168$. The mean steady-state value for the simulations initialised in $v_{1,k=0}$ are similar to the local minimum $v_{\mathrm{local},0}$, therefore resulting in almost the same power gain with respect to the DIC baseline simulation.

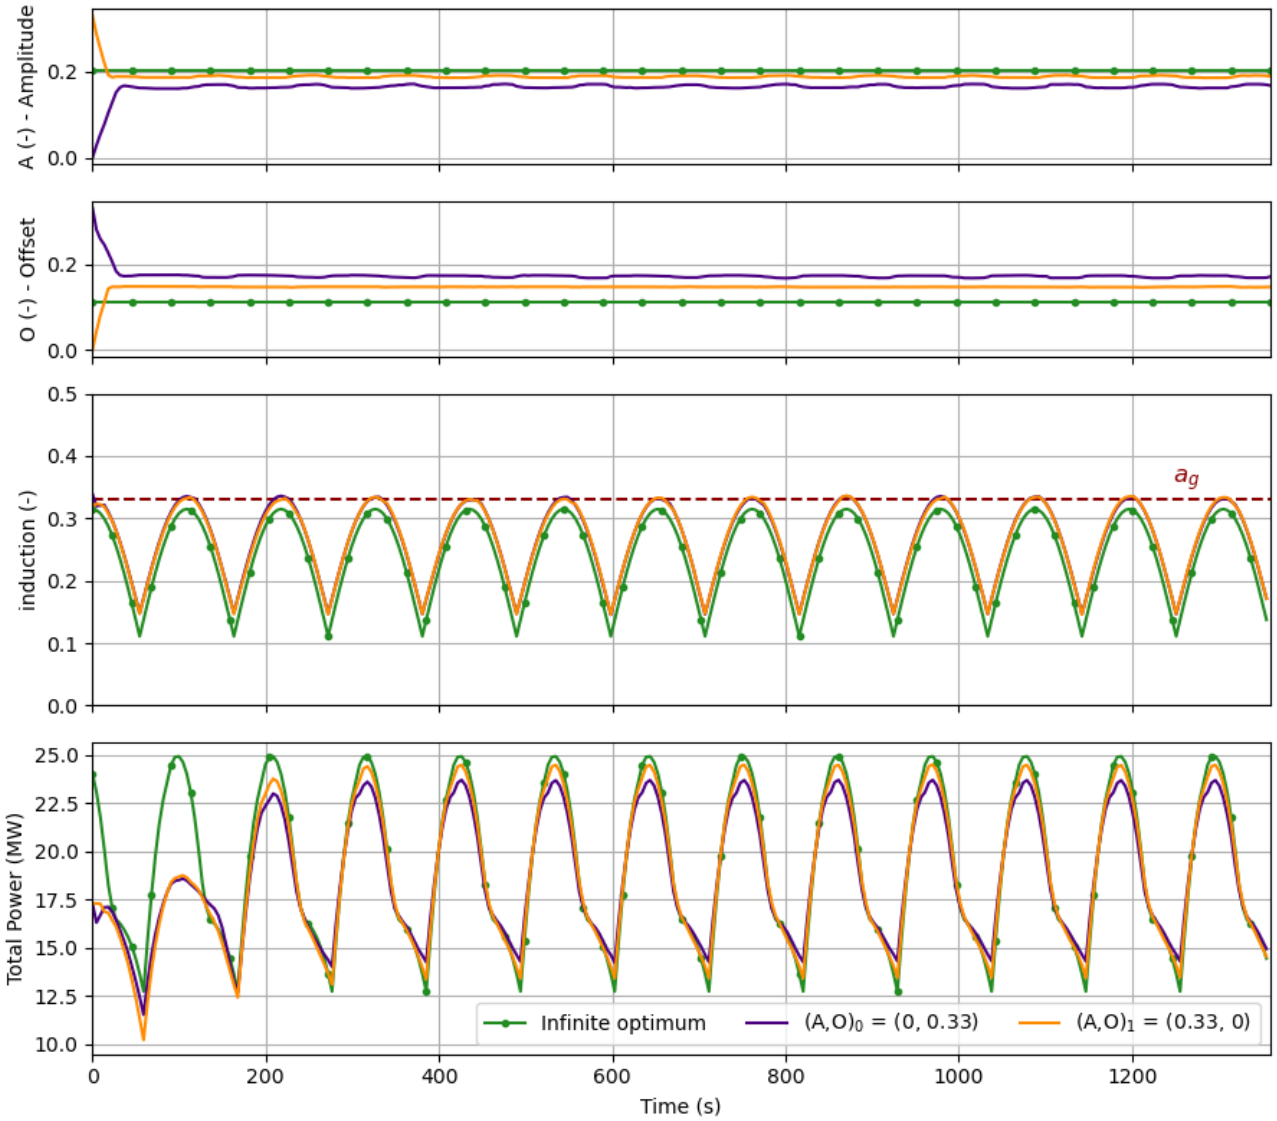


**Figure B-2:** All mean signals concerning the ten finite-horizon control simulations with the normalised objective function and a control horizon of $N_{\mathrm{h}} = 168$. The systems converge quickly towards their steady state, in which a slight periodicity remains.

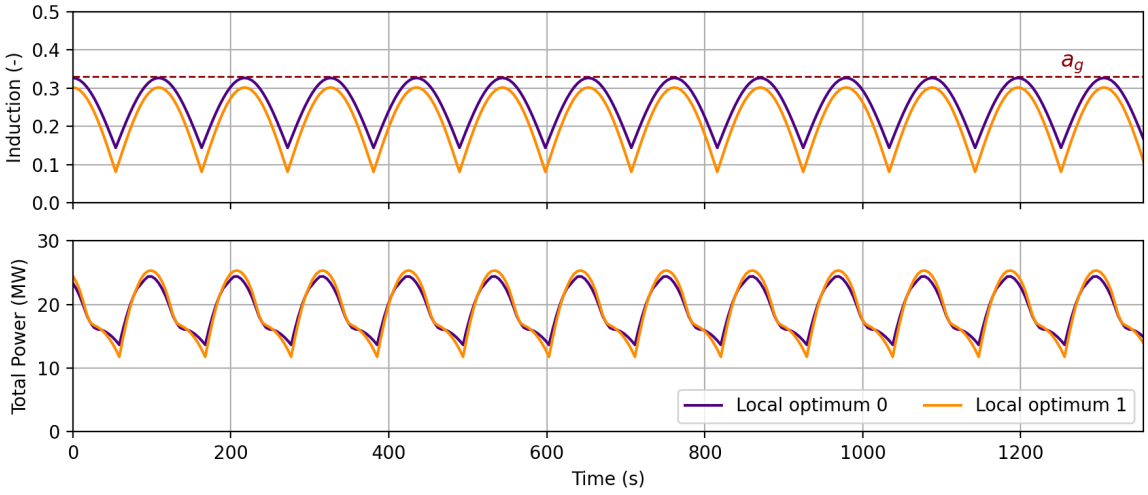


**Figure B-3:** The infinite horizon simulations for the absolute sine local minima configurations $v_{\text{local},0} = (A, O) = (0.183, 0.144)$ and $v_{\text{local},1} = (A, O) = (0.221, 0.081)$. The smaller mean induction factor value $\bar{a}_{\text{local},1}$ results in slightly better performance due to the better balance in power production emphasis between the two wind turbines.

# Bibliography


[1] M. van den Broek et al. "Adjoint optimisation for wind farm flow control with a free-vortex wake model". In: *Renewable Energy* 201 (Dec. 2022), pp. 752–765. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2022.10.120.

[2] C. Horowitz. "Paris Agreement". In: *International Legal Materials* 55.4 (Aug. 2016), pp. 740–755. ISSN: 0020-7829. DOI: 10.1017/S0020782900004253.

[3] C. Andrey et al. *Offshore renewable energy and grids: an analysis of visions towards 2050 for the Northern seas region and recommendations for upcoming scenario-building exercises*. Brussels, July 2021, p. 77. ISBN: 9789276472445.

[4] International Renewable Energy Agency. *RENEWABLE ENERGY STATISTICS 2022*. Tech. rep. 2022.

[5] O. Anaya-Lara et al. *Offshore wind energy technology*. 2018. ISBN: 1119097762.

[6] A. Kheirabadi and R. Nagamune. "A quantitative review of wind farm control with the objective of wind farm power maximization". In: *Journal of Wind Engineering and Industrial Aerodynamics* (2019). DOI: 10.1016/j.jweia.2019.06.015.

[7] B. Johnston et al. "Levelised cost of energy, A challenge for offshore wind". In: *Renewable Energy* 160 (Nov. 2020), pp. 876–885. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2020.06.030.

[8] J. Meyers et al. "Wind farm flow control: prospects and challenges". In: *Wind Energy Science* 7.6 (Nov. 2022), pp. 2271–2306. ISSN: 2366-7451. DOI: 10.5194/WES-7-2271-2022.

[9] J. Frederik et al. "Periodic dynamic induction control of wind farms: Proving the potential in simulations and wind tunnel experiments". In: *Wind Energy Science* 5.1 (Feb. 2020), pp. 245–257. ISSN: 23667451. DOI: 10.5194/WES-5-245-2020.

[10] J. Manwell, J. McGowan, and A. Rogers. *Wind Energy Explained: Theory, Design and Application*. 2nd. Wiley, 2009. ISBN: 978-0-470-01500-1.

[11] F. Bianchi, R. Mantz, and H. De Battista. *Wind turbine control systems : principles, modelling and gain scheduling design*. Advances in Industrial Control. London : Springer, 2007, 1 online resource (xix, 205 pages) : ISBN: 9781846284939. DOI: 10.1007/1-84628-493-7..

[12] T. Burton et al. *WIND ENERGY HANDBOOK*. 1981. ISBN: 0-471-48997-2.

[13] L. Andersson et al. "Wind farm control - Part I: A review on control system concepts and structures". In: *IET Renewable Power Generation* 15.10 (July 2021), pp. 2085–2108. ISSN: 1752-1424. DOI: 10.1049/RPG2.12160.

[14] M. Steinbuch et al. "Optimal control of wind power plants". In: *Journal of Wind Engineering and Industrial Aerodynamics* 27.1-3 (Jan. 1988), pp. 237–246. ISSN: 0167-6105. DOI: 10.1016/0167-6105(88)90039-6.

[15] Y. Su et al. "A coordinative optimization method of active power and fatigue distribution in onshore wind farms". In: *International Transactions on Electrical Energy Systems* 27.10 (Oct. 2017), e2392. ISSN: 2050-7038. DOI: 10.1002/ETEP.2392.

[16] M. Vali et al. "An active power control approach for wake-induced load alleviation in a fully developed wind farm boundary layer". In: *Wind Energy Science* 4.1 (Jan. 2019), pp. 139–161. ISSN: 23667451. DOI: 10.5194/WES-4-139-2019.

[17] M. Vali et al. "Model Predictive Active Power Control for Optimal Structural Load Equalization in Waked Wind Farms". In: *IEEE Transactions on Control Systems Technology* 30.1 (Jan. 2022), pp. 30–44. ISSN: 15580865. DOI: 10.1109/TCST.2021.3053776.

[18] M. Bastankhah and F. Porté-Agel. "A new analytical model for wind-turbine wakes". In: *Renewable Energy* 70 (Oct. 2014), pp. 116–123. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2014.01.002.

[19] B. Doekemeijer, D. van der Hoek, and J.W. van Wingerden. "Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions". In: *Renewable Energy* 156 (Aug. 2020), pp. 719–730. ISSN: 18790682. DOI: 10.1016/J.RENENE.2020.04.007.

[20] W. Munters and J. Meyers. "Towards practical dynamic induction control of wind farms: analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines". In: *Wind Energ. Sci* 3 (2018), pp. 409–425. DOI: 10.5194/wes-3-409-2018.

[21] J.W. van Wingerden et al. "Expert Elicitation on Wind Farm Control". In: *Journal of Physics: Conference Series* 1618.2 (Sept. 2020), p. 022025. ISSN: 1742-6596. DOI: 10.1088/1742-6596/1618/2/022025.

[22] T. Horvat, V. Spudic, and M. Baotic. "Quasi-stationary optimal control for wind farm with closely spaced turbines". In: *MIPRO, 2012 Proceedings of the 35th International Convention* (2012).

[23] J. Tian et al. "Active power dispatch method for a wind farm central controller considering wake effect". In: *IECON Proceedings (Industrial Electronics Conference)* (Feb. 2014), pp. 5450–5456. DOI: 10.1109/IECON.2014.7049333.

[24] A. Yılmaz and J. Meyers. "Optimal dynamic induction control of a pair of inline wind turbines". In: *Physics of Fluids* 30.8 (Aug. 2018), p. 085106. ISSN: 10897666. DOI: https://doi.org/10.1063/1.5038600.

[25] D. van der Hoek et al. “Effects of axial induction control on wind farm energy production-A field test”. In: (2019). DOI: 10.1016/j.renene.2019.03.117.

[26] J. Bartl and L. Sætran. “Blind test comparison of the performance and wake flow between two in-line wind turbines exposed to different turbulent inflow conditions”. In: *Wind Energy Science* 2.1 (2017), pp. 55–76. ISSN: 23667451. DOI: 10.5194/WES-2-55-2017.

[27] J. Annoni et al. “Analysis of axial-induction-based wind plant control using an engineering and a high-order wind plant model”. In: *Wind Energy* 19.6 (June 2016), pp. 1135–1150. ISSN: 1099-1824. DOI: 10.1002/WE.1891.

[28] P. Gebraad et al. “Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control”. In: *Wind Energy* (2016). DOI: 10.1002/we.1993.

[29] C. Bay et al. “Unlocking the Full Potential of Wake Steering: Implementation and Assessment of a Controls-Oriented Model”. In: *Wind Energy Science* (2019). DOI: 10.5194/wes-2019-19.

[30] P. Gebraad et al. “Wind plant power optimization through yaw control using a parametric model for wake effects - A CFD simulation study”. In: *Wind Energy* 19.1 (Jan. 2016), pp. 95–114. ISSN: 10991824. DOI: 10.1002/WE.1822.

[31] P. Fleming et al. “A simulation study demonstrating the importance of large-scale trailing vortices in wake steering”. In: *Wind Energy Science* 3.1 (Jan. 2018), pp. 243–255. ISSN: 23667451. DOI: 10.5194/WES-3-243-2018.

[32] B. Doekemeijer et al. “Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy”. In: *Wind Energy Science* 6.1 (Jan. 2021), pp. 159–176. ISSN: 23667451. DOI: 10.5194/WES-6-159-2021.

[33] P. Fleming et al. “Initial results from a field campaign of wake steering applied at a commercial wind farm – Part 1”. In: *Wind Energy Science* 4.2 (2019), pp. 273–285. ISSN: 23667451. DOI: 10.5194/WES-4-273-2019.

[34] P. Fleming et al. “Continued results from a field campaign of wake steering applied at a commercial wind farm - Part 2”. In: *Wind Energy Science* 5.3 (July 2020), pp. 945–958. ISSN: 23667451. DOI: 10.5194/WES-5-945-2020.

[35] M. Bastankhah and F. Porté-Agel. “Wind farm power optimization via yaw angle control: A wind tunnel study”. In: *Journal of Renewable and Sustainable Energy* 11.2 (Mar. 2019). ISSN: 19417012. DOI: 10.1063/1.5077038.

[36] F. Campagnolo et al. “Wind tunnel testing of wake steering with dynamic wind direction changes”. In: *Wind Energy Science* 5.4 (Oct. 2020), pp. 1273–1295. ISSN: 23667451. DOI: 10.5194/WES-5-1273-2020.

[37] P. Gebraad, P. Fleming, and J.W. van Wingerden. “Comparison of actuation methods for wake control in wind plants”. In: *Proceedings of the American Control Conference* 2015-July (July 2015), pp. 1695–1701. ISSN: 07431619. DOI: 10.1109/ACC.2015.7170977.

[38] J. Frederik et al. “The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms”. In: *Wind Energy* 23.8 (Aug. 2020), pp. 1739–1751. ISSN: 1099-1824. DOI: 10.1002/WE.2513.

[39] E. Bossanyi. "Individual blade pitch control for load reduction". In: *Wind Energy* 6.2 (2003), pp. 119–128. ISSN: 10954244. DOI: 10.1002/WE.76.

[40] A. Croce et al. "A CFD-based analysis of dynamic induction techniques for wind farm control applications". In: *Wind Energy* (Mar. 2022). ISSN: 10991824. DOI: 10.1002/WE.2801.

[41] L. Grüne and J. Pannek. *Nonlinear Model Predictive Control*. 2011. ISBN: 978-3-319-46024-6. DOI: 10.1007/978-3-319-46024-6.

[42] D. Simon. "Model Predictive Control in Flight Control Design : Stability and Reference Tracking". In: *Model Predictive Control in Flight Control Design : Stability and Reference Tracking* (Jan. 2014). DOI: 10.3384/LIC.DIVA-103742.

[43] J. Rawlings, D. Mayne, and M. Diehl. *Model Predictive Control: Theory, Computation, and Design 2nd Edition*. 2020. ISBN: 2020942771.

[44] S. Mulders et al. "Preventing wind turbine tower natural frequency excitation with a quasi-LPV model predictive control scheme". In: *Wind Energy* 23.3 (Mar. 2020), pp. 627–644. ISSN: 10991824. DOI: 10.1002/WE.2447.

[45] C. Bay et al. "Active Power Control for Wind Farms Using Distributed Model Predictive Control and Nearest Neighbor Communication". In: *Proceedings of the American Control Conference* 2018-June (Aug. 2018), pp. 682–687. ISSN: 07431619. DOI: 10.23919/ACC.2018.8431764.

[46] T. Faulwasser, L. Grüne, and M. Müller. *Economic nonlinear model predictive control*. Vol. 5. 1. Now Publishers Inc, 2018, pp. 224–409. ISBN: 978-1-68083-392-8. DOI: 10.1561/2600000014.

[47] J. Martins and S. Ning. *Engineering design optimization*. 2021, p. 637. ISBN: 9781108833417. DOI: doi:10.1017/9781108980647.

[48] N. Gionfra et al. "Wind farm distributed PSO-based control for constrained power generation maximization". In: *Renewable Energy* 133 (Apr. 2019), pp. 103–117. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2018.09.084.

[49] M. Vali et al. "Adjoint-based model predictive control for optimal energy extraction in waked wind farms". In: *Control Engineering Practice* 84 (2019), pp. 48–62. DOI: 10.1016/j.conengprac.2018.11.005.

[50] W. Munters and J. Meyers. "Optimal dynamic induction and yaw control of wind farms: effects of turbine spacing and layout". In: *Journal of Physics: Conference Series* 1037.3 (June 2018), p. 032015. ISSN: 1742-6596. DOI: 10.1088/1742-6596/1037/3/032015.

[51] Z. Wang et al. "A comparative study of common nature-inspired algorithms for continuous function optimization". In: *Entropy* 23.7 (July 2021), p. 874. ISSN: 10994300. DOI: 10.3390/E23070874/S1.

[52] P. Korošec and T. Eftimov. "Insights into Exploration and Exploitation Power of Optimization Algorithm Using DSCTool". In: *Mathematics 2020, Vol. 8, Page 1474* 8.9 (Sept. 2020), p. 1474. ISSN: 2227-7390. DOI: 10.3390/MATH8091474.

[53] S. Boersma et al. "A tutorial on control-oriented modeling and control of wind farms". In: *Proceedings of the American Control Conference* (June 2017), pp. 1–18. ISSN: 07431619. DOI: 10.23919/ACC.2017.7962923.

[54] G. Deskos, S. Laizet, and R. Palacios. "WInc3D: A novel framework for turbulence-resolving simulations of wind farm wake interactions". In: *ResearchGate* (2019). DOI: 10.13140/RG.2.2.19081.80487.

[55] J. Sørensen and W. Shen. "Numerical Modeling of Wind Turbine Wakes". In: *Journal of Fluids Engineering* 124.2 (June 2002), pp. 393–399. ISSN: 0098-2202. DOI: 10.1115/1.1471361.

[56] L. Martínez-Tossas, M. Churchfield, and S. Leonardi. "Large eddy simulations of the flow past wind turbines: actuator line and disk modeling". In: *Wind Energy* 18.6 (June 2015), pp. 1047–1060. ISSN: 1099-1824. DOI: 10.1002/WE.1747.

[57] R. Stanly et al. "Large-Eddy Simulation of a wind turbine using a Filtered Actuator Line Model". In: *Journal of Wind Engineering and Industrial Aerodynamics* 222 (Mar. 2022), p. 104868. ISSN: 0167-6105. DOI: 10.1016/J.JWEIA.2021.104868.

[58] L. Lignarolo et al. "Tip-vortex instability and turbulent mixing in wind-turbine wakes". In: *Journal of Fluid Mechanics* 781 (Oct. 2015), pp. 467–493. ISSN: 0022-1120. DOI: 10.1017/JFM.2015.470.

[59] D. van der Hoek et al. "Experimental analysis of the effect of dynamic induction control on a wind turbine wake". In: *Wind Energy Science* 7.3 (June 2022), pp. 1305–1320. ISSN: 23667451. DOI: 10.5194/WES-7-1305-2022.

[60] U. Piomelli. "Large-eddy simulation: achievements and challenges". In: *Progress in Aerospace Sciences* 35.4 (1999), pp. 335–362. ISSN: 03760421. DOI: 10.1016/S0376-0421(98)00014-1.

[61] J. Goit and J. Meyers. "Optimal control of energy extraction in wind-farm boundary layers". In: *Journal of Fluid Mechanics* 768 (2015), pp. 5–50. ISSN: 0022-1120. DOI: 10.1017/JFM.2015.70.

[62] W. Munters and J. Meyers. "An optimal control framework for dynamic induction control of wind farms and their interaction with the atmospheric boundary layer". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (Apr. 2017). ISSN: 1364503X. DOI: 10.1098/RSTA.2016.0100.

[63] P. Liu. "A General Theory of Fluid Mechanics". In: *A General Theory of Fluid Mechanics* (Jan. 2021), pp. 1–649. DOI: 10.1007/978-981-33-6660-2.

[64] J. Schluter. *Wind Turbine Wake Vortex Propagation (1) - YouTube*. URL: https://www.youtube.com/watch?v=p0n-LrCjokE&ab_channel=JorgSchluter.

[65] M. Glessmer and C. Janßen. "Using an Interactive Lattice Boltzmann Solver in Fluid Mechanics Instruction". In: *Computation 2017, Vol. 5, Page 35* 5.3 (July 2017), p. 35. ISSN: 2079-3197. DOI: 10.3390/COMPUTATION5030035.

[66] K. Brown et al. "Rapidly recovering wind turbine wakes with dynamic pitch and rotor speed control". In: *AIAA Scitech 2021 Forum* (2021), pp. 1–16. DOI: 10.2514/6.2021-1182.

[67] J. Katz and A. Plotkin. *Low-Speed Aerodynamics*. 2nd. Cambridge University Press, 2001. ISBN: 9780511810329. DOI: https://doi.org/10.1017/CBO9780511810329.

[68] D. Houck, N. Develder, and C. Kelley. "Comparison of a mid-fidelity free vortex wake method to a high-fidelity actuator line model large eddy simulation for wind turbine wake simulations". In: *Journal of Physics: Conference Series* 2265.4 (May 2022), p. 042044. ISSN: 1742-6596. DOI: 10.1088/1742-6596/2265/4/042044.

[69] T. Berdowski et al. "Three-dimensional free-wake vortex simulations of an actuator disc in yaw". In: *Wind Energy Symposium, 2018* (2018). DOI: 10.2514/6.2018-0513.

[70] L. Martínez-Tossas et al. "The aerodynamics of the curled wake: A simplified model in view of flow control". In: *Wind Energy Science* 4.1 (Jan. 2019), pp. 127–138. ISSN: 23667451. DOI: 10.5194/WES-4-127-2019.

[71] T. Lauß et al. "The Discrete Adjoint Gradient Computation for Optimization Problems in Multibody Dynamics". In: *Journal of Computational and Nonlinear Dynamics* 12.3 (May 2017). ISSN: 15551423. DOI: 10.1115/1.4035197/473875.

[72] E. Gaertner et al. *Definition of the IEA Wind 15-Megawatt Offshore Reference Wind Turbine Technical Report*. Tech. rep. 2020.

[73] A. Zaslavski. *Turnpike Phenomenon and Infinite Horizon Optimal Control*. 2014. ISBN: 978-3-319-08827-3. DOI: 10.1007/978-3-319-08828-0.

[74] T. Faulwasser and L. Grüne. "Turnpike Properties in Optimal Control: An Overview of Discrete-Time and Continuous-Time Results". In: *Handbook of Numerical Analysis* 23 (Nov. 2020), pp. 367–400. ISSN: 15708659. DOI: 10.1016/bs.hna.2021.12.011.

[75] D. Kingma and J. Lei Ba. "ADAM: A METHOD FOR STOCHASTIC OPTIMIZATION". In: *ArXiv* (2015).

[76] X. Jiang and S. Li. "BAS: Beetle Antennae Search Algorithm for Optimization Problems". In: *International Journal of Robotics and Control* 1.1 (Oct. 2017), p. 1. ISSN: 2577-7742. DOI: 10.5430/ijrc.v1n1p1.

[77] U. Ciri, M. Rotea, and S. Leonardi. "Model-free control of wind farms: A comparative study between individual and coordinated extremum seeking". In: *Renewable Energy* 113 (Dec. 2017), pp. 1033–1045. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2017.06.065.