MPC for the Slow Orbit Feedback Control at MAX IV

Emory Gassheld

My Karlsson

LUND
UNIVERSITY

Department of Automatic Control

MSc Thesis
TFRT-6197
ISSN 0280-5316

Department of Automatic Control
Lund University
Box 118
SE-221 00 LUND
Sweden

Printed in Sweden by Tryckeriet i E-huset
Lund 2023

Abstract

The MAX IV synchrotron radiation facility in Lund is designed to produce bright and high-energy X-ray radiation. One important component of achieving this is to keep the electron beam inside the storage rings centered and stable. The aim of this thesis was to develop an MPC (Model Predictive Control) that can accomplish this task more robustly than the currently commissioned I-controller. A controller was implemented in pyTango, a python module of Tango controls, using the Tango events system. By testing the controller on both a simulator and the real machine the result was a controller that fulfills its goal of keeping the beam centered while respecting the constraints of the system. Two additional features, offloading of another controller and controlling the radio frequency of the synchrotron to mitigate the impact of the MPC's control signals on the beam's energy, were also developed but due to a limited time frame testing was not finished. The results of the thesis show that using an MPC for this task has good potential. The MPC solves the problem at hand, and with more tuning and verification MAX IV should be able to use the MPC in production given that the issue of offloading of the other controller is resolved.

Acknowledgements

We would like to express our gratitude to our supervisor, Carla Takahashi, at MAX IV for her guidance and support throughout the project, as well as for providing us with all the necessary materials. We are also grateful to Magnus Sjöström for his expert help with the challenging physics concepts, which were particularly difficult for us as control engineering students. In addition, we would like to thank our supervisor at LTH, Pontus Giselsson, and our examiner, Bo Bernhardsson. Finally, we are thankful to MAX IV for providing us with the opportunity to work on such an interesting facility for our thesis. We would like to specifically thank the KITS department at MAX IV for hosting us and for the enjoyable conversations we had during lunch.

Contents

List of acronyms

BPM - Beam Position Monitor
FOFB - Fast Orbit Feedback
IPOPT - Interior Point Optimizer
KITS - Kontroll och IT-system (Controls and IT)
MPC - Model Predictive Control
RF - Radio Frequency
SOFB - Slow Orbit Feedback

1

Introduction

Producing highly brilliant X-ray radiation is a complex process. At MAX IV, electrons are accelerated to speeds close to the velocities of light and then deflected, which, in turn, makes them emit the sought-after X-ray radiation. The facility is a remarkable example of what can be achieved in the space where advanced physics and many different disciplines of engineering converge. The result is a world-leading light source in terms of brightness.

1.1 Problem formulation

The aim of this thesis is to implement a control system using model predictive control (MPC) for the Slow Orbit Feedback (SOFB) System at MAX IV. The primary function of the SOFB is to prevent the electron beam in the storage rings from drifting. The MPC will replace the current implementation of an I-controller, which uses singular value decomposition to invert the so-called response matrix in the SOFB. The main challenges with the current implementation are:

- Corrector magnet saturation. This leads to orbit drifts which means performance degradation and beam loss if left without action for too long. This if further discussed in Section 2.1.
- Manual intervention. It depends on operators to manipulate singular values to address temporary saturation issues, for which they are not trained, and this places an unnecessary workload on them.
- Undefined transitions. The state machine is not well defined, i.e. it is unclear what should occur during transitions. The state machine is further discussed in Section 5.6.

These problems are the reason why MAX IV now seeks a solution that better respects the system's constraints to avoid correction magnet saturation and enhance system robustness. The thesis aims to answer the questions

- Is it possible/suitable to use MPC for the SOFB?
- Is it possible to reduce or entirely get rid of the corrector magnet saturation problems?
- Can the robustness of the system be improved by using MPC?

1.2 Requirements

For the thesis, a number of requirements for the implementation of an MPC have been put forward. The implementation should:

- Be written in PyTango, see Section 5.2.
- Take advantage of the Tango event system, see Section 5.2.
- Include three simple states for the internal state machine: OFF, STANDBY and RUNNING, see Section 5.6.
- Take into consideration the connection to the Fast Orbit Feedback (FOFB) and how to offload it, see Section 4.3.
- Take into consideration the Radio Frequency (RF) adjustment, see Section 3.2.

2

Background

MAX IV is a synchrotron radiation facility located in Brunnshög, Lund, Sweden. MAX IV produces X-ray radiation for a wide variety of applications, including life sciences, biology, material sciences, and physics. Synchrotron radiation is produced when fast-travelling electrons are forced to change direction. Electrons are ejected from an electron gun at speeds close to the speed of light into a linear accelerator (linac) where they are further accelerated. The electrons are then sent to one of the two storage rings. Electrons of lower energy are sent to the smaller 1.5 GeV ring, with a circumference of 96 m, while higher energy electrons are sent to the 3 GeV ring, which has a circumference of 528 m. The storage rings consist of a number of achromats, which are arrays of bending and focusing magnets. In the storage rings, the electrons pass through insertion devices, which are arrays of magnets with alternating polarity. These magnets force the electrons to oscillate, and they emit energy in the form of light in the travelling direction. This light is directed into a beamline, where the sample that is to be studied is placed. The beamlines are all specialized for different uses, and the optics and detectors vary from beamline to beamline [MAX IV, 2022].

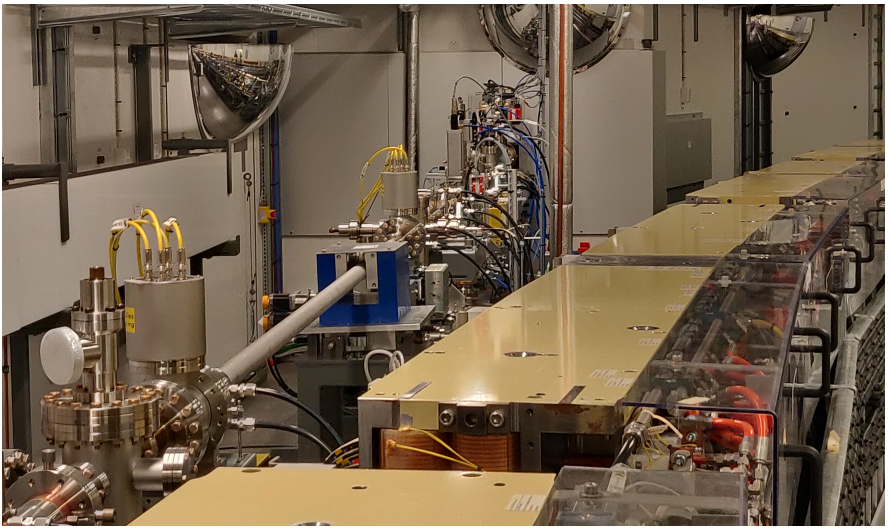

Figure 2.1 A picture from inside the tunnel of the 3 GeV ring at MAX IV. Curving off to the right is an achromat, which consists of an array of assorted magnets and other equipment. Inside the achromat is the pipe through which the electrons are traveling. The pipe seen on the left side of the achromat is the beginning of the beamline. The light emitted is enclosed in the pipe, which exits the accelerator tunnel through the wall in the middle of the image. Behind the photographer on the left side is the insertion device, where the electrons are made to oscillate and emit light. Photograph taken by one of the authors.

2.1 Orbit correction system at MAX IV

The orbit feedback correction system at MAX IV controls the trajectory of the electron beam inside of the storage rings. To achieve the high brightness goals of the produced synchrotron light, the low emittance and transverse stability of the electron beam is of great importance [Bell et al., 2017]. The electron beam has a transverse size of 1 to 4 μm and has an elliptical shape, with the longer axis in the horizontal plane [Sjöström et al., 2011]. The orbit correction system consists of two separate but connected systems: the Slow Orbit Feedback (abbreviated as SOFB) and the Fast Orbit Feedback (abbreviated as FOFB). The SOFB is software-based [Bell et al., 2017] and handles low-frequency disturbances like misalignments and drifts [Sjöström et al., 2011]. It is designed to operate at a 10 Hz rate [Bell et al., 2017]. The FOFB is hardware-based [Bell et al., 2017] and handles high-frequency disturbances like beam jitter [Sjöström et al., 2011]. It operates at a much higher rate, at 10 kHz. The two systems use the same set of sensors, called Beam Position Monitors (or BPMs), to keep track of the beam, but use different sets of actuator magnets to correct the beam. The actuator magnets of the FOFB have much higher bandwidth, but are much weaker than those of the SOFB, and the SOFB therefore

periodically off-loads the FOFB to prevent saturation of the FOFB magnets. This means that even though the subject of this thesis is the SOFB, the connection to the FOFB will also have to be considered [Sjöström et al., 2011].

The rings are equipped with Beam Position Monitors (BPMs), 200 for the 3 GeV ring, and 36 for the 1.5 GeV ring [Bell et al., 2017]. The BPMs output the vertical and horizontal position of the beam at a frequency of 10 kHz [Sjöström et al., 2011]. Therefore 400, respectively 72, position readings are outputted at each time instance.

The 1.5 GeV ring has 36 slow actuator magnets in the horizontal plane with a strength of $\pm$11.5 A and 36 in the vertical plane with a strength of $\pm$ 10.5 A, while the 3 GeV ring has 200 in the horizontal plane and 180 in the vertical plane, both with the strength $\pm$ 5 A. The actuators work by applying a current that changes the magnetic field produced by the actuator magnets that deflects the beam. The current applied corresponds to a specific beam deflection angle. The maximum beam deflection angle is in the order of $10^{-}4\mu$rad . Due to spacial limitations, one actuator in the vertical plane had to be removed from each of the 20 achromats during the design phase to leave room for the synchrotron light exiting towards the beamline. Since the SOFB deals with drifts and misalignments, the slow set of actuators were chosen for strength over bandwidth [Sjöström et al., 2011].

The current SOFB is implemented as a multiple input, multiple output I-controller which uses singular value decomposition to invert the so called response matrix, see Section 3.1. This design was chosen since the controller mostly has to deal with stationary errors and very infrequently set-point changes. The main problem with the current implementation is that the controller does not take the limits of the corrector magnets into account which often leads to saturation of the corrector magnets. If this is left unattended for to long it will lead to orbit drifts and in the worst case losing the beam. Currently, the only way to get the controller out of a saturated state is to manually manipulate the singular values of the decomposed response matrix. Apart from being able to periodically offload the FOFB, the SOFB also features the ability to correct the energy of the beam by changing the radio frequency of the rings. This is described in greater detail in Section 3.2. The current SOFB does not feature any anti-windup.

2.2 Model Predictive Control

Model Predictive Control, abbreviated as MPC, is a control system that uses a model of the system's dynamics to predict its future behavior. The prediction is based on the current states, disturbances, and current and future control signals. In each sample, the optimal control signals that minimize the cost function and respect the

constraints on states, control actions, rate of change of states and control signals are found. The model may be a linear discrete model such as

$$\begin{aligned} x_{k+1} &= \Phi x_k + \Gamma u_k \\ y_k &= C x_k, \end{aligned} \tag{2.1}$$

where $x_k \in \mathbb{R}^n$ are the states at time k and $u_k \in \mathbb{R}^m$ are the control signals at time k. If x_0 is the initial state, the k-th state is then given by

$$x_k = \Phi^k x_o + \sum_{n=0}^{k-1} \Phi^n \Gamma u_{k-1-n} \tag{2.2}$$

The cost function at time $t = k$, can be of the form

$$J = \sum_{t=k+1}^{k+H_p-1} e_t^T Q_1 e_t + \sum_{t=k}^{k+H_u-1} \Delta u_t^T Q_2 \Delta u_t + e_{k+H_p}^T Q_f e_{k+H_p} \tag{2.3}$$

where $e_t = x_{\text{ref}} - x_t$ is the predicted state error for time t, $\Delta u_t = u_t - u_{t-1}$ is the difference in control signal between time t and $t-1$ and e_{k+H_p} is the terminal state error at time $k+H_p$. The matrices Q_1, Q_2 and Q_f are positive semi-definite matrices that penalize the state errors, control signals and terminal state. The idea is that the contribution to the cost function is weighted differently depending on the state or control signal. For example, small changes in the control signal is achieved by a large penalty on the control signal differences, i.e., a large Q_2 relative to Q_1 and Q_f. The variable H_p is called the prediction horizon and determines how many steps ahead the MPC will simulate the system. Similarly H_u is the control horizon and determines how many steps ahead the optimal control signal will be computed. It is important to for H_p to be large enough to capture the dynamics of the system, but it is also important to not have H_p too large since it is computationally expensive to predict too far ahead. After time $t = k + H_u$, the control signal is assumed to be constant. The optimal control signal is then applied to the system. Typically, only the first computed control signal is applied and then in the next sample the optimal control signal H_u steps ahead is computed again. The optimization problem can be infeasible, meaning that there does not exist a solution that satisfies all the constraints. To avoid having the controller crash if the problem is infeasible, soft constraints can be introduced that relaxes the requirements on the constraints. for One of the big advantages of MPC is that the controller handles constraints which is important because there can be physical limits or safety limits on states and control signals. The formulation of the optimization problem can look like this [Cervin, 2022]:

$$\begin{aligned} \min_{x,u} \quad & J = \sum_{t=k+1}^{k+H_p-1} e_t^T Q_1 e_t + \sum_{t=k}^{k+H_u-1} \Delta u_t^T Q_2 \Delta u_t + e_{k+H_p}^T Q_f e_{k+H_p} \\ \text{s.t} \quad & x_{k+1} = \Phi x_k + \Gamma u_k \\ & \bar{x} = x_0 \\ & |x| \leq x_{max} \\ & |u| \leq u_{max} \end{aligned} \tag{2.4}$$

where $\bar{x}$ is a measurement or estimate of the current state. The entire state x needs to be available for measurements. To solve the optimization problem an optimization algorithm is used. There are several toolboxes in Matlab, Python and other languages that solve optimization problems. In Figure 2.2, a simple overview is provided of how the MPC works.

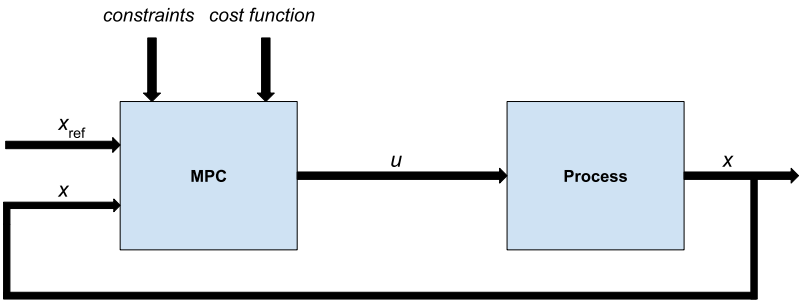

Figure 2.2 Schematic overview of the MPC process.

3

Model

While this is not a thesis on physics, it would be helpful to have some background knowledge of the accelerator physics that govern the dynamics of the system relevant to our control problem. Therefore, this section provides a basic explanation of some of the key concepts.

3.1 Basic accelerator optics

A point charge (in this case an electron) with charge, q, travelling through an electric field, $\mathbf{E}$, and a magnetic field, $\mathbf{B}$, with velocity, $\mathbf{v}$, will experience a force

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

called the *Lorentz force*. Therefore the trajectory of an electron can be manipulated by applying a magnetic field [Wolski, 2014].

The electron beam is composed of so-called bunches of electrons that need to be continuously focused as they repel each other due to their negative charge. In the transverse plane this is achieved by placing quadrupole magnets in the beam path. These produce a field which focuses the beam along one axis and spreads it out in the other. By adding another quadrupole magnet in series with the first one, but rotated 90 degrees, a focusing effect in both directions is achieved. Sextupole and octupole magnets, serving similar purposes, are also installed on the rings. Dipole magnets are used to deflect the beam to travel in a circular path by bending the beam trajectory [Wolski, 2014].

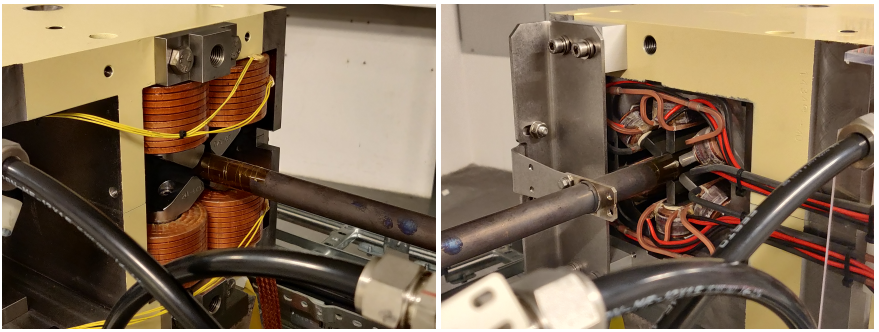

(a) Quadrupole magnet. **(b)** Sextupole magnet.

Figure 3.1 Photographs from inside of the tunnel of the 3 GeV ring showing a quadrupole and a sextupole magnet. The magnets surround the vacuum pipe in which the electrons are travelling.

Closed orbit distortion

The beam will encounter some static dipole errors when orbiting in the ring. These alter the equilibrium orbit of the electron beam from the reference orbit in what is called a closed orbit distortion. The effect of one such dipole error on the equilibrium orbit can be measured and assuming small enough orbit distortion, the non-linear effects of the sextupoles and octupoles can be ignored and the effect from all dipole errors can be added in a linear fashion [Wolski, 2014, p. 198]. The actuator magnets can be considered as dipole errors and the relationship between actuator magnet currents and BPM readings can be expressed in matrix form in what will be referred to as the response matrix, $\mathbf{R}$ [Bell et al., 2017]. The response matrix is in reality a dynamic mapping, but is considered static for simplicity. The Response matrix can be measured by changing the current of the actuator magnets one at a time and recording the sensor readings. This gives the system

$$\mathbf{x} = \mathbf{R}\mathbf{u} \tag{3.1}$$

where $\mathbf{x} \in \mathbb{R}^n$ is a vector of BPM readings, i.e., the states of the system, and $\mathbf{u} \in \mathbb{R}^m$ is the current to the corrector magnets, i.e., the control signal. It is worth noting that the sensors in the vertical plane are almost completely unaffected by the actuators in the horizontal plane and vice versa. This is good news, since it means that the control problem can be divided into two problems, decreasing the computational load significantly [Bell et al., 2017]. The horizontal and vertical plane will be treated separately for the remainder of this thesis.

Radio frequency cavities

The storage rings are fitted with so called radio frequency cavities, or RF cavities for short. These accelerate the electrons in the rings to compensate for the loss via

emission of synchrotron light, so that the energy of the electrons is maintained while they circulate the rings. They also focus the electron bunches in the longitudinal direction. The RF cavities work as such that when an electron bunch arrives in the cavity, an electric field is applied in such a way that low energy electrons arrive at a phase where they gain more than high energy electrons. A perfectly timed electron with the desired energy will be unaffected by the field in the RF cavity while lower energy electrons will be accelerated and higher energy electrons will be decelerated. The variations in the electric field need to be precisely timed to match the arrival of the electron bunches. This is where radio frequency comes in. The electromagnetic field in the RF cavities vary sinusoidally according to a set frequency of about 100 MHz (which is in the FM radio band) at MAX IV [CERN, 2023] [Lindvall, 2022]. The synchronization between the field switches in the RF cavities and the incoming beam is the origin of the name synchrotron [Gavaghan, 2001].

3.2 Compensation of energy shift using RF adjustment

When the actuator magnets are affecting the position of the beam, they also slightly change the energy of the beam. This effect occurs in the horizontal plane but not in the vertical plane. The model in the horizontal plane therefore needs to be expanded to include this effect. There are several models to describe and predict this change in energy, but the one used in this thesis is called the orbit lengthening model and describes the energy shift under the assumption that a change in the actuator currents change the length of the beam orbit [Wenninger, 1997]. This change in energy comprises a linear and a quadratic component where the quadratic component is negligible compared to the linear component. The change in energy is then given by the formula

$$\frac{\Delta E}{E} = -\frac{1}{L_0 \alpha} \eta_{Actuator}^{T} \mathbf{u} \tag{3.2}$$

where η_{Actuator} is given by the so called dispersion function, L_0 is the circumference of the storage ring and α is the momentum compaction factor [Wenninger, 1997]. The dispersion function describes how much each actuator affects the beam orbit length and is measured by the accelerator physicists. The momentum compaction factor describes how the momentum of the particle is affected by the path length and is also provided by the accelerator physicists.

A change in the energy of the beam leads to a change in the energy of the synchrotron light emitted into the beamlines. If the energy of the synchrotron light changes, the operators at the beamlines need to recalibrate their equipment to get the proper measurements. It is therefore undesired for the energy of the beam to change during operation. The change in energy caused by the actuator magnets can however be counteracted by adjusting the orbit frequency, f_o.

The momentum compaction factor is given by

$$\alpha = \frac{\Delta L/L_0}{\Delta p/p} \tag{3.3}$$

where $\Delta L/L_0$ is the change of the path length relative to the ring circumference and $\Delta p/p$ is the relative change of the momentum [Wiedemann, 2019, equation 8.124]. At speeds close to the speed of light, i.e. when $v \approx c$, the relative change of the momentum can be approximated to be equal to the relative change in beam energy, i.e

$$\frac{\Delta p}{p} \approx \frac{\Delta E}{E} \tag{3.4}$$

[Wiedemann, 2019]. From the relationship between particle velocity and orbit frequency the relationship between relative change of path length and the relative change in orbit frequency can be derived according to

$$\frac{L}{v} = \frac{1}{f_o} \tag{3.5}$$

Differentiation gives

$$\frac{\Delta L}{v} = -\frac{1}{f_o^2}\Delta f_o \tag{3.6}$$

Dividing (3.6) with (3.5) gives

$$\frac{\Delta L}{L} = -\frac{\Delta f_o}{f_o}$$

Inserting this result and (3.4) into (3.3) and rearranging gives how the orbit frequency affects the energy shift

$$\frac{\Delta E}{E} = -\frac{1}{\alpha}\frac{\Delta f_o}{f_o} \tag{3.7}$$

In this way the energy shift can be brought to zero. The radio frequency, f, needs to be an integer multiple of the orbit frequency. This integer multiple is called the harmonic number. At MAX IV, the radio frequency is chosen to be around 100 MHz, which gives a harmonic number of 176 for the 3 GeV ring and 32 for the 1.5 GeV ring. Since

$$\frac{\Delta f_o}{f_o} = \frac{\Delta f/n_{\text{harmonic}}}{f/n_{\text{harmonic}}} = \frac{\Delta f}{f}$$

(3.7) can be expressed in terms of the radio frequency instead of the orbit frequency

$$\frac{\Delta E}{E} = -\frac{1}{\alpha}\frac{\Delta f}{f} \tag{3.8}$$

However, the change of the radio frequency also affects the position of the beam, which means that the states will be affected. This relationship is

$$\Delta \mathbf{x} = -\frac{1}{\alpha}\eta_{\text{BPM}}\Delta f \tag{3.9}$$

where η_{BPM} is given by the dispersion function for the BPMs which describes how much a change in the frequency affects the sensor readings. A change in frequency, Δf will result in a change in sensor readings, $\Delta \mathbf{x}$ [Wiedemann, 2019, equation 10.83].

3.3 Known weaknesses in the model

The model used by our MPC has room for improvement. The current model does not account for any delays in the system, which is a questionable assumption. The model assumes that any change in the control signal immediately yields a response in the sensor readings, and that this response is consistent every time.

In regards to the delays, they can be expected to appear both on the actuator magnet side and on the BPM side of things. The actual process, i.e. the response of the position of the beam to the magnetic field changing can be assumed to be instantaneous, since the time scale on which this occurs (electrons travelling at close to the speed of light) is so fast compared to the time scale on which the SOFB operates (1 to 10 Hz). For the actuators, the desired control signal is sent from the server to the power supplies for each of the corrector magnets spread out around the circumference of the storage rings. The corrector magnets then have a certain rise time from the point in time where the new current is applied to when the new magnetic field strength is in place. This rise time is due to the dynamics of the power supplies as well as the conductive characteristics of the metallic pipe surrounding the electron beam. From the time that the BPMs make a readout there is a delay before the signal reaches the server and from there a delay while the server collects data from all of the BPMs and synchronizes them. These delays have been measured to be approximately 0.4 s. This is a considerable delay for a controller which is meant to run at 10 Hz.

The second assumption, that the response is the same every time a certain current is applied is not neccessarily true. This is due to something called magnetic hysteresis. Magnetic hysteresis is a phenomenon where a ferromagnet's dipoles align with an external magnetic field, and some of this alignment remains even after the field is removed [Chikazumi, 1997]. This means that in a way, all of the magnets' history is stored within it. It is therefore difficult and sometimes impractical to precisely predict how the magnets will react to a certain current.

4

Control

In this chapter the state space model of the control system is set up, and the control dynamics are motivated in Section 4.1. The optimization problem and control parameters are set up in Section 4.2.

4.1 Assembling the state space model

Under the assumption that all dynamics of the system are fast enough compared to the time constant of the Slow Orbit Feedback, the state space model will only contain the relationship of (3.1) and whatever control dynamics we choose to include in our controller. Since the controller's main objective is to handle drifts and stationary errors, it was decided to introduce integral action in the state space model. This is consistent with the old implementation of the SOFB. Using integral action and the relationship between change in actuator currents and change in sensor readings defined in (3.1), the system can be set up in a state space form as

$$\mathbf{x}_{k+1} = \mathbf{I}\mathbf{x}_k + \mathbf{R}\Delta\mathbf{u}_k \tag{4.1}$$

where $\mathbf{I} \in \mathbb{R}^{n\times n}$ is the identity matrix of dimensions corresponding to the length of the state vector, $\mathbf{x}_k \in \mathbb{R}^n$. The states of the system, $\mathbf{x}_k$ are the sensor readings from the BPMs while the control signals, $\mathbf{u}_k \in \mathbb{R}^m$, are the actuator currents. $\Delta\mathbf{u}_k \in \mathbb{R}^m$ is set up as the difference between the new control signal and the previous one, i.e.,

$$\Delta\mathbf{u}_k = \mathbf{u}_k - \mathbf{u}_{k-1} \tag{4.2}$$

and $\mathbf{R} \in \mathbb{R}^{n\times m}$ is the response matrix as defined in Section 3.1. The size of $\mathbb{R}$ is 36×36 in the horizontal and vertical plane respectively for the small ring, and 200×200 in the horizontal plane and 200×180 for the vertical plane for the large ring. This state space model is used when running the MPC without attempting to correct the energy of the beam via the radio frequency.

Adding RF to the state space model

The dynamics of the energy shift and the RF adjustment can be added to the state space model by introducing $\Delta E/E$ as a state, and the RF frequency, f, as a control signal. This means that the relationships presented in (3.2), (3.8) and (3.9) can be included in the state space model by concatenating the matrices as follows

$$\begin{bmatrix} \mathbf{x} \\ \frac{\Delta E}{E} \end{bmatrix}_{k+1} = \mathbf{I} \begin{bmatrix} \mathbf{x} \\ \frac{\Delta E}{E} \end{bmatrix}_{k} + \begin{bmatrix} \mathbf{R} & -\frac{1}{\alpha}\eta_{\text{BPM}} \\ -\frac{1}{L_0\alpha}\eta^T_{\text{Actuator}} & -\frac{1}{\alpha f} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{u} \\ \Delta f \end{bmatrix}_{k} \tag{4.3}$$

where $\mathbf{I} \in \mathbb{R}^{n+1\times n+1}$ is the identity matrix with size corresponding to the length of the new state vector which now includes the energy shift and where Δf is the difference between the new radio frequency and the previous one, i.e.

$$\Delta f_k = f_k - f_{k-1} \tag{4.4}$$

This means that the RF adjustment can be incorporated in the optimization problem solved by the MPC. The new state and control signal are also using integral action.

4.2 MPC Optimization

The optimization problem that is solved by the MPC is as follows:

$$\begin{aligned} \min_{x,u} \quad & J = \sum_{t=k+1}^{k+H_p-1} e_t^T Q_1 e_t + \sum_{t=k}^{k+H_u-1} u_t^T Q_2 u_t + e_{k+H_p}^T Q_1 e_{k+H_p} \\ \text{s.t} \quad & x_{k+1} = \mathbf{I}x_k + \mathbf{R}\Delta u_k \\ & \bar{x} = x_0 \\ & |x| \leq x_{\max} \\ & |u| \leq u_{\max} \end{aligned} \tag{4.5}$$

i.e. the MPC will minimize the cost function, J, subject to the constraints given by the state space model, initial values ($\bar{x}$), as well as maximum and minimum values for the states ($\pm x_{\max}$) and the control signals ($\pm u_{\max}$). The error vector is defined as $e_t = x_{\text{ref}} - x$. This is the optimization problem when running the MPC without the RF adjustment. When running the MPC with the RF adjustment, i.e., with the correction of the beam energy via radio frequency, the state space model from 4.3 is used instead.

There is a collection of parameters that need to be chosen to achieve the desired performance of the MPC, the first of which is the cost matrix for the states, Q_1. The positioning of the beam is more critical in certain places than others, particularly near the insertion devices. This is accounted for by giving the states corresponding to different sensors penalties that reflect their importance so that an error in the

more critical places is more costly than an error in a place where the beam position is of less importance. In the previous version of the SOFB, there exists a vector of weights which was used for a different purpose but contains the necessary information for sensor penalties. That is, more important sensors have a higher weight than less important sensors. This weights vector therefore serves as a starting point for the sensor penalties. In the small ring, the weights for the less important sensors are 1, while the weights for important sensors are 1000. In the large ring, the less important sensors are weighted 1, while the important sensors are weighted 100.

When the controller runs in with the RF-adjustment, a penalty for the energy shift, $\Delta E/E$ is required as the last diagonal element in Q_1. Since the energy shift is somewhere around 10^9 times smaller than the sensor values, using the current set of units, the penalty needs to be chosen very large for it to contribute to the cost function. The initial guess is to put the penalty to 10^{18} to get the contribution to the cost function from the energy shift to be in the same order of magnitude as the contribution from the other states.

No actuator magnets are more or less costly to use, and therefore the cost matrix for the control signal, Q_2, will have the same value on all diagonal elements. The relative costs of the control signals and states, rather than their absolute values, are important. Therefore, the costs for the control signals are initially set to one.

Additionally, the last element in Q_2 when running with the RF-adjustment determines the penalty on the radio frequency. Since it is more important to keep $\Delta E/E$ as close to zero than it is to keep the radio frequency constant, therefore the element in Q_2 corresponding to the RF is set to zero.

The actuator limits, $u_{\max}$, are determined by the maximum output of the actuator magnets. For the 3 GeV ring this is ± 5 A in both the horizontal and vertical plane. For the 1.5 GeV ring this is ± 11.5 A in the horizontal plane and ± 10.5 A in the vertical plane. If the beam deviates more than 500000 nanometers (or 0.5 mm) from the reference orbit when the machine is run at full energy, the beam is dumped to protect the machine. Therefore, the state limits, $x_{\max}$ are set to 500000 nanometers. For the beam to be considered stable, the beam needs to be kept within 10 % of the beam size, which in the vertical plane (where the beam size is the smallest) corresponds to 200 to 300 nm.

The state space model only includes dynamics from the controller and none from the process. This means that theoretically, the desired change in states can be achieved in one iteration, and there is no need to predict more than one step ahead. Thus the prediction horizon, H_p, and the control horizon, H_u, are both set to 1. If the measurements are noisy, it might be an idea to instead use a slower controller that corrects the beam in several steps instead.

4.3 FOFB offloading through mid ranging

The Fast Orbit Feedback (FOFB) actuators have a smaller working range compared to the SOFB actuators. For instance, in the horizontal plane of the 3 GeV ring the SOFB actuators can accommodate a deflection of 380 μrad (corresponding to the maximum current of $\pm$ 5 A), while the FOFB actuators only can handle a deflection of 10 μrad (corresponding to the maximum current of $\pm$ 2 A). This discrepancy puts the FOFB actuators at risk of saturating. This can be solved by having the SOFB periodically offload the FOFB through a method called mid ranging.

Mid ranging is a control structure which is useful for processes with two inputs and one measurement. It is particularly effective when one of the inputs has high precision but a small working range, while the other has low precision but a larger working range. The goal is to avoid saturation of the input with the smaller working range by having it operate in the middle of its working range. This can be achieved by letting the input with the larger working range offload the input with the smaller working range [Soltesz, 2021]. There are several ways to implement a mid ranging controller and for the offloading of the FOFB, two different structures were suggested, see Figure 4.1 and Figure 4.2.

The high precision, small working range input in the Orbit Feedback System is the Fast Orbit Feedback, and the low precision, large working range input is the Slow Orbit Feedback. The way the currently used Orbit Feedback System is implemented, there is a clear definition of what roles the SOFB and the FOFB should play. The SOFB handles reference tracking while the FOFB handles noise attenuation. The way that this is implemented in the implementation currently in production is that the SOFB makes a prediction of the states after its next control move and sends this prediction, $x_{\text{predicted}}$, to the FOFB. The FOFB then uses this prediction as its reference. Therefore it only handles noise and corrects to the orbit determined by the SOFB. The FOFB records its control signals over time and calculates the average control signal over the last five seconds. The FOFB then converts these average control signals into the average sensor reading error these control signals correspond to using the response matrix of the FOFB. Every five seconds the SOFB reads out these average errors, $x_{\text{average errors}}$. The implementation currently used in production, these average errors were converted to actuator signals using the SOFB's response matrix and added to the actuator signal before sending the signal to the corrector magnets. This further contributed to the saturation problems.

In the first suggestion for the MPC implementation of the SOFB, the average error is subtracted from the actual sensor readings to force the SOFB to add more control

signal and therefore offload the FOFB. This structure is presented in Figure 4.1. This solution is suggested since it most closely resembles the current implementation of the FOFB offloading while still respecting the constraints of the slow correctors. However, this solution somewhat blurs the distinction between the tasks assigned to each of the two controllers, as the FOFB now influences reference tracking. This is the case for the current implementation that is in production as well.

The second suggestion is to view $x_{\text{average errors}}$ as an indication of how much $x_{\text{predicted}}$ differs from the actual sensor readings. $x_{\text{average errors}}$ is added to the $x_{\text{predicted}}$ to compensate for the difference between prediction and reality, causing the FOFB to correct to a reference closer to the real position of the beam, thereby working in the middle of its range. This structure is presented in Figure 4.2.

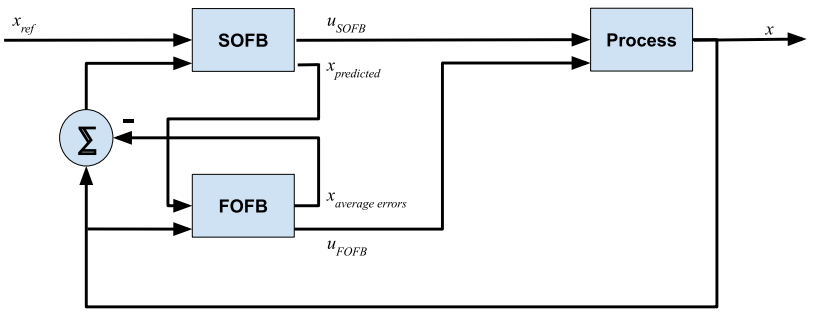

Figure 4.1 Schematic overview of the first suggestion of how the SOFB and the FOFB are connected to the process. The two controllers each have their own feedback loop meaning that the SOFB can operate when the FOFB is turned off and vice versa. The SOFB sends $x_{\text{predicted}}$, which is a prediction of the sensor readings with the planned control signal, to the FOFB. The FOFB returns $x_{\text{average errors}}$ to the SOFB where it is subtracted from the sensor readings, x. The SOFB block contains the MPC as well as the calculation of the predicted orbit.

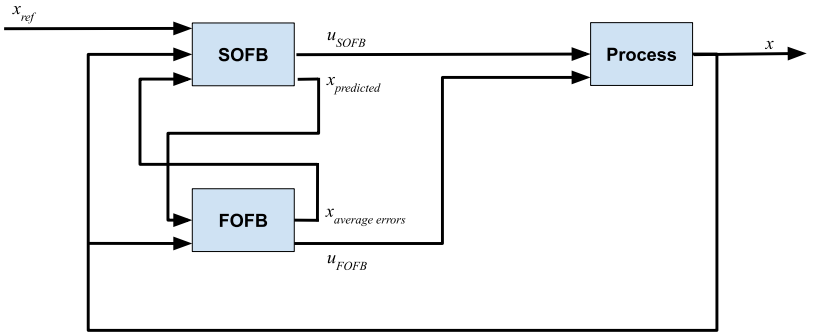

Figure 4.2 Schematic overview of the second suggestion of how the SOFB and the FOFB are connected to the process. The two controllers each have their own feedback loop meaning that the SOFB can operate when the FOFB is turned off and vice versa. The SOFB sends $x_{\text{predicted}}$, which is a prediction of the sensor readings with the planned control signal, to the FOFB. The FOFB returns $x_{\text{average errors}}$ to the SOFB. The SOFB block contains the MPC as well as the calculation of the predicted orbit.

5

Implementation

Initially, a simple MPC using the Python toolbox do-mpc, further described in Section 5.1, was implemented and tested on a basic model from the lecture notes of the Advanced course in control systems at Lunds Tekniska Högskola (LTH). Later, this was extended to an MPC for the SOFB. To facilitate communication with sensor devices and actuator devices, code was ported from the old SOFB implementation and tested on a test plant that mimicked the sensors and actuators. After successfully testing this, the MPC was tested on the small (1.5 GeV) ring. The code for the FOFB and RF was brought in from the old SOFB with necessary modifications to make it compatible with the MPC.

5.1 Implementation of MPC in Python using do-mpc

To implement the MPC in Python, the toolbox do-mpc was chosen for its simplicity and for containing all the necessary tools required [do-mpc, 2021]. It is an open-source toolbox for Model Predictive Control that has tools to set up and formulate control problems. It contains intuitive functions for setting up the state-space model of the problem, formulating the cost function and setting constraints and a prediction horizon. In the function make_step, the prediction and optimization is performed and the optimal control signal is returned. The toolbox allows for time-varying parameters in to incorporate disturbances or update setpoints while running. The MPC can be easily tested using do-mpc since the toolbox permits simulation of the system with real-time plots displaying the control signal, the states and their predicted path, and the cost function. The toolbox allows the user to save the control signals and states to a file for later analysis. The license for do-mpc is GNU LESSER GENERAL PUBLIC LICENSE which allows for use in software. The optimizer used by do-mpc to find the solution which minimizes the cost function is IPOPT (Interior Point Optimizer), a software for nonlinear optimization. IPOPT is an open-source license available for commercial purposes [Waechter and Laird, n.d.] This allows us to use do-mpc and IPOPT at MAX IV.

How do-mpc was used for the MPC

To structure the creation and configuration of the MPC, three functions were created: **setup_mpc**, **setup_model** and **setup_simulator**. The function **setup_model** sets up the state-space model, creates the variables for the states and control signals and sets the expression for the cost function. A time-varying parameter was created for the sensor references to enable updating the reference values. Another time-varying parameter was also created to set the current control signal u_k to be the old control signal u_{k-1} in the next iteration. The function returns a model-object containing this information. In the model-object the time-varying parameters are updated at each iteration. The function **setup_mpc** sets the parameters for the prediction horizon, the step size and sets the constraints for the sensor signal and the control signal. Finally, the function returns an mpc-object containing all this information. The last function, **setup_simulator**, receives the model given by **setup_model** as a parameter. This function is used exclusively for the simulations. The built-in function **make_step** is then used by the main program to compute the control signals given the current states and current control signal.

5.2 Communication using Tango control systems

The communication between hardware and software is facilitated by Tango Controls. Tango is a toolkit for controlling hardware and software using devices and building SCADA systems. It was originally developed at the European Synchrotron Radiation Facility (ESRF) in France [Tango Controls, 2015]. For this project, PyTango, a Python module that exposes the Tango C++ API to Python, was used [PyTango, n.d.] Each device can communicate with each other by sending and receiving data. Devices have a data type called attributes which can be read and written to by other devices [The TANGO Team, 2016]. Attributes can be several different data types, such as boolean, float, and string. Devices can subscribe to events in other devices such as a change in value for a certain attribute, and specify what should happen when such events occur. A client can configure and control a device using commands, which are functions in the device executing a sequence of actions. Devices have state information, which enables them to take different actions depending on their own state or other device states. The state machine is implemented as a function called dev_state which is then exposed as commands called State and Status in Jive (see next paragraph) [The TANGO Team, 2016]. Pressing State returns the current state of the device, while pressing Status gives an information message about the state. Several devices can be grouped together in a Tango Group, which is a collection of devices that share an attribute with the same name. This grouping can be used to write to several devices with the same attribute at the same time.

Jive is an application implemented in Java using Swing to browse and edit the

Tango database. In Jive, users can access the commands to a Tango device and read and write attributes [Tango Controls, 2020]. Jive is used by the operator to configure and run the MPC Tango device.

How Tango was used for the MPC

To configure and setup the MPC the operator has commands to load the sensors, load the actuators, the response matrix, load sensor weights, load sensor references and other parameters. When loading the actuator devices, the constraints are also read. There are functions that check that the MPC can read each device's parameters and raise an error if a device cannot be read. The operator can read the current value of each attribute in the MPC, such as the current prediction horizon, current actuator settings and so forth.

Each Beam Position Monitor (BPM) and corrector magnet is a device that communicates with the device running the MPC. Every time a BPM updates its position, an event is raised, which the MPC responds to by calling the function **on_sync_event**. This function stores each BPM's sensor reading, time stamp and value quality in a buffer. When all BPM's have sent their readings for that time iteration, the function on_sync_event can continue. To ensure that each sensor device has sent its value the time stamp of the oldest event and newest event are compared. If all events are from the same iteration, the time difference should be within a predefined tolerance. The buffer values are then transferred to the local variables that store the sensor quality and sensor values. After this, on_sync_event reads the BPM average errors from the FOFB and the make_step-function described above is called.

When the MPC has computed the new control signal, it gets sent to the actuators through a Tango group since all actuators share the same name for their attributes. Through the function on_stop_event the MPC can be stopped if one of the connected devices enter a state which is forbidden. The MPC then stops applying the correction and the variable external_interlock is set to true. The device that raised the stop event is then added to the list of active stop devices. In Figure 5.1, the communication scheme between the SOFB and the other devices are illustrated.

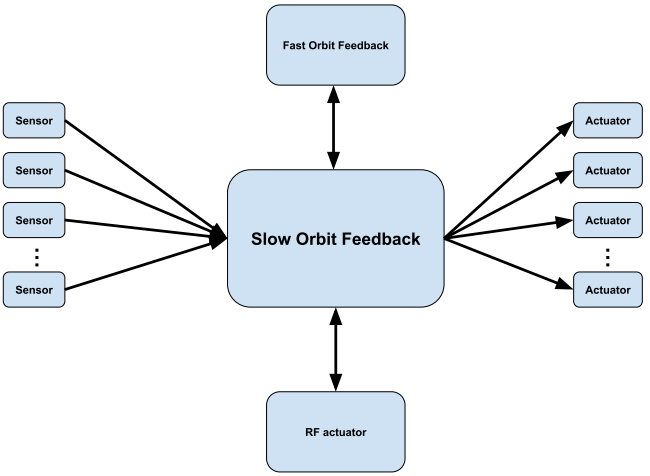

Figure 5.1 The communication scheme between the slow orbit feedback device and other devices.

5.3 Implementation of a test plant for simulations

To test the MPC and communications via Tango, a simple program was provided to us by our supervisor at MAX IV. The program simply used the model from the MPC to receive a control signal from the MPC, simulate one iteration, and via Tango send the new states to the MPC. The MPC subscribed to the output of the test plant via the function subscribe_sensor_devices. To emulate the real process as closely as possible, the simulations included the correct sensor references, sensor weights, and response matrix. However, the simulations did not allow for testing of the FOFB offloading or the RF adjustment since this was determined to be too challenging to simulate. Because of this, the simulations were mainly important for real-time development and testing of the MPC and communication via Tango. Additionally, the program could generate disturbances to test the MPC's ability to handle them by generating a value that was added as a disturbance to the output of the sensor values based on a specified time-dependent function.

5.4 Interaction with the Fast Orbit Feedback

The implementation described here is the first suggestion described in Section 4.3. The Fast Orbit Feedback runs on its own Tango device separate from the SOFB. Its primary role is to attenuate high frequency disturbances. The SOFB interacts with the FOFB in two ways: by offloading the FOFB at a set time interval and by sending the predicted orbit in each iteration to the FOFB. Some of the implementation for the interaction with the FOFB was ported from the old code.

The FOFB is offloaded every 5 seconds by the SOFB. This is done by checking if more than five seconds have passed since the last offloading. If this is true, then it is checked if the actuator settings for the FOFB has changed more than the allowed tolerance since the last offloading. This tolerance is given by the device property fofb_ps_tolerance and has the value 0.0001 A. If the change is larger than the tolerance, then finally the $x_{\text{average errors}}$ are loaded from the FOFB and gets added as a disturbance to the sensor values.

In each time iteration of the SOFB, the SOFB computes the predicted next step of the electron beam and sends to the FOFB. This works the same way as the model, i.e by taking the current sensor values and adding the difference between the new and old actuator value multiplied by the response matrix. The FOFB then uses this as its reference value for the orbit.

5.5 RF adjustment

The device property rf_correction_enabled in the MPC is used to toggle the RF adjustment on or off. When the RF adjustment is on, several attributes are initialized. These are attributes to load in sensor dispersions, actuator dispersions and the RF proxy. The sensor dispersions and the actuator dispersions are described in Section 3.2. The RF frequency proxy also has a function to load the limits for the RF, as computing an RF adjustment that deviates too much from the normal RF can dump the beam. The model used for the MPC is extended with an additional state for the energy shift and an additional control signal for the RF frequency. The response matrix is extended with an additional row and column according to equation 4.3 to account for the RF. Similarly, the sensor references and the penalty matrices are extended to account for the added state and control signal. In each iteration, the MPC tries to optimize the control signal so that the energy shift caused by the corrector magnets is zero while still correcting the beam. Since the energy shift is a state not measured by the BPMs, the MPC prints the computed energy shift in each iteration, allowing the user to track it while the MPC is running. The computed RF frequency is sent to the RF actuator via the proxy in the MPC.

5.6 State handling

The MPC can be in six states: MOVING, ON, STANDBY, INIT, FAULT and ALARM. To determine the current state of the MPC, several boolean flags are used to verify that specific conditions are met before entering a certain state. The default state is INIT when the program is started. When in INIT, the necessary matrices, lists and devices to run the MPC are not loaded. Once all the prerequisites needed to run the MPC are properly configured, the state switches to STANDBY. The operator can then turn on the MPC which sets up the MPC using the previously loaded matrices, lists and devices. When turned on the current actuator settings are also read once and stored in the variable actuator_current_values. This is to ensure that the model has the correct initial value for the actuators. In the ON state the control signal is computed, but not transmitted to the actuators. Finally, in the MOVING state, the control is being computed and sent to the actuators. The MPC can transition to the FAULT state if a device, list or matrix is loaded incorrectly or an error occurs in the process. The MPC can enter the ALARM state if any of the following conditions are met: an error is encountered while reading a sensor value, a sensor quality is invalid, there is a calculation error, an external device interlocks, there is an error in the FOFB, or the incoming event frequency from the sensors is less than 10 Hz. If the MPC is in the FAULT or ALARM state, the control signal is not being sent to the actuators. Figure 5.2 illustrates the different states and connections between states in the MPC. To operate, the MPC must transition from INIT to MOVING via STANDBY and ON.

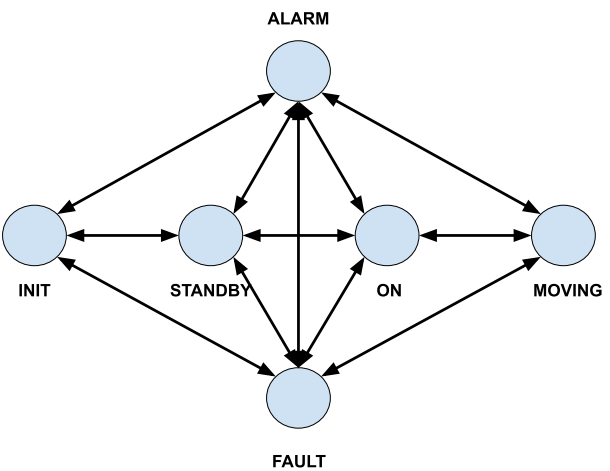

Figure 5.2 The different states and connections between states in the MPC. The MPC has to transition from INIT to MOVING via STANDBY and ON.

6

Results

In this chapter, the results of the simulations and experiments on the real rings will be presented. The experiments on the real rings was conducted during the period of March and April of 2023. Because MAX IV is a facility in operation, the ability to run tests was limited to time slots during maintenance Mondays and during study weeks, where operations on one of the rings were paused for internal research. Due to these constraints, most of the tests were conducted on the 1.5 GeV ring, which was the most readily available ring.

6.1 Simulations

The simulations of the MPC were conducted on both the horizontal and vertical planes on the 1.5 GeV ring and the 3 GeV ring. To simulate the real process as much as possible, we used the real constraints on the actuators together with the sensor references and sensor weights. For the simulations, the MPC and the simulation plant ran at 1 Hz. First the MPC was set to ON where the correction was calculated, but not applied. After approximately 20 seconds the mode was set to MOVING and the MPC started sending the correction to the test plant. Approximately 20 seconds later, a load disturbance of 10000 nm was added to the sensor readings to see how well the MPC could handle a large disturbance. This disturbance scenario is not very realistic. The magnitude of the disturbance is large compared to any disturbances we might expect on the real rings. It is also unlikely that a disturbance would affect all of the sensors the same. This disturbance scenario was chosen to get some interesting results rather than to simulate a physically plausible scenario. No sensor noise or model errors was added to the simulations. All the simulations were run without FOFB offloading or RF adjustment.

1.5 GeV ring

Figure 6.1 and Figure 6.2 show the sensor readings and the control signal for the simulation of the control in the horizontal plane of the 1.5 GeV ring. For this sim-

ulation, the controller was set to MOVING after 14 seconds and the disturbance added after 34 seconds.

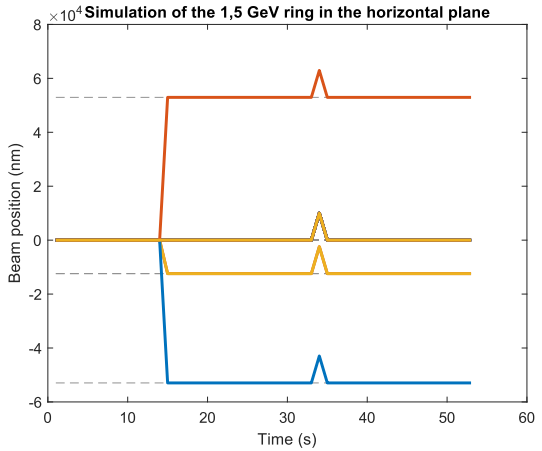

Figure 6.1 The colored lines represent the sensor values for the simulation of the 1.5 GeV ring the horizontal plane. The dashed gray lines represent the references for the sensor values. The controller manages to track the reference adequately. Since most of the sensors have the reference 0 many of the lines overlap. The controller was set to MOVING after 14 seconds and a load disturbance of 10000 nm added after 34 seconds.

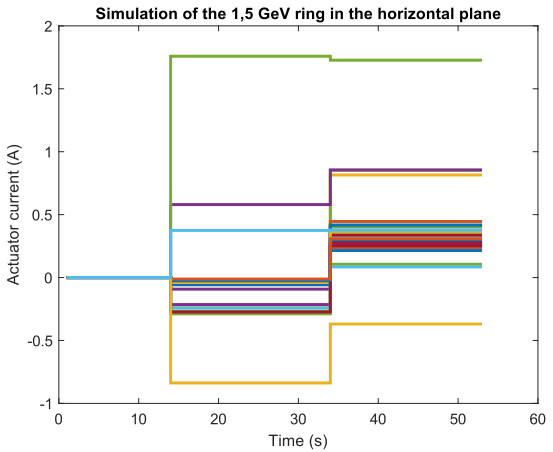

Figure 6.2 The colored lines represent the actuator values for the simulation of the 1.5 GeV ring in the horizontal plane. All actuators are working well within their working ranges (-11.5 A to 11.5 A). The controller was set to MOVING after 14 seconds and a load disturbance of 10000 nm added after 34 seconds.

Figure 6.3 and Figure 6.4 show the sensor readings and the control signal for the simulation of the control in the vertical plane of the 1.5 GeV ring. For this simulation the controller was set to MOVING after 14 seconds and the disturbance added after 34 seconds. In the vertical plane one actuator is saturating and the control signals are larger than for the horizontal plane. The MPC also fails to correct one of the sensors to its correct value.

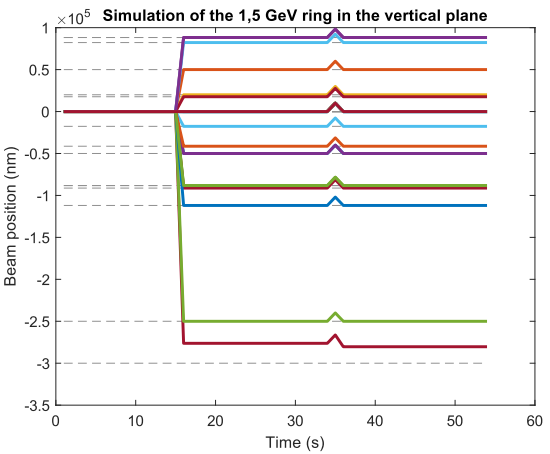

Figure 6.3 The colored lines represent the sensor values for the simulation of the 1.5 GeV ring the vertical plane. The dashed gray lines represent the references for the sensor values. The controller manages to track the reference adequately for all sensors except one. The controller was set to MOVING after 14 seconds and a load disturbance of 10000 nm added after 34 seconds.

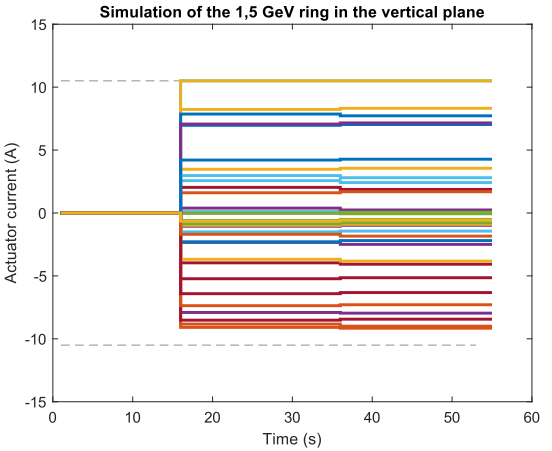

Figure 6.4 The colored lines represent the actuator values for the simulation of the 1.5 GeV ring in the vertical plane. The dashed gray line represent the limits of the working ranges of the actuators (-10.5 A to 10.5 A). The actuators are using most of their working ranges and one actuator is saturating. The controller was set to MOVING after 14 seconds and a load disturbance of 10000 nm added after 34 seconds.

3 GeV ring

In Figure 6.5 the sensor values for the simulation of the 3 GeV ring in the horizontal plane are shown. In Figure 6.6 the actuator values are shown. At 18 seconds the MPC is set to MOVING and sends the control signal to the plant. The sensor values then move to their reference values which are denoted by the dashed gray lines. At 37 seconds the load disturbance of 10000 nm is added to all sensor values. The MPC responds by computing a new control signal and the sensor values return to their reference values. The constraints on the actuators are ± 5 A.

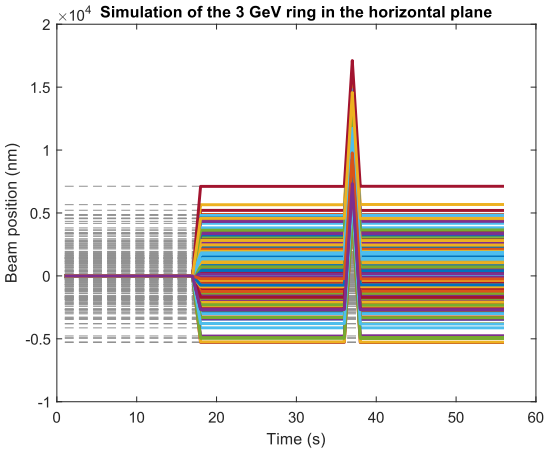

Figure 6.5 The colored lines represent the sensor values for the simulation of the 3 GeV ring the horizontal plane. The dashed gray lines represent the references for the sensor values. The controller was set to MOVING after 18 seconds and a load disturbance of 10000 nm added after 37 seconds. The controller manages to track the reference adequately.

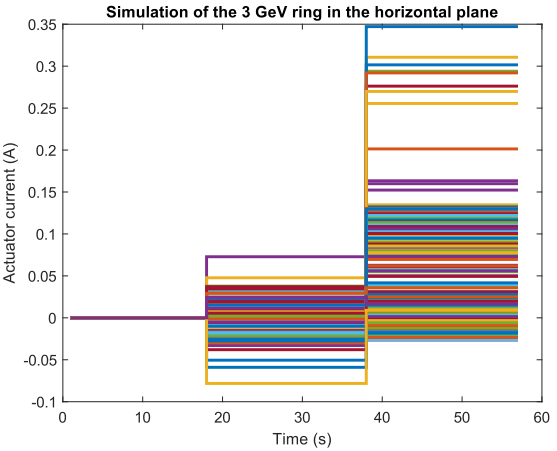

Figure 6.6 The colored lines represent the actuator values for the simulation of the 3 GeV ring in the horizontal plane. All actuators are working well within their working ranges (-5 A to 5 A). The controller was set to MOVING after 18 seconds and a load disturbance of 10000 nm added after 37 seconds.

In Figure 6.7 the sensor values for the simulation of the 3 GeV ring in the vertical plane are displayed. In Figure 6.8 the actuator values are shown. At 20 seconds the

MPC is set to MOVING and sends the control signal to the plant. A few of the sensor values have difficulty placing themselves at their correct reference values. This is especially noticeable for the references that are very large and distant from the rest of the references. This might be connected to the fact that some of the actuators are saturating and can not apply the magnitude required to move the beam to their correct reference. After applying the load disturbance, the MPC quickly handles it and returns the sensors to their previous values, but is still unable to place all sensors at the correct reference. It is worth noticing that one of the actuators stop saturating when the load disturbance is applied. One possible explanation for this result might be that the references are not realistic.

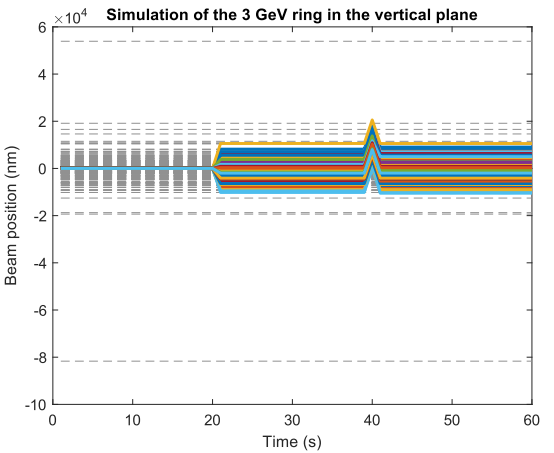

Figure 6.7 The colored lines represent the sensor values for the simulation of the 3 GeV ring the vertical plane. The dashed gray lines represent the references for the sensor values. The controller does not manage to track many of the references. The controller was set to MOVING after 20 seconds and a load disturbance of 10000 nm added after 40 seconds.

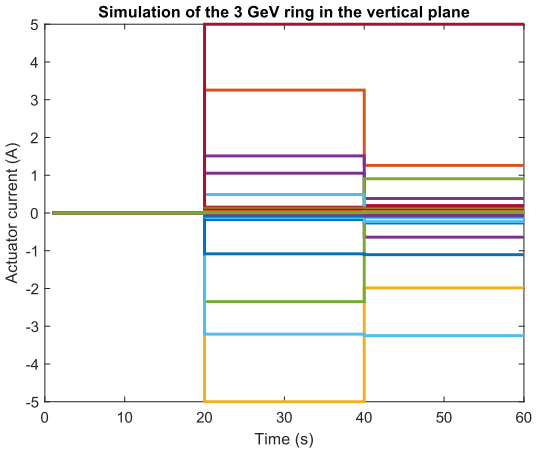

Figure 6.8 The colored lines represent the actuator values for the simulation of the 3 GeV ring in the vertical plane. Most actuators are working within their working ranges (-5 A to 5 A) but a few are saturating. The controller was set to MOVING after 20 seconds and a load disturbance of 10000 nm added after 40 seconds.

6.2 Results of running the MPC on the real rings

The experiments on the real rings were conducted by turning off the currently operating controller and then waiting a short period before turning on the MPC. The MPC ran at a 1 Hz rate. The choice of sampling rate is further discussed in Section 7.3. For the 3 GeV ring, BPM measurements were collected from the data archive at MAX IV that stores the last two weeks of measurements. The actuator measurements were saved from the Taurus window showing the plots. Taurus is the program used to plot real-time data at MAX IV. For the 1.5 GeV ring, BPM measurements were saved from storing the sensor data that was received by the MPC every second. This approach reduced the workload of data collection and eliminated the need to depend on our supervisor to obtain the sensor data from the archives. However, the drawback was that the MPC only received sensor data once per second compared to the archiver, which stored 10000 samples per second. This was deemed acceptable because the control signal was only updated once per second. For the 1.5 GeV ring the actuator values were saved using a script provided to us by our supervisor. This was because the data from Taurus could only be polled every third second which missed several updates to the actuator values. All the tests were run on a low current beam of 3 mA to avoid damaging the rings if something went wrong. When the facility is working at full capacity, the current is 300 mA in the 3 GeV ring and 50 mA in the 1.5 GeV ring. However, using a low current beam increased the signal-to-noise ratio of the BPMs.

1.5 GeV

In Figure 6.9 the sensor values for the 1.5 GeV ring in the horizontal plane are shown. In Figure 6.10 the actuator values are shown. At 16 seconds the MPC is set to MOVING and sends the control signal to the plant. The sensor values all return to their reference values. Before the MPC is started the sensor values can be seen drifting, but the MPC is able to quickly correct the beam. The actuator signals are very stable and only one of the actuators are saturating.

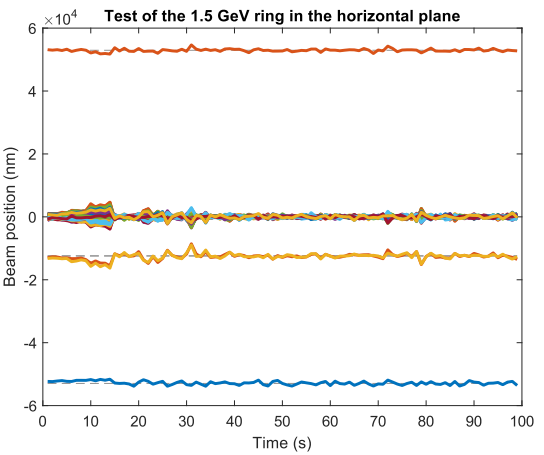

Figure 6.9 The colored lines represent the sensor values for the test of the 1.5 GeV ring in the horizontal plane. The dashed gray lines represent the references for the sensor values. The controller manages to track all of the references.

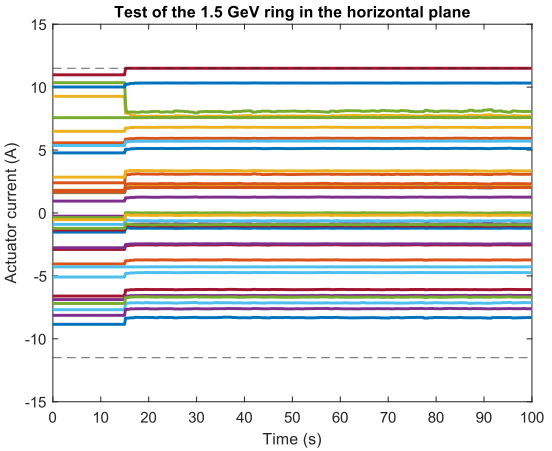

Figure 6.10 The colored lines represent the actuator values for the test of the 1.5 GeV ring in the horizontal plane. The actuators are working within their range (-11.5 A to 11.5 A) except one actuator which is saturating.

In Figure 6.11 the sensor values for the 1.5 GeV ring in the vertical plane are shown. In Figure 6.12 the actuator values are shown. At 18 seconds the MPC is set to MOVING and sends the control signal to the plant. Since the beam was already in the correct position when the MPC was turned on, the effect of the MPC is not as pronounced as in the other tests. The MPC is still shown to be able to control the beam since when it is turned on, the beam is still following its reference and none of the actuators are saturating. The MPC introduces some noise in the beam and the actuator signals are more noisy than in the horizontal plane. The reason for this is unclear.

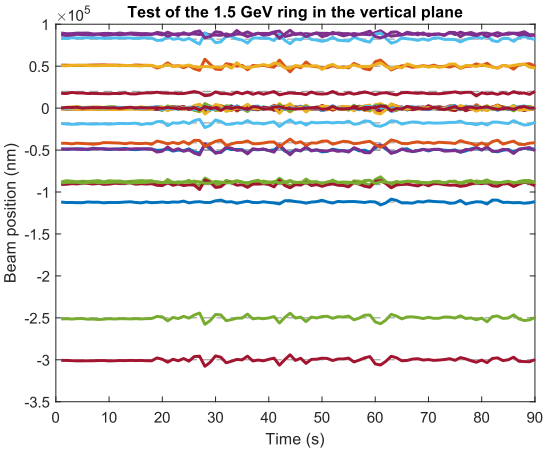

Figure 6.11 The colored lines represent the sensor values for the test of the 1.5 GeV ring in the vertical plane. The dashed gray lines represent the references for the sensor values. The controller manages to track all of the references.

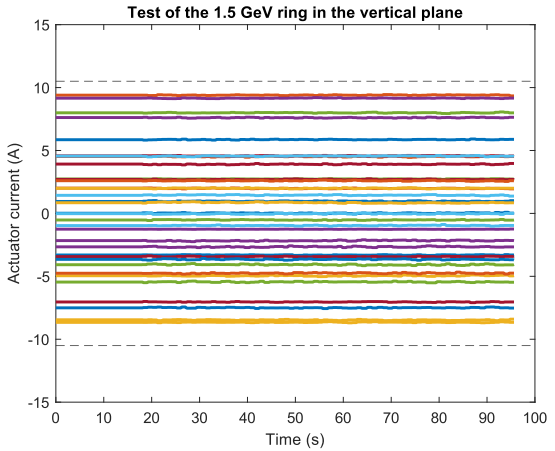

Figure 6.12 The colored lines represent the actuator values for the test of the 1.5 GeV ring in the vertical plane. None of the actuators are saturating (-10.5 A to 10.5 A).

3 GeV

In Figure 6.13 the sensor values for the 3 GeV ring in the horizontal plane are shown. In Figure 6.14 the actuator values are shown. At 16 seconds the MPC is set to MOVING and sends the control signal to the plant. The change of actuator signal is the difference between the old controllers computed optimal control signal and

the MPC's optimal control signal. The sensor values immediately react where some sensor values make a large jump. It is difficult to see the reference values because of the amount of sensors. It was afterwards discovered that the references used for the plot was not the same as the one used during the tests. The experience during the tests was however that the controller followed the references. The constraints on the actuators are ± 5 A. A few of the actuator values are close to saturating.

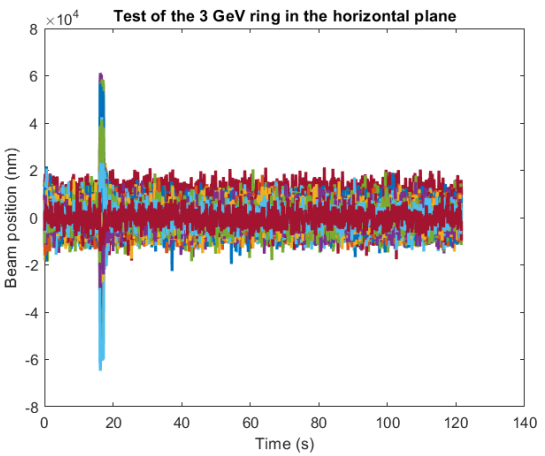

Figure 6.13 The colored lines represent the sensor values for the test of the 3 GeV ring in the horizontal plane.

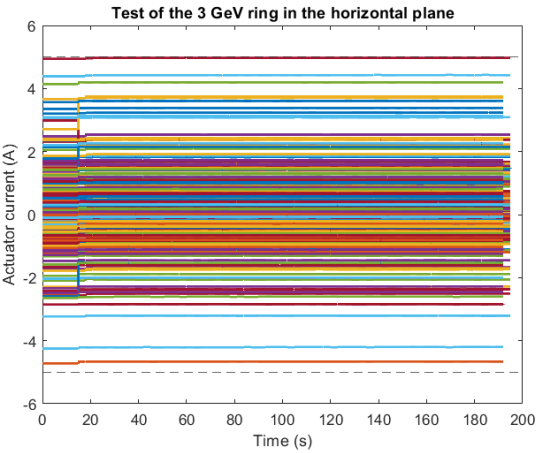

Figure 6.14 The colored lines represent the actuator values for the test of the 3 GeV ring in the horizontal plane. All actuators except for one are working within their working ranges (-5 A to 5 A).

In Figure 6.15 the sensor values for the 3 GeV ring in the vertical plane are shown. In Figure 6.16 the actuator values are shown. At 71 seconds the MPC is set to MOVING and sends the control signal to the plant. The change of actuator signal is the difference between the old controllers computed optimal control signal and the MPC's optimal control signal. The sensor values immediately react, but do not make any large jumps. It is difficult to see the dashed gray lines representing the reference values because of the amount of sensor values. According to the figure, the controller does not manage to track all of the references. This was afterwards discovered to be because we were provided with the wrong references when making the plots. The experience when running the tests was however that the references were being followed. The references used during test were unfortunately not saved and are changed quite frequently at MAX IV. The constraints on the actuators are ± 5 A. None of the actuators are saturating. Most of the actuators do not change their signal significantly when the MPC is turned on. There is one actuator that becomes jittery. The reason for this is unclear.

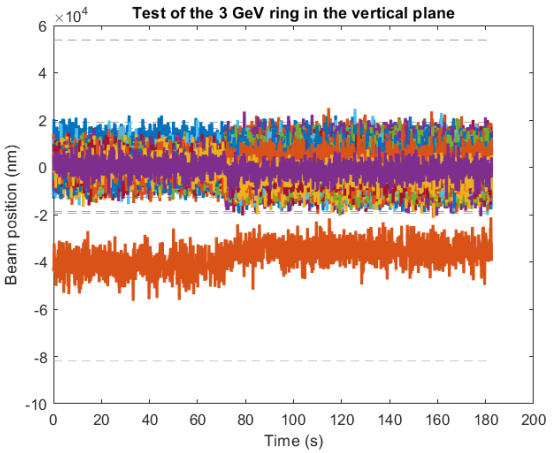

Figure 6.15 The colored lines represent the sensor values for the test of the 3 GeV ring in the vertical plane. The dashed grey lines represent the references for the sensor values. It turned out that these references were not the same as had been used during the test.

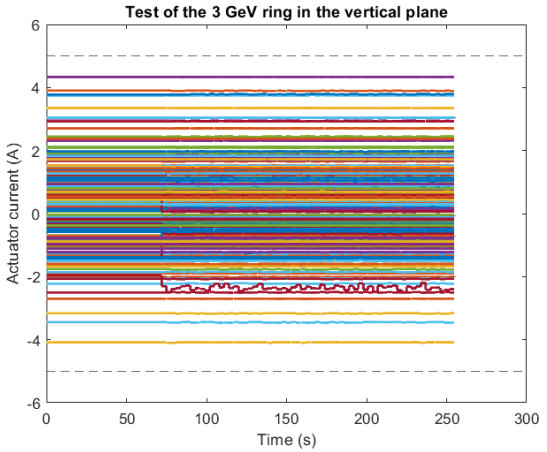

Figure 6.16 The colored lines represent the actuator values for the test of the 3 GeV ring in the vertical plane. All actuators are working within their working ranges (-5 A to 5 A).

6.3 RF adjustment

The simulator did not include any features to simulate the RF, therefore the RF adjustment could not be tested in simulation. The RF adjustment was tested on both

the 3 GeV ring and the 1.5 GeV ring but with no confirmed successful result. A working RF adjustment was expected to correct the beam to a stable solution with no dispersion patterns when the controller was turned on with the RF adjustment activated. The radio frequency would be corrected to a frequency close to the design frequency of 99 931 000 Hz (the deviation being in the magnitude of tens of Hertz).

Although a stable solution was found by the MPC during early tests on the big ring, the result was deemed invalid due to several theoretical inaccuracies. During several tests, the MPC computed a frequency that deviated largely from the set RF frequency and the MPC had to be turned off. Several solutions were attempted with different penalties on the energy shift but this did not affect the results. A penalty on the rate of change on the RF frequency was introduced, which had an effect on the rate of change but ultimately did not result in a successful control of the energy shift. It was also attempted to change the model so that the average energy shift caused only by the actuator magnets was set to zero i.e excluding equation 3.8 from the model. This did however not work at all. After many different attempts a small but significant mistake was found in that equation 3.9 incorrectly included a f in the denominator. This f had made its way into the derivation of the equations but had through a typo not been included in the model during the first tests. After the mistake was corrected, a stable solution was once again found by the MPC during tests on the small ring, suggesting that the solution may have been close all along. Unfortunately, there was no more test time available to record, recreate, or verify this result.

6.4 FOFB offloading

The testing of the FOFB offloading was conducted on both the small ring and large ring by setting to the MPC to MOVING and allowing the MPC offload the FOFB every 5 seconds, sending the next predicted orbit to the FOFB every iteration. However, the SOFB was never able to offload properly, instead the FOFB quickly saturated and the machine was forced to be turned off to avoid dumping the beam. The offloading was tested with both methods suggested in Section 4.3. Since the offloading could not be tested on the simulator, a lot of the debugging had to be done against the Tango devices of the real machine. Unfortunately, all attempts resulted in unstable offloading where the FOFB saturated after varying amounts of time. At best, the SOFB ran with FOFB offloading for about 45 seconds before the FOFB became unstable. If the offloading had worked properly we would have expected to see the FOFB working around zero and BPM average errors to be low. However, the average errors increased over time instead.

7

Discussion

In this chapter, the results presented in the previous chapter will be discussed. The results will be evaluated and compared to the goals set up in the introduction.

7.1 Comparison of simulations and real tests

The simulations, although simple, provided us with a lot of valuable information. First and foremost, they provided an environment for testing and debugging the code. This project, apart from being a control project for a physics application, was in practice a lot about us learning Python and PyTango. Therefore, having an environment to test against was immensely helpful. Secondly, the simulations gave us some indications of what to expect when running on the real rings.

In the horizontal plane of the 1.5 GeV ring, the simulation suggested that good reference tracking and load disturbance handling could be expected. The control signals were well within their working range, using only circa the range -1 A to 1.8 A, see Figure 6.1 and Figure 6.2. The test on the real ring demonstrated good reference tracking (Figure 6.9), while the control signals were spread out over a larger portion of the working range, with one actuator at its limit, as seen in Figure 6.10. The results for the 1.5 GeV ring in the vertical plane were similar, with the expectation of using a larger portion of the working range of the actuators, which is also the result of the run on the real rings. Reference tracking in this plane was also good. For the 1.5 GeV ring in the vertical plane, compare Figure 6.3 and Figure 6.4 with Figure 6.11 and Figure 6.12.

For the 3 GeV ring, the simulations led us to expect good reference tracking in the horizontal plane (see Figure 6.5 and Figure 6.6), while we expected to encounter some trouble in the vertical plane (see Figure 6.7 and Figure 6.8). In the tests on the actual rings, the reference tracking can not be determined from the figures (see Figure 6.13, Figure 6.14, Figure 6.15 and Figure 6.16). This is because the reference values provided to us to make the plots were not the same that were used during the

tests. The results still show that the MPC finds a feasible solution and manages to keep the beam stable.

7.2 Performance of the MPC

The MPC was repeatedly able to find stable solutions which respected the actuator limits while removing the state error. This implies that the robustness of the controller is good, and even when subjected to load disturbances and changing the radio frequency to produce dispersion patterns, the MPC performed well. Compared to the old I-controller, the MPC showed a very good ability to manage disturbances. The response of the MPC was also significantly faster than that of the old I-controller. The I-controller has a longer integration time, which makes it slow which is evident when the beam starts in a bad position, requiring a long time for the controller to correct the beam to a better position. However, this also makes the I-controller less sensitive to noise. The MPC, on the other hand, is fast but also sensitive to noise, as discussed in greater detail later. The results show that MPC is not only a possible method for controlling the beam inside the storage rings of a synchrotron facility but also a suitable one. While there is much work to be done, the results obtained through this thesis are very promising.

7.3 Rate of operation

The question of which rate to run the SOFB at is one that followed us through the whole project. The current SOFB is said to be intended to be run at 10 Hz. However, it is currently only running at 5 Hz. Due to the known flaws in the model presented in Section 3.3, our MPC is not running at 10 Hz either. To ensure that the delay does not disturb the MPC, we opted to run the controller at 1 Hz, providing roughly a 2.5-fold safety margin to the delay. It should be possible to run the controller at 2 Hz, and still be unaffected by the delay. For rates higher than this it cannot be guaranteed that the delay would not pose any problems. A Tango attribute called "throttle" has been implemented to allow for easy adjustment of the rate, should future development require it. It should be noted that running the MPC at 1 Hz has proven effective, and higher rates may not be necessary.

7.4 Noise handling

During testing it was discovered that the sensor readings were more noisy than with the current I-controller. This may be because the current I-controller attenuates its

signal before applying it, thereby not reacting as much to noise in the sensors. It may be of interest to add a noise model to the MPC to try to reduce the noise. This would however require extensive research and development and is outside the scope of this thesis. The effect of the noise was considered minor and is not imperative for keeping the beam stable.

7.5 Prediction horizon and computational cost

It was decided to use a prediction horizon of one because there are no dynamics in the model longer than one time step. All the dynamics in the model come from changing the control signal. Reducing the prediction horizon to one raised doubts about the purpose of using Model Predictive Control when there is no predictive element. However, the MPC's ability to incorporate system constraints into its optimization process was the primary motivation for using it in this project, particularly to address the saturation issues with the SOFB. As such, the use of MPC was intended to compute a control signal that respects the constraints. A more sophisticated model including more of the dynamics of how the beam moves as well as of the delays could be developed and in that case it might be warranted to have a prediction horizon larger than one. However, that would require significant research and development. During simulations it was discovered that having a prediction horizon larger than 1 was computationally heavy, such that the MPC could not compute the new control signal at a 10 Hz rate. This was when running on the KITS lab-machine which is slower than the real machine. It may be worth exploring other optimization algorithms besides IPOPT, which do-mpc currently uses, to reduce computational demands. However, the program seems to be running fine on the server, not exceeding limits neither in terms of memory usage or processing.

7.6 The choice of do-mpc

During planning, there were a number of alternatives for how to implement the MPC, but do-mpc was the final choice. This proved to be a satisfactory choice and the toolbox was instrumental in implementing the MPC. Setting up the model, the constraints, the cost function was smooth and the optimization step was fast and reliable. It worked well together with Tango.

7.7 FOFB offloading

The FOFB offloading is critical for beam delivery during operations and must work for the MPC to be deployed. There are a few theories on why the FOFB offloading failed. One theory is that the FOFB has a stationary error, which has not previously been noticed. The implementation of the SOFB currently used in production is believed to compensate for this stationary error, but since the FOFB has never run together with anything but the implementation of the SOFB currently used in production, this stationary error has never been detected. The first suggested solution for the FOFB offloading was suggested since it most closely resembles the solution used in the implementation of the SOFB currently used in production. However, it did not make sense for us to disturb the MPC in this way for the sake of the FOFB and the second suggestion was developed. The need for offloading is a product of a discrepancy between the predicted orbit that the FOFB uses as its reference and the actual orbit. Therefore, it was suggested to focus the efforts there, at the source of the problem. The second suggestion aims to make the predicted orbit more like reality, while the first suggestion aims to make reality more like the predicted orbit. It is now up to the engineers at MAX IV to decide which suggestion to move forward with.

7.8 RF adjustment

More testing is needed to verify that the RF adjustment works as desired. The fact that the MPC finds a stable solution is a promising result. However, it remains to be proven that the MPC actually manages to keep the energy shift at zero. In the implementation of the MPC, $\Delta E/E$ exists as an internal variable rather than a measurable variable. This means that we might have come up with a solution which brings this internal variable to zero, but not the actual energy shift. The RF adjustment part of the thesis was probably the most theoretically challenging part and against that background, it is hard to say with confidence that we actually solved the right problem. During the first tests of the RF adjustment, we did not know what to expect, which of course made it hard to draw conclusions on the results we were getting. The RF adjustment is not as critical for the system as the FOFB offloading. The MPC implementation of the SOFB could be used in production without the RF adjustment, it is however desired for this part to work as well. For numerical stability, it might have been better to scale the magnitude of the energy shift, rather than the penalty in the penalty matrix.

7.9 Difficulties with complex systems

As stated by our supervisor at MAX IV, we are dealing with complex real-world real-time systems that come with their unique sets of challenges. This implies that there is a likelihood that a factor believed to be outside the scope of our work could be influencing our problem. The challenges encountered in identifying the issues with the FOFB offloading, and to some degree with the RF adjustment, are possibly symptoms of this. In other words, there may be more elements influencing the system than previously thought.

7.10 State machine

It was decided to use a different state machine than the one proposed in the beginning of the project. When investigating, the device required four states to describe the device better. The INIT, STANDBY, ON and MOVING states provided better and more intuitive descriptions than three states would. It was decided to use the MOVING state instead of the RUNNING state to be better aligned with the definition and usage of states by the rest of the Tango devices in the MAX IV library. The ALARM and FAULT states are necessary to ensure the safety of the system.

8

Conclusions

An MPC has been developed for stabilizing and controlling the electron beam in both the large and small ring at MAX IV, as well as RF adjustment to keep the beam energy constant and midranging to offload the FOFB and keep it within its working range. The project was implemented using PyTango, taking advantage of the Tango event system for all communication with the rings. The control of the beams have been successfully demonstrated in the small ring and large ring in both the vertical and horizontal planes. More tuning might be needed to reduce the noise that can be seen in the sensors when the MPC is used on the real rings. This noise did not occur in the simulations, likely due to the simplicity of the simulations. The tests of the controller showed that it handles events correctly, and that the computational cost of using an MPC did not seem to exceed the limits of the server, neither in terms of memory usage nor processing. The RF adjustment showed promise of working but more testing is required to verify that it works, as described in Section 6.3. In contrast, the FOFB offloading was not able to offload correctly in testing and also here more time is needed to for testing. The design of the FOFB offloading might need to be reconsidered to make it work. The simulations showed that the model and constraints were feasible in all cases but for the vertical plane of the 3 GeV ring. The implementation of the state machine was changed during development to include six states instead of three with better state names that reflect their actual behavior. The controller has a well-defined state machine and a well-defined command sequence for configuration and start up. The results of the project can in many ways be viewed as a success.

8.1 Challenges

It has been very special to have gotten the chance to work on such a high technological machine such as the MAX IV synchrotron. It took a while to get used to the thought that two masters students were even allowed to run their code on the synchrotron and not just be able to develop against a simulated environment. The

fact that the machine that our control system is controlling is so advanced, big and valuable makes the experience memorable but also poses a few challenges. One challenge was the limited time to test our code on the machine. We could only test during maintenance windows and since that is the case for all other development at MAX IV, those time slots during maintenance windows were scarce. All in all we got nine opportunities of one to a few hours each to run tests and collect data for the thesis, often with a week or more between opportunities. Therefore, we had to develop large amounts of code before testing or troubleshooting, hindering an agile approach. However, the small simulator we developed in the beginning helped a lot and we were able to do some code validation because of it. When we got to the first test on the machine our code ran without any bugs which was a big relief. After this we decided that it would be too complex and time consuming to try to expand the simulator to be able to simulate the FOFB-offloading and the RF-correction and from there on we worked more in the dark and needed time to debug at each test opportunity.

Another challenge was the lack of documentation for the old SOFB I-controller that was used as a starting point, as it is technically still in development and not adequately documented. The definition between development and production at MAX IV is a bit blurry since it is not just a laboratory in the sense that it delivers light to be used in other research, but just as much a laboratory in all things engineering, IT and accelerators, constantly improving.

8.2 Outlook

Further work is required before our MPC can be used in production. First and foremost, the FOFB-offloading needs to be functional so that the SOFB can be used together with the FOFB. While the MPC seems to be doing a sufficient job and the RF adjustment also seems to be working, both require more validation before they can be considered reliable. The MPC is running with the same parameters that were initially set, so there might be room to improve the performance of the MPC by tuning the control parameters. Our initial plan included that we would have the time and opportunity to try out different sets of parameters but unfortunately we never got the chance to do this. The MPC is noisier than the I-controller and it could be of interest to add a noise model to the MPC to see if the noise level can be reduced. The MPC seems to introduce this noise which, while not fatal for the MPC's performance, is still undesired. Another development that could be useful is to develop a more advanced model of the orbit system which takes into consideration delays and the dynamics of how the beam behaves.

Bibliography

Bell, P., M. Sjöström, M. Lindberg, V. Martos, and V. Hardion (2017). "A general multiple-input multiple-output feedback device in Tango for the MAX IV accelerators". In: *Proceedings of ICALEPCS2019*. DOI: 10.18429/JACoW-ICALEPCS2019-WEPHA012.

CERN (2023). *Accelerating: radiofrequency cavities*. Last accessed 22 February 2023. URL: https://home.cern/science/engineering/accelerating-radiofrequency-cavities.

Cervin, A. (2022). *Automatic Control, Advanced Course: Lecture 12*. Lund University.

Chikazumi, S. (1997). *Physics of ferromagentism*. 2nd ed. Oxford University Press. ISBN: 9780191569852.

Gavaghan, H. (2001). "What is a synchrotron?" *Nature* **410**:722. DOI: 10.1038/35070715.

Lindvall, R. (2022). *Radio frequency*. Last accessed 22 February 2023. URL: https://www.maxiv.lu.se/beamlines-accelerators/accelerators/radio-frequency/.

MAX IV (2022). *How MAX IV works*. Last accessed 7 February 2023. URL: https://www.lunduniversity.lu.se/research-innovation/max-iv-and-ess/how-max-iv-works.

do-mpc (2021). *Model predictive control python toolbox*. Last accessed 20 February 2023. URL: https://www.do-mpc.com/en/latest/%20do_mpc.

PyTango (n.d.). *Welcome to PyTango documentation*. Last accessed 20 February 2023. URL: https://pytango.readthedocs.io/en/stable/.

Sjöström, M., J. Ahlbäck, M. Johansson, S. Leemann, and R. Lindvall (2011). "Orbit feedback system for the MAX IV 3 GeV storage ring". In: *Proceedings of IPAC2011*. URL: https://accelconf.web.cern.ch/IPAC2011/papers/mopo010.pdf.

Soltesz, K. (2021). *Systems Engineering/Process Control: Lecture 10*. Lund University.

Tango Controls (2015). *What is tango controls?* Last accessed 20 February 2023. URL: https://www.tango-controls.org/what-tango-controls/.

Tango Controls (2020). *Jive*. Last accessed 24 april 2023. URL: https://tango-controls.readthedocs.io/en/latest/tools-and-extensions/built-in/jive/index.html.

The TANGO Team (2016). *The tango control system manual, version 9.2*. Last accessed 20 February 2023. URL: http://ftp.esrf.fr/pub/cs/tango/tango_92.pdf.

Waechter, A. and C. Laird (n.d.). *Ipopt documentation*. Last accessed 23 May 2023. URL: https://coin-or.github.io/Ipopt/.

Wenninger, J. (1997). *Orbit Corrector Magnets and Beam Energy*. CERN - SL/Note 97-06 (OP).

Wiedemann, H. (2019). *Particle Accelerator Physics*. 4th ed. Springer International Publishing AG Switzerland. ISBN: 978-3-319-18317-6.

Wolski, A. (2014). *Beam Dynamics In High Energy Particle Accelerators*. Imperial College Press, London. ISBN: 9781783262786.

Lund University
Department of Automatic Control
Box 118
SE-221 00 Lund Sweden

Document name
MASTER'S THESIS

Date of issue
June 2023

Document Number
TFRT-6197

Author(s)
Emory Gassheld
My Karlsson

Supervisor
Carla Takahashi, MAX IV, Sweden
Magnus Sjöström, MAX IV, Sweden
Pontus Giselsson, Dept. of Automatic Control, Lund University, Sweden
Bo Bernhardsson, Dept. of Automatic Control, Lund University, Sweden (examiner)

Title and subtitle
MPC for the Slow Orbit Feedback Control at MAX IV

Abstract

The MAX IV synchrotron radiation facility in Lund is designed to produce bright and high-energy X-ray radiation. One important component of achieving this is to keep the electron beam inside the storage rings centered and stable. The aim of this thesis was to develop an MPC (Model Predictive Control) that can accomplish this task more robustly than the currently commissioned I-controller. A controller was implemented in pyTango, a python module of Tango controls, using the Tango events system. By testing the controller on both a simulator and the real machine the result was a controller that fulfills its goal of keeping the beam centered while respecting the constraints of the system. Two additional features, offloading of another controller and controlling the radio frequency of the synchrotron to mitigate the impact of the MPC's control signals on the beam's energy, were also developed but due to a limited time frame testing was not finished. The results of the thesis show that using an MPC for this task has good potential. The MPC solves the problem at hand, and with more tuning and verification MAX IV should be able to use the MPC in production given that the issue of offloading of the other controller is resolved.

Keywords

Classification system and/or index terms (if any)

Supplementary bibliographical information

ISSN and key title
0280-5316

ISBN

Language
English

Number of pages
1-56

Recipient's notes

Security classification