



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

PAUL SCHERRER INSTITUT



START-TO-END MODELLING OF THE AWA MICROBUNCHED ELECTRON COOLING POP-EXPERIMENT

MASTER THESIS

Department of Physics

ETH Zurich

written by

ARNAU ALBÀ

supervised by

Dr. A. Adelman

Collaborators:

Dr. A. Fallahi

Dr. G. Ha

Dr. J. P. Power

Dr. A. Zholents

June 7, 2020

Abstract

The Electron-Ion Collider being built at Brookhaven National Laboratory requires cooling of hadron beams to achieve and maintain a high luminosity. Hence a new method known as Micro-Bunched Electron Cooling is being considered for reducing energy spread and emittance of hadrons. This technique employs space charge effects in drifts to induce energy modulation in electron bunches starting from density modulation. In this thesis we discuss and perform a numerical study on the possibility of replacing said drifts with much shorter magnetic wigglers, and guide an upcoming experiment at the Argonne National Laboratory (ANL) that will investigate the feasibility of this replacement. To simulate the experiment, an electromagnetic solver is added to the existing 3D accelerator code OPAL, such that it can now simulate undulators and wigglers by means of the external library MITHRA, which uses FDTD and PIC to solve Maxwell's equations. The new code is tested and benchmarked, and subsequently used to analyse the effects of a wiggler on an electron bunch with different characteristics. It is found that the reduced speed of a bunch passing through a wiggler yields strong space-charge effects, but also radiation effects which improve the gain in energy modulation. The radiation effects are found to be highly dependent on the transverse size of the beam, as predicted by the theory. Start-to-end simulations of the ANL experiment are carried out to aid in its preparation, and to choose the optimal parameters to obtain measurable impact from the wiggler.

Contents

List of Figures	iii
------------------------	------------

List of Tables	vi
-----------------------	-----------

1	Introduction	1
1.1	Acknowledgements	2
2	Numerical Approach	2
2.1	OPAL Static Solver	2
2.2	OPAL-MITHRA Full-Wave Solver	4
2.3	Convergence Tests	10
2.3.1	Static Solver	10
2.3.2	Full-Wave Solver	14
2.4	Code Benchmarking	18
3	Experiment and Simulations	20
3.1	Micro-Bunched Electron Cooling	20
3.1.1	Wiggler as a Drift Replacement	23
3.2	Experiment Overview	24
3.3	Simulations	27
3.3.1	Wiggler Simulations	27
3.3.2	Surrogate Model by R. Bellotti	35
3.3.3	Start-to-End Simulations	36

4	Conclusions and Future Work	40
	Bibliography	41
A	Comparison Wiggler Simulations and Drift at Reduced Gamma	44
B	Comparison Wiggler Simulations and Drift of Same Length	51

List of Figures

1	<i>Schematics of a Free-Electron-Laser [14].</i>	4
2	<i>Comparison between the particle trajectories in the laboratory (left) and the co-moving frame (right). Particles are not static in the comoving frame, despite their average speed being zero.</i>	7
3	<i>The Lorentz transformation shrinks the undulator, and elongates the bunch and the resonant wavelength [10].</i>	8
4	<i>Solver distribution for the start-to-end simulation of the AWA beamline. (image from [5]).</i>	9
5	<i>Simulated parts of the beamline for the convergence analysis of the full-wave solver (a), and of the static solver (b).</i>	11
6	<i>Slicing the phase-space plot and obtaining slice-energy and slice-energy-spread values, for the comparison between simulations. For the sake of clarity, only 30 slices are drawn in this image, although the convergence analysis was done with 100 slices.</i>	12
7	<i>Static solver convergence with respect to number of particles N and number of particles-per-cell NpC. The rest of parameters are fixed at $MX = MY = 64$, $MT = 512$.</i>	13
8	<i>Static solver convergence with respect to number of longitudinal grid-points MT. The rest of parameters are fixed at $NpC = 8$, $MX = MY = 64$.</i>	13
9	<i>Static solver convergence with respect to number of transverse grid-points $MX = MY$. The rest of parameters are fixed at $NpC = 8$, $MT = 512$.</i>	14
10	<i>Full-wave solver convergence with respect to number of particles N and number of particles-per-cell NpC. The rest of parameters are fixed at $\Delta = (25, 25, 4) \mu\text{m}$, $L = (9, 9, 4) \text{mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.</i>	15
11	<i>Full-wave solver convergence with respect to longitudinal grid-spacing Δz. The rest of parameters are fixed at $\Delta x = \Delta y = 17 \mu\text{m}$, $L = (9, 9, 4) \text{mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.</i>	15

12	<i>Full-wave solver convergence with respect to transverse grid-spacing $\Delta x = \Delta y$. The rest of parameters are fixed at $\Delta z = 2 \mu\text{m}$, $L = (9, 9, 4) \text{ mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.</i>	16
13	<i>Full-wave solver convergence with respect to length of computational domain L_z. The rest of parameters are fixed at $\Delta = (25, 25, 4) \mu\text{m}$, $L_x = L_y = 9 \text{ mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.</i>	16
14	<i>Full-wave solver convergence with respect to transverse size of computational domain $L_x = L_y$. The rest of parameters are fixed at $\Delta = (25, 25, 4) \mu\text{m}$, $L_z = 4 \text{ mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.</i>	17
15	<i>Initial current in the simulation of the LCLS experiment described in [21].</i>	18
16	<i>Longitudinal phase-space of the electron bunch (a) 80 cm before and (b) 10 cm after the wiggler in the LCLS experiment used for benchmarking.</i>	19
17	<i>Comparison of the final slice-energy between MacArthur's paper [21] and an OPAL simulation of the same wiggler and bunch.</i>	20
18	<i>Basic idea behind the Coherent electron Cooling (CeC) scheme [2].</i>	21
19	<i>Schematics of the EIC infrastructure, including the electron beamline for MBEC [26].</i>	22
20	<i>Different improvements on the basic CeC scheme. (a) Micro-Bunched Electron Cooling with one amplification section, (b) MBEC with two amplification sections, (c) MBEC with one amplification section, with a wiggler replacing the drift [27].</i>	23
21	<i>Example of longitudinal distribution used in the AWA experiments.</i>	25
22	<i>Detailed plan of the experimental setup used at AWA, provided by G. Ha.</i>	26
23	<i>Longitudinal phase-space and density distribution after the wiggler for simulations at 44 MeV. The bunch head is on the right side of the plots.</i>	30
24	<i>Longitudinal phase-space and density distribution after the wiggler for simulations at 63 MeV. The bunch head is on the right side of the plots.</i>	31
25	<i>Ratio between the energy modulation at the centre of the bunch induced in a wiggler and in a drift of the same length, at 44 MeV (left) and at 63 MeV (right). The individual phase-space comparisons can be found in the appendix.</i>	32
26	<i>Two bunches with $\lambda/c = 0.4 \text{ ps}$ and $\sigma_r = 1.8 \text{ mm}$ simulated through a 1.1 m drift, to be compared with the results from the wiggler. Case (a) corresponds to $E = 44 \text{ MeV}$, and case (b) to $E = 63 \text{ MeV}$.</i>	32
27	<i>Comparison between a wiggler simulation and a 1.1 m drift at reduced energy to imitate the wiggler effects. The comparison for the rest of simulations can be found in the appendix.</i>	34

28	<i>Comparison of energy modulation at the centre of the bunch between simulations through a wiggler and through a low-energy drift, at 44 MeV (left), and at 63 MeV (right).</i>	34
29	<i>Comparison of the energy variation over the whole bunch between simulations through a wiggler and through a low-energy drift, at 44 MeV (left), and at 63 MeV (right).</i>	35
30	<i>Dashboard running the surrogate model that can quickly find the required accelerator settings.</i>	36
31	<i>The bunch longitudinal phase-space and density distribution, as measured on two screens at the AWA beamline: (a) right after the gun $s = .11$ m and (b) at $s = 5.22$ m in between linacs 2 and 3. This was an initial test with 8 pulses separated by 2 ps and a high charge of 40 nC. Both parameters were subsequently reduced for the final simulations.</i>	37
32	<i>Beam statistics along the beamline, with parameters chosen such that the energy spread remains low.</i>	38
33	<i>Longitudinal phase-space and density distribution on a YAG monitor before (plots at the top) and after (plots at the bottom) the wiggler. (a) refers to the simulation at 66 MeV, (b) to the simulation at 44 MeV case 1, and (c) to the simulation at 44 MeV case 2,</i>	39

List of Tables

1	Parameters used in the simulations for the convergence analysis of the static solver.	11
2	Parameters used in the simulations for the convergence analysis of the full-wave solver.	14
3	Summary of the approximate values to be used in the simulation of the AWA beamline.	17
4	Values used in the OPAL simulation of the LCLS experiment used for benchmarking.	20
5	Tunable parameters during the experiment at the AWA facility.	25
6	Parameters used for the OPAL full-wave solver in this section.	27
7	Conditions satisfied for each set of parameters in the simulations of the wiggler at 44 MeV. The case coloured in blue is the only one that, à priori, satisfies the constraints from Geloni's paper [4] such that the wiggler effects are the same as in a low-energy drift.	27
8	Conditions satisfied for each set of parameters in the simulations of the wiggler at 63 MeV. The case coloured in blue is the only one that, à priori, satisfies the constraints from Geloni's paper [4] such that the wiggler effects are the same as in a low-energy drift.	28
9	Parameters used for OPAL's solvers for the start-to-end simulations.	38
10	Beamline settings used for the 4 and 6 cavity start-to-end simulations from figures 32 and 33.	39

1 Introduction

The Electron-Ion Collider is a unique new machine currently being designed at Brookhaven National Laboratory (BNL), reusing the infrastructure already in place for the Relativistic Heavy Ion Collider (RHIC) [1]. Its main goal is to give further insight into the physics of gluons and quarks, and how they give rise to the mass and spin of protons and neutrons, the building blocks of the universe.

One of the requirements for this accelerator is a high luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for a wide range of centre-of-mass energies, from ~ 20 to 100 GeV . Such high luminosities require very small phase-space volumes, which are hard to obtain and to preserve, since intra-beam scattering and small defects in the infrastructure degrade the beam over time. This is especially true for the hadron or ion beam since, unlike lighter particles such as leptons and electrons, they have low radiation damping, and hence require some type of external cooling. To this effect, a yet to be tested method has been proposed for the upcoming EIC, known as Micro-Bunched Electron Cooling (MBEC) [2], which consists in copropagating an electron beam with the hadrons in a short section of the ring, and then having them exchange energy through a series of chicanes and drifts, effectively cooling the hadron beam.

More recently, based on a series of publications [3,4] that analyse the effect of wigglers, undulators and FELs, Zholents suggested to replace the drifts from the MBEC scheme with wigglers, that could have the same effect but in a shorter distance, and thus reducing the total length of the cooling section [5]. To test the feasibility of this option, a Proof-of-Principle (POP) experiment will be carried out at the Argonne Wakefield Accelerator (AWA), where electrons will be brought to relativistic speeds and through a one metre wiggler. By experimenting with different particle distributions and energies we expect to obtain a better understanding of the effects of a wiggler, and a validation of the theory.

One of the main goals and motivations of this thesis was to study and guide the preparation of this POP-Experiment. In this document we present this study through simulations of the AWA accelerator and the wiggler in different regimes, and compare it to the theory.

The second motivation for this thesis was to develop a code that would allow for the start-to-end 3D simulation of Free Electron Lasers (FELs). A full electromagnetic simulation of such machines can be very computationally expensive, and that is why nowadays these simulations are done with codes that work only for certain regimes, and can thus make assumptions that speed-up the code [6–8]. Recent developments have made it possible for a handful of codes to simulate undulators and wigglers starting from first principles, but to our knowledge no single code can simulate the entire machine including cavities, bunch compressors, and undulators. For this thesis we use OPAL [9], an existing C++ code for the 3D simulation of particle accelerators, and enhance it such that it can include undulators in its simulations. This is achieved by use of an external library, MITHRA [10], which is a recently published code for the full-wave simulation of undulators. It is with this enhanced version of OPAL that we can model the entire AWA beamline with the wiggler, and simulate electron bunches from start-to-end. The wiggler experiments will also serve as a test for the code.

The following text is divided into two main parts. In the first half the OPAL-MITHRA integration is thoroughly discussed and benchmarked, and a convergence analysis is carried out. In the second part the MBEC and the physics of a wiggler are explained in detail, followed by the simulations

to study the AWA experiment and the effects of a wiggler for beams with different characteristics, and comparison to the predictions from the literature.

1.1 Acknowledgements

I would firstly like to thank my supervisor Dr. Andreas Adelman for giving me the opportunity to participate and do my small part in such interesting and exciting research. Furthermore, I would like to thank him for always being available and open to discuss any matter during the whole duration of the thesis, and offering important advice and guidance when it was needed. Working with him and with the Accelerator Modelling and Advanced Simulations group at PSI has been a true pleasure, and has helped me grow as a researcher and as a person. In general, I am thankful for all my colleagues at PSI, with whom every encounter has been pleasant, and from whom I have asked and received much help and advice on several occasions. A special thanks goes to my colleague and fellow student Renato Bellotti, who's work concerning the AWA MBEC experiment was of great help in this project.

I am also very grateful to our colleagues at ANL Dr. John Power, Dr. Gwanghui Ha, and Dr. Alexander Zholents, who have influenced this work through weekly discussions, and who will be carrying out the experiments here described. Their input has been essential, and this work could not have been done without it.

Finally, I would like to thank Dr. Arya Fallahi, who wrote the code MITHRA, an essential tool used in this work. He was open to discuss and accept modifications on the code as required for this thesis, and helped on many occasions to understand results or problems with the simulations. His help has been much appreciated.

2 Numerical Approach

2.1 OPAL Static Solver

The Object-Oriented Parallel Accelerator Library (OPAL) is an open source C++ code developed at the Paul Scherrer Institut (PSI) for the simulation of particle accelerators on large multicore HPC clusters [9]. It has the capability of simulating a variety of linear and circular machines in high resolution and with a large number of particles, and includes 3D space-charge effects. It is an open source software currently used and tested in several research institutes, and is in constant development and improvement.

OPAL takes as an input an initial distribution of particles and a beamline defined by the user, and then integrates each particle's trajectory in 3D space through discrete time-steps. At each step, the electric and magnetic fields from the beamline elements and the self-fields from the bunch are computed and used to push and kick the particles. Computing the self-fields of a bunch is often the performance bottle-neck for 3D electromagnetic simulation codes, and it is hence the part of the algorithm that can most benefit from tricks and approximations that could speed it up. In the case of OPAL the electrostatic and magnetostatic approximations are exploited,

which greatly simplify Maxwell's equations to

$$\begin{cases} \vec{\nabla} \wedge \vec{E} = 0 \Rightarrow \vec{E} = -\vec{\nabla}\phi, \\ \vec{\nabla} \wedge \vec{B} = \mu_0 \vec{j}, \end{cases} \quad (1)$$

where \vec{E}, \vec{B} are the electric and magnetic fields respectively, \vec{j} the current density, and ρ the charge density [11].

In addition, a Lorentz transformation into a frame of reference moving at the bunch's average speed $\vec{\beta}$ is performed on the particle positions before solving the equations, which means that the particles are quasi-static and the current and the magnetic field vanish $\vec{B} = \vec{j} \simeq 0$. This leaves us with a single equation to be solved in the comoving frame, Poisson's equation:

$$\nabla^2 \phi = -\frac{\rho}{\epsilon_0}. \quad (2)$$

Once the scalar potential ϕ has been computed, it can be used to compute the electric field in the comoving frame \vec{E} , followed by a Lorentz transformation back into the laboratory frame that yields both magnetic and electric fields:

$$\begin{cases} \vec{E}_{lab} = \gamma \vec{E} - (\gamma - 1) E_z \hat{z}, \\ \vec{B}_{lab} = \gamma \frac{\vec{\beta} \wedge \vec{E}}{c}. \end{cases} \quad (3)$$

To solve Poisson's equation efficiently OPAL uses a Particle-In-Cell (PIC) scheme. At each time-step the charge density ρ is interpolated to a grid of points that tightly encloses the whole bunch. The grid is then Lorentz boosted into the static frame, which is computationally less demanding than transforming the particle positions. Then Poisson's equation is solved on each grid-point, from which the field can be interpolated back to the particle positions. Finally the Boris pusher [12] scheme is used to update the positions and momenta of the particles.

An additional advantage of using a discretised mesh to compute the fields is that a Fast Fourier Transform (FFT) can be used to speed up the process. A solution to Poisson's equation is a convolution between the charge density ρ and the Green's function subject to the appropriate boundary conditions [11]. In Fourier space a convolution becomes a simple multiplication, and hence the bottle-neck of the whole algorithm is the Fourier Transform itself. Using an FFT represents a speed-up over more common methods that scale quadratically with the number of particles; this method reaches the solution in a time of the order $O(2N \log(2N))$, where N is the number of grid-points.

As a summary, the following lines go over the main points of the algorithm in each time-step:

1. Resize grid to enclose bunch
2. Charge deposition on discrete grid-points $\rho_{i,j,k}$
3. Lorentz transform to static reference frame

4. Solve Poisson $\nabla^2 \phi = -\rho/\epsilon_0$
5. Evaluate self-field $\vec{E} = -\vec{\nabla}\phi$
6. Lorentz transform to laboratory frame
7. Interpolate fields at particle positions $\vec{E}_{i,j,k} \rightarrow \vec{E}(\vec{x}), \vec{B}_{i,j,k} \rightarrow \vec{B}(\vec{x})$
8. Add external fields $\vec{E} += \vec{E}_{ext}, \vec{B} += \vec{B}_{ext}$
9. Update $\vec{x}_{n+1}, \vec{p}_{n+1}$ one time-step

This method of computing the space-charge fields is efficient and reliable for many cases, but it has its downfalls too. When doing the electrostatic and magnetostatic approximations, coupling between \vec{E} and \vec{B} fields is completely lost, which means that electromagnetic waves cannot propagate through space. In addition, if the particles have a large energy spread or velocity spread, the current and the magnetic fields can no longer be neglected in the comoving frame. This is the case for example in an undulator, the core element of Free-Electron-Lasers.

2.2 OPAL-MITHRA Full-Wave Solver

An undulator or wiggler is formed by an array of alternating dipoles that cause charged particles going through them to oscillate left and right in a wiggly trajectory (fig. 1). Light particles such as electrons emit a non-negligible amount of synchrotron radiation when subject to this alternating acceleration, which can affect other electrons in the bunch, and lead to micro-bunching and coherent radiation in an FEL [13]. As seen in the previous section, OPAL's static solver cannot account for these effects, and it is for this reason that up until now OPAL could not include undulators in the beamlines that it can simulate.

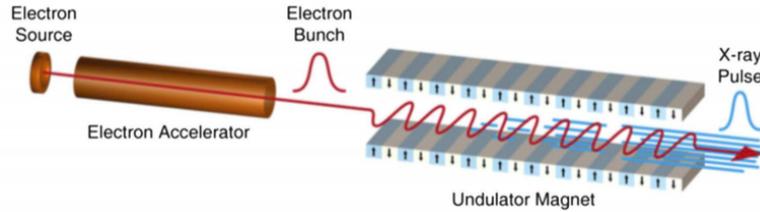


Figure 1: *Schematics of a Free-Electron-Laser [14].*

In order to fully capture the physics of an undulator one needs to compute the EM fields by solving the Maxwell equations [11]. By rearranging them and using the Lorenz Gauge we can reduce the problem to solving two wave equations:

$$\begin{cases} \vec{\nabla} \cdot \vec{E} = \frac{\rho}{\epsilon_0}, \\ \vec{\nabla} \wedge \vec{E} = -\frac{\partial \vec{B}}{\partial t}, \\ \vec{\nabla} \cdot \vec{B} = 0, \\ \vec{\nabla} \wedge \vec{B} = \mu_0 \vec{J} + \frac{1}{c^2} \frac{\partial \vec{E}}{\partial t}, \end{cases} \Rightarrow \begin{cases} \nabla^2 \vec{A} - \frac{1}{c^2} \frac{\partial^2 \vec{A}}{\partial t^2} = -\mu_0 \vec{J}, \\ \nabla^2 \phi - \frac{1}{c^2} \frac{\partial^2 \phi}{\partial t^2} = -\frac{\rho}{\epsilon_0}, \end{cases} \quad (4)$$

where \vec{A} and ϕ are the vector and scalar potentials respectively, related to the EM fields by

$$\begin{cases} \vec{E} = -\frac{\partial \vec{A}}{\partial t} - \vec{\nabla}\phi, \\ \vec{B} = \vec{\nabla} \wedge \vec{A}. \end{cases} \quad (5)$$

These wave equations and the equations of motion of the particles are coupled through the source terms ρ and \vec{j} , and through the Lorentz force that the fields exert on the particles

$$\vec{F} = q \left(\vec{E} + \vec{v} \wedge \vec{B} \right). \quad (6)$$

This means that the problem requires two separate time-marching algorithms, one to integrate the wave equations and one to integrate the particle positions, which communicate between themselves the charge distribution and the field values at each time-step. Thus OPAL needs a new solver that will be able to simulate the parts of the beam trajectory that include an undulator, by carrying out these two algorithms. In addition we would like OPAL to be able to switch back and forth between this solver and the static solver, in order to efficiently simulate entire beamlines or start-to-end FELs.

To achieve this, we decided to make use of the external C++ library MITHRA, developed by A. Fallahi et al. [10]. At the start of this project MITHRA was a standalone code for the simulation of electrons passing through an undulator, but during this thesis it was transformed into a library that can be used in other programmes, and was enhanced with new features that improved its accuracy and usability, and made it work well in conjunction with OPAL.

There are many codes that currently exist and are widely used for the simulations of Free-Electron-Lasers (e.g. [6–8]), but MITHRA was chosen for being the only one that uses Finite-Differences in Time-Domain (FDTD) to solve the 3D wave equations. Many codes work under certain assumptions to reduce the computational load, but which also render the code accurate only for certain use cases. We wanted OPAL to be able to simulate any undulator, and thus MITHRA seemed like a good option that had been thoroughly benchmarked.

The FDTD method [15] is a common and successful approach to computational electrodynamics, but it has some drawbacks that make it very computationally expensive. The main idea behind it is to discretise the wave equations (4) using finite differences in time and space to obtain an explicit update formula for the fields. The equations to solve at time $t = \Delta t(n + 1)$ are

$$\begin{aligned}
\phi_{i,j,k}^{n+1} &= -\phi_{i,j,k}^{n-1} + 2 \left[1 - \left(\frac{c\Delta t^2}{\Delta x} \right) - \left(\frac{c\Delta t^2}{\Delta y} \right) - \left(\frac{c\Delta t^2}{\Delta z} \right) \right] \phi_{i,j,k}^n \\
&\quad + \left(\frac{c\Delta t^2}{\Delta x} \right) (\phi_{i+1,j,k}^n + \phi_{i-1,j,k}^n) \\
&\quad + \left(\frac{c\Delta t^2}{\Delta y} \right) (\phi_{i,j+1,k}^n + \phi_{i,j-1,k}^n) \\
&\quad + \left(\frac{c\Delta t^2}{\Delta z} \right) (\phi_{i,j,k+1}^n + \phi_{i,j,k-1}^n) + (c\Delta t)^2 \rho_{i,j,k}^n, \\
\bar{A}_{i,j,k}^{n+1} &= -\bar{A}_{i,j,k}^{n-1} + 2 \left[1 - \left(\frac{c\Delta t^2}{\Delta x} \right) - \left(\frac{c\Delta t^2}{\Delta y} \right) - \left(\frac{c\Delta t^2}{\Delta z} \right) \right] \bar{A}_{i,j,k}^n \\
&\quad + \left(\frac{c\Delta t^2}{\Delta x} \right) (\bar{A}_{i+1,j,k}^n + \bar{A}_{i-1,j,k}^n) \\
&\quad + \left(\frac{c\Delta t^2}{\Delta y} \right) (\bar{A}_{i,j+1,k}^n + \bar{A}_{i,j-1,k}^n) \\
&\quad + \left(\frac{c\Delta t^2}{\Delta z} \right) (\bar{A}_{i,j,k+1}^n + \bar{A}_{i,j,k-1}^n) + (c\Delta t)^2 \rho_{i,j,k}^n,
\end{aligned}$$

where Δx , Δy , Δz are the grid spacings, and Δt the time-step for the field computation.

Unlike OPAL's static solver, which can stretch and rotate the mesh as necessary to tightly enclose the bunch at each time-step, FDTD requires the mesh to be fixed in space because each grid-point is updated using its previous values, and thus it needs to enclose the whole domain where there will be EM fields of interest during the simulation. In the case of an undulator this means that the computational domain should have the same length as the undulator, usually within a range $L_u \simeq 1\text{--}100$ m, but should also have a grid spacing small enough to resolve the bunch of length $\sigma_z \simeq 0.1\text{--}1$ mm and the resonance frequency of the radiation $\lambda_r \simeq 0.01\text{--}100$ μm , which represents an order of $\sim 10^6\text{--}10^{10}$ grid-points just for the longitudinal axis, extremely large even for the current state-of-the-art HPC clusters. In addition, FDTD algorithms need to satisfy the Courant-Friedrichs-Lewy condition [16] to be stable, which sets an upper limit on the time-step Δt proportional to the Δz , and would imply also $\sim 10^6\text{--}10^{10}$ steps for a single simulation.

It is for this reason that many codes make assumptions such as considering some periodicity in the beam to only simulate a slice, or the slowly-varying envelope approximation, which gives a parabolic problem instead of hyperbolic, and for which the time-step and grid spacing can be chosen independently. However, MITHRA greatly reduces the computational cost of the simulation by doing the entire computation in a Lorentz-boosted frame at a constant speed, as suggested in [17, 18].

It is well known that a bunch with an average energy $E = mc^2\gamma$ passing through an undulator gets slowed down to

$$\beta_0 = \sqrt{1 - \frac{1}{\gamma_0^2}}, \quad \text{where} \quad \gamma_0 := \frac{\gamma}{\sqrt{1 + \frac{K^2}{2}}}, \quad (7)$$

where K is the undulator strength parameter defined as

$$K = \frac{eB_{max}\lambda_u}{2\pi m_e c}. \quad (8)$$

with B_{max} the peak magnetic field in the undulator, and λ_u the undulator period [13].

MITHRA performs a Lorentz transformation from the laboratory frame to a moving frame at speed β_0 at the start of the simulation, and by doing so obtains two important advantages.

Firstly, since the grid-points are moving at a constant speed with the bunch, the computational domain only needs to enclose the bunch, instead of the whole undulator. Note that this is different from the electrostatic regime; individual particles are still moving at and accelerating to speeds close to the speed of light as they wiggle back and forth, only the average speed of the bunch is constant. The trajectory of the beam in the comoving frame can be seen in figure 2.

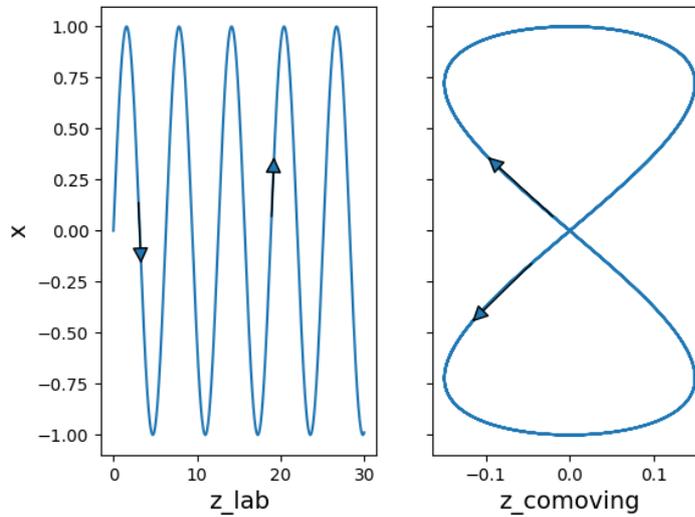


Figure 2: Comparison between the particle trajectories in the laboratory (left) and the comoving frame (right). Particles are not static in the comoving frame, despite their average speed being zero.

Secondly, using a boosted reference frame, the resonant wavelength of the radiation and the undulator period become of the same length (fig. 3). In the laboratory frame the radiation has a resonant wavelength given by

$$\lambda_r = \frac{\lambda_u}{2\gamma_0^2}, \quad (9)$$

which can be several orders of magnitude smaller than the undulator period, and hence requires a very small Δz . However, in the comoving frame the resonant wavelength is larger and is given by

$$\lambda_r \rightarrow \lambda'_r = \frac{\lambda_u}{\gamma_0}. \quad (10)$$

Contrarily, the undulator shrinks due to the Lorentz transformation, and its period becomes

$$\lambda_u \rightarrow \lambda'_u = \frac{\lambda_u}{\gamma_0}. \quad (11)$$

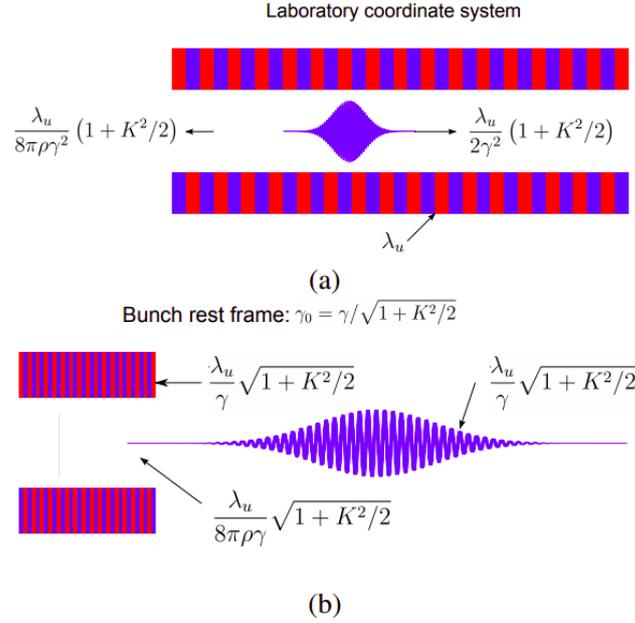


Figure 3: *The Lorentz transformation shrinks the undulator, and elongates the bunch and the resonant wavelength [10].*

With these changes in lengths the required number of longitudinal grid-points is reduced by $\sim 10^3$ – 10^5 depending on the exact parameters. The number of time-steps is also reduced by γ_0^2 .

By means of the Lorentz boost, and a few other techniques and details to ensure accuracy, MITHRA achieves an efficient full-wave 3D simulation of a bunch passing through an undulator. The final algorithm is roughly summed up in the following lines:

1. Lorentz transformation on the bunch, grid, and undulator, is applied only once at the start of the simulation
2. At each time-step:
 - (a) Charge and current deposition on discrete grid-points $\rho_{i,j,k}, \vec{j}_{i,j,k}$
 - (b) Update $\phi_{n+1}, \vec{A}_{n+1}$ one time-step
 - (c) Evaluate fields $\vec{E} = -\frac{\partial \vec{A}}{\partial t} - \vec{\nabla} \phi, \vec{B} = \vec{\nabla} \wedge \vec{A}$
 - (d) Interpolate fields at particle positions $\vec{E}_{i,j,k} \rightarrow \vec{E}(\vec{x}), \vec{B}_{i,j,k} \rightarrow \vec{B}(\vec{x})$
 - (e) Add external fields $\vec{E} += \vec{E}_{ext}, \vec{B} += \vec{B}_{ext}$

(f) Update $\vec{x}_{n+1}, \vec{p}_{n+1}$ one time-step

This code can be run on multiple CPUs in parallel using the MPI-protocol. For details on the algorithm, benchmarks, and examples, the reader is referred to the MITHRA manual [19].

As mentioned earlier, MITHRA was adapted into a static library such that OPAL could use its functions and classes for a new full-wave solver, and so a new OPAL class `UNDULATOR` has been defined as a beamline element, with its own full-wave solver attached to it. When a bunch being simulated by OPAL reaches the fringe-field of an undulator, a Lorentz boost is applied to all the particles, and the MITHRA functions take over from this point on, to integrate the trajectory and compute the fields as described above, along the whole undulator (see fig. 4). In this way the rest of the beamline, where the electrostatic approximation remains mostly valid, can be simulated faster with the static solver, and only the undulator makes use of the more computationally expensive FDTD full-wave algorithm.

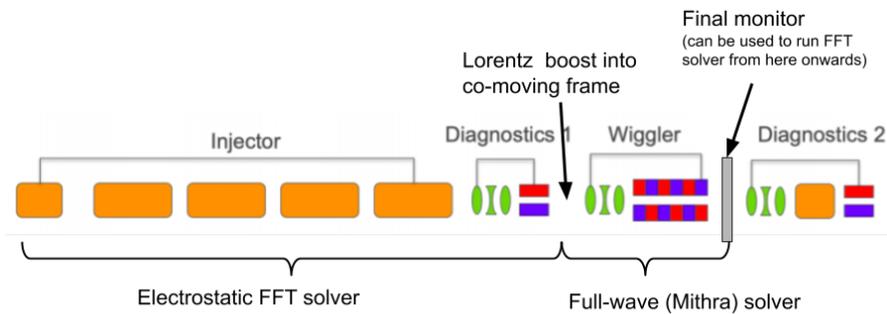


Figure 4: Solver distribution for the start-to-end simulation of the AWA beamline. (image from [5]).

The `UNDULATOR` element can be introduced to a beamline with the usual OPAL syntax:

```
UND: UNDULATOR, ELEMEDGE = 26.0, K = 10.81,
      LAMBDA = 8.5e-2, NUMPERIODS = 10,
      MESHLENGTH = { 20e-3, 20e-3, 8e-3 }, MESHRESOLUTION = { 30e-6, 30e-6, 2e-6},
      TRUNORDER = 2, TOTALTIME = 5e-9, FNAME = "output_params.job";
```

where

- `ELEMEDGE` defines the z -coordinate in the beamline where the fringe fields of the undulator start,
- `LAMBDA` is the undulator period λ_u ,
- `NUMPERIODS` the number of periods N_u , and hence the undulator length is `LAMBDA*NUMPERIODS`,
- `MESHLENGTH` is the size of the computational domain (L_x, L_y, L_z) , which should be slightly larger than the bunch size,

- **MESHRESOLUTION** is the grid-spacing ($\Delta x, \Delta y, \Delta z$), which should be small enough to resolve the resonant frequency,
- **TRUNORDER** can be 1 or 2, and gives the order of the Absorbing Boundary Conditions (ABC) used in the computational domain boundaries,
- **TOTALTIME** is the time that the full-wave simulation should last (in the future OPAL will calculate this automatically),
- **FNAME** is the name of the file that specifies specific outputs desired from the simulation. It is used to indicate where to place phase-space monitors, whether to record emitted radiation, and other possible outputs. This is the same as the *jobfile* in MITHRA, and one can refer to the manual [19] for the details.

At the moment of writing, an OPAL simulation stops after the undulator, and thus undulators can only be included at the end of a beamline. The possibility to change back from full-wave to static solver and continue the simulation further downstream is currently in development, and will soon be implemented.

2.3 Convergence Tests

In order to test the validity of the simulations, it is important to have an understanding of how the schemes perform and how they depend on computational parameters. A finer grid and a higher number of particles will yield results closer to reality, but will also require much more time and computational power.

One of the main interests in the experiments described in this thesis is to compare the longitudinal phase-space of electron beams in the Argonne Wakefield Accelerator (AWA), so we will use simulations on this beamline for the convergence analysis.

2.3.1 Static Solver

For the static solver we will compare simulations on the linear acceleration part of the AWA, where the static approximation is valid. This part consists of a photocathode (electron gun) that emits the electron bunch, 4 RF cavities each providing an increase in energy of ~ 10 MeV, and a YAG screen at 11 metres to capture the phase-space (fig. 5).

The electron bunch used for these simulations is similar to the one used in the experiments, consisting in a shifted set of Gaussian distributions, to give a bunched beam. The exact beamline and bunch parameters were chosen using the surrogate model of the AWA beamline [20] and can be found in table 1. The reader is referred to section 3.2 for a description of each parameter from the table.

IBF	467.90 A
IM	165.98 A
GPHASE	1.93 °
ILS1	1.93 A
ILS2	19.18 A
ILS3	1.87 A
Q	1.0 nC
λ	1.3 ps
σ_r	4.71 mm

Table 1: Parameters used in the simulations for the convergence analysis of the static solver.

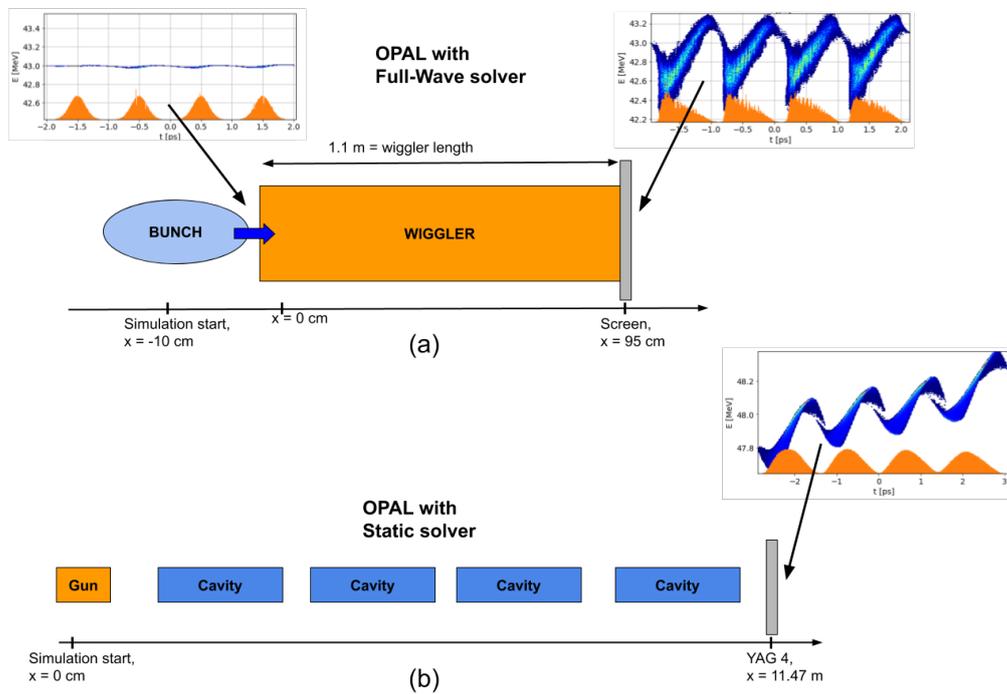


Figure 5: Simulated parts of the beamline for the convergence analysis of the full-wave solver (a), and of the static solver (b).

To compare two phase-space plots, the bunch's longitudinal phase-space is divided into 100 longitudinal slices, and for each one we measure the average slice-energy E_i and slice-energy-spread σ_{E_i} . From these values we can compute two errors:

$$\epsilon_1 := \sum_{i=0}^{N_{slices}} \frac{|E_i - \hat{E}_i|}{N_{slices}} = \langle |E_i - \hat{E}_i| \rangle, \quad (12)$$

$$\epsilon_2 := \sum_{i=0}^{N_{slices}} \frac{|\sigma_{Ei} - \hat{\sigma}_{Ei}|}{N_{slices}} = \langle |\sigma_{Ei} - \hat{\sigma}_{Ei}| \rangle, \quad (13)$$

where \hat{E} and $\hat{\sigma}_E$ refer to the reference simulation, which was done with the highest resolution, and is considered to be the *truth* when evaluating the error of a simulation. Since the error is defined as the average of the slice-errors, we will take the standard deviation of the slice-errors as the errorbars. This process is summarised in figure 6. Note that the longitudinal axis is time t instead of z as one might expect. This is because the phase-space is recorded on a screen, so each particle is recorded at the same z -coordinate, but with different times. As long as the bunch has a speed $\vec{\beta} \simeq \beta_z$, time is proportional to z and can be used interchangeably.

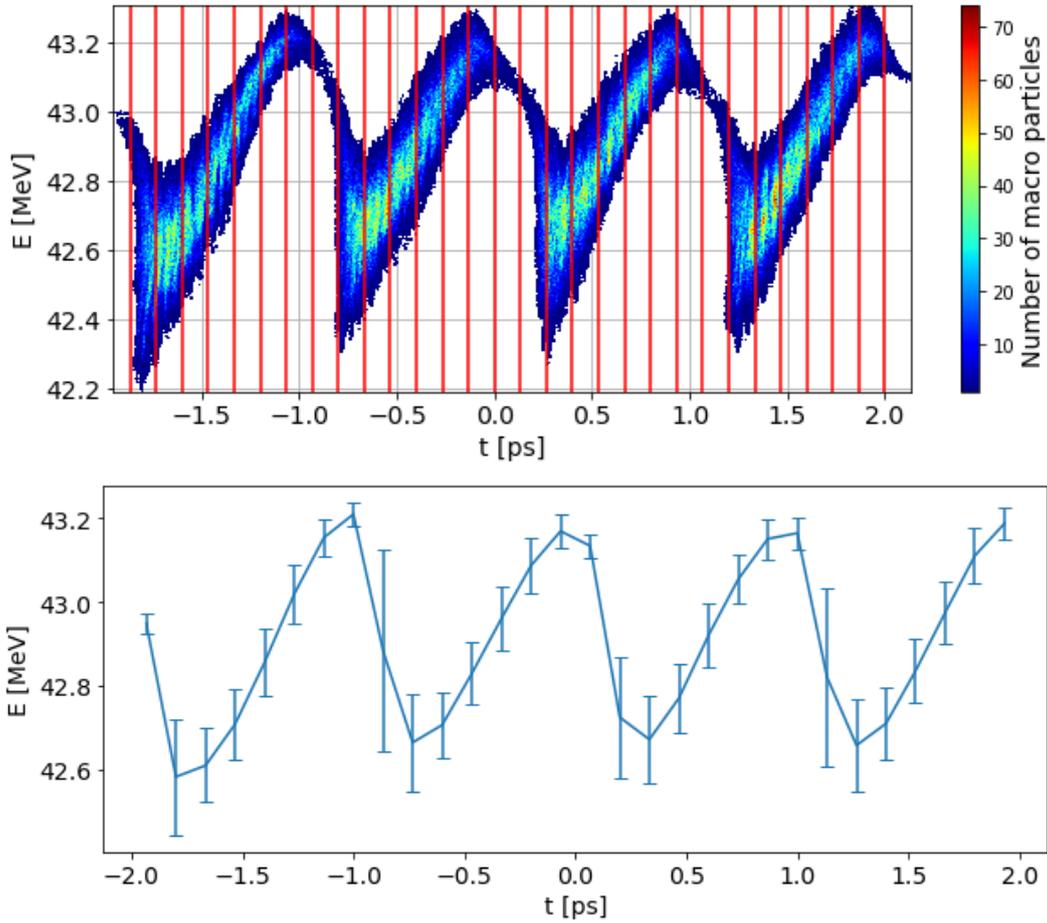


Figure 6: *Slicing the phase-space plot and obtaining slice-energy and slice-energy-spread values, for the comparison between simulations. For the sake of clarity, only 30 slices are drawn in this image, although the convergence analysis was done with 100 slices.*

Figures 7, 8, 9 show the convergence study with respect to the number of particles N , the Number of Particles-per-Cell NpC , the number of longitudinal grid-points MT , and the number of transverse grid-points MX, MY . The longitudinal and transverse grid-points are studied separately since, for the specific beam parameters that are studied in this thesis, a higher longitudinal resolution is required.

Note that the number of grid-points is used as a parameter instead of the grid-spacing $\Delta x, \Delta y, \Delta z$, because the static solver resizes the mesh at each time-step to fit the bunch, so the grid-spacing changes during the simulation. Additionally, note that for the grid-point convergence the Number of Particles-per-Cell NpC is kept constant, rather than the total number of particles N , since reducing the cell number alone would decrease NpC , and noise would increase, distorting the phase-space plot. It is well known that PIC schemes can show non-negligible noise if there are not enough particles per cell.

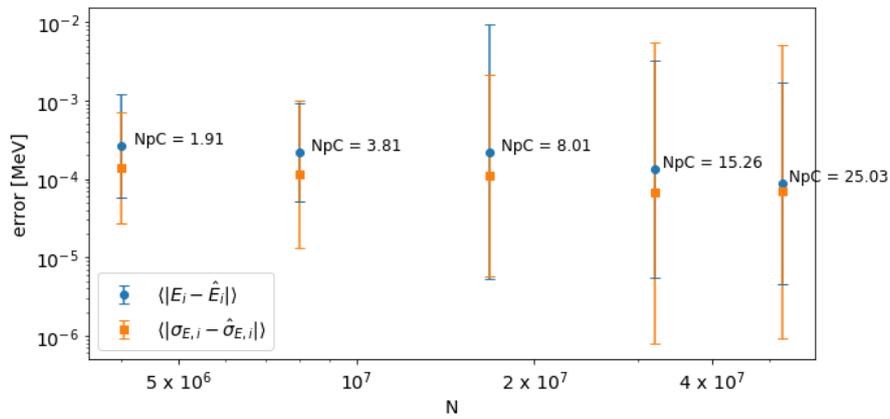


Figure 7: *Static solver convergence with respect to number of particles N and number of particles-per-cell NpC . The rest of parameters are fixed at $MX = MY = 64$, $MT = 512$.*

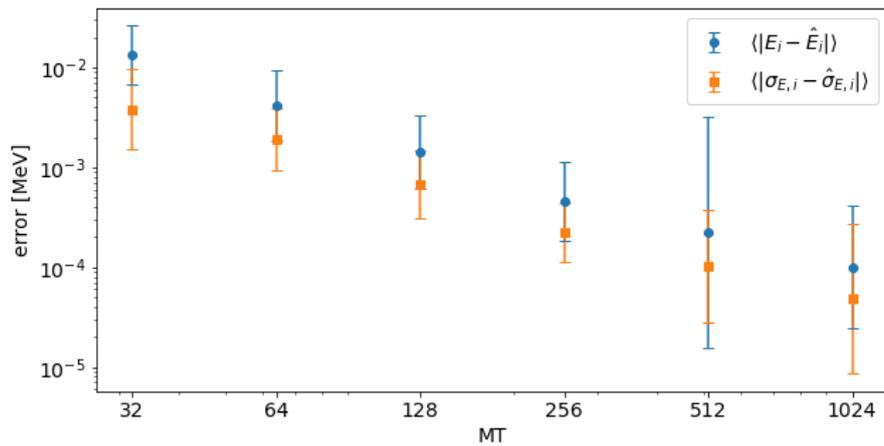


Figure 8: *Static solver convergence with respect to number of longitudinal grid-points MT . The rest of parameters are fixed at $NpC = 8$, $MX = MY = 64$.*

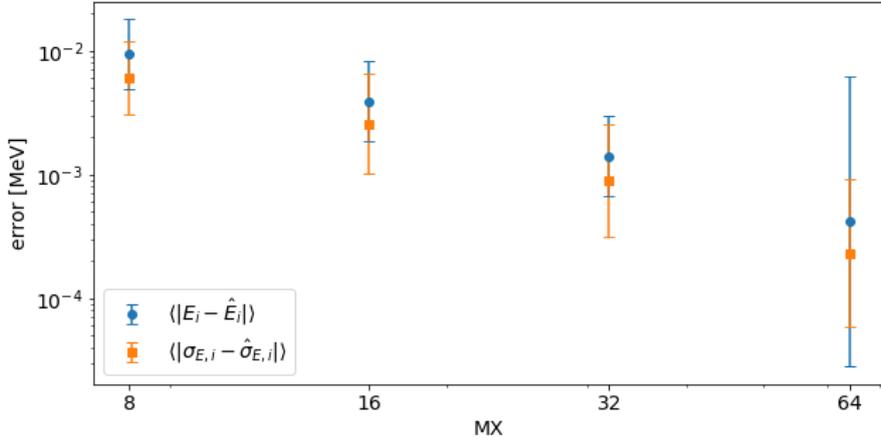


Figure 9: *Static solver convergence with respect to number of transverse grid-points $MX = MY$. The rest of parameters are fixed at $NpC = 8$, $MT = 512$.*

From the plots regarding MT , and MX, MY we can see that, although the error has not yet converged, the values tested clearly give an error below the resolution that can be obtained from a spectrometer, estimated to be under 10 keV. The number of particles has surprisingly hardly any effect on the error, which is constantly at ~ 0.1 keV. Since OPAL cannot run with less than 1 particle-per-cell ($N \geq N_{cells}$) we cannot test a lower number of particles, but we can speculate that the error might converge at $N < 4e6$, and what we are seeing is an error that has already converged. It is nonetheless strange that $NpC = 1$ doesn't generate noticeable noise. OPAL's static solver is usually run with 10 particles per cell or more [9]. A summary of the adequate parameters for the simulation of the AWA can be found in table 3.

2.3.2 Full-Wave Solver

For the full-wave solver we will inject a mono-energetic bunch at $E = 43$ MeV a few centimetres in front of the AWA wiggler, and capture the phase-space after the wiggler with another YAG screen (see figure 5). The full simulation parameters can be found in table 2, with each parameter explained in detail in section 3.2.

K	10.81
λ_u	85 mm
N_u	10
E	43 MeV
Q	0.3 nC
λ	1.0 ps
σ_r	681 μm

Table 2: Parameters used in the simulations for the convergence analysis of the full-wave solver.

Figures 10 to 14 show the results of the convergence study with respect to the number of particles N , the Number of Particles-per-Cell N_{pC} , the grid-spacing Δ , and the size of the computational domain L . Here it makes sense to use grid-spacing as a parameter as the grid has a fixed size throughout the whole simulation, and in this way it can be compared to other values such as the radiation wavelength or the bunch size. We also study the convergence with respect to the size of the computational domain, as we expect the reflection of waves on the boundaries to have a non-negligible effect. For these simulations N_{pC} is not constant, because the bunch can shrink and grow throughout the simulation, but the grid cannot. This is a disadvantage of using FDTD with respect to the static solver, since at some point during the simulation N_{pC} could be very low and induce noise.

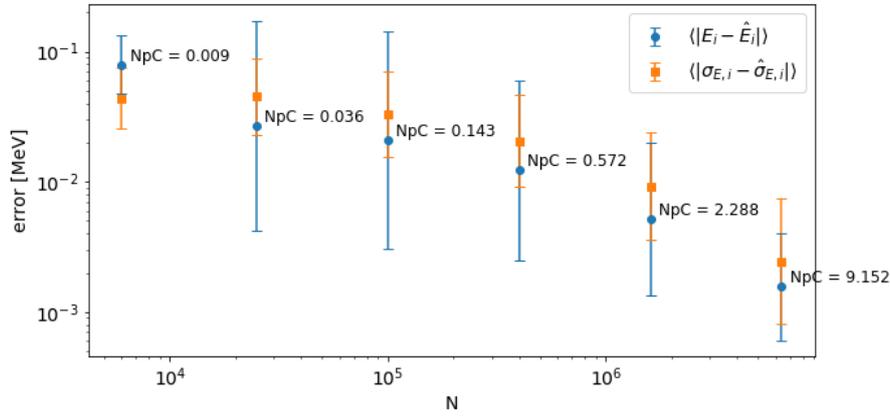


Figure 10: Full-wave solver convergence with respect to number of particles N and number of particles-per-cell N_{pC} . The rest of parameters are fixed at $\Delta = (25, 25, 4)$ μm , $L = (9, 9, 4)$ mm, and an initial bunch size $\sigma = (681, 681, 1200)$ μm .

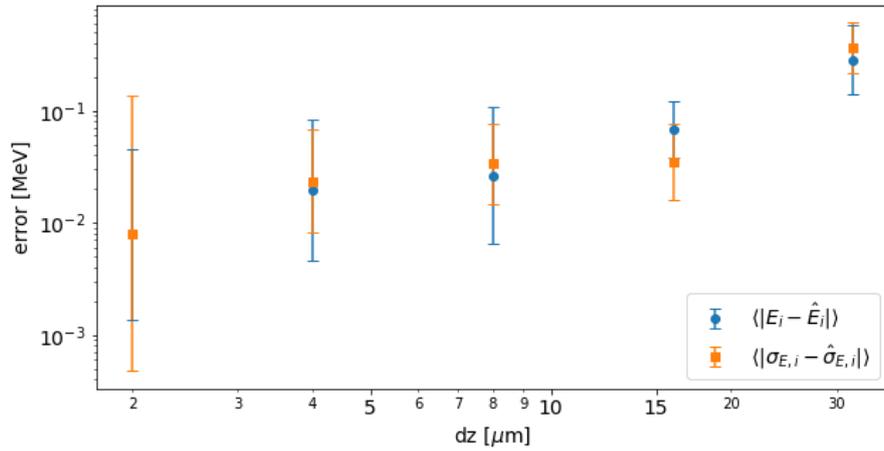


Figure 11: Full-wave solver convergence with respect to longitudinal grid-spacing Δz . The rest of parameters are fixed at $\Delta x = \Delta y = 17$ μm , $L = (9, 9, 4)$ mm, and an initial bunch size $\sigma = (681, 681, 1200)$ μm .

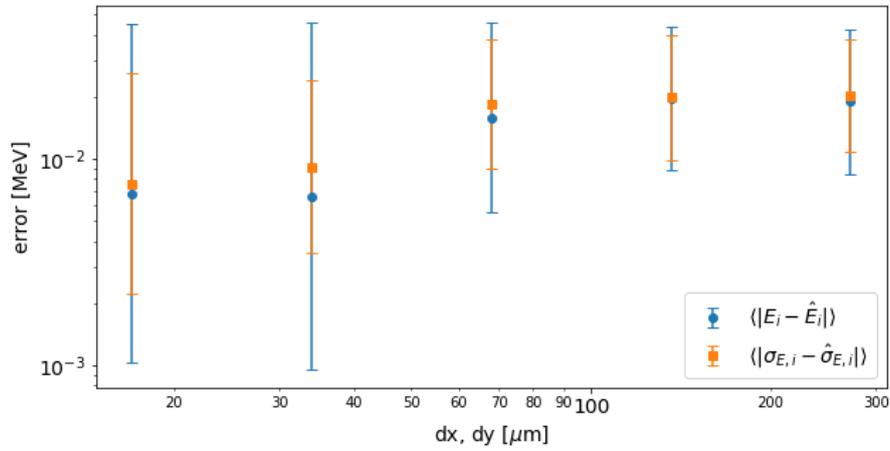


Figure 12: *Full-wave solver convergence with respect to transverse grid-spacing $\Delta x = \Delta y$. The rest of parameters are fixed at $\Delta z = 2 \mu\text{m}$, $L = (9, 9, 4) \text{ mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.*

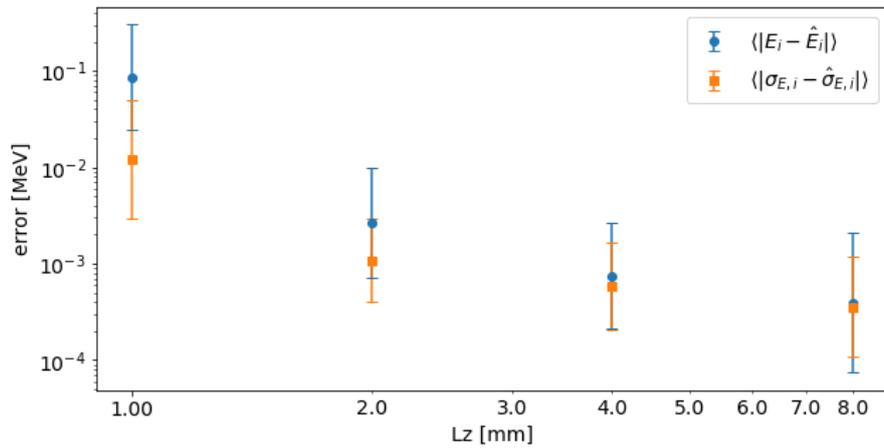


Figure 13: *Full-wave solver convergence with respect to length of computational domain L_z . The rest of parameters are fixed at $\Delta = (25, 25, 4) \mu\text{m}$, $L_x = L_y = 9 \text{ mm}$, and an initial bunch size $\sigma = (681, 681, 1200) \mu\text{m}$.*

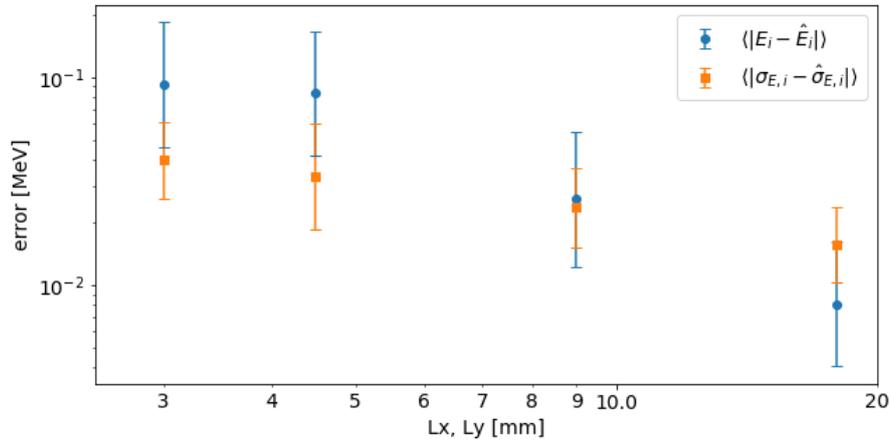


Figure 14: Full-wave solver convergence with respect to transverse size of computational domain $L_x = L_y$. The rest of parameters are fixed at $\Delta = (25, 25, 4)$ μm , $L_z = 4$ mm, and an initial bunch size $\sigma = (681, 681, 1200)$ μm .

For all the parameters we achieve an error below 10 keV, which is approximately the resolution that will be seen when doing the experiment. In table 3 the adequate parameters for the simulations are summarised. With these settings, the simulation of a full undulator on 64 CPUs takes approximately 12 hours.

Some of these parameters can be relaxed without affecting too much the outcome, reducing the time of computation by a factor of 4 or 5. The parameters that should not be relaxed are Δz and L_x, L_y , as from the plots they appear to have the largest impact on the error. This is to be expected for Δz , as it needs to be small to resolve the resonant wavelength, but also because it is proportional to the time-step (such that CFL is satisfied and to ensure the correct dispersion). Regarding L_x, L_y , this can be explained by the fact that the ABCs used in MITHRA only absorb perpendicularly incident waves, but can have unrealistic reflections for oblique incidence. The radiation from the undulator will be mostly forwards along the z-direction, but some of it will be at an angle, hitting the transverse boundaries. Thus the transverse boundaries need to be ~ 10 times larger than the bunch transverse size.

Static Solver		Full-Wave Solver	
NpC	10	NpC	10
MT	128	Δz	2 μm
MX, MY	32	$\Delta x, \Delta y$	40 μm
		L_z	4 mm
		L_x, L_y	20 mm

Table 3: Summary of the approximate values to be used in the simulation of the AWA beamline.

2.4 Code Benchmarking

MITHRA has been benchmarked against other codes through simulation of FELs in different regimes, as well as in free-space propagation. These tests can be seen in the original publication [10]. Here we provide an additional benchmark with a wiggler and bunch slightly similar to the cases that will be studied in the AWA experiment.

In [21], MacArthur et al. investigate the possibility to generate a single cycle energy modulation in an electron bunch, by means of a wiggler. The idea is that if a short energy modulation can be obtained in the centre of the bunch, a chicane can transform it into a short electron pulse that will increase the peak current. This bunch can then be used in a long undulator and undergo the usual Self-Amplified Spontaneous Emission (SASE) process of FELs, and has the potential to generate sub-femtosecond x-ray pulses [22]. This way of achieving the energy modulation is an alternative to other previously proposed methods, namely by means of a laser mixed with a wiggler [23].

To do this, a relativistic electron bunch with a high current in the tail (fig. 15) is passed through a strong 6-period wiggler. Several theoretical models discussed in the paper suggest that the radiation from the tail will generate the necessary single-cycle energy modulation in the bunch centre, and this is later tested with the 3D electromagnetic code OSIRIS [24] and with experiments on the Linac Coherent Light Source (LCLS) at SLAC. The bunch phase-space before and after the wiggler as simulated on OPAL with the full-wave solver can be seen in figure 16, and a comparison with the results from the paper in figure 17. Note that in these plots the bunch head is to the right, and bunch tail is to the left. Excellent agreement can be seen between the codes, and both cases show the sought energy modulation in the centre of the bunch. The specific parameters used for the OPAL simulation are shown in table 4.

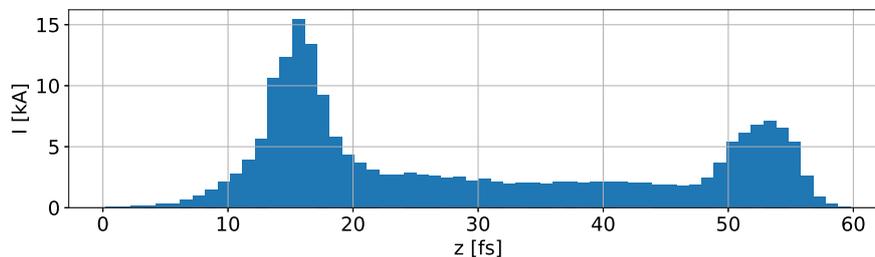


Figure 15: *Initial current in the simulation of the LCLS experiment described in [21].*

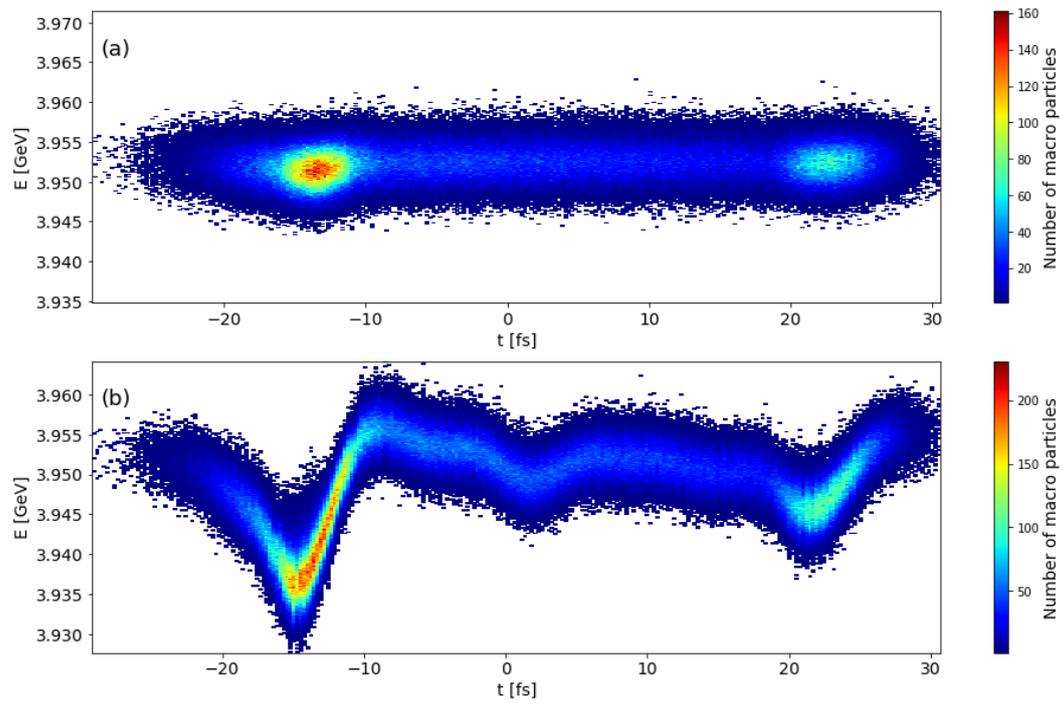


Figure 16: *Longitudinal phase-space of the electron bunch (a) 80 cm before and (b) 10 cm after the wiggler in the LCLS experiment used for benchmarking.*

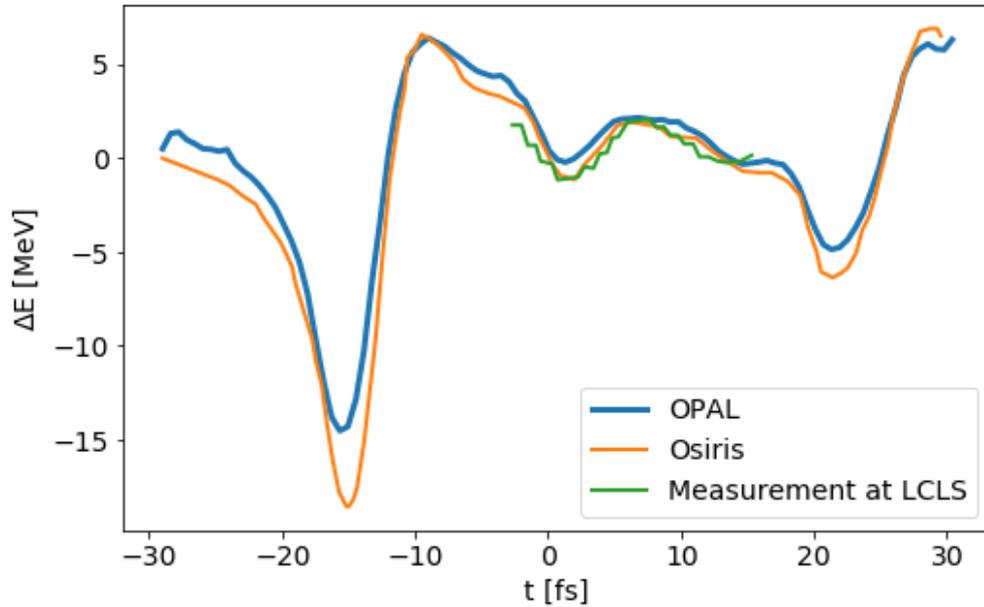


Figure 17: Comparison of the final slice-energy between MacArthur’s paper [21] and an OPAL simulation of the same wiggler and bunch.

Q	200 pC
E	3.95 GeV
σ_E	2 MeV
K	51.5
λ_u	35 cm
L_u	2.1 m
N	$5 \cdot 10^5$
$(\Delta x, \Delta y, \Delta z)$	(30, 30, 0.1) μm
(L_x, L_y, L_z)	(1800, 1800, 85) μm
σ_r	75 μm

Table 4: Values used in the OPAL simulation of the LCLS experiment used for benchmarking.

3 Experiment and Simulations

3.1 Micro-Bunched Electron Cooling

In the field of particle accelerators, cooling a particle bunch refers to decreasing the phase-space volume of the bunch. This can be achieved both by decreasing bunch size, or by decreasing its energy or momentum spread. Therefore a cool beam has particle velocities with a small standard

deviation, and a small transverse radius, which allows for higher luminosity and thus higher probability of interactions with the opposite beam in a collider, or a fixed target.

The high luminosity that is required in the future EIC [1] will need some form of cooling that maintains small beam sizes for hours. The electron beam in the collider will have *natural* cooling due to synchrotron radiation, but the hadron beam has negligible radiation due to the much larger mass of hadrons. The intra-beam scattering, and beamline defects that accumulate over turns will degrade the beam and increase phase-space volume, which will decrease luminosity, and might eventually lead to particle losses and machine activation or damage.

The method proposed to be used in the EIC is Micro-Bunched Electron Cooling (MBEC), which is based on an idea initially proposed by Derbenev in 1992 [25]. The scheme proposed by Derbenev was named Coherent electron Cooling (CeC), and consists in having an electron beam that copropagates with the hadrons for a short distance, and acts both as a pick up of the energy deviations and as a kicker/corrector of the deviation. The electrons and hadrons first meet at a point called the modulator, where the hadrons impose small energy variations in the electron beam due to Coulomb interaction. The two beams are then separated, and brought through dispersive sections. For the hadrons, this means that more energetic ones will travel a shorter distance, and so faster hadrons will shift ahead of the slower ones within the bunch. For the electrons, this dispersive section will have the effect of transforming the energy modulations into small current peaks or micro-bunches. The two beams are then brought back together in the kicker, where this time the electron micro-bunches will act back on the hadrons through Coulomb forces, giving a negative kick to those with a positive energy deviation, and a positive kick to those with less energy. The process is schematically explained in figure 18.

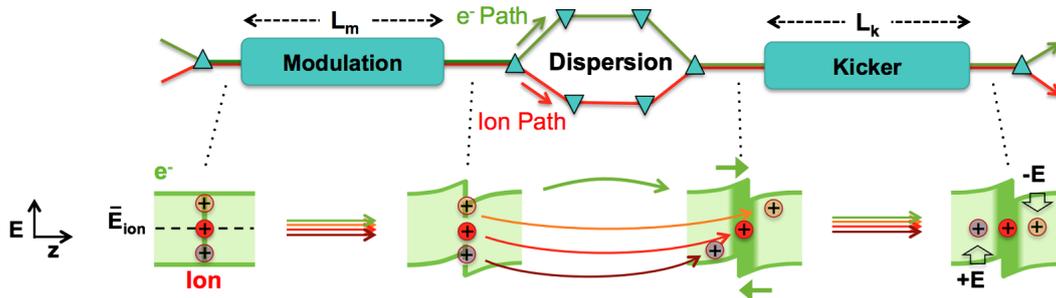


Figure 18: *Basic idea behind the Coherent electron Cooling (CeC) scheme [2].*



Figure 19: *Schematics of the EIC infrastructure, including the electron beamline for MBEC [26].*

Clearly, the key to efficiently cooling a hadron beam with this scheme is to get the largest density modulation possible in the electron bunch by the time it reaches the kicker, such that the correction kick on the hadrons will be as strong as possible. To this effect, Ratner proposed to exploit the Micro-Bunching Instability (MBI) as a means of amplification of the density modulations [2]. MBI is an effect usually seen as detrimental, where a small random energy perturbation in a bunch gets amplified through a series of chicanes and drifts. MBEC uses this effect by adding a series of amplification sections between the modulator and the kicker. Each amplification section consists of a drift, where the density modulations will go through a quarter plasma oscillation $\lambda_p/4$ and be transformed into energy modulations, and a chicane, where the energy modulations are turned back into density modulation, but with a stronger peak density.

Stupakov calculated the feasibility of using this scheme with a collider based on the values of the eRHIC in Brookhaven, and estimated that a cooling section with two amplification sections would be enough to reduce the cooling time from ~ 51 hours, to under 1 hour, and hence make it a viable option [27].

Figure 19 shows the infrastructure of the future EIC, and how the electron cooling section might be included to it. Figure 20 shows different setups for the MBEC scheme, including one with a wiggler instead of a drift, as will be explained in the next section.

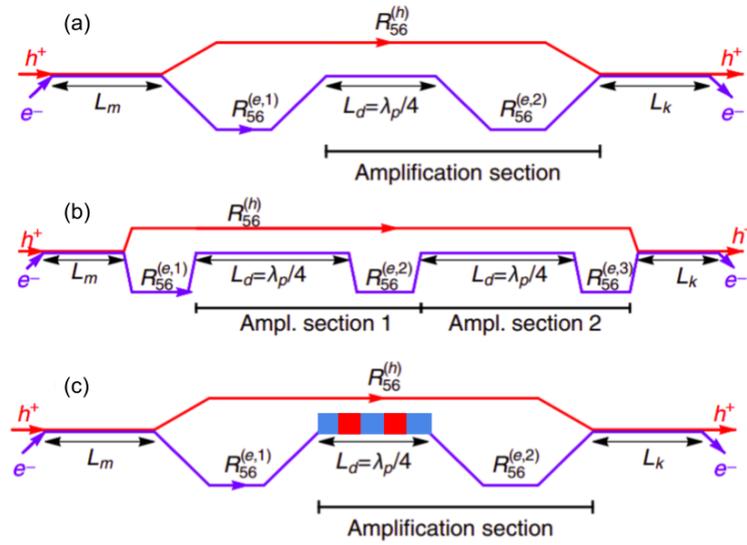


Figure 20: *Different improvements on the basic CeC scheme. (a) Micro-Bunched Electron Cooling with one amplification section, (b) MBEC with two amplification sections, (c) MBEC with one amplification section, with a wiggler replacing the drift [27].*

3.1.1 Wiggler as a Drift Replacement

In 2019 Zholents suggested replacing the drifts in the amplification sections of the MBEC scheme with short wigglers, as it could shorten the overall length of the cooling section [5]. The goal of the drifts in MBEC is to allow the density modulations to turn into energy modulations by means of longitudinal space-charge effects. As explained in section 2.3.2, particles entering a wiggler are slowed down to a speed equivalent to that of particle in a drift with a *reduced* gamma factor

$$\gamma_0 = \frac{\gamma}{\sqrt{1 + K^2/2}}. \quad (14)$$

Since space-charge effects are inversely proportional to γ^2 , one can intuitively assume that a wiggler will have stronger longitudinal space-charge fields, and thus will be able to transform density to energy modulation in a shorter length.

However, this is not so straightforward, as there are other effects at play in a wiggler:

- The non-zero dispersion due to dipoles will cause slower, less energetic electrons to lag and fall behind the faster, more energetic ones.
- Electrons emit synchrotron radiation due to the constant acceleration. Electrons within slippage length and diffraction transverse size can radiate coherently, and eventually lead to microbunching and SASE radiation.

- In addition to the wiggling in the horizontal plane, there is vertical focusing that will cause vertical β -oscillations.

Geloni et al. [4] give some guidance from theory on how to choose parameters and how they will effect the radiation and space-charge fields. In the paper, the effect of a wiggler is examined under the following conditions:

- i The bunch is much longer than the radiation wavelength, so the whole bunch will not radiate coherently: $\lambda \gg \lambda_r$, where $\lambda := \sigma_z$,
- ii A steady state of the wake-fields is achieved, but far from the saturation regime $L_s \gg 2\gamma_0^2\lambda$, where L_s is the saturation length,
- iii The effects of the metallic surroundings are negligible $a \gg \gamma_0\lambda$, where a is the radius of the beam pipe.

Under these constraints, it is theorised that if $\sigma_r \gg \sigma_{\text{diff}}$, where

$$\sigma_{\text{diff}} = \sqrt{\lambda\lambda_u}, \quad (15)$$

radiation fields can be neglected and the only fields left are the longitudinal space-charge fields, which are exactly the same as in a drift as in a drift with bunch energy $E = mc^2\gamma_0$.

To validate this theory, and to better understand the effects of a magnetic wiggler, experiments will be carried out at AWA.

3.2 Experiment Overview

The experiment at AWA will consist of a photo-cathode gun that emits electrons, followed by 4 or 6 cavities, each accelerating roughly by 10 MeV, and finally a 1.1 m wiggler designed specifically for these experiments. The beamline will also have focusing solenoids and quadrupoles, as well as several diagnostics. Namely, after the wiggler there will be a transverse deflecting cavity followed by a spectrometer that will allow a visualisation of the longitudinal phase-space, the main interest of the experiment. The gun and linacs occupy the first 11 m of the beamline, then there is a ~ 15 m drift that can be used for focusing or diagnostics, before reaching the wiggler placed at 26 m. A general schematic of the experimental setup can be seen in figure 4, and a detailed plan with all the elements in figure 22.

The AWA gun can produce trains of electron bunches with Gaussian or flattop longitudinal distributions by means of UV Birefringent Crystals [28], and thus can also mimic a beam with density modulation, just like the one that would come out of the dispersive section in the MBEC scheme. Recall that we are trying to see how and if a wiggler can transform a density modulation into an energy modulation. It was decided that the most feasible bunch to use in the experiments would be a train of 4 Gaussian pulses along the longitudinal axis, each with a Full-Width Half-Maximum of 0.3 ps ($\simeq 90 \mu\text{m}$). The transverse distribution is a uniformly filled circle. Each generated bunch will be defined by its radius σ_r (corresponding to the laser spot-size shone on the cathode), its charge Q , and the separation between Gaussians λ . These three parameters

can be chosen freely within some boundaries. An example of the longitudinal distribution of a bunch after the gun can be seen in figure 21. Note that here λ is proportional to the total bunch length σ_z and will be used when comparing to the theory exposed in the previous section.

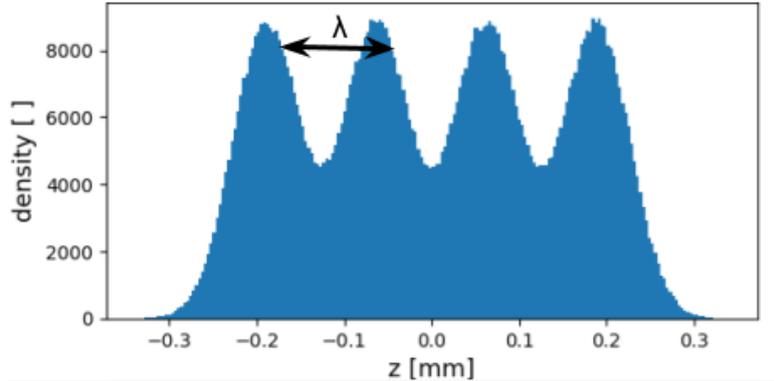


Figure 21: *Example of longitudinal distribution used in the AWA experiments.*

The wiggler built for this experiment has 10 periods of 8.5 mm, and a strength of $K = 10.81$, which is equivalent to a peak magnetic field $B_{max} = 1.37$ T. In the simulations of the beamline we shall study the case with 4 and with 6 cavities, which will be roughly equivalent to 40 MeV and 60 MeV respectively. However, at the time of writing only 4 cavities are working at the AWA, and it is probable that only the experiments at 40 MeV will be done.

The other beamline parameters that can be tuned differently for each simulation are shown in table 5.

The experiment was set to happen in April 2020, but due to the Covid-19 situation has been delayed, and is likely to happen sometime before September 2020.

Q	Bunch charge
λ	Pulse separation
σ_r	Laser spot radius = Initial bunch transverse radius
GPHASE	Phase of the RF cavity at the gun
IBF	Current on the Bucking and Focusing solenoids at the gun
IM	Current on the Matching solenoid right after the gun
ILS1,2,3	Current on the three Linac Solenoids

Table 5: Tunable parameters during the experiment at the AWA facility.

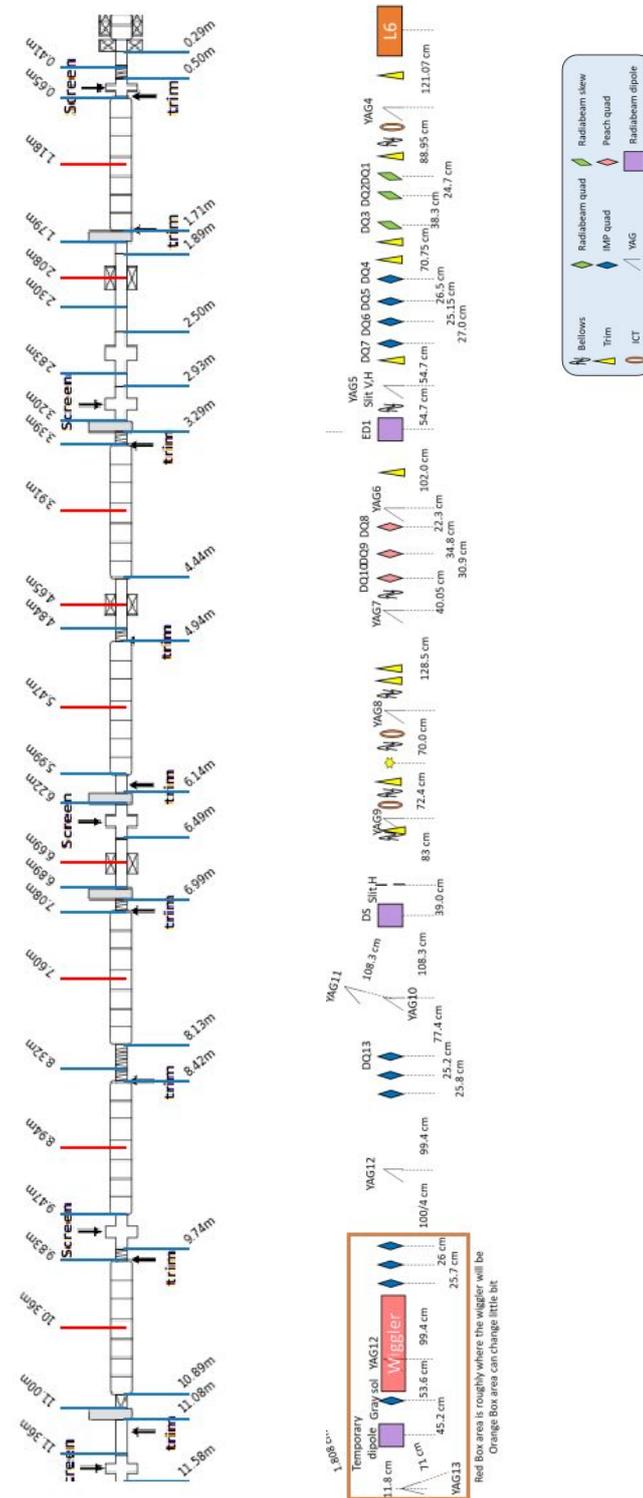


Figure 22: Detailed plan of the experimental setup used at AWA, provided by G. Ha.

3.3 Simulations

3.3.1 Wiggler Simulations

To get a better understanding we shall first simulate a bunch going through the wiggler only, rather than the whole accelerator. In this way we have complete control over the bunch at the entry of the wiggler, and it is easier to scan a range of bunch parameters. A cold bunch (zero energy spread) is inserted 18 cm before the wiggler, and then allowed to pass through it. A monitor at the end captures the resulting phase-space. These simulations are done entirely with the full-wave solver in OPAL. The simulation setup is as in figure 5.a, and the solver parameters are summarised in table 6. The number of particles needed to be increased from 3 million to 30 million for the simulations with the highest charge density, as otherwise unrealistic noise heavily polluted the results.

N	$3 \cdot 10^6$ or $30 \cdot 10^6$
$(\Delta x, \Delta y, \Delta z)$	$(30, 30, 2)$ μm
(L_x, L_y, L_z)	$(20, 20, 8)$ mm

Table 6: Parameters used for the OPAL full-wave solver in this section.

A set of specific parameters were chosen in order to explore different regimes, and to test the theory exposed in section 3.1.1. These consist of two energies $E = 44, 63$ MeV, two pulse-separation lengths $\lambda/c = .4, 2$ ps, and three transverse sizes $\sigma_r = 340, 680, 1800$ μm , giving in total 12 different simulations. All of these parameters satisfy being far from saturation $L_s > 2\gamma_0^2\lambda$, and having negligible effects from the beam pipe $a > \gamma_0\lambda$, but not all of them satisfy the two remaining conditions $\lambda > \lambda_r$ and $\sigma_r > \sigma_{\text{diff}}$ where the theory deems the radiation negligible. The 12 simulations, their parameters, and the conditions they satisfy are summarised in tables 7 and 8. All of the simulations are done with $Q = 300$ pC, since preliminary tests showed that higher charges, above ~ 5 nC, could not preserve density modulations for long enough.

$E = 44$ MeV, $\lambda_r = 340$ μm , $\gamma_0 = 11.2$	$\lambda/c = 0.4$ ps, $\lambda = 120$ μm , $\sigma_{\text{diff}} = 510$ μm	$\lambda/c = 2$ ps, $\lambda = 600$ μm , $\sigma_{\text{diff}} = 1.1$ mm
$\sigma_r = 340$ μm	$\lambda \ll \lambda_r, \sigma_r < \sigma_{\text{diff}}$	$\lambda > \lambda_r, \sigma_r \ll \sigma_{\text{diff}}$
$\sigma_r = 680$ μm	$\lambda \ll \lambda_r, \sigma_r > \sigma_{\text{diff}}$	$\lambda > \lambda_r, \sigma_r < \sigma_{\text{diff}}$
$\sigma_r = 1.8$ mm	$\lambda \ll \lambda_r, \sigma_r \gg \sigma_{\text{diff}}$	$\lambda > \lambda_r, \sigma_r > \sigma_{\text{diff}}$

Table 7: Conditions satisfied for each set of parameters in the simulations of the wiggler at 44 MeV. The case coloured in blue is the only one that, à priori, satisfies the constraints from Geloni's paper [4] such that the wiggler effects are the same as in a low-energy drift.

$E = 63$ MeV, $\lambda_r = 166$ μm , $\gamma_0 = 16$	$\lambda/c = 0.4$ ps, $\lambda = 120$ μm , $\sigma_{\text{diff}} = 510$ μm	$\lambda/c = 2$ ps, $\lambda = 600$ μm , $\sigma_{\text{diff}} = 1.1$ mm
$\sigma_r = 340$ μm	$\lambda < \lambda_r, \sigma_r < \sigma_{\text{diff}}$	$\lambda \gg \lambda_r, \sigma_r \ll \sigma_{\text{diff}}$
$\sigma_r = 680$ μm	$\lambda < \lambda_r, \sigma_r > \sigma_{\text{diff}}$	$\lambda \gg \lambda_r, \sigma_r < \sigma_{\text{diff}}$
$\sigma_r = 1.8$ mm	$\lambda < \lambda_r, \sigma_r \gg \sigma_{\text{diff}}$	$\lambda \gg \lambda_r, \sigma_r > \sigma_{\text{diff}}$

Table 8: Conditions satisfied for each set of parameters in the simulations of the wiggler at 63 MeV. The case coloured in blue is the only one that, à priori, satisfies the constraints from Geloni’s paper [4] such that the wiggler effects are the same as in a low-energy drift.

The resulting phase-spaces and particle density distributions after the wiggler can be seen in figures 23 and 24. All simulations yielded some energy modulation, which was to be expected. More importantly, if we compare these phase-space plots to those obtained by simulating a 1.1 m drift with the same beam characteristics, we see that in every case the energy modulation is stronger in the wiggler than in the drift (fig. 25).

As one might expect, larger transverse sizes yield smaller modulations due to the reduction in charge-density, and larger λ also gives larger modulation since each pulse in the bunch is almost only interacting with itself. The simulations at 44 MeV and 63 MeV show almost the same peak-to-peak energy modulations, but the lower energy case has a higher relative energy modulation $\Delta E/E$, larger by a factor of 1.4, which would give stronger cooling since the chicane effects depend on relative energy difference.

The amount of dispersion in the wiggler is also noticeable, especially for the lower energy simulations. The R_{56} of the whole wiggler is given by

$$R_{56} \simeq 2N_u \lambda_r, \quad (16)$$

thus ~ 7 mm at 44 MeV, and ~ 3.5 mm at 63 MeV. Additionally, the delay in time and position gained over a certain distance depends on the relative energy spread, and is given by

$$\Delta t = \frac{1}{c} \frac{\Delta\gamma}{\gamma} R_{56}, \quad (17)$$

which yields a time delay Δt almost three times larger for the simulations at 44 MeV than those at 63 MeV. It is unclear whether this effect is detrimental or not. If the bunch already has an energy modulation when entering the undulator, the large R_{56} could destroy the density modulation before it can be enhanced. However, since the MBEC scheme has a dispersive section following the drift/wiggler, there might be some way to profit from the dispersion already present in the wiggler.

When looking at the charge density distribution from the plots, it is clear that at lower transverse sizes the distribution becomes asymmetric between back and front of the bunch, especially for the simulations at 44 MeV. In a drift where only space-charge forces are present, we would expect the bunch length to grow symmetrically, as the energy gain in the bunch head should be exactly opposite to the energy loss in the bunch tail, and the bunch centre (the peak of the Gaussian distribution) should stay in the same position relative to the bunch. From this we conclude that the asymmetry comes from the radiation energy loss that is happening in the bunch centre,

where the current is the highest. As predicted by the theory, the larger transverse sizes where $\sigma_r > \sigma_{\text{diff}}$ show negligible radiation effects, as the bunch is still longitudinally symmetric after the wiggler.

The paper by Geloni et al. gives a formula (equation 94 from [4]) for calculating the gain in energy modulation over a certain distance when the four conditions mentioned above are satisfied. In our simulations, this is only satisfied by the two cases with $\lambda/c = 2$ ps and $\sigma_r = 1.8$ mm (coloured in blue in tables 7 and 8).

If we apply the formula, we find $\Delta E = 330$ KeV for the 44 MeV simulation, and $\Delta E = 162$ KeV for the 63 MeV simulation. Both of these results underestimate the energy modulation, which in the simulations is $\Delta E \simeq 500$ KeV for both energies. However, this could be due to the definition that we have chosen for λ . Here we are using λ to describe the separation between Gaussian pulses in the bunch, but in Geloni's paper it is $\lambda = \sigma_z$. In addition the formula given assumes a Gaussian transverse and longitudinal distribution. With these points in mind we could argue that in fact the simulations closest to the constraints from the paper are the two with $\lambda/c = 0.4$ ps and $\sigma_r = 1.8$ mm, since they are more similar to a Gaussian single bunch, and have a total length of $l_z \simeq 600$ μm , larger than the radiation wavelength λ_r . If we apply the formula to these two cases, and take $\lambda = \sigma_z = \frac{600}{2.355}$ μm , we find $\Delta E = 780$ KeV for the 44 MeV simulation, and $\Delta E = 380$ KeV for the 63 MeV simulation. Both these results are closer to the $\Delta E \simeq 550$ KeV seen in the plots, measured from head to tail of the bunch as if it were a single Gaussian.

For these two plots, that seem to better agree with the theory, we can use the same formula to estimate how long a drift should be in order to obtain the same difference in energy as in the wiggler. The formula predicts that for a 1.1 m drift (same length as wiggler), the energy modulation will be 15 and 8.5 KeV, for the 44 and 63 MeV simulations respectively. It also predicts, that a drift of ~ 50 m would be necessary to achieve the same modulations as the wiggler. We can compare these numbers by simulating these bunches through 1.1 metre drifts, as shown in figure 26. The head-to-tail modulations after the drift are approximately 80 KeV for the simulation at 44 MeV, and 40 KeV for the simulation at 63 MeV. This is much larger than the calculated values, but still one order of magnitude smaller than the energy modulations from the wiggler. Again, this formula is derived on the assumption of Gaussian distributions, which is not exactly the case here.

$E = 44 \text{ MeV}$,
 $\lambda_r = 340 \text{ }\mu\text{m}$,
 $\gamma_0 = 11.2$

$\lambda/c = 0.4 \text{ ps}$,
 $\lambda = 120 \text{ }\mu\text{m}$,
 $\sigma_{\text{diff}} = 510 \text{ }\mu\text{m}$

$\lambda/c = 2 \text{ ps}$,
 $\lambda = 600 \text{ }\mu\text{m}$,
 $\sigma_{\text{diff}} = 1.1 \text{ mm}$

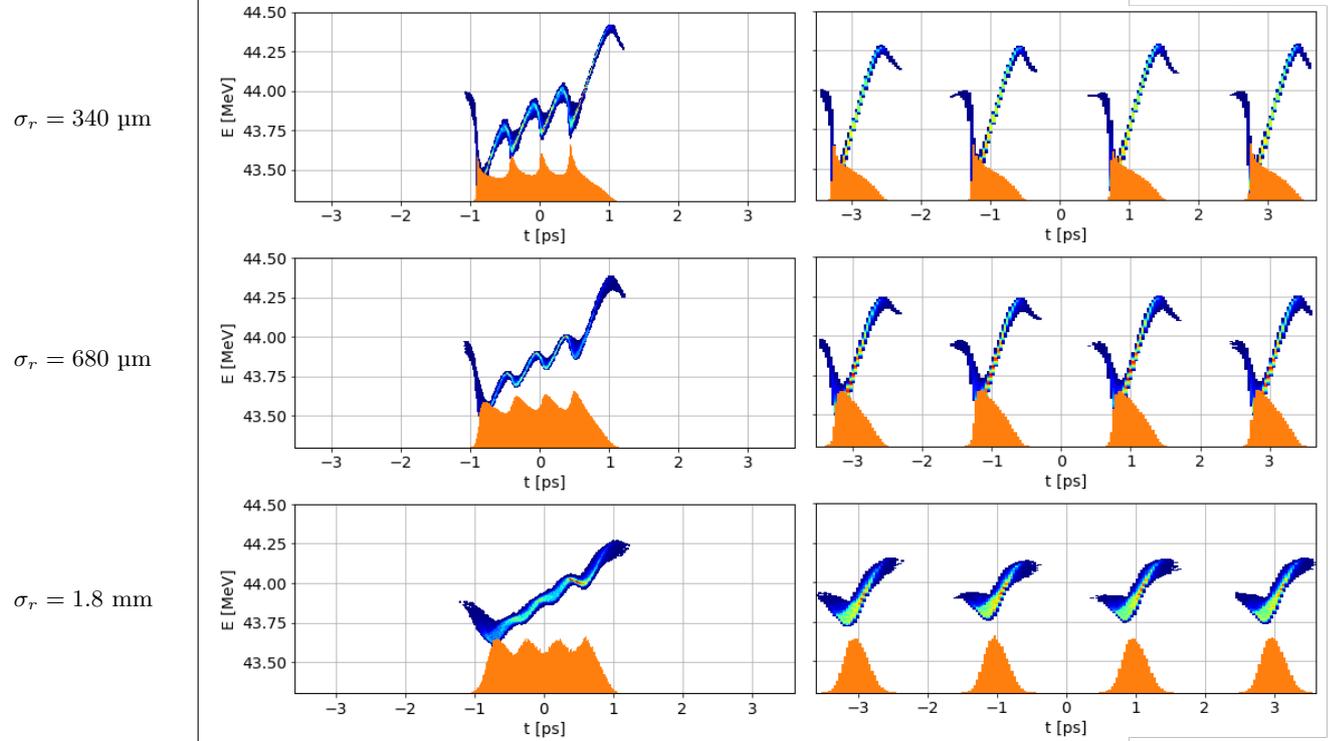


Figure 23: Longitudinal phase-space and density distribution after the wiggler for simulations at 44 MeV. The bunch head is on the right side of the plots.

$E = 63$ MeV,
 $\lambda_r = 166$ μm ,
 $\gamma_0 = 16$

$\lambda/c = 0.4$ ps,
 $\lambda = 120$ μm ,
 $\sigma_{\text{diff}} = 510$ μm

$\lambda/c = 2$ ps,
 $\lambda = 600$ μm ,
 $\sigma_{\text{diff}} = 1.1$ mm

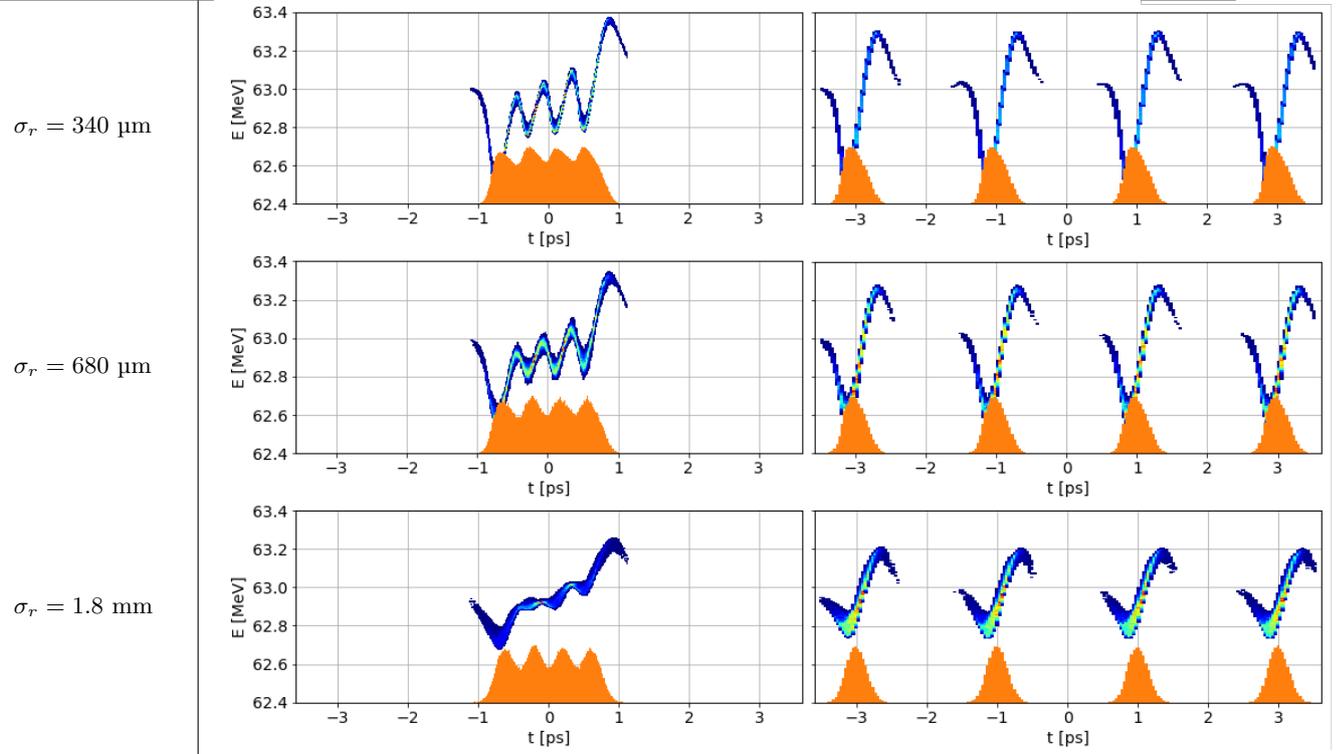


Figure 24: Longitudinal phase-space and density distribution after the wiggler for simulations at 63 MeV. The bunch head is on the right side of the plots.

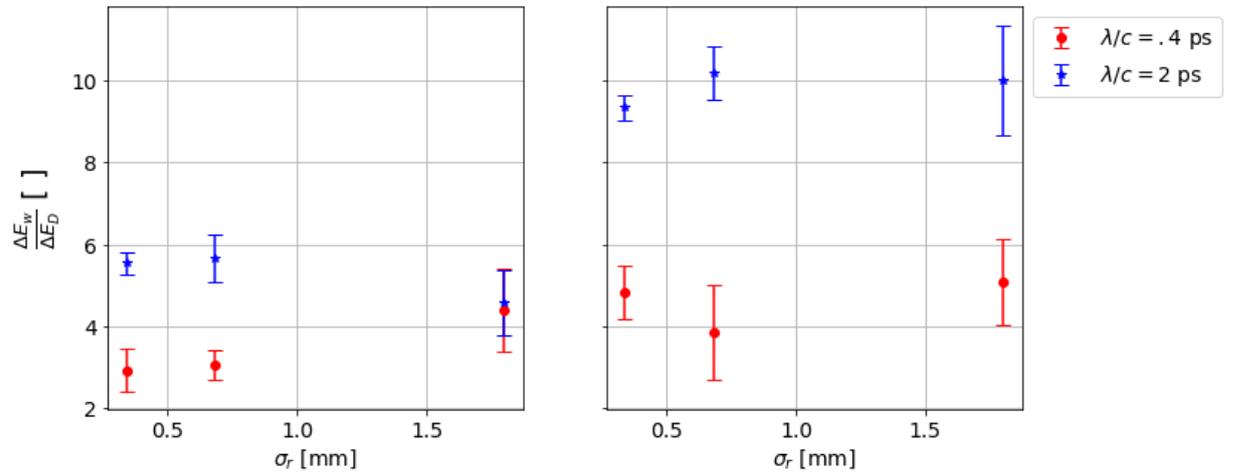


Figure 25: Ratio between the energy modulation at the centre of the bunch induced in a wiggler and in a drift of the same length, at 44 MeV (left) and at 63 MeV (right). The individual phase-space comparisons can be found in the appendix.

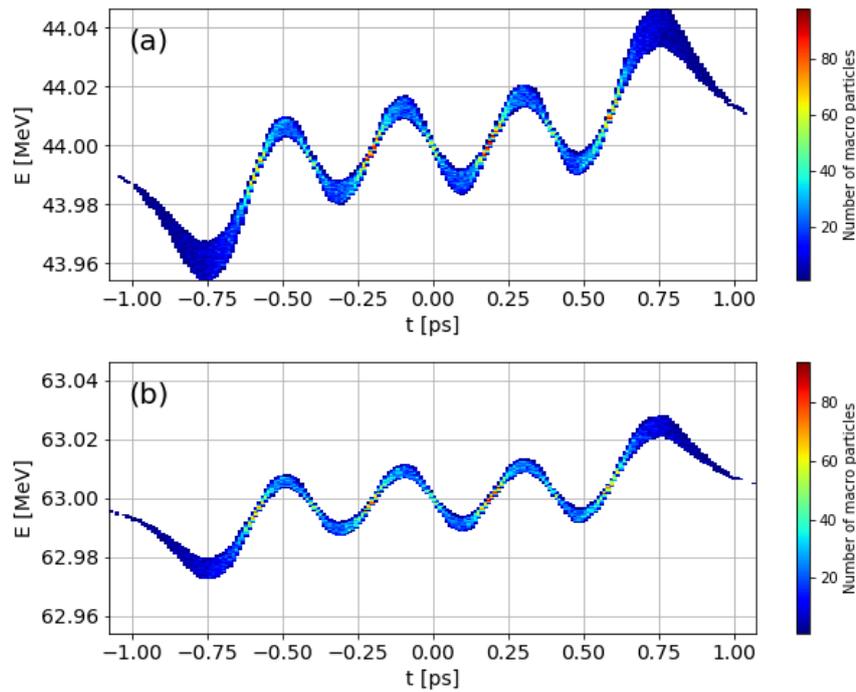


Figure 26: Two bunches with $\lambda/c = 0.4 \text{ ps}$ and $\sigma_r = 1.8 \text{ mm}$ simulated through a 1.1 m drift, to be compared with the results from the wiggler. Case (a) corresponds to $E = 44 \text{ MeV}$, and case (b) to $E = 63 \text{ MeV}$.

To see the effect of the *reduced* gamma in the simulations, we can compare them to a 1.1 m drift simulation, in which the bunch has a lower energy $E = mc^2\gamma_0$. A few examples can be seen in figure 27, and the rest can be found in the appendix. When comparing the longitudinal phase-space, it seems that only the energy modulation should be compared, as the length of the bunch is drastically different with and without the wiggler. This is most likely explained by the fact that at such low energies, energy modulations quickly drift from each other as if they were in a highly dispersive section. If we compute the time difference between two particles in the low-energy 1 metre drift, we get

$$\Delta t = \frac{1}{c} \frac{\Delta\gamma}{\gamma_0} 2N_u \lambda_r, \quad (18)$$

which is over 7 times larger than the dispersion in the wiggler. Hence if the same energy modulation is induced in the wiggler as in the reduced-gamma drift, the drift will quickly have its density modulation washed out through dispersion. This means that effects of a wiggler aren't exactly comparable to those of a low-energy drift, but could be considered better in some way: the bunch in the wiggler benefits of wake-fields that generate large energy modulations as if it were propagating in free-space at a lower energy, but can preserve the micro-bunching longer than a low-energy drift.

In order to have a better overview of the comparison between wiggler and low-energy drift, figure 28 shows the energy modulation at the centre of the bunch, for these two cases. We see that in most of the simulations the ratio between energy modulation in a wiggler and in a low-energy drift is close to 1, except for a few cases where the wiggler's energy modulation is an order of magnitude higher. These cases are those at 44 MeV and $\lambda/c = .4$ ps. As mentioned above, in the cases where λ is small, the comparison to the theory is more adequate if the micro-bunches are considered as a single long Gaussian bunch. In figure 29 we again compare the energy modulation from a wiggler and a low-energy drift, but in this case we use the total minimum-to-maximum energy difference over the whole bunch, as if there were no micro-bunches. In this case we see that all wiggler simulations show a similar energy modulation to a drift with γ_0 , and in addition we see that for the simulations at 44 MeV, the increase in beam radius reduces the ratio, and brings it closer to 1. This last remark shows the dependence of the radiation on the transverse size of the beam. When the beam becomes wide enough, radiation can be neglected.

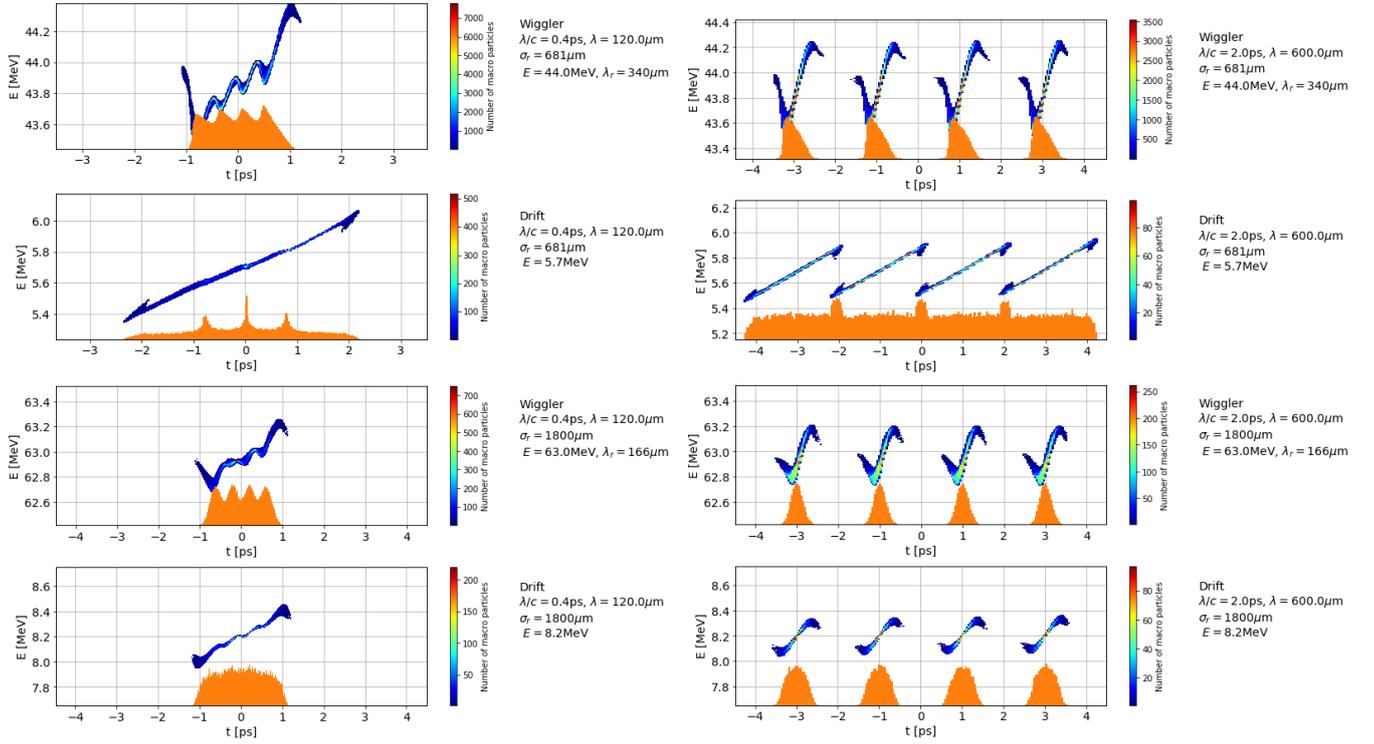


Figure 27: Comparison between a wiggler simulation and a 1.1 m drift at reduced energy to imitate the wiggler effects. The comparison for the rest of simulations can be found in the appendix.

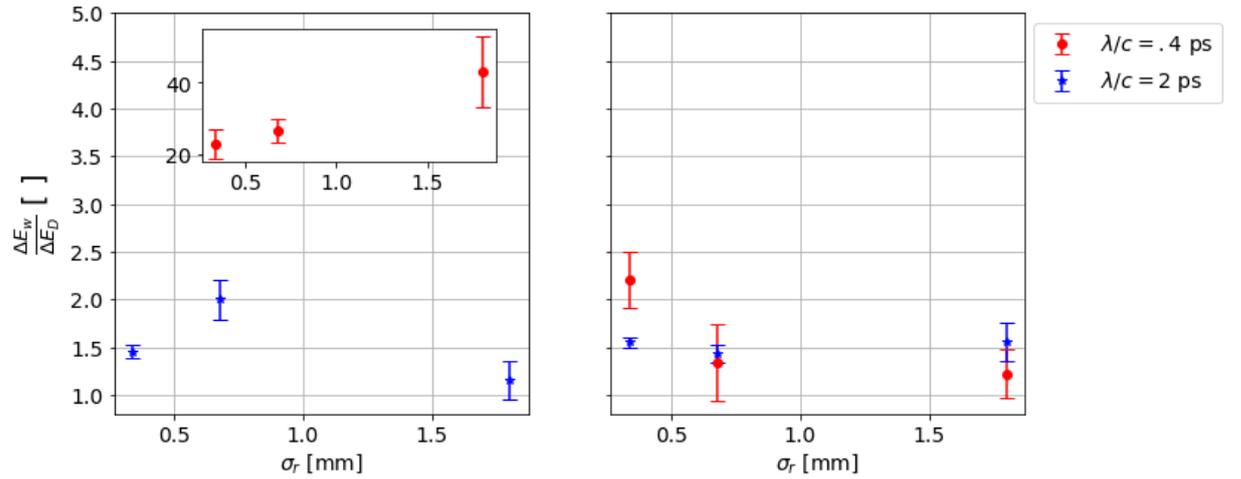


Figure 28: Comparison of energy modulation at the centre of the bunch between simulations through a wiggler and through a low-energy drift, at 44 MeV (left), and at 63 MeV (right).

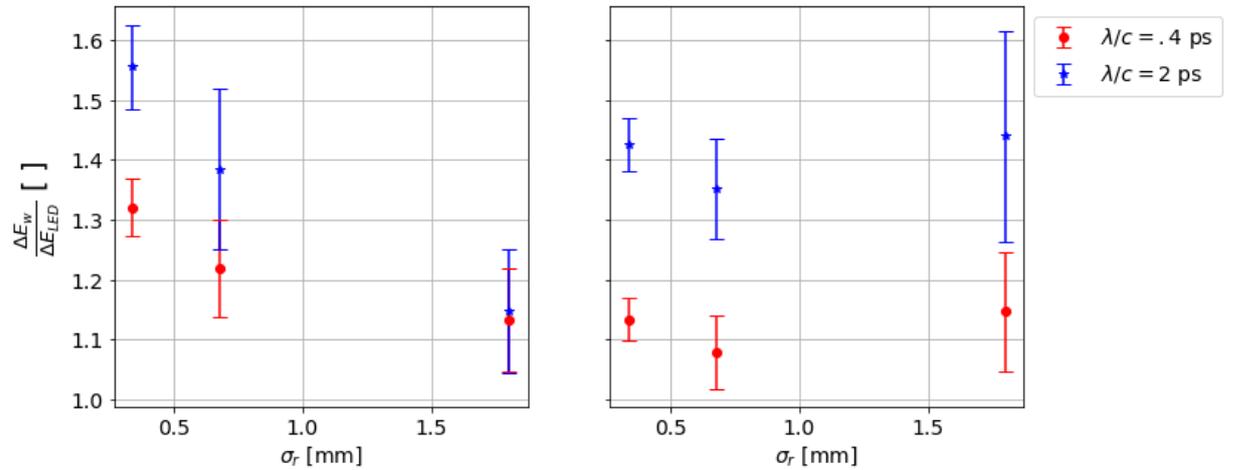


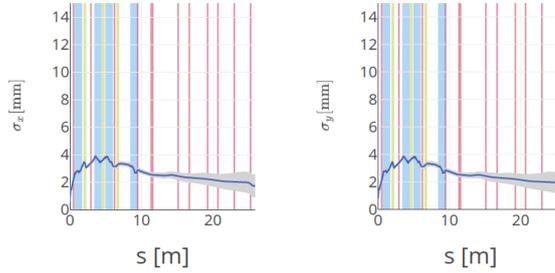
Figure 29: Comparison of the energy variation over the whole bunch between simulations through a wiggler and through a low-energy drift, at 44 MeV (left), and at 63 MeV (right).

3.3.2 Surrogate Model by R. Bellotti

This thesis has been carried out in close collaboration with another MSc student, who developed an excellent surrogate model for the first 26 metres of the AWA beamline, from gun to wiggler entrance [20]. This model was trained on thousands of OPAL simulations, and now can predict quantities of interest in a matter of seconds. He made a dashboard (fig. 30) that takes as input the settings described in table 5, and predicts the transverse size, emittance, energy and energy-spread, and momentum-coordinate correlations for every point along the beamline.

The main advantage of this model is that it can very quickly run a genetic algorithm and find an optimal beam configuration under given constraints. This proved very useful when specific beam parameters were required at the wiggler or some other point in the beamline.

Beam sizes



IBF [A]	450	[450, 550]	IM [A]	100	[100, 260]
GPHASE [°]	-50	[-50, 10]	ILS1 [A]	0	[0, 250]
ILS2 [A]	0	[0, 200]	ILS3 [A]	0	[0, 200]
Bunch charge [nC]	0.3	[0.3, 5]	lambda [ps]	0.3	[0.3, 2]
Laser radius [mm]	1.5	[1.5, 12.5]			

Quantile of the residuals to use as uncertainty:

0.95

cavity

YAG

solenoid

Figure 30: Dashboard running the surrogate model that can quickly find the required accelerator settings.

3.3.3 Start-to-End Simulations

We finally show some results of start-to-end simulations. These were all done as explained in figure 4; each example is a single simulation from the gun up to a monitor behind the wiggler, and OPAL automatically switches between solvers when necessary. In the examples that follow, the surrogate model was used to find an optimal configuration that would give the smallest possible energy-spread of the bunch when reaching the wiggler. This was necessary because it was noticed that, as the beam with initial density modulation was propagated to the wiggler, it acquired an energy chirp due to longitudinal space-charge forces. This chirp could “wash-out” the density modulation before reaching the wiggler, or right when entering the wiggler due to the combination of a large energy spread and a large R_{56} . Figure 31 shows one of these examples in which the longitudinal space-charge at low energies caused the density modulation to vanish. This figure is from an initial test simulation before the charge or number of pulses in the bunch had been decided.

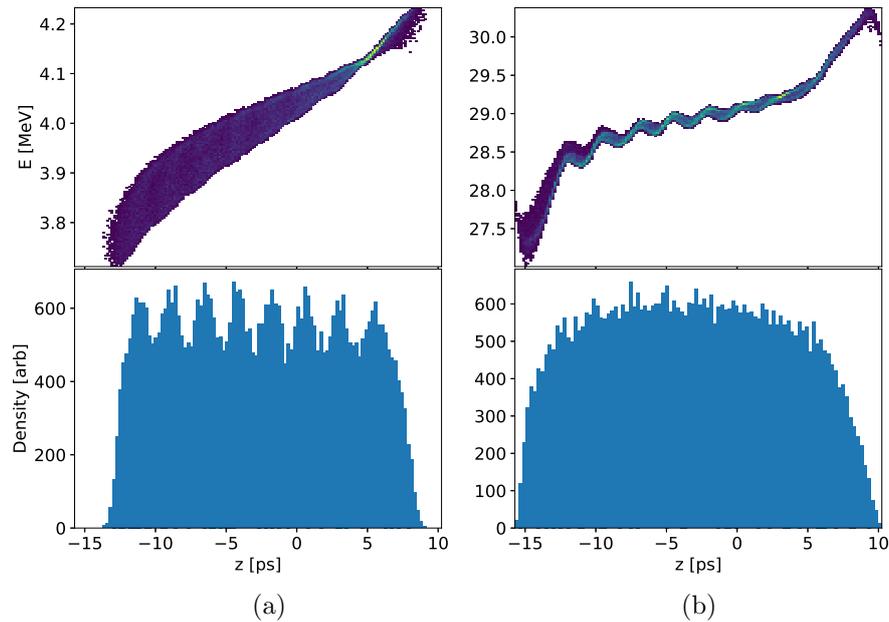


Figure 31: *The bunch longitudinal phase-space and density distribution, as measured on two screens at the AWA beamline: (a) right after the gun $s = .11$ m and (b) at $s = 5.22$ m in between linacs 2 and 3. This was an initial test with 8 pulses separated by 2 ps and a high charge of 40 nC. Both parameters were subsequently reduced for the final simulations.*

Figures 32 and 33 show three sets of parameters that were chosen from the surrogate model's optimisation. Two of the cases use 4 cavities and reach an energy $E \simeq 44$ MeV, and the other one runs with 6 cavities and reached over 60 MeV. The two 4-cavity cases only differ slightly in the emittance and beam-size when reaching the wiggler. No quadrupole focusing was used in these simulations, only solenoids provided focusing. The parameters used for the solvers in these simulations are displayed in table 9, and the three sets of parameters defining each simulation in table 10.

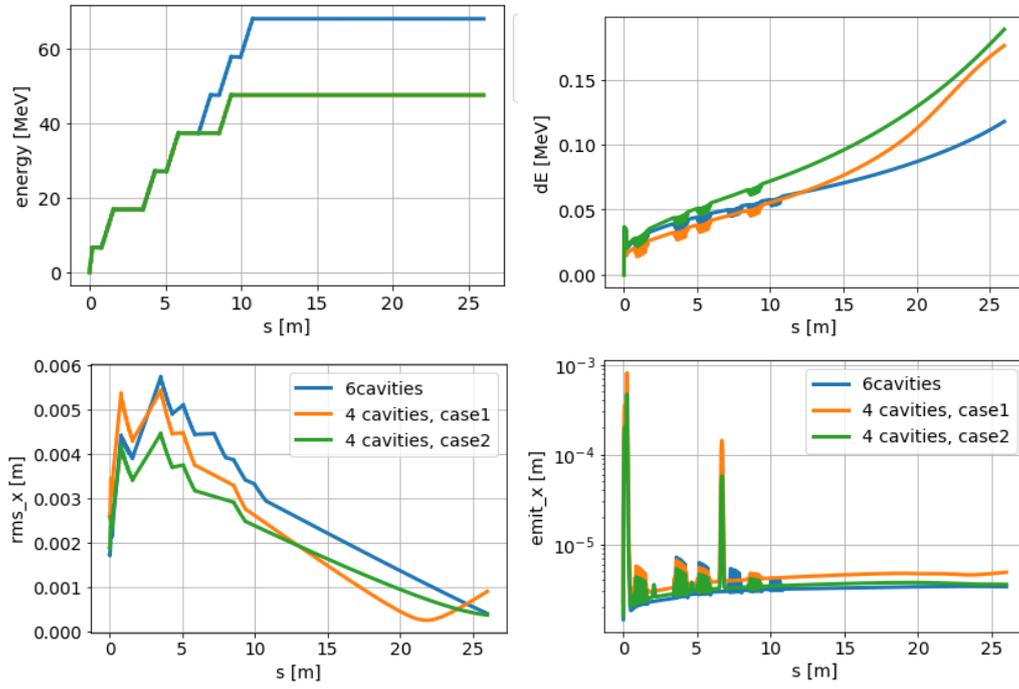


Figure 32: *Beam statistics along the beamline, with parameters chosen such that the energy spread remains low.*

N	$1.7 \cdot 10^7$
(MX, MY, MT)	$(64, 64, 512) \mu\text{m}$
DT	10^{-12}
$(\Delta x, \Delta y, \Delta z)$	$(40, 40, 2) \mu\text{m}$
(L_x, L_y, L_z)	$(14, 14, 7) \text{mm}$

Table 9: Parameters used for OPAL's solvers for the start-to-end simulations.

	4 cavs, case 1	4 cavs, case 2	6 cavs
Q [nC]	0.3	0.3	0.3
λ [ps]	2.0	2.0	2.0
σ_r [mm]	5.24	3.81	3.50
IBF [A]	512	511	500
IM [A]	150	150	120
GPHASE [°]	6.4	8.5	7.0
ILS1 [A]	0.0	0.4	0.0
ILS2 [A]	0.0	0.3	0.0
ILS3 [A]	26.0	14.3	0.0

Table 10: Beamline settings used for the 4 and 6 cavity start-to-end simulations from figures 32 and 33.

All three examples show an increase in energy modulation after the wiggler. As with the simulations on the wiggler alone, we see that the lower energy cases, in addition to the energy modulation, also show an asymmetry in the density distribution due to the stronger dispersion. The dependence on the transverse size can also be observed, as figures 33(b) and 33(c) have very similar beam parameters, their main difference being the transverse size upon entering the undulator (see fig. 32), and they accordingly show different increases in energy modulation.

These are promising results that suggest that the experiments at AWA will be able to observe the differences in energy modulation between a drift and a wiggler.

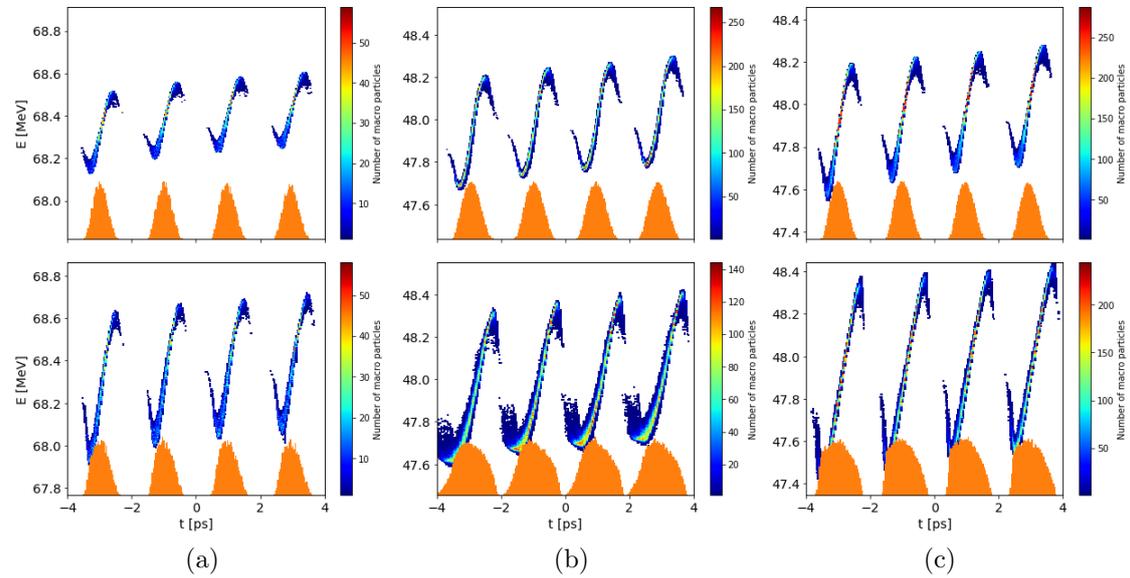


Figure 33: Longitudinal phase-space and density distribution on a YAG monitor before (plots at the top) and after (plots at the bottom) the wiggler. (a) refers to the simulation at 66 MeV, (b) to the simulation at 44 MeV case 1, and (c) to the simulation at 44 MeV case 2,

4 Conclusions and Future Work

In the first part of this thesis the codes OPAL and MITHRA were successfully merged into a new full-wave solver that OPAL can use for wigglers and undulators. A convergence analysis shows that the error can be decreased to an amount negligible for the experiments in this thesis, but we did not see the error converge. For convergence we expect that a finer grid and a higher number of particles are required, increasing the computational load significantly, as this solver is more expensive than the static one. For this reason it is essential that OPAL can switch back and forth between solvers when necessary. At the moment the static to full-wave transition is implemented, and full-wave to static is in development and will be available in the near future. A benchmark with the code OSIRIS and measurements at LCLS from [21] showed agreement regarding the phase-space modulation. From the start-to-end simulations of the AWA experiment, we conclude that OPAL can now simulate full FELs, including the linear accelerator, compression stages, and undulators in one single simulation. In addition, MITHRA and OPAL can accurately model a wider regime of beams and wigglers than most other FEL codes which use approximations to reduce computational load.

The code was used in the analysis of the upcoming experiments at AWA, where the feasibility of using wigglers in the MBEC scheme for cooling hadrons and ions will be assessed. Simulations on the wiggler with different beam regimes all showed a conversion from density to energy modulation stronger than in a drift, reaching in some cases an energy difference 10 times stronger in the wiggler than in free-space propagation. From the literature we can estimate that the energy modulation achieved in a 1 metre wiggler at these energies would require ~ 50 metres to be achieved in a drift. As predicted by the theory, radiation effects become negligible for wide beams, where the transverse size is larger than the radiation diffraction size. These cases could be compared to the energy modulation appearing in a low-energy drift, although the values did not exactly match. This is due to the fact that the Gaussian pulses forming the bunches considered in this experiment are shorter than the radiation wavelength, and thus don't entirely satisfy the constraints from the theory. The bunches with shorter overall length and large transverse size were found to more closely resemble a low-energy drift, since the internal micro-bunches were close enough that the whole bunch was similar to a single Gaussian. In thinner pencil-like beams, radiation lead to an energy-loss in peak current parts of the bunch, which enhances the energy modulation and yields better results that one would get from a low-energy drift. Measuring the emitted radiation from these simulations could give further insight into how it affects the modulation.

Finally, start-to-end simulations showed that the beam parameters at emission from the gun need to be chosen carefully in order to preserve the density modulation up to the wiggler, and such that the energy spread is not too large. A surrogate model of the beamline developed at PSI proved extremely useful when finding optimal beam parameters that could maintain density modulation at low energy spreads up until the wiggler. The full simulations of the accelerator show promising results, where a measurable increase in energy modulation from the wiggler can be observed. As expected from the simulations on the wiggler, transverse size plays an important role in the effects of the radiation, and thus focusing through solenoids and quadrupoles will be necessary in the experiment in order to get measurable changes in energy. Further start-to-end simulations are required before the experiment, if possible including focusing from quadrupoles.

Bibliography

- [1] A. Accardi et al. Electron-Ion Collider: The next QCD frontier. *The European Physical Journal A*, 52(9):268, 2016.
- [2] D. Ratner. Microbunched Electron Cooling for High-Energy Hadron Beams. *Physical Review Letters*, 111(084802), 2013.
- [3] E. L. Saldin, E. A. Schneidmiller, and M. V. Yurkov. Radiative Interaction of Electrons in a Bunch Moving in an Undulator. *Nucl. Instrum. Meth.*, A(417):158–168, 1998.
- [4] G. Geloni, E. Saldin, E. Schneidmiller, and Mikhail Yurkov. Longitudinal Impedance and Wake from XFEL Undulators. Impact on Current-Enhanced SASE Schemes. *Nucl. Instrum. Meth.*, pages 228–247, 2007.
- [5] A. Zholents et al. Wiggler Impact on the Electron Beam Self Energy Modulation in View of the Electron-Ion Collider R&D. *AWA Needs and Opportunities Workshop*, 2019.
- [6] S. Reiche. *Numerical Studies for a Single Pass High Gain Free-Electron Laser*. PhD thesis, Universität Hamburg, Hamburg, 1999.
- [7] E. L. Saldin, E. A. Schneidmiller, and M.V. Yurkov. FAST: a three-dimensional time-dependent FEL simulation code. *Nucl. Instrum. Meth.*, A(429):233–237, 1999.
- [8] T. M. Tran and K. S. Furtele. TDA - A Three-Dimensional Axisymmetric Code for Free-Electron-Laser (FEL) Simulations. *Computer Physics Communications*, 54:263–272, 1989.
- [9] A. Adelman et al. OPAL a Versatile Tool for Charged Particle Accelerator Simulations. 2019. [arXiv:1905.06654](https://arxiv.org/abs/1905.06654).
- [10] A. Fallahi, A. Yahaghi, and F. Kärtner. MITHRA 1.0: A Full-Wave Simulation Tool for Free Electron Lasers. *Computer Physics Communications*, 228:192–208, 2018.
- [11] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, 1962.
- [12] J. P. Boris. Relativistic Plasma Simulation-Optimization of Hybrid Code. In *Proceedings of Fourth Conference on Numerical Simulations of Plasmas*, 1970.
- [13] C. Pellegrini. The Free-electron Laser Collective Instability and the Development of X-ray FELs. In *Proceedings of the 2001 Particle Accelerator Conference*, Los Angeles, CA, USA, 2001.

- [14] A. Marinelli. Introduction to the Physics of Free Electron Lasers. [http://web.stanford.edu/group/pulse_institute/static/uxss2016/UXSS2016-Marinelli\(1\).pdf](http://web.stanford.edu/group/pulse_institute/static/uxss2016/UXSS2016-Marinelli(1).pdf), 2016. Accessed 30.05.2020.
- [15] B. Archambeault, O. M. Ramahi, and C. Brench. *EMI/EMC Computational Modeling Handbook*. Springer, Boston, MA, 1998. Chapter: The Finite-Difference Time-Domain Method.
- [16] R. Courant, K. Friedrichs, and H. Lewy. On the Partial Difference Equations of Mathematical Physics. *IBM Journal*, 1(2):215–234, 1967.
- [17] W.M. Fawley and J.-L. Vay. Use of the Lorentz-Boosted Frame Transformation to Simulate Free-Electron-Laser Amplifier Physics. *AIP Conference Proceedings*, 1086:346, 2009.
- [18] J.-L. Vay. Noninvariance of Space- and Time-Scale Ranges under Lorentz Transformation and the Implications for the Study of Relativistic Interactions. *Physical Review Letters*, 98(130405), 2007.
- [19] A. Fallahi. MITHRA 2.0, A Full-Wave Simulation Tool for Free Electron Lasers. <https://github.com/aryafallahi/mithra>. Accessed on 15.05.2020.
- [20] R. Bellotti. Accelerating Accelerators: Fast Surrogate Models for Beam Prediction. MSc Thesis, ETHZ, 2020.
- [21] J. P. MacArthur et al. Phase-Stable Self-Modulation of an Electron-Beam in a Magnetic Wiggler. *Physical Review Letters*, 123(214801), 2019.
- [22] E. L. Saldin, E. A. Schneidmiller, and M. V. Yurkov. Self-Amplified Spontaneous Emission FEL with Energy-Chirped Electron Beam and its Application for Generation of Attosecond X-ray Pulses. *Physical Review ST Accelerators and Beams*, 9(050702), 2006.
- [23] A. Zholents. Method of an Enhanced Self-Amplified Spontaneous Emission for X-ray Free Electron Lasers. *Physical Review ST Accelerators and Beams*, 8(040701), 2005.
- [24] R. A. Fonseca et al. OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for modeling Plasma Based Accelerators. In *Proceedings of the International Conference on Computational Science-Part III*, pages 342–351, 2002.
- [25] Ya. S. Derbenev. On Possibilities of Fast Cooling of Heavy Particle Beams. In *AIP Conference Proceedings*, volume 253, page 103, 1992.
- [26] The electron-ion collider. "<https://www.bnl.gov/eic/>". Accessed: 03.05.2020.
- [27] G. Stupakov and P. Baxevanis. Microbunched Electron Cooling with Amplification Cascades. *Physical Review Accelerators and Beams*, 22(034401), 2019.
- [28] J. G. Power and Chunguang Jing. Temporal Laser Pulse Shaping for RF Photocathode Guns: The Cheap and Easy way using UV Birefringent Crystals. In *AIP Conference Proceedings*, volume 1086, page 689, 2009.



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

START-TO-END MODELLING OF THE AWA
MICROBUNCHED ELECTRON COOLING POP-EXPERIMENT

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):

Albà Jacas

First name(s):

Arnau

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the ['Citation etiquette'](#) information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date

Zürich, 30.05.2020

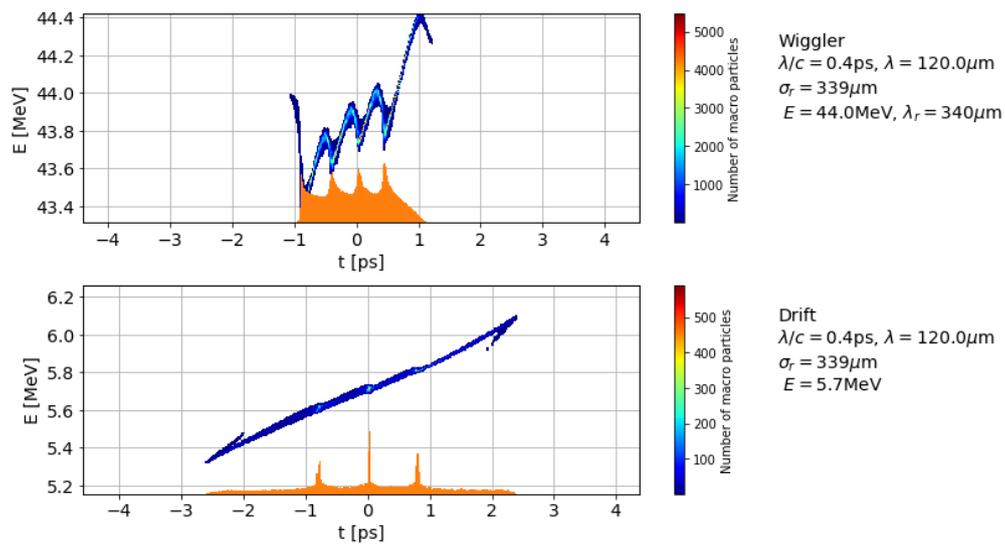
Signature(s)

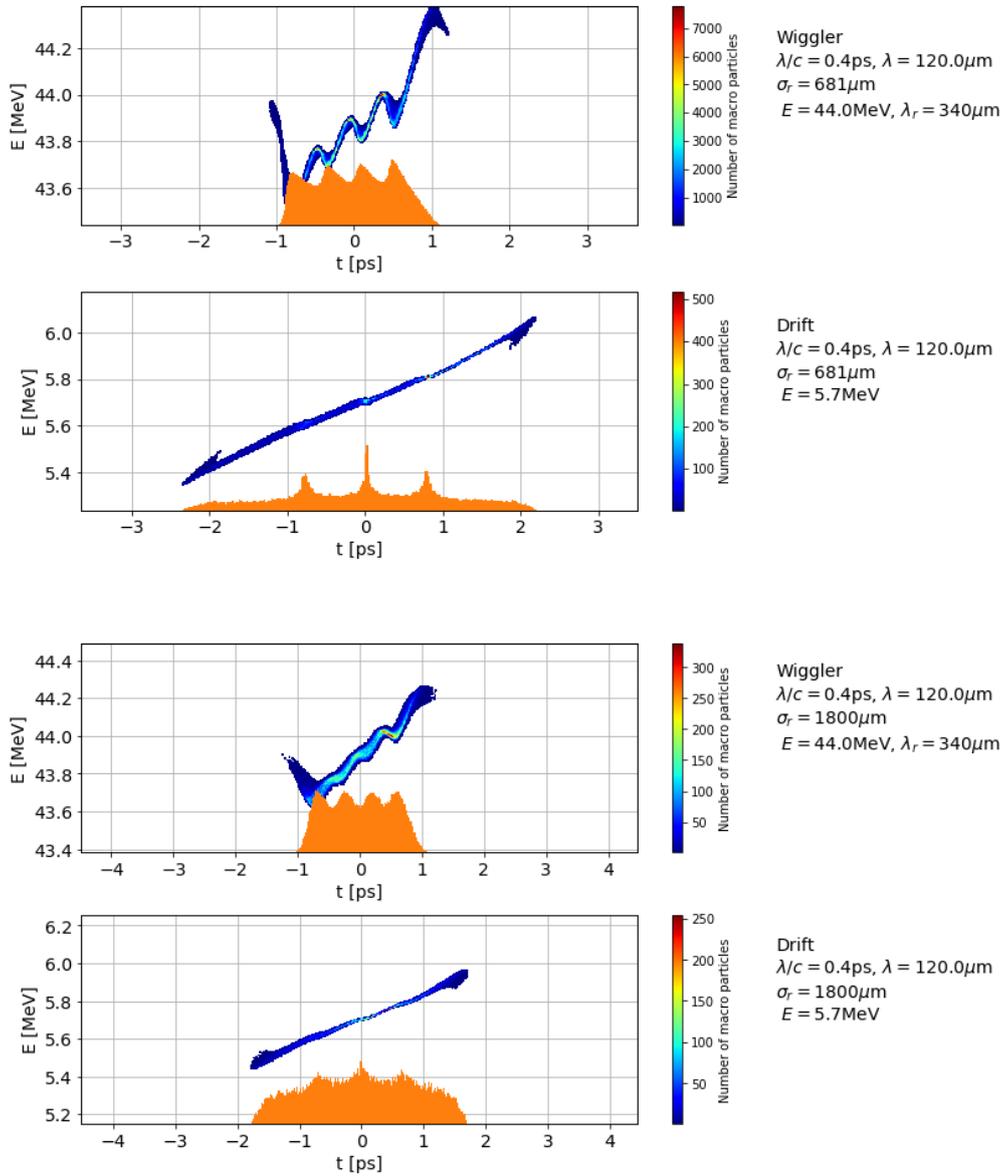
For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.

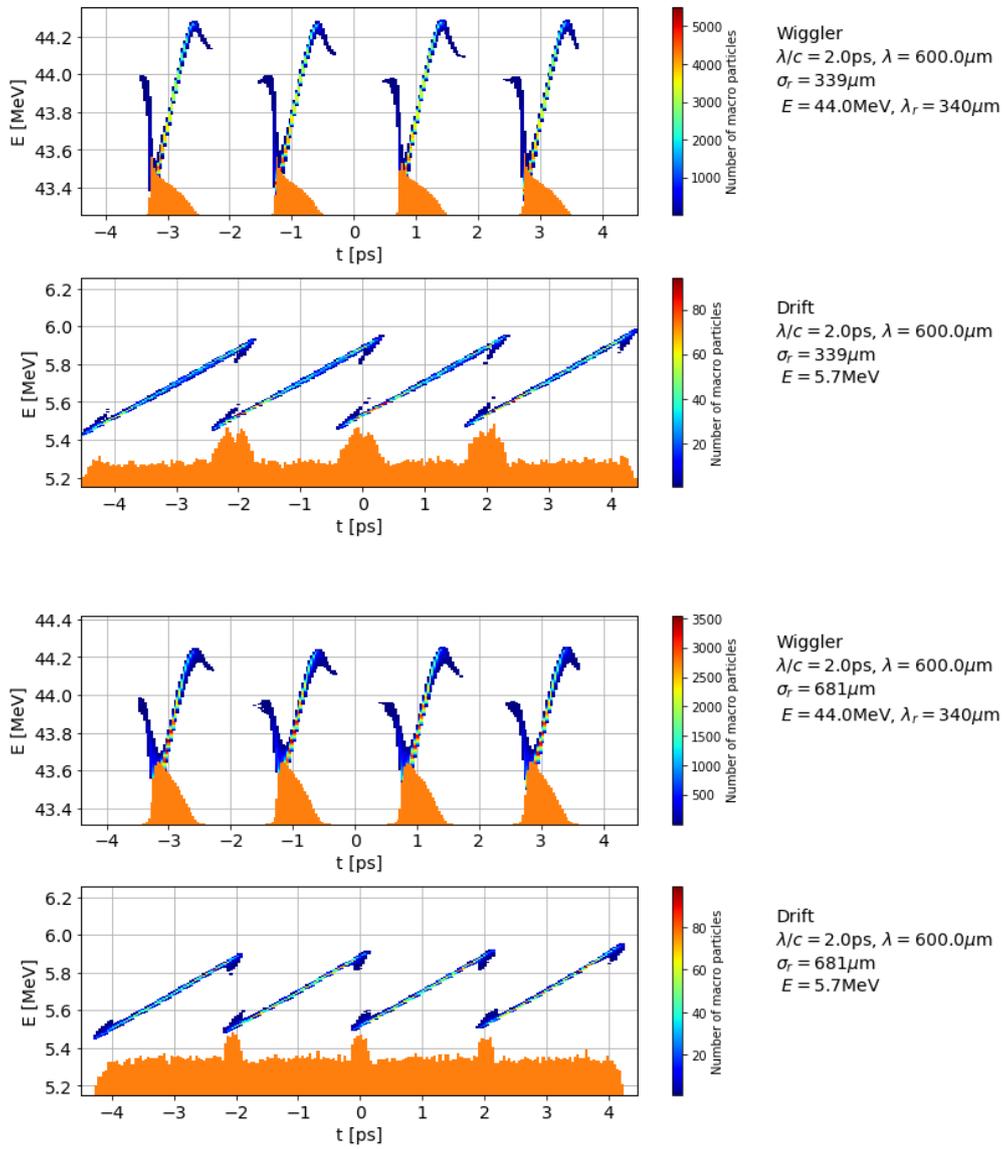
Appendix A

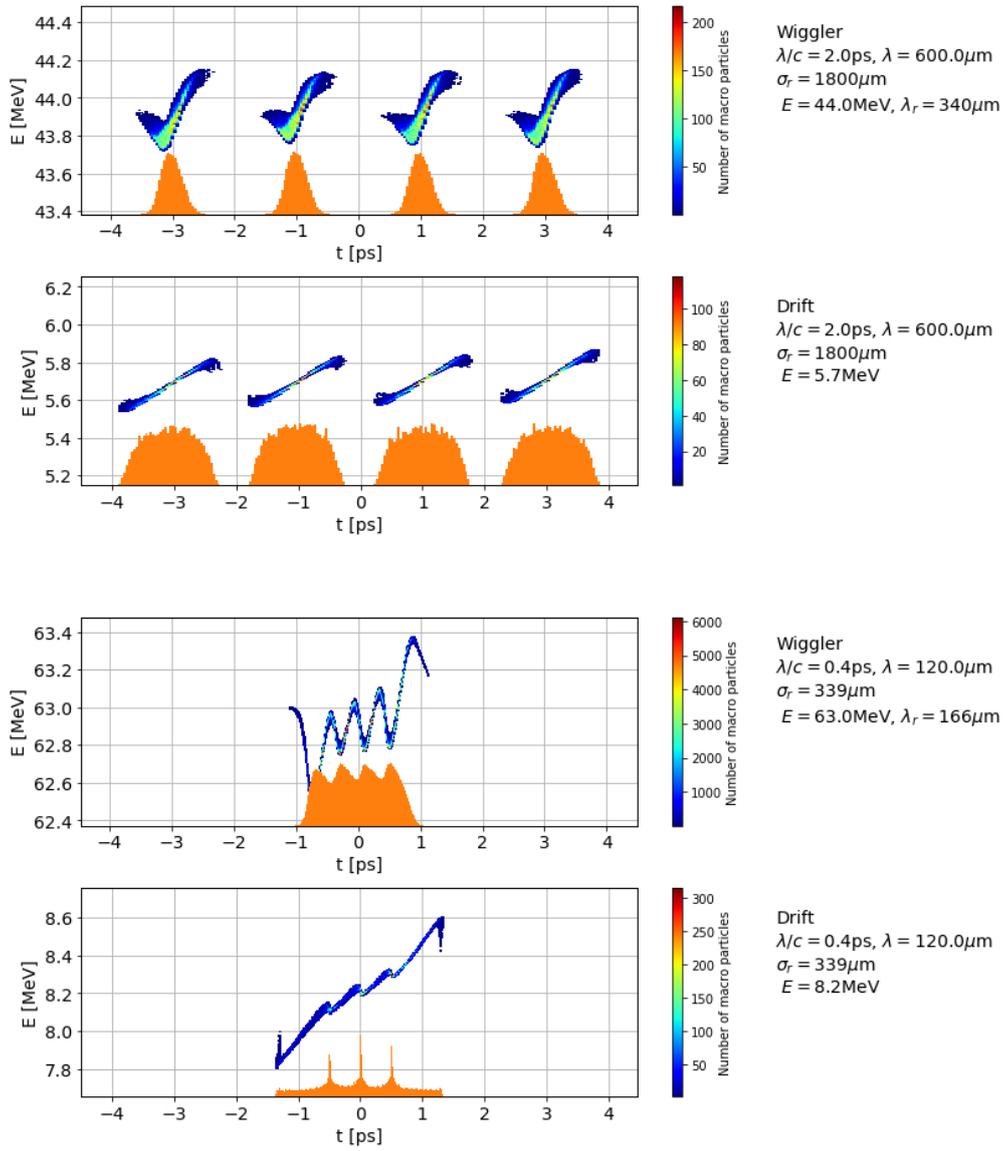
Comparison Wiggler Simulations and Drift at Reduced Gamma

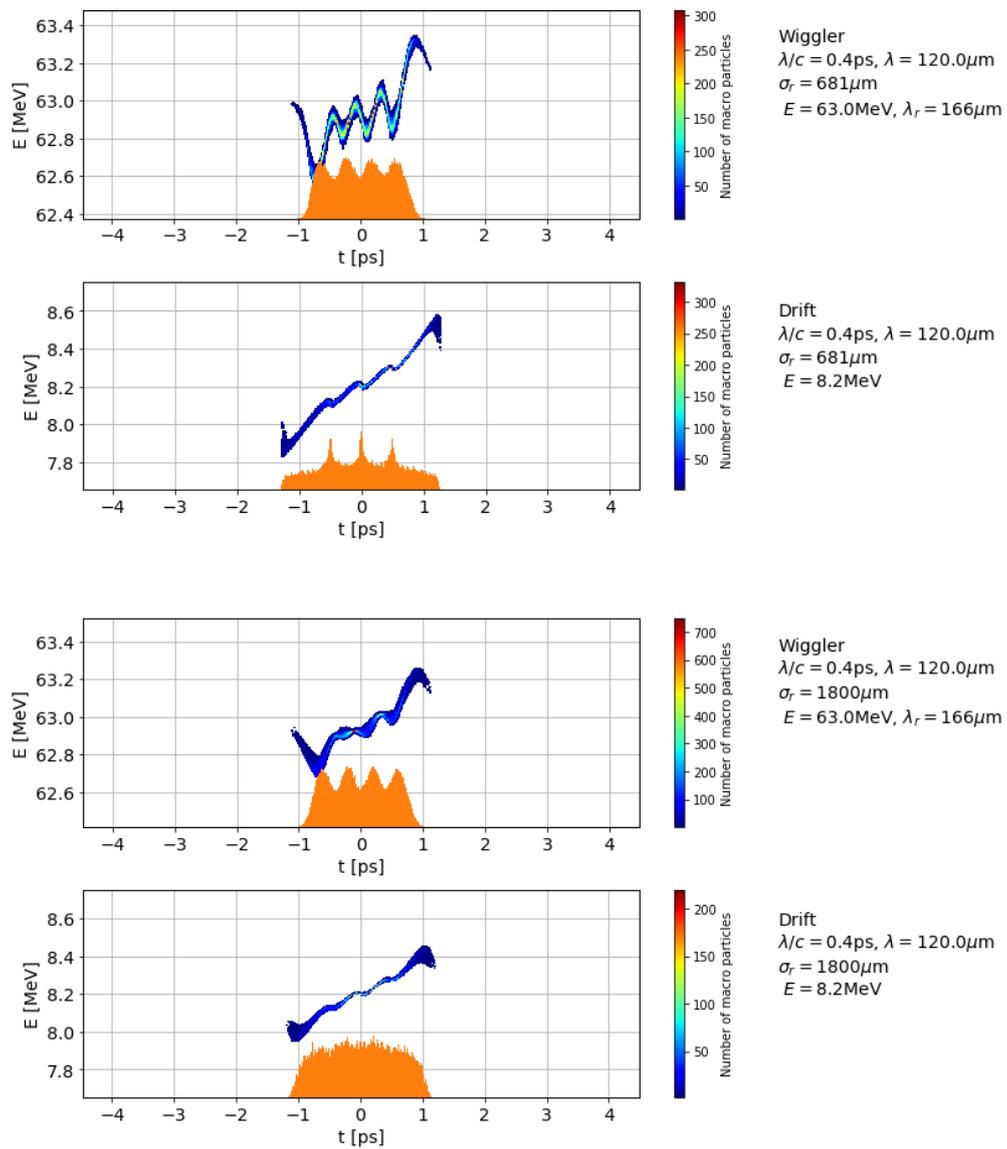
The following figures compare each simulation of the bunch passing through a wiggler to the bunch passing a low-energy 1 metre drift.

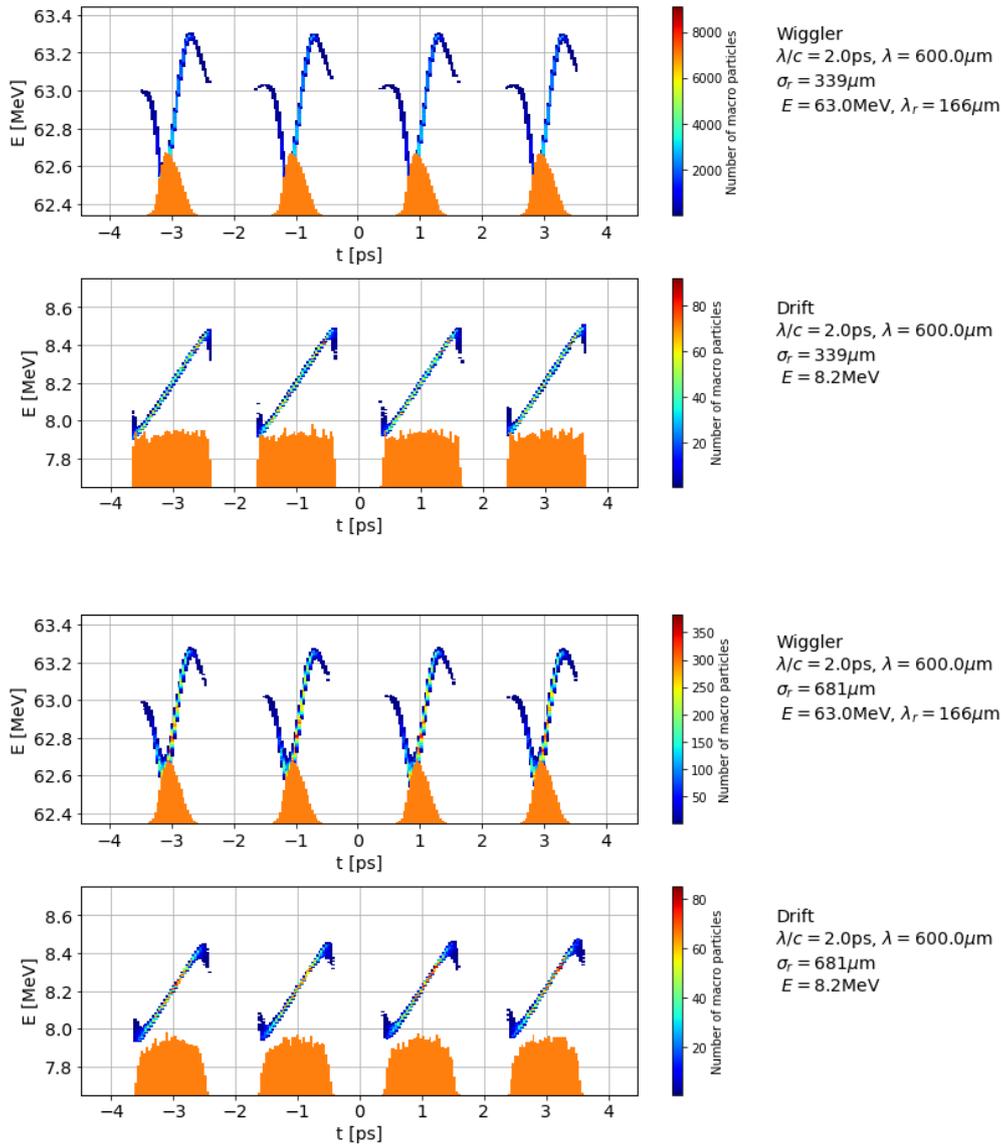


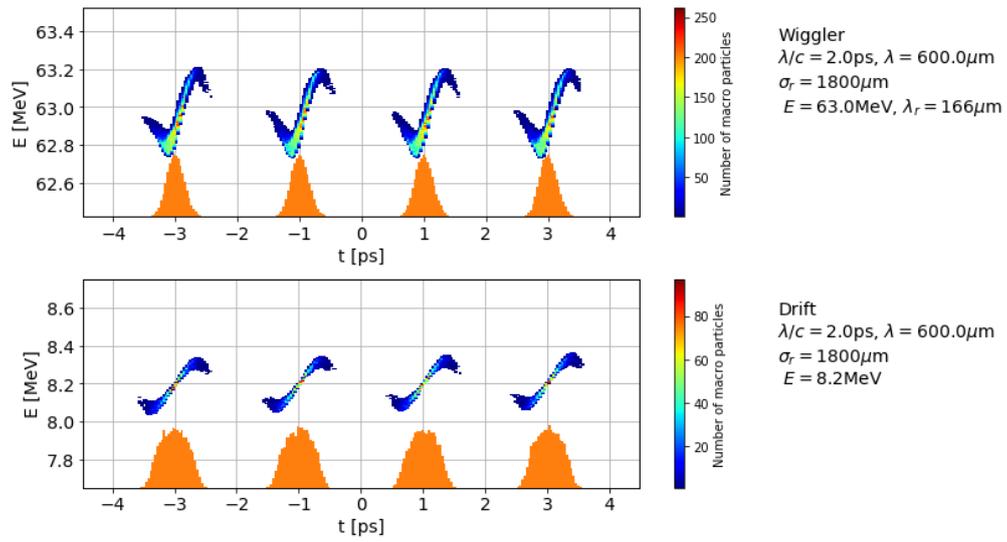












Appendix B

Comparison Wiggler Simulations and Drift of Same Length

The following figures compare each simulation of the bunch passing through a wiggler to the bunch passing a 1 metre drift.

