

MONTE CARLO SIMULATIONS OF ULTRA HIGH VACUUM AND SYNCHROTRON RADIATION FOR PARTICLE ACCELERATORS

Thèse no. 7063 présenté le 3 Mai 2016
à la Faculté des Sciences de Base
programme doctoral en Physique
École Polytechnique Fédérale de Lausanne

pour l'obtention du grade de Docteur ès Physique
par

Márton Ady

a être évalué par le jury:

Cécile Hébert, président du jury
Leonid Rivkin, directeur de thèse
Roberto Kersevan, co-directeur de thèse
Minh Quang Tran, rapporteur interne
Mikael Eriksson, rapporteur externe
Katsonobu Oide, rapporteur externe

Lausanne, 2016



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE



Résumé

Avec la préparation de Hi-Lumi LHC en cours, et les machines du FCC à l'étude, les accélérateurs atteindront des énergies sans précédent et donc une très grande quantité de rayonnement synchrotron (SR), provoquant la désorption de photoélectrons et de molécules depuis les parois de l'accélérateur. Ces désorptions contribuent à l'accumulation de nuage d'électrons et augmentent la pression résiduelle - les deux effets réduisant la durée de vie du faisceau. Dans les accélérateurs actuels, les deux effets sont parmi les principaux facteurs limitants. Le calcul précis du rayonnement synchrotron et de la pression est donc un enjeu important dès la phase de conception.

Ce projet de thèse décrit la modernisation et la mise à jour majeure de deux codes, Molflow et Synrad, écrits originellement par R. Kersevan dans les années 1990. Ils sont basés sur la méthode de Monte Carlo à particules de test, et permettent les calculs de la pression en condition de l'ultra-vide et du rayonnement synchrotron. Les nouvelles versions contiennent de la nouvelle physique, et sont construites comme un outil tout-en-un mis à la disposition du public.

Les méthodes de calcul de vide existantes sont résumées. Les algorithmes utilisés par Molflow relatifs au calcul des états stationnaire et évolutif sont ensuite présentés. Certaines des bonnes pratiques pour contrer les problèmes les plus communs que pose la simulation des systèmes de grandes dimensions sont également discutées. Les résultats issus de la simulation sont comparés à ceux issus de la théorie, et validés par deux expériences.

Les principales étapes de simulation de rayonnement synchrotron sont ensuite présentées. Les Propriétés du SR sont résumées, ainsi que les optimisations qui permettent de simuler une physique sous-jacente plutôt complexe à une vitesse supérieure. Les résultats de l'algorithme de génération de photons du logiciel sont comparés aux données publiées.

Le phénomène de la désorption stimulée par photons et sa littérature sont résumés, puis les deux expériences dédiées de photodésorption menées à KEK (Tsukuba, Japon) sont présentées: l'une réalisée à partir de six échantillons à température ambiante et l'autre à température de l'azote liquide.

Un calcul simple de rayonnement synchrotron est effectué pour la région d'interaction LHeC, permettant la comparaison des résultats de Synrad+ et des calculs analytiques publiés. Les calculs ont ensuite été répétés prenant en compte une description plus précise de la géométrie. Le profil de pression d'un absorbeur en Y de la source de lumière Max IV (démarrée récemment) est calculé en utilisant à la fois MolFlow+ et Synrad+.

Enfin, l'analyse de la pression dans la région d'interaction SuperKEKB est présentée. Elle comprend la modélisation de la chambre à vide et du système optique, le calcul du rayonnement synchrotron, puis les simulations de pression. Il est confirmé que la pression ainsi calculée lors du fonctionnement de la machine répond aux exigences de conception.

Abstract

With preparation of Hi-Lumi LHC fully underway, and the FCC machines under study, accelerators will reach unprecedented energies and along with it very large amount of synchrotron radiation (SR). This will desorb photoelectrons and molecules from accelerator walls, which contribute to electron cloud buildup and increase the residual pressure - both effects reducing the beam lifetime. In current accelerators these two effects are among the principal limiting factors, therefore precise calculation of synchrotron radiation and pressure properties are very important, desirably in the early design phase.

This PhD project shows the modernization and a major upgrade of two codes, Molflow and Synrad, originally written by R. Kersevan in the 1990s, which are based on the test-particle Monte Carlo method and allow ultra-high vacuum and synchrotron radiation calculations. The new versions contain new physics, and are built as an all-in-one package - available to the public.

Existing vacuum calculation methods are overviewed, then the steady-state and time-dependent algorithms behind the ultra-high vacuum simulator Molflow are presented. Some practices to tackle the most common problems that arise when simulating large systems are also discussed. Results are compared to theory, and validated through two experiments.

Next the the main steps of synchrotron radiation simulations are presented. Properties of SR are summarized, along with optimizations that allow simulating the rather complex underlying physics at a higher speed. The resulting software's photon generation algorithm is benchmarked against published data.

The phenomenon of photon stimulated desorption and its literature is overviewed, then two dedicated photodesorption experiments carried out in KEK (Tsukuba, Japan) are presented: one with six room-temperature samples and an other at liquid nitrogen temperature.

A simple synchrotron radiation calculation is performed for the LHeC interaction region, allowing to compare Synrad+ results with published analytic calculations. Then the calculations are repeated for a more precise geometry description.

The pressure profile of a crotch absorber of the recently started Max IV light source is calculated using Molflow+ and Synrad+ together.

Finally the pressure analysis of the SuperKEKB interaction region is presented, consisting of modeling the vacuum chamber and the optics, calculating synchrotron radiation, then performing vacuum simulations. It is confirmed that pressure is expected to meet the design requirements during operation of the machine.

Keywords:

Ultra High Vacuum, Synchrotron Radiation, Photon Stimulated Desorption, LHeC, Max IV, SuperKEKB, Molflow, Synrad

Acknowledgements

Doing a PhD is much more fun when one is not alone: I had the chance to be part of CERN's Vacuum and Surface Coatings group during this project, whose all members I owe many thanks for the help I got through the years.

I'd like to acknowledge a few people who have been specially helpful to me and mention them by name:

First of all, *Roberto Kersevan* for guiding me through this thesis as my supervisor. For all the work presented below, he must be credited for having a brilliant idea twenty years ago and creating - much ahead of the times - the first versions of the Monte Carlo codes I built upon. None of my recent work would have been possible without the good foundations and the continuous support he provided - thank you!

Leonid Rivkin for supervising my work from the university side, and providing strategic advice at crucial moments of the project.

Our group leader *Paolo Chiggiato* for offering the opportunity to pursue a PhD at CERN in the first place, and then going out of his way making it possible that I participate in workshops, schools and experiments in the UK, US and Japan - making my doctoral studies much more colorful than expected.

Yasunori Tanimoto for the warm welcome at KEK and the excellent preparation of our experiments that were a very important milestone of my doctoral studies.

Jason Carter from APS for giving a helping hand with the SuperKEKB case study, and the abundant feedback that helped making Synrad better.

Alan Van Drie from Tri Alpha Energy for his excellent handnotes about gas dynamics and helpful Molflow feedback.

Pedro C. Pinto, *Antonios Sapountzis*, *Benjamin Holliger*, *Wilhelmus Vollenberg*, *Berthold Jennings* and *Mario Armin Hermann* for all the help with the sample preparations and the experiments.

Christina Yin Vallgren for showing the way to finishing a PhD and giving advice so many times next to a coffee, *Ákos Hencz* and *Thibaut Richard* for their entertaining company throughout the years and *Ida Aichinger* for being the best Mathematica and Latex expert in the office.

Finally, my family: they were there all the time, helping throughout the years.

Geneva, 3 Mars 2016

Márton Ady

Contents

Abstract	iii
Acknowledgements	v
Introduction	1
1 Vacuum calculations	3
1.1 Introduction	3
1.1.1 Analytic methods	4
1.1.2 Analogy to electrical networks	7
1.1.3 Angular coefficient method	8
1.1.4 Commercial simulators	9
1.1.5 Monte Carlo methods	10
1.2 Molflow	11
1.2.1 Steady-state simulations	13
1.2.2 Time-dependent Molflow+ simulations	25
1.2.3 Post-processing in Molflow+	29
1.2.4 Sectioning the geometry	33
1.2.5 Theoretical validation	39
1.2.6 Experimental validation (CERN)	43
1.2.7 Experimental validation (LIGO)	47
2 Synchrotron radiation simulations	51
2.1 Introduction	51
2.2 Defining magnetic regions	51
2.2.1 Location, curvature and beam direction	54
2.2.2 Beam width and divergence	55
2.3 Synchrotron radiation properties	56
2.3.1 Spectrum	56
2.3.2 Representation	58
2.3.3 Vertical distribution and polarization	59
2.3.4 Horizontal distribution	61
2.4 Photon generation process	64
2.4.1 Source point choice	65
2.4.2 Choosing photon energy	65
2.4.3 Choosing vertical angle	65
2.4.4 Choosing horizontal angle	65

Contents

2.4.5	Reverse interpolating in a cumulative distribution function	66
2.4.6	Validity	67
2.5	Photon tracing	70
2.5.1	Photon reflection	70
2.5.2	Rough surface scattering	79
2.5.3	Validation	82
2.5.4	Low-flux mode	86
3	Photon stimulated desorption	89
3.1	PSD data in literature	90
3.1.1	Pressure calculation methods	91
3.1.2	Conversion between linear and surface dose	92
3.2	Experiments at KEK	93
3.2.1	Samples	94
3.2.2	Experimental setup	95
3.2.3	Experiment	98
3.2.4	Results	100
3.2.5	Experiment summary	108
3.3	Experiments at liquid N ₂ temperature	109
3.3.1	Samples	109
3.3.2	Experimental setup	111
3.3.3	Results	114
3.3.4	Discussion	116
4	Case studies	119
4.1	Simulation of SR power for the LHeC interaction region quadrupoles	119
4.1.1	Introduction	119
4.1.2	The IR region of the LHeC	120
4.1.3	Magnetic regions	120
4.1.4	Geometry	121
4.1.5	Results (approximate geometry)	121
4.1.6	Precise solution	123
4.1.7	Conclusion	124
4.2	The MAX IV crotch absorber problem	126
4.2.1	Introduction	126
4.2.2	Synchrotron radiation	126
4.2.3	Vacuum simulations	128
4.3	The SuperKEKB interaction region	133
4.3.1	The Belle experiment	133
4.3.2	Vacuum chamber modeling	134
4.3.3	The interaction region optics	135
4.3.4	Synchrotron radiation simulation	137
4.3.5	PSD conversion	143
4.3.6	The vacuum system	145
4.3.7	Vacuum simulations	148
4.3.8	Effects of roughness	150
4.3.9	Discussion of results	151

Summary and future plans	153
A High-pressure experiments	157
A.1 Introduction	157
A.2 HIE-ISOLDE and SPIRAL2 safety systems	158
A.3 Experimental setup	160
A.4 Tests and results	161
A.5 Additional test	162
A.6 Conclusion	164
B Miscellaneous simulations	165
B.1 Thermal radiation simulations - McCryoT	165
B.2 Lighting calculations	167
B.3 Moving parts	168
B.3.1 Gaede pump	168
B.3.2 Time-dependent problems	171
B.3.3 Space shuttle	172
Bibliography	181
Curriculum Vitae	183

Introduction

This thesis is written at an interesting time for the accelerator physicist: with preparation of Hi-Lumi LHC fully underway, and the FCC machines under study, in the near future, particles are expected to be accelerated to energies never seen before.

Increasing the energy poses several technological challenges, though. A charged particle in a circular machine will emit synchrotron radiation in the bends. The power of the radiation scales with the fourth power of the energy. Therefore, once the new machines start operating, an unprecedented amount of radiation will be generated. Apart from safety considerations, the high radiated power will desorb photoelectrons and molecules from the accelerator walls, which will contribute to electron cloud buildup and increase the residual pressure - and through both effects the beam lifetime would be reduced. In current accelerators these two effects are among the principal limiting factors, contributing substantially to the cost of the machines.

Because of the problems above, it is more important than ever to precisely calculate synchrotron radiation and pressure properties, desirably in the early design phase. There are a handful of tools already available for this task, however each has its own limitations and the learning curve to use them is not negligible.

This PhD project shows the modernization and a major upgrade of two codes, Molflow and Synrad, originally written by R. Kersevan in the 1990s, which are based on the Monte Carlo method and allow ultra-high vacuum and synchrotron radiation calculations. The new versions contain new physics, and are built as an all-in-one package with emphasis on compatibility both with CAD programs and also with each other. Available to the public, they are versatile tools for engineers in the accelerator domain.

To write them, several challenges had to be solved, experimental data had to be collected, and the ready versions needed to be constantly benchmarked through comparison with literature and through case studies where they were put to use in solving real problems. This thesis shows the development process and the validation, going along the following structure:

Chapter 1 overviews existing vacuum calculation methods. It presents the steady-state and time-dependent algorithm behind Molflow, our ultra-high vacuum simulator. It also shows some practices to tackle the most common problems that arise when simulating large systems. Results are compared to theory, and finally two experiments are presented where Molflow+ is benchmarked against measured data.

Chapter 2 presents the main steps of setting up a synchrotron radiation study. While doing so, it summarizes the main properties of the synchrotron light, and shows optimizations I undertook to allow simulating rather complex underlying physics at an acceptable speed. The resulting software, Synrad's photon generation algorithm is benchmarked against published data. The chapter continues with the description of the reflection and scattering model used, and a simple radiation protection example demonstrates its ability to treat surfaces in a custom, user-defined way. Finally, the scattering model is compared to published data, and an optimization idea is mentioned.

Chapter 3 familiarizes the reader with the phenomenon of photon stimulated desorption, and gives a hint on how to use existing data in literature for vacuum calculations. It then presents two dedicated photodesorption experiments carried out at the KEK laboratory in Tsukuba, Japan, one with six room-temperature samples and an other at liquid nitrogen temperature which is of interest for the HiLumi LHC and the FCC-hh.

Chapter 4 puts the codes presented in chapters 1 and 2 to use.

First a simple synchrotron radiation calculation is performed for the LHeC interaction region, allowing to compare my results with published analytic calculations. Then the capabilities of Synrad+ are demonstrated by changing the geometry to a more precise case which the analytic calculations avoided due to its complexity.

Next the coupled vacuum / synchrotron radiation calculation capabilities of the codes are showcased by using Moflow+ and Synrad+ together to calculate the pressure properties of a crotch absorber of the recently started Max IV light source.

In the final section of the chapter a full analysis of the SuperKEKB interaction region is presented, that consists of modeling the vacuum chamber and the optics, simulating synchrotron radiation and discussing the effect of surface roughness, then performing vacuum simulations with some simplifications applied to the pump regions. It is confirmed that pressure is expected to meet the design requirements during operation of the machine.

Appendix A gives an outlook to gas dynamics outside the ultra-high vacuum regime, and since simulations with mixed pressure ranges are very difficult, we measure gas propagation velocity on a shock tube experimental setup at CERN.

Appendix B hints the reader at other potential uses of the Monte Carlo methods shown so far, by highlighting the similarities between vacuum, thermal radiation and lighting calculations. Finally, an add-on of the original Moflow+ code allowing to include moving parts in the simulated system is presented.

1 Vacuum calculations

1.1 Introduction

When we design particle accelerators, our goal is to collide the accelerated particles - either with a fixed target or an other accelerated particle -, or to store them and use their synchrotron radiation for various experiments. In both cases, as particles interact with matter, machines have to be designed in a way that the beam of accelerated particles is contained in a chamber which is as clean as possible.

From an operational point of view, residual gas in an accelerator causes various undesirable phenomena, most importantly beam losses through various effects. One of the main causes of such losses is Bremsstrahlung, sometimes referred to as braking radiation, where the accelerated particles interact with the nuclei of the residual gas. The beam lifetime due to the effect is expressed as [1]

$$\tau = \frac{X_0}{W\rho c}$$

where X_0 is the radiation length of the residual gas, ρ is the density and c is the speed of light, and W is a constant that depends on the energy spread of the stored particles¹.

From the above expression that defines the lifetime, with the density in the denominator, we can see that keeping the chamber "clean" means reducing the residual gas pressure. This is achieved by various (primary, turbomolecular, ion, getter, cryogenic) pumps², and in a modern accelerator, the pressure must be kept at a level in the order of 1×10^{-8} mbar (for the LHC, see chapter 12 of [5]).

At such a low pressure, the gas behaves differently from standard conditions: the dynamics can be quantified through the Knudsen number, defined as

$$Kn = \frac{\lambda}{L}$$

¹ $W = \frac{4}{3} \ln\left(\frac{\Delta E}{E}\right) - \frac{5}{6}$

²Although the technological overview of vacuum is out of the scope of this chapter, I can recommend several good introductions to the topic, such as [2, 3, 4].

Where L is the characteristic dimension of the vacuum chamber, and

$$\lambda = \frac{k_B T}{\sqrt{2} \pi d^2 p}$$

is the mean free path, where k_B is the Boltzmann constant [J/K], T is the temperature [K], p is pressure [Pa], and d is the diameter of the gas particles [m]. Based on the Knudsen number, the gas dynamics can be classed into different regimes³:

- **Continuum (viscous) flow** ($Kn < 0.001$): This is the what we experience in standard conditions, such as with atmospheric air: gas behavior is determined by intermolecular collisions and internal friction (viscosity). The regime can be further divided into turbulent and laminar flow, based on the Reynolds number. Simulations are usually carried out using fluid dynamics equations, and several commercial finite-element solvers can handle the task.
- **Transitional flow** ($0.001 < Kn < 1$): At this regime the mean free path is comparable to the chamber dimensions, viscosity still plays a role but is less important
- **Free molecular flow** ($1 < Kn$): Intermolecular collisions can be neglected; the gas molecules are to be treated individually.

The free molecular flow condition is typically achieved in an accelerator (with chamber diameter of a few centimeters) below 1×10^{-3} mbar. In terms of pressure, we call this range High Vacuum (HV, between 1×10^{-3} and 1×10^{-7} mbar) and Ultra High Vacuum (UHV, below 1×10^{-7} mbar), as opposed to Rough Vacuum (above 1 mbar). Calculation methods presented in this thesis assume free molecular flow, and are therefore valid up to the pressure limit defined above.

In UHV conditions, as there is no viscosity, fluid dynamics equations do not hold, and gas behavior must be modeled as a statistics of a very large number of individual particles doing random walk within the walls. When particles collide with the wall, their impulse changes thus they exert a force on it. This impulse change rate is the origin of the pressure. There are several ways to calculate the pressure in a vacuum system or for a whole accelerator in the UHV range:

1.1.1 Analytic methods

Most analytic methods⁴ apply two fundamental equations of UHV theory. In a vacuum chamber, where gas enters with an outgassing rate Q (in this thesis I'll use mbarl/s as units), and exits with a pumping speed S [l/s], the equilibrium pressure will be

$$p = \frac{Q}{S} \tag{1.1}$$

³Regime limits are not strict and vary across literature

⁴An extensive summary of these methods can be found at [6]




Figure 1.1: Analytic vacuum problem - first example

In a vacuum system gas flows from higher to lower pressure areas⁵, and the flow rate is limited by the geometry. If between two points the pressure difference is dP , then the gas flow rate from the high pressure point towards the lower will be

$$Q = C \cdot dP \quad (1.2)$$

where C is the conductance [l/s], which in free molecular flow depends only on the geometry and on the velocity of the gas molecules. Using these two equations, many simple vacuum problems can be solved analytically. For example, in the setup shown in Fig.1.1, gas flows to a chamber, which in turn is connected to a second chamber equipped with a pump. Using the notations of the figure, we can write equation 1.2 and 1.1 for the two chambers:

$$Q_{1 \rightarrow 2} = C \cdot (p_1 - p_2)$$

and

$$p_2 = \frac{Q_{1 \rightarrow 2}}{S_2}$$

$$p_1 = \frac{Q}{S_1}$$

Here S_1 is the local pumping speed in chamber 1, that we have to calculate. We expect it to be lower than S_2 because the pump's performance is reduced by the finite conductance C . Now, because of the conservation of mass, we can notice that $Q = Q_{1 \rightarrow 2}$. Solving the equation system, we can express the pressures in the two chambers as:

$$p_2 = \frac{Q}{S_2}$$

$$p_1 = Q \left(\frac{1}{C} + \frac{1}{S_2} \right) = \frac{Q}{S_1}$$

obtaining $\frac{1}{S_1} = \frac{1}{C} + \frac{1}{S_2}$ as the local (or effective) pumping speed of S_2 in chamber 1. Using this new result (local pumping speed of a distant pump separated by conductance C) we can calculate continuous pressure profiles, resembling more to a real vacuum problem in an accelerator (Fig.1.2).

⁵Assuming there is no temperature gradient




Figure 1.2: Analytic vacuum problem - second example




Figure 1.3: Pressure profile of a periodic vacuum system with pumps placed at 5 m intervals

In this case the outgassing comes from the wall of the chamber uniformly, and we place pumps of speed S at an interval of d . The conductance between two points – similar to admittance (the reciprocal of resistance) in an electrical network - will scale with the inverse of the distance: $C(L) = C'/L$ where L is the distance between the two points and C' is the specific conductance of a unit length. The pressure at location x between two pumps will be (neglecting all but the two closest pumps):

$$p(x) = \frac{Q_{\text{linear}}}{S(x)}$$

$$S(x) = \frac{1}{\frac{x}{C'} + \frac{1}{S}} + \frac{1}{\frac{d-x}{C'} + \frac{1}{S}}$$

Substituting $S = 1001/\text{s}$, $Q_{\text{linear}} = 1 \times 10^{-7} \text{ mbar l/s/m}$, $d = 5 \text{ m}$, $C' = 10 \text{ ml/s}$, we obtain a parabolic pressure profile, with the local minima at the pump locations ($x = 0$ and $x = 5 \text{ m}$), as in Fig.1.3.

This method is suitable for modeling large linear structures. Many "in-house" codes have been written using the principles above, such as Vasco [7] and IdaVac at CERN, VacCalc[8] at SLAC, etc. These codes usually describe the vacuum system as a succession of linearly connected vacuum elements, each having its own outgassing, conductance and pumping speed. They also tend to include dynamic outgassing (molecules desorbing due to synchrotron radiation impinging on the walls) and even electron cloud effects (accumulation of electrons in the space through secondary emission), and aggregate all vacuum equations (one for each element) into a matrix representing the whole system.

Table 1.1: Electrical circuit analogy of vacuum systems

Vacuum concept	Electrical analogue
C (conductance)	$Y=1/R$ (admittance)
P (pressure)	U (potential)
dP (pressure difference)	U (voltage)
Q (gas flow)	I (current)
V (volume)	C (capacity)
Pump (with finite pumping speed)	Ground (through resistance)
Gas source	Current generator

The problem with these methods is that the system must be linear, and the conductance must be known for all elements. The former requirement is not too constraining (most accelerators are linear systems), however determining conductances analytically can be practically impossible in certain cases, as I'll explain below.

The conductance of a straight pipe of length L and diameter d was first calculated by Knudsen[9]:

$$C = \frac{8}{3} \sqrt{\frac{2}{\pi}} \cdot \frac{A^2}{sL} \cdot \sqrt{\frac{1}{d}}$$

where A is the cross-section area and s is the circumference. It was shown by Clausing[10] that if the tube is not circular, the formula is incorrect and a double integral over the tube length must be used instead. Over the following years, a catalogue was built containing analytic conductance expressions for the most common shapes and vacuum elements.

The problem with this approach is that it's very difficult to apply for complicated vacuum structures (such as in Fig.1.9), and even for simple pipes, it does not take into account the beaming effect: in a pipe with a predominant flow direction the molecules will deviate from the isotropic direction distribution, distorting any analytic conductance formula. Nowadays, analytic methods are usually used hand-in-hand with finite element solvers or Monte Carlo calculations to obtain the conductance of all elements, which in turn can be plugged in to the analytic vacuum equations.

1.1.2 Analogy to electrical networks

In the previous section I mentioned that the conductance of a vacuum system is similar to the admittance of electrical networks. In fact, similarities are much more extensive: almost every vacuum concept can be translated to an electrical analogue, as shown in Table1.1.

This means, for example, that the fundamental vacuum equation $Q = C \cdot dP$ is in fact the analogue of Ohm's law, $I = U/R$. The importance of the analogy is that we can use electronic network solvers (such as the the freeware LTSpice, used in CERN's vacuum group) to treat our vacuum problems, even time-dependently. An example is shown in Fig.1.4.

Although such solvers are powerful and can usually treat very large systems which are not necessary linear, the conductance of each element must still be calculated externally, moreover




Figure 1.4: Fig: simple vacuum system and its electrical network analogue. Source: [11]




Figure 1.5: View factor method overview, redrawn based on Fig. 2 of [12]

some effects (like temperature differences and beaming) can't be included.

1.1.3 Angular coefficient method

Originally applied for radiative heat transfer calculation, the view factor or angular coefficient method can be also used for vacuum. Similar to the thermal radiation from Lambertian surfaces, the directional distribution of particles reflecting from a surface is proportional to the cosine of the surface normal and the direction[9].

We can section all surfaces in our system to N small elements. Then, according to Knudsen's law, element i will receive

$$\Phi_{in,i} = \sum_{N, i \neq j} \phi_{out,j} \cdot \frac{\cos(\theta_i) \cos(\theta_j)}{r^2 2\pi \cdot dF_i}$$

flux from all other elements. Φ_{out} is the total flux emitted by element j , and $\theta_i, \theta_j, dF_1, dF_2, r$ are notations as shown in Fig.1.5.

Each element can have a sticking factor s_i (similar to the emissivity of radiative heat transfer), and a proper source of particles e_i , so if $\Phi_{in,j}$ is the total incident flux on element j , then

$$\phi_{out,j} = e_j + (1 - s_j) \cdot \phi_{in,j}$$




Figure 1.6: ANSYS simulations of a pressure profile based on the radiative heat transfer module

This way all surface elements can be connected through a series of algebraic equations, and the system behavior can be described by an $N \times N$ matrix, where element i, j will be the view factor

$$\frac{\cos(\theta_i) \cos(\theta_j)}{r^2 2\pi \cdot dF_i}$$

between elements i and j . The solvability (time and memory requirement) of the system will largely depend on the matrix size determined by the mesh resolution. Such a method (next to the default Monte Carlo algorithm) is applied in Molflow+ and compared[12] to analytic and Monte Carlo solutions.

1.1.4 Commercial simulators

Traditionally simulating UHV was a narrow field with a limited user base, and as such it was outside of the world of commercial programs. Some finite element solvers, such as ANSYS (Fig.1.6) could be "hacked" to use their radiative heat transfer module to vacuum calculations.

The situation has changed very recently, when Comsol (Fig.1.7) came out with version 4.3 of their multiphysics simulator in 2014. That version, for the first time, included a free molecular flow module, based on the angular coefficient method, which in fact was a "rebranding" of its already existing radiative heat transfer module⁶. Initial testing at CERN showed that the computation time is very long compared to our in-house methods, probably due to using a matrix that becomes huge for a finer mesh. However, at the time of writing, the most recent version, 5.2 has been published. It seems to have evolved the module and added new features, including iterative solving of UHV problems and an interface with the *Mathematical Particle Tracing* module, allowing transmission probability calculations with the Monte Carlo method.

⁶I base this on the governing equations of the two modules (visible to the user), as the calculation method is proprietary

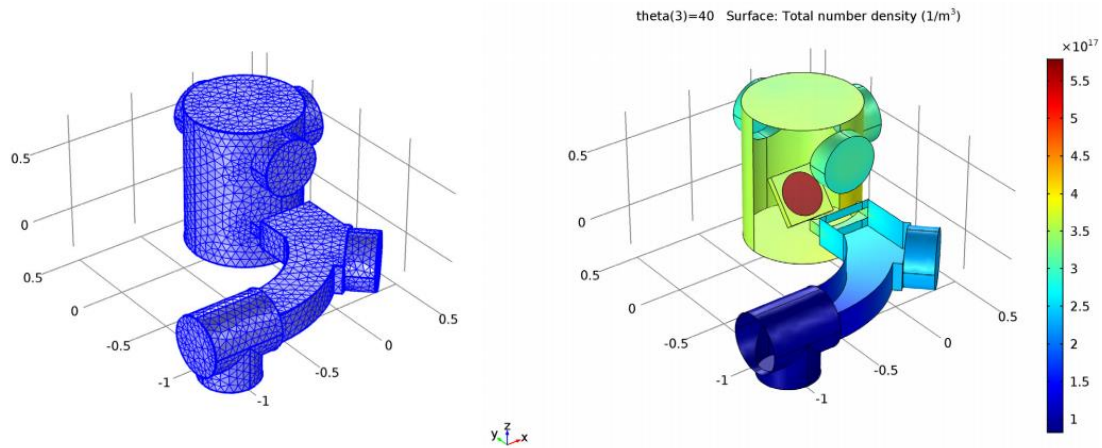


Figure 1.7: Density simulation of an ion implanter model, as found on the public website of Comsol Multiphysics' free molecular flow module




Figure 1.8: Molflow+ simulation (right) of the transmission probability from source to exit (left, in red) of the LINAC4 H- extraction region. Source: [11]

Having an advanced interface with CAD programs and a large development team behind the product, it is certainly interesting to keep an eye on future versions.

1.1.5 Monte Carlo methods

As written above, the problem with many analytic methods is that the conductance cannot be calculated for complex geometries therefore it must come from an external source. The Monte Carlo based vacuum simulators tackle the problem by launching a large number of “test” particles (Fig.1.8) in a vacuum system which is described either parametrically or more often by a set of polygons. When a collision is detected between a test particle and the wall, it is relaunched based on Knudsen's cosine law. After a succession of hits, the test particle will leave the system at one of the pumping locations.

For a complicated geometry, the transmission probability α_{tr} from source S to target T can be defined as the ratio of test particles leaving the system at exit T to all generated particles. For a given average molecule velocity $\langle v \rangle$ (governed by the gas mass and temperature), we can obtain the conductance of the element:

$$C_{S \rightarrow T} = \alpha_{TR} \frac{1}{4} \nu A_S$$

where A_S is the source area.

As explained later, based on statistics from the test particles, we can calculate – apart from the conductance – physical quantities like the pressure and density distribution of the system, and accurately model beaming effects. Computer implementations of the method have been published by Luo, Giegerich, Day (Provac3D, [13]), Malyshev [14], Suetsugu [15], Pace and Poncet[16] and Garion[17].

1.2 Molflow

Molflow, dating back to the 1990s, is a Monte Carlo simulator intended to calculate pressure profiles and conductances in ultra-high vacuum. It uses the test particle Monte Carlo method where – as introduced above - a limited number of virtual, “test” particles represent a larger amount of physical molecules and the derived quantities of the test particles are scaled up to match the physical numbers.

In Molflow the geometry is described by polygons called facets. All facets have user-defined physical properties (temperature, opacity, sticking, etc.). To perform a vacuum simulation, we insert molecules in the system and trace them until they are pumped. In between, collisions with the wall (referred to as hits) are registered, and where necessary, stored in memory. We do not account for collisions between molecules as the mean free path of molecules in the accelerator pressure ranges is shown before to be above the characteristic length of the vacuum chambers.

This is in contrast with direct simulation Monte Carlo (DSMC) methods, where the vacuum chamber volume is divided into 3D cells, and a set of virtual particles is kept in memory and their positions are evolved as time advances.

Historically the DSMC method [18] supersedes the TPMC, and is generally considered more advanced as treatment of inter-molecular collisions allows simulating pressures in the viscous regime (where the Knudsen number is below 1). However, in particle accelerator applications, where pressures are, with few exceptions, well within the free molecular flow regime, I argue that TPMC methods are still a good choice because of the following advantages:

- As intermolecular collisions are omitted, only one test particle has to be kept in memory. Its trajectory, assumed to be straight between the chamber walls, can be determined by ray tracing, from the point of insertion (outgassing point) to the pump location. This is in contrast with DSMC methods where each 3D cell has to store a statistically significant number of particles, increasing the memory requirement substantially.
- The TPMC method allows simulations to be event-driven: Molflow+ first looks for collisions with the walls, and once they happen, based on the position and the particle




Figure 1.9: Complex geometry and its Molflow+ equivalent

velocity, it calculates the time of the hits. This is in contrast with time-driven simulations, where time is advanced periodically by a given step and the new molecule locations are iterated, checking at each step if the new position is within the chamber's volume. The problem with that method is that the time step must be scaled to the smallest feature in the system, therefore in case of a straight flight path between two walls, possibly hundreds of time steps are required in time-driven algorithms vs. only one in event-driven ones.

- Since test particles are independent, the vacuum system's behavior is linear: the algorithm is particularly suitable for parallelization: in Molflow's case each CPU core traces one test particle, and hits are summed when results are evaluated. For larger parallelization (cluster, GPU), only the geometry with the physical properties have to be passed for each thread, and simulation state synchronization, usually a bottleneck on GPUs, is not necessary.
- While 3D meshing is straightforward in simple geometries, for complex ones (grids, screws, probes, RF fingers) it becomes more error-prone (see Fig.1.10): small geometry details can increase the number of cells used, and the mesh size, if not chosen correctly, can introduce non-converging solutions. In TMPC simulations only surfaces are described, imported from CAD programs in the STL format.

Due to the lower memory requirement and the faster algorithm explained above, most vacuum chambers can be simulated with the TMPC method on an average desktop computer. As an example of a complex geometry, the tank with electrodes and correctors in Fig.1.9 can be described by 37.000 facets and a pressure solution with acceptably low statistical scattering is reached within an hour (corresponding to 20 million hits) in Molflow+. Volumetric meshing of the geometry can take over a day and the consequent DSMC simulation of the system can typically be achieved only with high-performance computation clusters due to the large memory and computational requirement.

The original Molflow code – developed outside of this PhD project – calculated results as




Figure 1.10: 3D tetrahedral meshing failures with Autodesk Inventor 2016 – a 3D mesh can take days to generate and the memory requirement is well outside the capability of current desktop PCs

numbers of Monte Carlo hits – which could be converted into pressures or density in case of a steady-state simulation where the whole system was isothermal. Based on the ray-tracing engine and the custom GUI framework⁷, I have extended the algorithm to include the distribution of the molecule velocities (so pressure can be calculated in non-equilibrium scenarios), and once this is achieved, make simulations time-dependent. This new code is presented below.

1.2.1 Steady-state simulations

First Molflow's steady-state algorithm is presented: in this case, there is a continuous influx of gas particles to the system, with constant pumping speeds, and our goal is to determine the (constant) pressure, density and other physical quantities on the facets. It is worth mentioning that unless the user explicitly defines time-dependent parameters, this steady-state algorithm is used. The algorithm's main steps, with an overview on Fig.1.11, are identical to almost all TPMC vacuum calculators. Details for each step are explained below.

1.2.1.1 Particle generation

In the steady-state case the outgassing rate $Q = d(pV)/dt$ is constant. Gas may come from the walls through thermal desorption, or from injection points. In both cases the source points will be represented by facets where for each facet f the user defines the outgassing Q_f . From

⁷Programmed by of J.L.Pons, ESRF in 2007




Figure 1.11: Main steps of Molflow's algorithm

this quantity, the influx rate of physical particles, $dN_{f,real}/dt$ can be calculated: using the ideal gas equation:

$$pV = N_{real}k_B T \quad (1.3a)$$

$$\frac{dN_{f,real}}{dt} = \frac{pV}{dt} \frac{1}{k_B T_f} = \frac{Q_f}{k_B T_f}, \quad (1.3b)$$

where T_f is the temperature of the gas source that facet f represents.⁸

For each generated virtual particle, Molflow+ chooses a starting location. The probability of choosing a certain facet as source is proportional to the local flux dN_{real}/dt of molecules on the facet, derived above. Once the source facet is decided, the start position of the test particle is chosen randomly with uniform distribution on its surface.

The total number of particles entering the system every second, $\sum \frac{dN_{real}}{dt}$ is the sum of the particle fluxes on the source facets. Even for a low outgassing rate, for example 1×10^{-8} mbar l/s, it is in the order of 1×10^{12} particles/second, and our computing power is insufficient to simulate all of them. This is why we represent a large number of real particles with a smaller number of virtual (test) particles. In this case, if we generate a total $N_{virtual}$ test particles, each one will represent a number of real particles per second equal to:

$$K_{\frac{real}{virtual}} = \frac{\sum dN_{real}/dt}{N_{virtual}}$$

⁸For clarity, I would emphasize that in the steady-state mode – similar to the SR flux in Synrad, presented in chapter 2 – a test particle doesn't represent an absolute number of physical molecules, but instead – as the system is in equilibrium – a certain influx or rate of particles (incoming particles per second). Consequently, a virtual hit on the wall will represent a certain number of real hits per second, contributing to the (constant) pressure.




Figure 1.12: Pressure profiles in a cube with uniform desorption, based on the directional distribution of emitted and reflected particles

1.2.1.2 Choosing particle direction

For particle generation and surface rebounds, Knudsen's cosine law is applied, which states [9] that the probability ds that a molecule leaves a surface in solid angle $d\omega$ is

$$ds = d\omega \pi \cdot \cos \theta$$

where θ is the angle between the particle's direction vector and the surface normal⁹.

Using results from [20], the ϕ azimuth and θ angles can be generated as follows:

$$\phi = rnd \cdot 2\pi$$

$$\theta = \sin^{-1} \sqrt{rnd}$$

where rnd is a pseudo-random number generated uniformly between 0 and 1. We generate these random numbers using the Mersenne-Twister algorithm[21].

1.2.1.3 Choosing particle speed

For simplicity, several Monte Carlo vacuum simulator implementations calculate either with the average molecule speed [15], or a constant speed that remains the same through the collisions [17]. In a real system the speed distribution of molecules in a given volume of ideal gas at equilibrium is described by the Maxwell-Boltzmann distribution[22], and after each collision the particle spends a certain sojourn time on the wall during which it thermalizes

⁹If we have Monte Carlo tools available, Knudsen's cosine law can be "proven" by a simple thought experiment. If we have a cube with uniform desorption and uniform sticking on all sides, we expect its pressure to be uniform at every point. In Fig.1.12 I model three cubes, one with cosine desorption obeying Knudsen's law, an other one with "uniform" desorption, meaning that each outgoing direction is chosen with equal probability, and one with a "collimated" desorption, where the probability of choosing a direction is proportional to $\cos^5(\theta)$, therefore preferring directions close to the facet's normal vector.

On the pressure values, visualized by textures, we can see that any deviation from the cosine law results in particles tending to get stuck in corners or on the contrary, tending to avoid them. Only a random walk following the cosine law will result in uniform pressure regardless of the geometry shape.

For a more formal proof that the cosine distribution must be chosen see [19]

Chapter 1. Vacuum calculations

hence obtaining a new speed. Implementing this accurately can simulate effects where after a gas injection some faster particles can arrive to distant locations faster than that predicted by the average gas speed. To include these effects, I choose therefore a new molecule speed at every collision, as it happens in a real vacuum system.

In my event-driven algorithm, test particles aren't generated volume-wise, but - as explained above - representing the flux of incoming particles during a given time through source facets. This is important as the speed distribution is different in the latter case: fast molecules of the gas cross or hit a surface more frequently than slow ones.

If the probability density function (p.d.f.) of molecular speeds in a volume is $f(v)_{gas}$, then it is possible¹⁰ to calculate the p.d.f. of molecules colliding with a wall during a period of time, $f(v)_{coll}$:

$$f(v)_{coll} = v^3 \exp\left(-\frac{v^2}{2a^2}\right) \frac{1}{2a^4} \quad (1.4)$$

where

$$a = \sqrt{\frac{k_B T}{m}}$$

Now that we have the speed distribution of molecules crossing or colliding on a facet, we can generate molecules according to it by the numerical inversion method [23]: we calculate

¹⁰We start from the fact that fast particles will collide more often, or more precisely, the collision frequency will scale linearly with the particle speed v :

$$f(v)_{coll} = f(v)_{gas} \cdot v \cdot C \quad (1.5)$$

where C is a normalizing factor, that can be obtained since the integral of the collision p.d.f. must by definition be 1:

$$\begin{aligned} \int_0^\infty f(v)_{coll} \cdot dv &= 1 \\ \int_0^\infty f(v)_{gas} \cdot v \cdot C \cdot dv &= 1 \\ C \cdot \int_0^\infty f(v)_{gas} \cdot v \cdot dv &= 1 \\ C \cdot \langle v \rangle_{gas} &= 1 \end{aligned}$$

$$C = \frac{1}{\langle v \rangle_{gas}} \quad (1.6)$$

where $\langle v \rangle_{gas}$ is the average molecule speed in the gas. In case of the Maxwell-Boltzmann distribution, $\langle v \rangle_{gas} = \sqrt{\frac{8RT}{\pi m}}$ with T the gas temperature, R the ideal gas constant and m the molar mass (in kg). Assuming that we have a gas with Maxwell-Boltzmann distribution, where the speed p.d.f is known[22]:

$$f(v)_{gas(M.B.)} = v^2 \exp\left(-\frac{v^2}{2a^2}\right) \frac{1}{a^3} \sqrt{\frac{2}{\pi}}$$

and substituting the value of the normalizing factor, we can express $f(v)_{coll}$, as eq.1.4 in the main text.

Table 1.2: The Maxwell-Boltzmann and the modified speed distribution

	Maxwell-Boltzmann	Surface collisions
P.D.F	$f(v)_{gas} = v^2 \exp\left(-\frac{v^2}{2a^2}\right) \frac{1}{a^3} \sqrt{\frac{2}{\pi}}$ $a = \sqrt{\frac{k_B T}{m_{[kg/particle]}}} = \sqrt{\frac{RT}{M_{[kg/mole]}}}$	$f(v)_{coll} = v^3 \exp\left(-\frac{v^2}{2a^2}\right) \frac{1}{2a^4}$
C.D.F	$CDF_{gas} = \text{erf}\left(\frac{v}{\sqrt{2}a}\right) - \sqrt{\frac{2}{\pi}} \frac{v \cdot \exp\left(-\frac{v^2}{2a^2}\right)}{a}$	$CDF_{coll} = -\exp\left(-\frac{v^2}{2a^2}\right) \left[1 + \frac{v^2}{2a^2}\right] + 1$
avg.	$\langle v \rangle_{gas} = \int_0^\infty v \cdot f(v)_{gas} dv = \sqrt{\frac{8}{\pi}} a$	$\langle v \rangle_{coll} = \int_0^\infty v \cdot f(v)_{coll} dv = 3\sqrt{\frac{\pi}{8}} a$
speed	$\langle v \rangle_{gas, 300K, 28g/mole} = 476.2 \text{ m/s}$	$\langle v \rangle_{coll, 300K, 28g/mole} = 561 \text{ m/s}$
$\langle \frac{1}{v} \rangle$	$\langle \frac{1}{v} \rangle_{gas} = \int_0^\infty \frac{1}{v} \cdot f(v)_{gas} dv = \frac{\sqrt{2}}{a}$ $\langle \frac{1}{v} \rangle_{gas, 300K, 28g/mole} = 0.0027 = \frac{1}{374} \text{ m/s}$	$\langle \frac{1}{v} \rangle_{coll} = \int_0^\infty \frac{1}{v} \cdot f(v)_{coll} dv = \frac{\sqrt{\pi}}{a}$ $\langle \frac{1}{v} \rangle_{coll, 300K, 28g/mole} = 0.0021 = \frac{1}{476.2} \text{ m/s}$
Com- parison	$\langle v \rangle_{coll} / \langle v \rangle_{gas} = \frac{3\pi}{8} \approx 1.178$ $\langle \frac{1}{v} \rangle_{gas} / \langle \frac{1}{v} \rangle_{coll} = \frac{4}{\pi} \approx 1.273$	

the cumulative distribution function (C.D.F) for the speed distribution, which will assign for every given speed v a probability (between 0 and 1) that a particle's velocity is lower than v . Then we generate pseudo-random numbers uniformly between 0 and 1 (using again the Mersenne-Twister algorithm), then interpolate in the C.D.F. the speed v corresponding to the generated probability. For this, the C.D.F. for our modified speed distribution must be calculated, which - unlike for the Maxwell Boltzmann distribution - can be expressed in closed form:

$$F(v)_{coll} = \int_{-\infty}^v f(v')_{coll} dv' = 1 - \exp\left(-\frac{v^2}{2a^2}\right) \left[1 + \frac{v^2}{2a^2}\right]$$

For speedup of the simulation, at the beginning, this C.D.F. function's values are calculated for 100 equidistant speed values between 0 and $4 \cdot v_{\text{most_probable}} = 4 \cdot \sqrt{\frac{2RT}{M_{kg/mole}}}$ for every facet temperature present in the system (this interval contains most of the distribution), and the algorithm - instead of interpreting the analytic expression at each collision - performs an inverse lookup: by reverse interpolation it looks up the speed that belongs to a generated random number.

An additional speedup could be achieved by calculating the inverse of the C.D.F. and doing forward interpolation (interpolate the Y value at a known X value) instead of reverse, but due to the long tail, i.e. non-zero probability for high speeds in the Maxwell-Boltzmann distribution this would introduce sampling artifacts at high speeds.

A summary of the Maxwell-Boltzmann and the modified speed distributions are presented in Table 1.2, and the two p.d.f. functions are shown in Fig.1.16.

1.2.1.4 Finding next collision location

Once the particle's direction and velocity are calculated, Molflow+ uses a ray-tracing algorithm [6] to find the next collision point with a facet. Then, depending on the sticking coefficient (see below) of the hit facet, the particle is either pumped and the algorithm proceeds to generate the next particle from source, or it rebounds from the facet:

1.2.1.5 Pump or bounce

The user can define pumps in the system by assigning sticking factors s (probability of absorption for the impinging particle) or pumping speeds S to certain facets. Assuming equilibrium, the two quantities can be converted to each other:

$$S[m^3/s] = s \cdot \frac{1}{4} v_{avg}[m/s] \cdot A[m^2]$$

where v_{avg} is the average molecule speed and A is the facet area. When a collision happens with a pump facet, a uniformly distributed pseudo-random number is generated, and if it is smaller than s , the test particle is removed from the system and a new one is generated at a source facet. Otherwise the test particle rebounds.

1.2.1.6 Registering hits

Pressure on a surface comes from the momentum change rate of particles colliding with it. Every facet (or texture or profile slice, see later) has three counters in memory that we increment by the following quantities:

- We increment N_{hit} , the "number of Monte Carlo hits" counter by 1. Knowing the number of MC hits on a facet will allow us to calculate the impingement rate, moreover it provides important information on the statistical accuracy of our results.
- To $\sum I_{\perp}$, the "total orthogonal momentum change" counter we add the orthogonal momentum change, $mv_{\perp} = mv \cos \theta$ of the incoming and outgoing particle (θ is the test particle's incident angle, v is its speed and m is its mass)¹¹. We will need this counter to calculate the pressure.
- We also use a third counter, $\sum \frac{1}{v_{\perp}}$ to store the sum of the reciprocals of the orthogonal speed components, $\frac{1}{v_{\perp}} = \frac{1}{v \cos \theta}$. We need to store this quantity for the calculation of particle density near the facet, as explained later.

1.2.1.7 Rebound

Once we have registered the hit by incrementing the counters, unless the particle is pumped, the algorithm proceeds to the next iteration: it assigns a new velocity and a new direction for the outgoing particle. The direction is chosen according to the user setting: by default, a facet

¹¹in case of desorption, the outgoing component is added, whereas the incoming for absorption and both for reflection

is diffuse so a new random direction is generated following Knudsen's cosine law, but the user can also choose specular reflection.

A particle can interact with the wall which results in energy and momentum exchange. Its new speed thus depends on the thermal accommodation coefficient, A_{acc} . By default, total thermalization ($A_{acc} = 1$) is assumed, in which case - regardless of the incident velocity - a new velocity v_{wall} is generated according to the frequency-modified speed distribution, which takes into account the facet's temperature.

The accommodation coefficient is defined differently across literature (for example, [24] uses my notation comparing the incident and outbound energies as opposed to [25] comparing the incident and outgoing velocities whereas [26] is distinguishing an orthogonal and tangential accommodation factor). I consider the thermalization coefficient as the ratio of kinetic energy change compared to that of total thermalization (definition above). In this case, the new velocity is calculated as:

$$v_{new}^2 = v_{old}^2 + A_{acc}(v_{wall}^2 - v_{old}^2)$$

where v_{old} is the incident velocity and the thermal accommodation coefficient, A_{acc} must be between 0 and 1.

Once the new speed and direction are chosen, then the orthogonal momentum change and speed reciprocal counters of the facet are incremented again to take into account the outgoing particle's impulse and reciprocal speed.

It is worth mentioning that this algorithm is executed parallel on every CPU core present in the system (one core taking care of one test particle at a time), and taking advantage of the linearity of UHV systems, the counters are summed in the end.

1.2.1.8 Calculating physical quantities

Using the counters above, the following physical quantities can be calculated:

1.2.1.8.1 Impingement rate The impingement rate on a wall is the number of collisions, N_{real} per second per unit area. To calculate it, we need the following quantities:

- The number of incoming physical molecules per second each test particle represents, $K_{\frac{real}{virtual}}$
- The number of test particle hits on the given facet, N_{hit}
- The area of the facet, A

With all the above, the impingement rate can be expressed:

$$z = \frac{N_{real}}{dt \cdot A} = \frac{N_{hit} \cdot K_{\frac{real}{virtual}}}{A}$$




Figure 1.13: For a given position within the volume, only particles with high enough orthogonal velocity will hit the surface in the next Δt period. Redrawn based on [27]

Molflow+ applies a few more coefficients to correct for double-sided and half-transparent facets¹², but for simplicity they are omitted here.

1.2.1.8.2 Pressure When particles desorb, absorb or bounce on a facet, their rate of momentum change exerts a force on the wall, which - divided by the area - is the pressure. Since we have previously stored the sum of momentum change rates in a dedicated counter, to obtain the pressure all we have to do is to scale it up by the $K_{\frac{real}{virtual}}$ ratio (which already includes the time normalization to one second) and then divide by the facet area:

$$p = \frac{\sum dI}{dt \cdot A} = \frac{\sum I_{\perp} \cdot K_{\frac{real}{virtual}}}{A}$$

1.2.1.8.3 Density It is not straightforward to calculate density directly in Molflow+: in an event-driven simulator, the main idea is to simulate events only on walls and ignore everything in between, whereas for a direct density measurement we would need to count the number of test particles in a test volume. It is, however, achievable through a simple derivation¹³ that connects the density in a volume with the number of particles crossing a surface nearby:

As shown in Fig.1.13, a molecule traveling with velocity v_{\perp} will hit a surface within time Δt if it is closer than $v_{\perp} \Delta t$ to the wall. Therefore, the number of collisions for molecules traveling with velocity v_{\perp} during the next Δt period is:

$$\# \text{collisions, } v_{\perp} = \frac{N}{V} A \Delta t v_{\perp}$$

where N/V is the density of molecules and A is the surface area. Since the molecular speeds are distributed, we integrate this expression over the velocity distribution function $f(v_{\perp})$ ¹⁴:

$$\# \text{collisions} = \frac{N}{V} A \Delta t \int_0^{\infty} v_{\perp} f(v_{\perp}) dv_{\perp} = \langle n \rangle_{\text{volume}} A \Delta t \langle v_{\perp} \rangle_{\text{gas}}$$

¹²these half-transparent facets serve to represent periodic structures, for example grids, where – regardless of the hit location – there is a certain, constant probability of a particle being able to cross to the other side

¹³This derivation can be found in many places in literature, one example is [27]

¹⁴For clarity, here I ignore the molecules that move away from the surface ($v_{\perp} < 0$). Otherwise I would integrate from $-\infty$ to $+\infty$ and divide the result by 2

And since the impingement rate is the number of collisions per unit time and area, we can write:

$$\langle n \rangle_{\text{volume}} = \frac{\# \text{collisions}}{A \Delta t \langle v_{\perp} \rangle_{\text{gas}}} = \frac{z_{\text{surface}}}{\langle v_{\perp} \rangle_{\text{gas}}}$$

Here z_{surface} is the impingement rate on the surface, and $\langle v_{\perp} \rangle_{\text{gas}}$ is the average of the orthogonal speed components of those particles that move toward the surface. We've already calculated the impingement rate above. Obtaining $\langle v_{\perp} \rangle_{\text{gas}}$ requires more steps, though: unless we use test volumes or other tools, we won't be able to compute it directly. What we can calculate, however, is the speed distribution on an adjacent facet. Using the definition $\int_0^{\infty} f(v)_{\perp \text{gas}} dv = 1$ we can write:

$$\langle v_{\perp} \rangle_{\text{gas}} = \frac{\langle v_{\perp} \rangle_{\text{gas}}}{\int_0^{\infty} f(v)_{\perp \text{gas}} dv}$$

$$\langle v_{\perp} \rangle_{\text{gas}} = \frac{\langle v_{\perp} \rangle_{\text{gas}}}{\int_0^{\infty} f(v)_{\perp \text{gas}} v \frac{1}{v} dv}$$

Then using eq.1.6:

$$\langle v_{\perp} \rangle_{\text{gas}} = \frac{1}{\int_0^{\infty} f(v)_{\perp \text{gas}} v C \frac{1}{v} dv}$$

Then, using eq.1.5 we can write:

$$\langle v_{\perp} \rangle_{\text{gas}} = \frac{1}{\int_0^{\infty} f(v)_{\perp \text{coll}} \frac{1}{v} dv} = \frac{1}{\langle \frac{1}{v_{\perp}} \rangle_{\text{coll}}}$$

At this point we moved from using $\langle v_{\perp} \rangle_{\text{gas}}$ to $\langle \frac{1}{v_{\perp}} \rangle_{\text{coll}}$ which can be sampled by Molflow+. Therefore the density on a facet, using the $\sum \frac{1}{v_{\perp}}$ counter, will become:

$$\langle n \rangle_{\text{volume}} = \frac{z_{\text{surface}}}{\langle v_{\perp} \rangle_{\text{gas}}} = z_{\text{surface}} \cdot \langle \frac{1}{v_{\perp}} \rangle_{\text{coll}} = \frac{N_{\text{hit}} \cdot K_{\text{virtual}}^{\text{real}}}{A} \cdot \frac{\sum \frac{1}{v_{\perp}}}{N_{\text{hit}}} = \frac{\sum \frac{1}{v_{\perp}} K_{\text{virtual}}^{\text{real}}}{A}$$

1.2.1.8.4 Average velocity As described in the previous point and also shown in Table 1.2, the average speed in the gas can be calculated:

$$\langle v_{\perp} \rangle_{\text{gas}} = \frac{1}{\langle \frac{1}{v_{\perp}} \rangle_{\text{coll}}} = \frac{N_{\text{hit}}}{\sum \frac{1}{v_{\perp}}}$$




Figure 1.14: Solid angles for small and large incident angles: even in uniform directional distribution, more particles would reach the target facet (in the center of the hemisphere) from large incident angles




Figure 1.15: Incident angle distribution on a facet in equilibrium state, for a gas of 28 g molar mass and 300 K temperature

1.2.1.8.5 Angular profiles In Molfow+ we can also calculate the directional distribution of incident molecules on a surface. A 100-value histogram counts the absolute value of the incident angle (ie. two particles coming from opposite sides with the same incident angle will count the same, since facets don't have default directions). The incident angle range of 0..90 degrees is distributed over the 100 profile values. One important note is that - similar to velocities - there is a difference between angular distribution in a volume of gas and the angular distribution of particles that cross or hit a surface during a given time. This is because particles with large incident angles (incoming almost parallel to the surface) might originate from a larger solid angle than those that arrive perpendicularly. This is demonstrated in Fig.1.14.

Therefore a normalization is included (as a user option) to correct for this effect: the value for each bin of the angular distribution is normalized by the corresponding solid angle (the solid angle for incident angle θ is $d\omega(\theta, \phi) = \sin\theta d\theta d\phi$, therefore the normalization consists of a division by $\sin\theta$, then a second division to assure that the sum of the distribution is 1). In equilibrium, without this correction, the distribution would be $p(\theta) \approx \sin(\theta) \cos(\theta)$, and with the correction, it becomes $p(\theta) \approx \cos(\theta)$ which is what Knudsen's cosine law stipulates. This is shown in Fig.1.15.




Figure 1.16: Incident speed and orthogonal velocity distribution on a facet in equilibrium state

1.2.1.8.6 Speed profiles We can also sample the distribution of the incident velocities of test particles on a facet. The 100 histogram bins correspond to the speed range $0.4 \cdot v_{\text{most_probable}} = 0.4 \cdot \sqrt{\frac{2RT}{M_{\text{kg/mole}}}}$. As with the angles, a surface-volume conversion is possible: in this case each value of the distribution is divided by the speed it belongs to, then the distribution is renormalized (effectively converting from $f(v)_{\text{coll}}$ to $f(v)_{\text{gas}}$ of Table 1.2). There are two types of speed profiles (see Fig. 1.16): one for the absolute value of the incident speed (giving information on the gas temperature), and one for the orthogonal component (which determines the pressure).

1.2.1.9 Error of the results

When assessing the statistical error of Molflow's results, we can apply a scientific (quantitative) or a practical (qualitative) approach. I first present the quantitative. In mathematical terms, a Monte Carlo simulator is approximating a physical result by carrying out a given number of experiments and counting the number of positive outcomes.

In case of calculating the transmission probability of a tube, for example, each traced test particle would be one experiment, and the positive outcome would be that the particle leaves the system on the side which is opposite to where it came from. The transmission probability p would correspond to the probability of positive outcomes: $p = n/N$ where n is the number of test particles leaving on the opposite side, and N the total number of traced test particles.

In case of the pressure at a certain location, for simplicity we'll assume that the gas has a Maxwellian behavior (i.e. the system is isothermal and beaming effects are not important). In this simpler case the pressure is directly proportional to the number of hits at the location. So – omitting some normalizing constants – the pressure can be expressed as a function of $p = h/N$, where h is the number of hits at that point, N is the total number of hits and thus p is the probability that a hit happens at the given location.

The two examples above can be translated to probability theory: the discrete probability distribution of the number of successes in a sequence of N independent yes/no experiments, each of which yields success with probability p , is the binomial distribution.

The expected value of positive outcomes of such a series of experiments is $N \cdot p$, and the

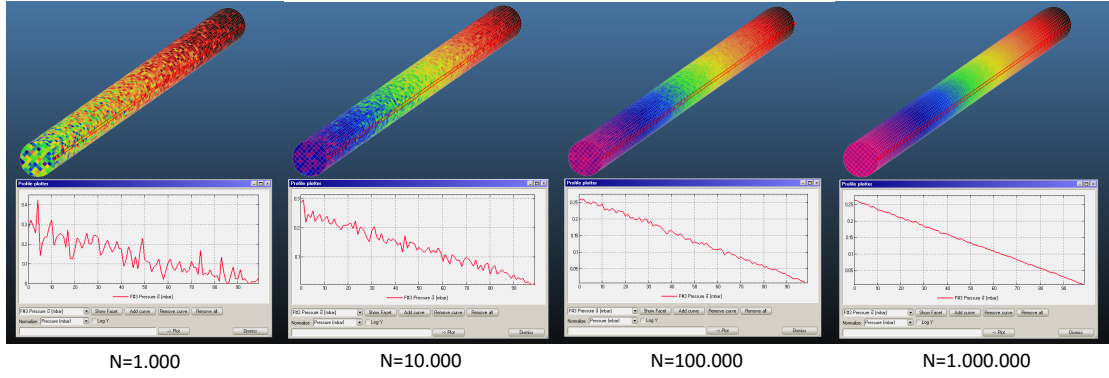


Figure 1.17: Texture and profile results for a vacuum tube with desorption at the left and pumping at both sides, as a function of the number of traced particles.

variance of the result is $N \cdot p \cdot (1 - p)$.

As suggested by Suetsugu [15], for Monte Carlo simulations we can assume a very large number of experiments, where using the notation of the two examples above, $p \sim n/N$ or $p \sim h/N$. In this case the normalized standard deviation, expressed as a percentage of hits (positive outcomes), can be derived from the variance formula above as:

$$\sigma_n (\%) = \frac{\sigma}{n} \cdot 100 = \sqrt{\frac{1}{n} \left(1 - \frac{n}{N}\right)} \cdot 100 \quad (1.7)$$

Suetsugu also correctly points out that as we trace a particle through consecutive hits, not each experiment is independent from the others – statistical fluctuations in adjacent locations are correlated distorting the above estimation. I would argue that due to this, the exact error depends on the geometry of each simulation and as such there is no universally valid formula to estimate it. However, from eq.1.7 we can retain that the error of $p \sim n/N$ scales with $1/\sqrt{N}$: for example, to reduce the error of the calculated pressure to its half requires to run the simulation four times longer.

Since the quantitative formula only gives an approximation, it is important to mention the qualitative approach that most Molflow+ users go along with when simulating a vacuum system. As mentioned above, post-processing visualization tools are updated every second as simulation runs, which – as shown in Fig.1.17 – means in practice that profiles and textures become smoother. In most cases the user has an expectation of the pressure profile, so statistical scattering can be visually identified through inhomogeneities of drawn textures.

Staying at the example in Fig.1.17 for example, the engineer using the software either knows that the profile from desorption to the pump should be linear, thus he/she can fit a straight line on the obtained plot and isolate the statistical error. Otherwise the user can argue that the system is axisymmetric, therefore texture cells at the same longitudinal coordinate should have the same color, and deviation from it must be due to statistical scattering. The dependence of the error can be verified to vary with $1/\sqrt{N}$. In the above example I calculated




Figure 1.18: Normalized standard deviation of Molflow+ results as a function of the simulated test-particles

the deviation of the number of hits ¹⁵ from the theoretical solution (which in our example is a pressure profile of a straight line) and plotted it as a number of the traced test particles in Fig.1.18.

The plot, containing a fit with the analytic formula confirms that the error squares approximately with the inverse of the square root of the number of hits.

1.2.2 Time-dependent Molflow+ simulations

Once we keep track of the molecule speeds, apart from the possibility to calculate the pressure in a direct way (based on particle momentum change), it also becomes possible to determine the time when each hit happens, which opens the way to simulations in the time domain. On every facet, instead of calculating one global (steady-state) statistics using every registered hit, we take into account only those that happen at a certain moment - thus we can obtain the physical quantities for that instant.

The algorithm is still event-driven, and is essentially the same as that of the steady-state simulations explained above, but some corrections and fine-tuning is required.

1.2.2.1 Parameters

When our simulation treats time-dependent problems, both the outgassing and the facets' physical parameters, such as sticking factor, opacity, etc. can vary over time. This is included in Molflow+ through parameters (see Fig.1.19): the user references a parameter (by name) at the given facet's property list, then he proceeds to define that parameter by entering a list of time moments and the parameter's values at those instants. For every event on that facet, Molflow+ will use the actual, time-dependent value, obtained by linear interpolation within the user-defined defined time range, and by constant extrapolation outside. This is done directly for sticking factors and opacity, and with an intermediate step (integration) for time-dependent outgassing, see later.

¹⁵I extracted the number of Monte Carlo hits instead of the pressure because the latter is already normalized by the number of traced particles

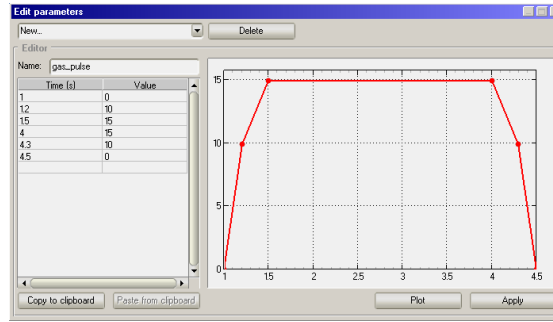


Figure 1.19: Time-dependent parameter definition for a gas pulse in Molflow+

1.2.2.2 Moments and time window

We aim to simulate a large number - possibly billions - of hits. In the original, steady state algorithm, this doesn't cause a memory problem since these hits aren't stored individually: they only increment counters on facets.

In the time-dependent mode, however, the direct way to gain statistics would be to store a list of hits along with the time they occur for every facet¹⁶. This would, consequently, mean that the longer the simulation runs, the more memory we'd need. But limiting the number of hits, and along with it the statistical accuracy is not an option, therefore we have to implement an indirect approach.

We use binning: we ask the user to decide in advance at what moments he is interested in the pressure and density profiles. This is usually a series of moments t_1, t_2, \dots, t_N with equal or logarithmically increasing differences, ie. the time range when the dynamics of the system is interesting. We will then create three counters, for every moment and every facet:

- $N_{hit,i}$, the number of MC hits on the facet at moment i
- $\sum I_{\perp,i}$, the sum of orthogonal momentum changes of particles hitting the facet at moment i
- $\sum \frac{1}{v_{\perp,i}}$, the sum of the reciprocals of the orthogonal speed components of particles hitting the facet at moment i

Consequently, if a hit happens at t_i , it will increment the i th counter. That way the pressure and other physical quantities at that moment can be calculated similarly to the steady state using values of the counter with the same index, and the user has several tools to visualize the dynamic behavior of the system.

One problem to tackle though is that no hit will happen exactly at t_i with infinite precision. Therefore we introduce the term time window, serving as a tolerance: we will treat all hits between $t_i - \frac{t_{window}}{2}$ and $t_i + \frac{t_{window}}{2}$ to have happened at the i th moment and accordingly, increase the i th counter.

¹⁶As my simulations are event-driven, there is no default "time step" to store results for.




Figure 1.20: Desorption rate and number of desorbed molecules during injection

The time window can be set by the user. In practice, it has to be small enough to reveal the fast pressure changes in the system, but large enough to contain a number of hits that is sufficient for the desired statistical accuracy. The good choice is one that matches the pressure change speed of the simulated system, and sometimes it needs to be found by trial and error across multiple simulations. This is demonstrated during the validation in section 1.2.5.1.

1.2.2.3 Particle generation

When an outgassing is time-dependent, we must:

- Generate particles so that the distribution of their time of insertion follows the user-defined desorption parameter.
- Know how many real molecules one test particle represents.

Both tasks can be solved by integrating the user-defined desorptions on the range $t = 0..t_{\text{last}}$, for all outgassing facets, where t_{last} is the last user-defined moment where the pressure needs to be calculated (after which we are not interested in the behavior of the system). We store the integrated desorption function - as illustrated in Fig.1.20. -, which will also tell us the total number $N_{f,real} = \int_0^{t_{\text{last}}} \frac{dN_{f,real}(t)}{dt} dt$ of desorbed particles, on facet f .

For each generated virtual particle, Molflow+ then chooses a starting location. The probability of choosing a certain facet f as source is proportional to $N_{f,real}$, just like in steady-state. Once the source facet is chosen, the start position of the test particle is generated randomly with uniform distribution on its surface. To decide the desorption time, we generate a random number uniformly between 0 and $N_{f,real}$, then interpolate the desorption time in the integrated desorption function corresponding to that random number. That way, the particles will enter the system following the user-defined distribution, and when the simulation traced N_{virtual} test particles, each one will represent

$$K_{\frac{real}{virtual}} = \frac{\sum_f N_{f,real}}{N_{\text{virtual}}}$$

real molecules.

1.2.2.4 Radioactive decay

As we already keep track of the flight time for every particle, it is straightforward to add radioactive decay to the algorithm: upon particle generation, the MC algorithm assigns the instant of decay, t_{decay} for each particle of user-defined half-life $T_{1/2}$ entering the system, with the method described hereafter. The probability that the particle's lifetime will be t is given by the exponential probability density function:

$$P(t) = \frac{1}{\tau} \exp\left(-\frac{t}{\tau}\right)$$

where $\tau = T_{1/2} / \ln(2)$. Random numbers following this distribution are then generated by the inversion method: a random number rnd , generated uniformly between 0 and 1 must equal the cumulative distribution function:

$$rnd = 1 - \exp(-t/\tau)$$

The solution of this equation for t is:

$$t = -\tau \cdot \ln(1 - rnd)$$

Using rnd in place of $1 - rnd$, we can express the time of decay:

$$t_{\text{decay}} = t_{\text{des}} - \tau \cdot \ln(rnd)$$

where t_{des} is the time at which the particle enters the system. Each time a particle hits a facet, the time of the hit t_{hit} is compared to t_{decay} , and the particle is eliminated at the first hit when $t_{\text{hit}} > t_{\text{decay}}$. As Molflow+ is event-driven, the exact location of the decay (in the volume, between two hits) has no interest because statistics is gathered only on surfaces. One exception is if the user sets up counter facets: these transparent facets in the volume execute the lifetime check of all passing particles, therefore - if needed - they allow visualization of the decay en route between the vacuum chamber walls.

1.2.2.5 Wall sojourn time

Once the time when the particle enters the system is known, for every subsequent hit the flight time is incremented by the time elapsed since the last collision (which is calculated by dividing the distance from the previous hit location with the particle's speed). That way for every hit the time it occurs is known.

One phenomenon to take into account for making these hit times more accurate is the residence time τ of a molecule on the surface.

The physical process is essentially the same as in case of the radioactive decay: each molecule on the surface has a certain, constant probability of escape, resulting in an exponential probability distribution of the residence time. Its average value has the expression[28]:

$$\tau = \tau_0 \cdot \exp(-Q/RT_f)$$

where T_f is the facet temperature, τ_0 is the molecule vibration period, and Q is the molar adsorption heat. If the user defines these two parameters (two, because T_f is already defined) then each bounce time is delayed with a sojourn time that is generated following the exponential probability density function.

1.2.2.6 Calculating physical quantities

We can calculate impingement rate, pressure and density similar to the methods in the steady-state simulations, with one remarkable difference: while at steady state, each test particle represented a certain number of physical molecules entering the system *per second*, this time $K_{\frac{real}{virtual}}$ stands for a certain - absolute - number of physical molecules that enter the system during the whole simulated process. Therefore, in time-dependent calculations, the dt in the pressure and impingement rate formulas is not included anymore in $K_{\frac{real}{virtual}}$, and has to be replaced by the time window, as that represents the duration during which hits are considered to happen at the same moment and increase the same counter. Therefore, for the i th moment, physical quantities can be calculated as follows:

1.2.2.6.1 Impingement rate Per definition, the impingement rate is the number of collisions per second per unit area, however, our counter contains hits for the period of the time window, so to have the per second hit rate, we divide by t_{window} to get:

$$z_i = \frac{N_{real,i}}{dt \cdot A} = \frac{N_{hit,i} \cdot K_{\frac{real}{virtual}}}{t_{\text{window}} A}$$

1.2.2.6.2 Pressure Similar to the impingement rate, we need to divide by t_{window} which represents dt in the pressure formula at the i th moment:

$$p_i = \frac{\sum dI}{dt \cdot A} = \frac{\sum I_{\perp,i} \cdot K_{\frac{real}{virtual}}}{t_{\text{window}} A}$$

1.2.2.6.3 Density To calculate the density at moment i , we start from the impingement rate and the sum of reciprocal orthogonal velocities at that moment, and apply the correction by t_{window} :

$$\langle n \rangle_{\text{volume},i} = \frac{z_{\text{surface},i}}{\langle v_{\perp} \rangle_{\text{gas},i}} = z_{\text{surface},i} \cdot \langle \frac{1}{v_{\perp}} \rangle_{\text{coll},i} = \frac{N_{hit,i} \cdot K_{\frac{real}{virtual}}}{t_{\text{window}} A} \cdot \frac{\sum \frac{1}{v_{\perp,i}}}{N_{hit,i}} = \frac{\sum \frac{1}{v_{\perp,i}} K_{\frac{real}{virtual}}}{t_{\text{window}} A}$$

1.2.2.6.4 Average velocity Like at steady-state, the average speed at moment i can be calculated:

$$\langle v_{\perp} \rangle_{\text{gas},i} = \frac{1}{\langle \frac{1}{v_{\perp}} \rangle_{\text{coll},i}} = \frac{N_{hit,i}}{\sum \frac{1}{v_{\perp,i}}}$$

1.2.3 Post-processing in Molflow+

Although results of Molflow+ simulations can be exported for external processing, the design idea was to create an all-in-one vacuum simulation tool, therefore some post-processing is provided within the graphical user interface. Since these post-processing tools are often used in the thesis' images, they are briefly presented here.




Figure 1.21: local U, V coordinate system of a facet and the longitudinal profile slicing

1.2.3.1 Profiles

Each facet has a local, 2D coordinate system, with the origin being one of the facet's vertices, and its local U, V vectors delimiting the bounding box, orthogonal to each other. By default, the U vector is aligned with the facet's longest side, though this is adjustable by the user.

As shown in Fig.1.21, a profile splits the facet to a given number of slices (currently 100) along the U axis, and samples the pressure, density, number of Monte Carlo hits or the impingement rate independently for each slice. The goal is to plot – within the user interface – the evolution of a physical quantity along a direction.

There are three special profiles, presented earlier which count a distribution of a quantity for the entire facet:

- Angular profiles sample the distribution of the incident angles
- Speed profiles sample the distribution of the incident molecular speed
- Orthogonal velocity profiles are similar to speed profiles, but sample only the component of the molecule speed orthogonal to the facet

As an example, for a simple tube with a length/radius ratio of 10, I desorb gas at one side with $Q = 1 \times 10^{-4}$ mbarl/s and plot the pressure, speed and angular distributions on a wall facet in Fig.1.22, after 500.000 desorbed test particles.

1.2.3.2 Textures

Textures – similar to profiles – can sample a physical quantity within a facet, but split it in not one but two directions (along the local U and V vectors on the plane of the facet). The resolution (cell size) is chosen by the user. Results can be read with a Texture Plotter tool numerically, or visualized by a linear or logarithmic color scale, with parameters adjustable by the user.

As most physical quantities depend on the cell area, it must be calculated carefully for cells on the edges. The area of these partial cells can be very tiny compared to full cells, and therefore they might receive only a very small number of MC hits during the whole simulation. For plotting and autoscaling purposes, it is the best to ignore their values – this can be done

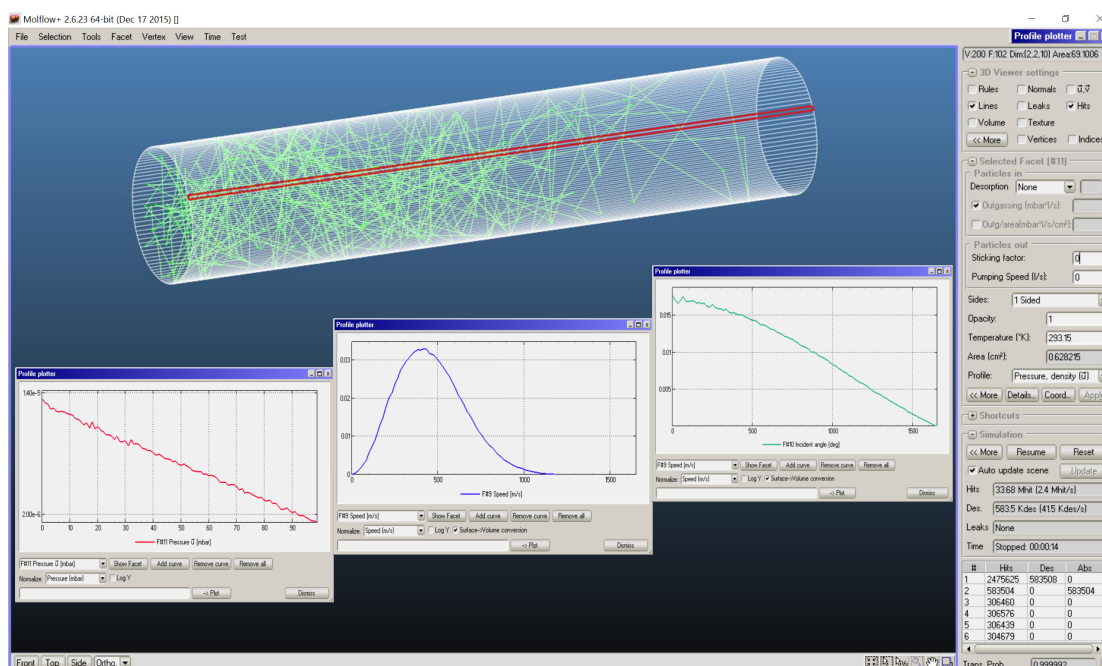


Figure 1.22: Pressure, speed and angular profiles for a simple tube within Molflow's interface

automatically in Molflow+ if their area is smaller than a given percent of whole cells (as in Fig.1.23).

Visualization is supported for the pressure, the impingement rate and the density.

1.2.3.3 Direction vectors

For each texture cell, the user has the option to evaluate the prevailing molecular direction. This is the average¹⁷ of all velocity vectors incident to the surface (or crossing it – in case of transparent facets). This is shown in Fig.1.24 for a simple effusion example, where a gas is injected to an infinite, empty volume (modelled by a box with sticking walls) with an outgassing rate of $Q = 1 \times 10^{-4}$ mbar l/s.

The importance of these direction vectors is emphasized when simulating non-isothermal systems, such as the hot filament / cold walls example in Fig.1.25, where we can clearly observe a deviation from the isotropic directional distribution of equilibrium gases.

1.2.3.4 Formulas

We also let the user get more insight into the system by allowing him to define formulas, which are symbolic expressions referring to facet parameters. One can refer to the number of hits, the pressure, impingement rate, etc. of a facet or a group of facets, and even extract global

¹⁷averaging is done vectorially: two velocity vectors in the opposing directions cancel each other. Consequently, in equilibrium, this average velocity vector is zero.

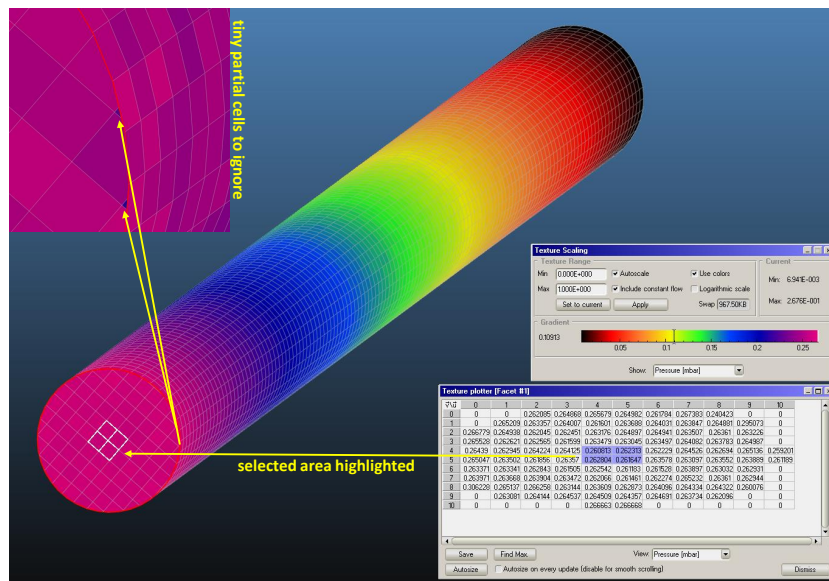


Figure 1.23: Textures and mesh cells within Molflow+

simulation parameters, such as the mean free path (in this sense, the average distance of flight between two wall hits), mean pumping path (average flight distance from desorption to pumping or to radioactive decay point), total particle influx of the system or the gas mass.

As common mathematical operators, trigonometric functions and physical constants are also included, among others, these formulas are extensively used for conductance calculations, as in Fig.1.26.

1.2.3.5 Time-dependent visualization tools

When performing time-dependent simulations, all the post-processing tools above are available. The user can step through the predefined moments and the on-screen visualizations are updated for that given moment. Succession of these visualized system states can be used as frames to compile short videos showing the system behavior.¹⁸

In the time-dependent mode, two additional tools are available:

- A *Timewise Plotter* (Fig.1.27) that graphs the pressure profile evolution in a certain direction across different moments on a chosen surface.
- The *Pressure Evolution* (Fig.1.28) that shows how the pressure on a profile changes over time: the X scale contains all the user-defined moments, and the Y value is the pressure, density or velocity at that moment.

¹⁸Two examples of such videos are shared on these permalinks:
<https://goo.gl/fgTQR> (Gas expansion and pumpdown following an RF cavity electric breakdown)
<https://goo.gl/Pd31Q6> (Gas expansion and pumpdown in a straight tube)




Figure 1.24: Direction vectors and density texture in an effusion example

1.2.4 Sectioning the geometry

The complexity limits of the systems that we can simulate depend on the computation speed. In general, the more facets we have, the longer the ray-tracing algorithm takes to determine the closest collision point and thus the slower the simulations are. There are some ways, though, to include large or complex vacuum systems and at the same time mitigate the speed impact of the large number of facets.

1.2.4.1 Teleports

Molflow+ allows the definition of periodic boundary conditions. When the ray-racing algorithm finds a collision with such a boundary, the virtual particle is moved to an other, referenced facet (the opposite side in case of a periodic B.C.) – hence the name "teleport". If we can replace a large geometry with only one pattern that is repeated, we reduce the number of facets accordingly.

In the example below, a simple RF structure is simulated with gas injection on one side and pumping on the other. Two models are built for comparison: one containing the full revolution, and one with only a 45° slice (Fig.1.29). In the latter case, the walls of the slice are set as teleports referencing each other. The total desorption of the slice is one-eighth of the full revolution (1×10^{-5} mbarl/s vs. 8×10^{-5} mbarl/s) resulting in the same influx density, and the pumping speed is also one-eighth (in practice, this results in the same sticking factor of 1 for both surfaces). As the surface normals are tilted by 45 degrees compared to each other, when a teleport happens, the particle direction is also rotated (Fig.1.30) – to correspond to the real case.

The resulting simulation speed improvement is summed in Table1.3, and the resulting pressure profile (Fig.1.31) is the same.

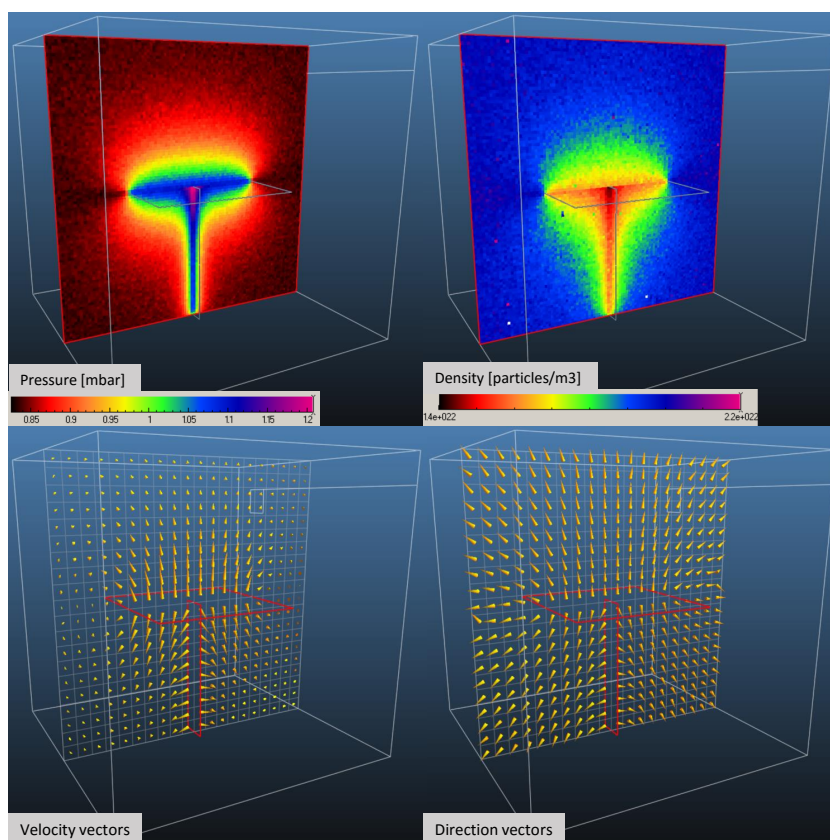


Figure 1.25: Pressure, density, velocity and direction vectors in a non-isothermal system at equilibrium; with a hot filament at $T = 800$ K at the center and walls at room temperature

One other use of teleports is to include a reusable geometry feature only once. This can be for example a pumping group. Molflow+ can memorize where test particles were teleported from, and if molecules find their way back to the teleporting surface, they are transported back to the origin. This – as shown on the example in Fig1.32 is essentially a single-to-multiple logical assignment, once again reducing the number of facets.

1.2.4.2 Structures

As stated before, the ray-tracing checks for collisions with all facets and then among the collisions found, it selects the shortest (although due to optimizations the exact implementation

Table 1.3: Periodic B.C. performance results

	Full revolution	45° slice with teleports
Number of facets	25491	3492
Traced molecules/sec	10.5	36
Hits/sec	8200	28000

$$\frac{\text{Outgassing rate [mbar}\cdot\text{l/s]}}{\text{Conversion factor mbar}\cdot\text{l/s} \rightarrow \text{Pa}\cdot\text{m}^3/\text{s}} \cdot \frac{\text{Boltzmann constant}}{100} \cdot \frac{\text{Temperature on facet 5 (Gas source)}}{(\text{Kb} \cdot \text{T5})} \cdot \frac{\text{All absorbed molecules on the exits}}{\text{SUM(A, S1)}} \cdot \frac{\text{Total number of desorbed molecules}}{\text{SUMDES}}$$

Conversion factor mbar.l/s → Pa.m³/s Temperature on facet 5 (Gas source) Facet group 1 (containing all exits)

Figure 1.26: Molflow+ formula for calculating the outgoing flux of particles on a vacuum part

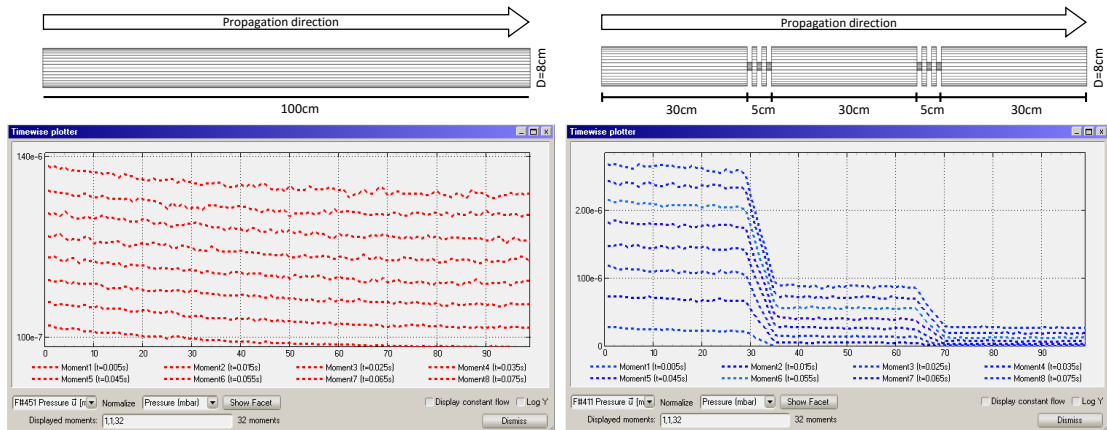


Figure 1.27: Timewise Plotter in Molflow+ 2.6 showing the pressure evolution at eight different moments in a tube without (left) and with (right) an acoustic delay device

is somewhat different). In Molflow, it is possible to speed up the process by dividing a – usually linear – geometry to distinct sections, and defining them as separate structures. In this case collisions are only sought with facets in the actual structure. Connection between the structures is done by link facets. When a test particle crosses such a link facet, its actual structure is updated.

A very simple example is shown in Fig.1.33, where a short pipe is split in half, with a link facet in the middle. The speed advantage is almost linear: in theory splitting to N structures would yield an N times faster simulation, but the speedup is mitigated by the extra collisions with link facets, the fact that not all structures have the same number of facets, and the non-negligible user time spent with setting up the sectioning and the links.

To give a real-life example, the method has been applied for a static analysis of the radioactive

Table 1.4: Teleport performance results

	Four pumping ports	One p.p. with teleport
Number of facets	701	601
Traced molecules/sec	1.600	2.200
Hits/sec	560.000	776.000




Figure 1.28: Pressure evolution plotter window showing the pressure response signal on a target after a desorption pulse at $t = 0$




Figure 1.29: Periodic structure (left) and its full equivalent (right)

ion-beams facility ISOLDE. Working together on the problem with my colleague M. Maietta, as explained in detail in her MSc thesis [29], the model of the ISOLDE vacuum line was split to 9 different structures, as shown in Fig.1.34, making the simulation approximately four times faster, as shown in Table 1.4.

1.2.4.3 Simulating by parts (Convolution method)

When simulating linear vacuum systems with powerful pumps, a relatively frequent problem is the large pressure difference between different parts. When Molflow+ traces molecules, a test particle always represents the same number of physical particles (or in case of steady-state simulations, a flux of a certain number of particles per second). This means that if there is a part of the geometry where the pressure is orders of magnitude higher than at other parts, most of the ray-tracing will "focus" on that high-pressure part.

This is shown on an example in Fig.1.35 showing a linear vacuum system with a desorption of 1×10^{-3} mbarl/s on one side and a series of pumps with sticking factor 1 placed downstream. Most of the test particles are caught by the pumps before they can reach the low-pressure region on the right side. Since the pressure difference is of 5 orders of magnitude (1×10^{-3} mbar

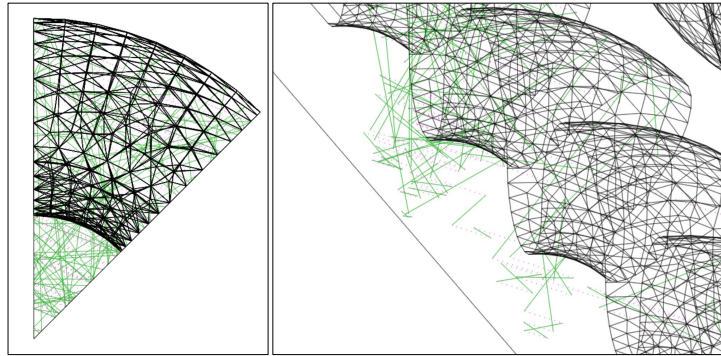


Figure 1.30: Particle trajectories rotating as following teleports




Figure 1.31: Pressure results for the periodic and the full RF model

vs. 1×10^{-8} mbar), to simulate a single hit on the low-pressure side, 100.000 hits must be calculated on the high-pressure side. This results in an inefficient distribution of test particles, and the statistical scattering will be very high on the low-pressure side.

An idea to tackle this issue is to do two (or even more) different simulations instead of one: If we determine that – as shown in Fig.1.35 – only 0.1% of the particles desorbed from point *A* will reach the middle (point *B*) of the system (we will refer to this ratio as transmission probability), then we could launch a second simulation, desorbing particles only from point *B* to the right side, with a physical desorption rate which is 0.1% of the original. That way the speedup on the low-pressure side is of a factor of 1000. We could apply this method repeatedly, launching a new simulation from every point where the pressure drops significantly (launching a third simulation at point *C* would improve statistics downstream by an additional factor of 100).

When applying this method, we have to take into account the beaming effect, though: in a tubular geometry where there is a dominant molecular direction (from the desorption point to the right, in our case), the molecular motion is not isotropic. In Molflow, we can calculate the distribution of the incident angles at the desired relaunch point through angular profiles, and – if significantly different from isotropic – we can adjust the directional distribution of the new simulation's desorption to one that is similar to what we sampled.




Figure 1.32: The same system modeled with 4 distinct pumping groups (left) and a reused structure with teleports (right). Teleports are marked with yellow dashed lines. The resulting pressure profile (color code, on the nearest side of the pipes) is the same.




Figure 1.33: A simple system split to two structures, with the link facets (red) in the middle

Staying at the ISOLDE vacuum system, this method was applied to calculate the time of flight of radioactive isotopes from point A to C, as shown in Fig.1.36. The method is described in detail in Chapter 6 of [29]. To summarize it, the idea is the following:

If the time of flight from point A to B is described by the (discrete) distribution f , and the T.O.F from B to C is described by g , then the T.O.F from A to C will be a sum of two random variables, and its distribution will be described by the convolution of f and g [30]:

$$(f * g)[n] = \sum_{m=-\infty}^{\infty} f[m] g[n - m]$$

Since due to the pumps upstream of point B, there is a significant pressure drop between points A and B, it is a good idea to – as described before – perform two separate simulations.

- The first simulation will determine the T.O.F. distribution f and the transmission probability from point A to B (distance: 661 cm)
- The second will launch test particles from point B, with a physical outgassing rate




Figure 1.34: Sectioned system of the ISOLDE transfer line, drawn based on geometry from Maddalena Maietta[29]

Table 1.5: Sectioning details

Structure	Number of facets
1	9180
2	699
3	959
4	321
5	192
6	682
7	143
8	741
9	1049
SUM	13966

adjusted by the previously determined transmission probability, and give the T.O.F. distribution g between points B and C (of 690 cm distance)

- The resulting T.O.F. distribution for the whole system will be the convolution of f and g , as shown in Fig.1.37.

1.2.5 Theoretical validation

The original, steady state algorithm of Molflow has been validated and used on many occasions in accelerator design since its first publication over the last 20 years. We concentrate therefore on the new algorithm capabilities:

Table 1.6: Sectioning performance

	No sectioning	9 structures
Traced molecules / sec	2.2	8.2
Hits / sec	4000	17000




Figure 1.35: Linear system with desorption from the left and pumps at locations marked by red. The pressure profile, as calculated after 1 million desorbed test particles, is shown on a logarithmic scale under the geometry, with transmission probabilities marked in yellow at three locations. Statistical scattering is high at low-pressure parts.




Figure 1.36: ISOLDE transfer line used for Time of Flight calculations. Redrawn based on Figure 6.1 of [29]

1.2.5.1 Time-dependent simulations

First we compare a simulation in the time domain against textbook results. We simulate injection and pump-down processes in a cylinder of 10 cm length and 10 cm diameter (thus a volume of $V = 0.785\text{l}$). We inject nitrogen from one side with a rate of $Q = 1 \times 10^{-6} \text{ mbarl/s}$ between $t_{\text{start}} = 0.005 \text{ s}$ and $t_{\text{stop}} = 0.020 \text{ s}$ while constantly pumping at the other side with $S = 5001/\text{s}$. The gas and the walls are kept at 20°C , and the cylinder's length/radius ratio is small enough to treat it as a single volume where conductance doesn't limit the flow of the gas.

While the pump is running, every second it extracts S/V ratio of the N molecules in the system. Including the gas injection, the governing differential equation will be:

$$\frac{dN}{dt} = \frac{dN_{\text{inject}}}{dt} - \frac{S}{V}N$$

Knowing that the pressure - in this case - is proportional to the number of molecules in the system, and using vacuum theory to express the steady-state pressure as $p_{\text{const}} = Q/S$, we can




Figure 1.37: Results after the same computing time, with and without the convolution method. Redrawn based on Figure 6.2 of [29]




Figure 1.38: The 3D model representing the cylinder. The green lines show the test particle trajectories; the transparent facet in the middle (shown in red) samples the pressure.

solve the equation for the injection and the pump-down durations:

$$p(t)_{\text{injection}} = \frac{Q}{S} (1 - \exp(-\frac{t - t_{\text{start}}}{\tau}))$$

$$p(t)_{\text{pumpdown}} = \frac{Q}{S} \exp(-\frac{t - t_{\text{stop}}}{\tau})$$

To simulate the above process, we create a model in Molflow+, presented in Fig.1.38. It consists of 100 sides and 2 cap facets of the cylinder, and an additional transparent facet in the middle to measure the average pressure in the system. The wall facets use diffuse reflection, one cap facet serves as gas source (governed by a time-dependent desorption function, which is $Q = 1 \times 10^{-6}$ mbarl/s between t_{start} and t_{stop} and 0 elsewhere), and the other as the 500 l/s pump, which converts to a sticking factor of 0.54.

Three simulations are performed with a time window of 1×10^{-3} s, 1×10^{-4} s and 1×10^{-5} s respectively. We record the system state at 40, 400 and 4000 equidistant moments between $t = 0$ and $t = 0.04$ s, therefore the distance between the recorded moments corresponds to the




Figure 1.39: Pressure profiles for an injection-pumpdown process with too large, optimal and too small time windows.

time window length in all three cases. We run the simulation with 1 million test particles. This results in approximately 11 million hits, and for reference, the computing takes around 10 seconds on a standard desktop PC. The pressure values for all moments are plotted in Fig.1.39.

Looking at the results we can make two observations. The first is that choosing the time window correctly is important: on the left plot of Fig.1.39 we can see that a large time window will smooth out thus hide the beginning of the injection and the pumpdown. The optimal time window (middle), $t_{\text{window}} = 1 \times 10^{-4}$ s shows a very good match between theory and simulation. On the other end, choosing a too small time window will introduce excessive statistical scattering: to gain the same number of hits (i.e. accuracy) for every moment with a ten times smaller time window, we would need to run the simulation ten times longer.

There is no universally correct size for the time window, but the system's dynamics can give us the order of magnitude: in this case, the characteristic pump-down time of the system is $\tau = \frac{V}{S} = 1.57 \times 10^{-3}$ s, and the optimal time window size was approximately one-tenth of it.

1.2.5.2 Non-isothermal systems

Results of the new algorithm for non-isothermal systems can be demonstrated and checked against theory through a simple vacuum part with two vessels connected by a narrow pipe. One vessel is hot ($T_1 = 450$ K), the other is cold ($T_2 = 100$ K), and the pipe's temperature changes linearly in between. The cubes have a volume of 2 l, the connecting pipe has a 3x3 cm rectangular cross-section and a length of 30 cm.

Gas in this system will exhibit a phenomenon called thermal transpiration: even at equilibrium, due to higher molecular velocities, pressure will be locally higher on the hot side, while the molecule density will be higher on the colder side. The theoretical[31] pressure and density ratios are given by the formulas

$$p_1/p_2 = \sqrt{T_1/T_2}$$

$$n_1/n_2 = \sqrt{T_2/T_1}$$

The Molflow+ model shown in Fig.1.40 consists of 36 facets. The pipe connection is split to six 5 cm long sections to model the temperature gradient in 50 K steps. All facets have textures




Figure 1.40: Pressure and density at equilibrium in a non-isothermal system

Table 1.7: Analytic and Molflow+ results for the thermal creep problem

Temperature	Pressure	Density
$T_1 = 100 \text{ K}$	$p_1 = 2.89 \times 10^{-5} \text{ mbar}$	$n_1 = 2.10 \times 10^{18} \text{ m}^{-3}$
$T_2 = 450 \text{ K}$	$p_2 = 6.13 \times 10^{-5} \text{ mbar}$	$n_2 = 9.91 \times 10^{17} \text{ m}^{-3}$
Theoretical ratio:	Molflow+ ratio:	Molflow+ ratio:
$\sqrt{T_2/T_1} = 2.121$	$p_1/p_2 = 2.122$	$n_2/n_1 = 2.115$

with 1x1 cm cell size to sample and visualize the pressure, impingement rate and density. One side facet has a $Q = 1 \times 10^{-7} \text{ mbar l/s}$ outgassing rate and an $s = 10^{-6}$ sticking factor, so low that the system is in quasi-equilibrium (each inserted test-particle has an average of 12 million hits all around the system before being pumped). 2 billion hits were simulated, taking about 10 minutes of computing time. The pressure and density ratios in table 1.7 show a good match between simulation and theory.

1.2.6 Experimental validation (CERN)

In 2014, when the new time-dependent code was first introduced internally at CERN, one of its immediate practical applications was the simulation of the AWAKE plasma cell opening process. In the AWAKE experiment[32], a plasma cell must be periodically opened to allow a laser pulse to pass through, which in turn allows the pressurized Rubidium in the cell to propagate down the vacuum line. The pressure gradient created by this process has an undesired defocussing effect, therefore precise modeling of the system is necessary. The new code was used to simulate the process, and subsequently an experimental station was assembled at CERN to test the prototype of the system.

1.2.6.1 Experimental setup

The test station, as shown in Fig.1.41, consists of a 4 m long, 40 mm diameter cylinder that represents the plasma cell. It is separated from the vacuum line by a VAT gate valve. The vacuum line consists of three approx. 7 m long, 80 mm diameter stainless steel tubes. Pumping is performed at the vacuum side of the gate valve by a 500 l/s turbo-molecular pump. There were additional pumps installed on the system which helped the initial pump-down process but were separated during our tests.




Figure 1.41: Overview of the experimental bench at CERN (not to scale)




Figure 1.42: Parts of our experimental bench at CERN

The experiment is performed in a way to simulate the actual cycle the AWAKE plasma cell goes through:

- We pressurize the plasma cell prototype to 1×10^{-2} mbar, the reference pressure of the AWAKE plasma cell, through the variable leak valve at $P6$
- By pressing a button on its controller, we open the gate valve (actuated by compressed air), allowing gas to flow out of the pressurized cell
- Approximately one second later, we close the gate valve
- During a 30 second pause, the gas that entered the vacuum line is pumped down, and the cell is pressurized to 1×10^{-2} mbar again

During the cycles, pressure is measured at each interconnection of the 7 m pipes (referred to as $P1$, $P2$ and $P3$), on each side of the gate valve ($P4$ and $P5$), and also at the middle of the plasma cell prototype ($P6$). We have placed IKR070 (Penning) gauges of range $10^{-8}..10^{-3}$ mbar to $P1$, $P2$ and $P3$ where pressure is expected to be low at all times, IMR260 (all-range) gauges of range $10^{-6}..1$ mbar at $P5$ and $P6$ since those points are pressurized to 1×10^{-2} mbar, and both types of gauges at $P4$ which are at low pressure for most of the cycle but reach a high-pressure peak after the opening of the valve. Sampling of the gauge voltages is performed at 19.2 kHz

and converted to pressures during post-processing, taking into account the different gauge sensitivities for the different types of gases used: we performed our tests with either air or Krypton (choosing the latter because its mass is close to that of Rubidium).

For each test run we repeated the cycle four times. System behavior differed significantly between the first and second cycles (at the first cycle we started from a clean state), and differed only slightly between subsequent cycles (manual valve operation couldn't guarantee the same valve opening duration). For the comparison with our simulations below, we have selected a representative cycle, using air as injected gas.

1.2.6.2 Simulation

A CAD model representing the vacuum part of the real system was created and imported to Molflow+. After merging coplanar surfaces, it was reduced to 1700 facets. The fast-opening valve, as seen in red in Fig.1.43 is represented by a facet with time-dependent opacity, whereas the 567 l/s turbo-molecular pump is modeled with a single facet (of sticking factor 1) at the aperture of the real pump. The experiment cycle was modeled with the following time-dependent parameters:

- Gas desorption was set at point *P6* to be 0.321 mbarl/s between $t = 0.1..0.3$ s and 0 outside of that time range. This pressurizes the valve-delimited cell volume to 1×10^{-2} mbar.
- Valve opacity was set to 1 (closed state) until $t = 2$ s. Then it was set to decrease to 0 (fully open state) within 100 ms. It is left open until $t = 3.2$ s, after which its opacity changes back to 1 during 50 ms. With these values the simulated valve opening/closing times coincide with those measured electronically.
- Pressure is recorded in 600 equidistant steps between $t = 0$ and $t = 6$ s, and the time window was chosen to match the time step: $t_{\text{window}} = 0.01$ s.

30 billion test particles were desorbed to reach a good statistical accuracy. For reference, the calculation took 11 hours on a standard desktop PC. This relatively long simulation time comes from the fact that all hits up to the flight time of 6 s need to be calculated, taking 48.000 collisions per desorbed test particle on average. Some speedup was achieved by removing particles once their flight time reached 6 s, as opposed to tracing them until they are pumped. As the simulation doesn't account for the residual gases in the real system, the measured background pressure was added to all calculated pressure values.

1.2.6.3 Results

For comparison, the measured values were offset in time so that the valve opening happens at $t = 2$ s for both the simulated and the measured datasets. Results (measured vs. simulated pressures) are shown in Fig.1.44. Error of the measured values can be estimated to be 15% (precision of an average vacuum gauge whose temperature is not controlled), and the error of the simulated values, the statistical scattering depends on the pressure (the lower the pressure, the smaller the number of hits thus the larger the scattering). The following can be observed:




Figure 1.43: Molflow+ model of the system, with the facet representing the fast valve in red




Figure 1.44: Pressure evolution - simulation vs. measurement

- At $P5$ and $P6$ before $t = 1\text{ s}$ (red and blue markers), simulation over/undershoots the initial pressure of 1×10^{-2} mbar. This is because the initial state was achieved in the simulation with a short, pulsed gas injection at $P6$ (overshoot), which then spreads in the plasma cell reaching equilibrium with $P5$.
- At $P4$, the measured pressure from both the IKR070 and the IMR060 gauges is displayed (dark and light gray): in the high pressure region, the Penning gauge is over range, therefore the IMR060 signal should be trusted; as opposed to the low pressure region, where the latter is under range and the real pressure is indicated by the IKR070 gauge. The transition can be well observed at the valve closing moment, $t \approx 3.3\text{ s}$.
- Most importantly, the measured pressures have delays, originating from the gauges. We can approximate it by using the valve opening moment, $t = 2\text{ s}$ as reference: whereas the simulation predicts an immediate pressure change at $P4$ and $P5$ (which is the physical



Figure 1.45: Aerial shot of LIGO Livingston observatory with the Y-arms visible.
Photo: Stefan Ballmer (LIGO), used with permission

solution as gas will flow from the plasma cell as soon as the valve is open), looking at the measurement curves, we can see a significant delay at both points.

- Correcting for this approx. 0.2 s of delay would result in a significantly better match between the measured and simulated values, however it was deliberately not done to benchmark the accuracy of Molflow+ simulations "out of the box".

Apart from the gauge delays and the trick necessary to measure in both low and high pressures at P_4 , the simulated values match reasonably well with the measured results. Differences can be attributed to physical phenomena not included in the simulation, like different components of air propagating with different velocities, pumping speed inexactitude, etc. It is also important to consider that Molflow+ assumes molecular flow condition, which is not strictly the case in the 1×10^{-2} mbar region.

1.2.7 Experimental validation (LIGO)

As opposed to the previous validation, where we carried out both the experiment and the simulation in person, Molflow+ was also tested against measurements externally, and results were shared with us. As this experiment is a nice example of the code's capabilities, the test is included in this thesis, but the credit goes to the LIGO team for its execution.

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is a multi-kilometer-scale scientific experiment located at Hanford, Washington and at Livingston, Louisiana, for the detection of gravitational waves. It is a cooperation of many scientific institutes, including Caltech and MIT¹⁹. It is essentially an interferometer hosted in a vacuum tube. In 2015, an upgrade was carried out, changing the name to Advanced LIGO. Data collection began in September 2015, and on the 11th of February 2016, the LIGO scientific collaboration announced the first direct observation of gravitational waves.

¹⁹Introduction from LIGO website: <https://www.ligo.caltech.edu/page/about>




Figure 1.46: Molflow+ model of a 2km long Y-arm with the injection point and the detector. Geometry by Dr. Yulin LI

Table 1.8: LIGO leak problem simulation parameters

Number of facets	64
Molecules traced	385000
Hits	383 billion
Recorded moments	0, 1, 2 ... 2000 s, $t_{\text{window}} = 1$ s
Computation time	69.4 hours ²³

The two observatories feature an L-shaped vacuum system with 4 km long legs. The optical components are installed in a vacuum chamber with 1×10^{-9} torr²⁰ pressure, which – with its 10.000 m³ volume – takes more than a month to pump down. This vacuum environment filters vibrations from the outside world, eliminates air currents and provides thermal insulation.

Before commissioning, due to a large leak in the middle of a 2 km long Y-arm²¹, a series of leak detections were performed to find the source. As part of the process, a helium injection took place in the middle and engineers measured the He signal at a detector placed 550 m from the injection point.

A Molflow+ simulation was carried out by Dr. Yulin LI²² to compare the detector signal with that predicted by simulations and a diffusion model developed by the LIGO team to explain the long response time.

As shown in Fig.1.46, the 2 km long, 1.2 m diameter section was modeled by 64 facets, with desorption at the middle point lasting for 64 seconds. Simulation parameters are summed in table 1.8.

Measurement data, courtesy of Dr. Mike Zucker²⁴, was compared with Molflow+ results. The simulated pressure values were scaled to match the initial background and the peak values (since we don't have all parameters to convert between an apparent He leak rate on the detector

²⁰More about the vacuum system: <https://www.ligo.caltech.edu/page/vacuum>

²¹The full, 4 km long arms are split in two by a middle station

²²Cornell Laboratory for Accelerator-Based Sciences and Education (CLASSE), Cornell University

²³Simulation took considerable time because of the large number of random walk hits necessary for test particles to propagate a 2km-long section

²⁴LIGO/MIT




Figure 1.47: LIGO leak problem: measured detector signal compared to simulated values. Redrawn based on a plot of Dr. Yulin LI

and real pressure), and plotted on the same graph (Fig.1.47).

The scaled simulated data matches measured values almost perfectly. One subtle difference is that the simulated pressure rise is somewhat faster than the measured one. The reason might be either a delay in the detector's signal, or the fact that the wall sojourn time – not included in the simulations – becomes important since each desorbed particle's random walk consists of an average of 1 million hits.

2 Synchrotron radiation simulations

2.1 Introduction

This chapter presents Synrad+, a tool to simulate synchrotron radiation (SR) in accelerators. The first version was created by R. Kersevan in the 1990s[33], building on the codebase of MOLFLOW¹: instead of tracing molecules, it traced virtual photons and could visualize the flux distribution in the 3D geometry. Using the original codebase, I have modernized it, both in terms of graphical interface and by optimizing the photon generation algorithm. While doing this, priority was given to ensure compatibility with Molflow+ presented in the previous chapter: both programs can read each other's file formats, and this new simulator - renamed to Synrad+ - is intended to be used as part of a package helping to calculate the SR-induced dynamic gas load for vacuum simulations.

Our synchrotron radiation simulations go along the following steps, each detailed later:

- The user defines the optics and the beam parameters of the machine
- The source of the synchrotron radiation, *magnetic regions* are calculated
- Virtual photons are generated from these sources
- Ray-tracing is performed until photons leave the system or are absorbed:
 - Collisions are found with the chamber wall and photons are absorbed or reflected
 - Treating the wall chamber as a smooth or rough surface, reflected photons get a new direction
- Results are computed, then visualized within the program or exported for post-processing

2.2 Defining magnetic regions

To generate photons, we require at least one source. We could choose from several possibilities for defining it, depending on what we request as input from the user and what we calculate

¹MOLFLOW is the original UHV simulator written by R.Kersevan in the 1990s, and the modern version is distinguished by being renamed to Molflow+

internally.

Usually an engineer performing a SR simulation knows the beam parameters and the optics of a machine, therefore in Synrad+ – instead of the complete beam trajectory - we ask the user to set up a so-called magnetic region: a volume where the magnetic field is known at every point. This magnetic field can be defined either directly (coordinate-wise, using either absolute or beam length coordinates), or by describing common optics elements. Supported elements include:

- Dipoles
- Quadrupoles
- Analytic magnetic fields (for helicoidal, sinusoidal or other periodic structures)

When using these elements, their macroscopic properties (size, strength, orientation) are input by the user and the magnetic field is calculated within Synrad+.

- In case of dipoles, the magnetic field is constant.
- For periodic structures, such as wigglers, helicoidal and sinusoidal elements, the description is analytic.
- As for quadrupoles, to determine a point's magnetic field, we transform the real-world X , Y , Z coordinates of the point into the quadrupole's own coordinate system. If Z' is the quad's axis, and X' and Y' point from the center of its entrance to the standard quadrupole focusing/defocusing directions (in center between the southern and northern magnet poles), we define the parameters in Table 2.1, shown in Fig.2.1:

First we transform the P point coordinates to the quadrupole's own coordinate system.²

Then, with the P'_x and P'_y local coordinates obtained, we can express the magnetic field in the quad's coordinate system: if the point is inside the quadrupole of length L ($0 < P'_z < L$), then

$$B'_x = -K \cdot P'_y$$

²We calculate the offset of P compared to the quadrupole's entrance:

$$dX = x_P - x_{\text{center}}$$

$$dY = y_P - y_{\text{center}}$$

$$dZ = z_P - z_{\text{center}}$$

Then we transform P coordinates to those of the quadrupole:

$$P''_x = dX * \cos(\beta) + dZ * \sin(\beta)$$

$$P''_y = -dX * \sin(\alpha) * \sin(\beta) + dY * \cos(\alpha) + dZ * \sin(\alpha) * \cos(\beta)$$

$$P'_z = -dX * \cos(\alpha) * \sin(\beta) - dY * \sin(\alpha) + dZ * \cos(\alpha) * \cos(\beta)$$

$$P'_x = P''_x * \cos(rot) + P''_y * \sin(rot)$$

$$P'_y = -P''_x * \sin(rot) + P''_y * \cos(rot)$$

Table 2.1: Quadrupole parameters in Synrad+

α	Z' angle with real-world XZ plane (pos. towards Y)
β	Z' angle with real-world YZ plane (neg. towards X)
rot	quadrupole focusing plane rotation around Z' respective to X
K	Dipole strength, in T/m




Figure 2.1: Left: quadrupole coordinate system's α , β , rot angle definition
Right: point definition in quadrupole's coordinates

$$B'_y = -K \cdot P'_x$$

$$B'_z = 0$$

otherwise the magnetic field is 0.

Then we get the real-world magnetic field coordinates by an inverse transformation.³

Apart from the magnetic field, we need the following properties for a fully defined source:

- Start position and starting direction of the beam at the first point of the region
- Region limits
- Beam properties, such as:

³To get the B_x, B_y, B_z coordinates from B'_x, B'_y, B'_z :

$$B''_x = B'_x \cdot \cos(rot) - B'_y \cdot \sin(rot)$$

$$B''_y = B'_x \cdot \sin(rot) + B'_y \cdot \cos(rot)$$

$$B''_z = B'_z$$

$$B_x = B''_x \cdot \cos(\beta) - B''_y \cdot \sin(\alpha) \cdot \sin(\beta) - B''_z \cdot \cos(\alpha) \cdot \sin(\beta)$$

$$B_y = B''_y \cdot \cos(\alpha) - B''_z \cdot \sin(\alpha)$$

$$B_z = B''_x \cdot \sin(\beta) + B''_y \cdot \sin(\alpha) \cdot \cos(\beta) + B''_z \cdot \cos(\alpha) \cdot \cos(\beta)$$

- Particle type
- Energy
- Current
- We can assume that a beam is ideal (has no width), otherwise we define:
 - Emittance
 - Horizontal/vertical coupling
 - Energy spread
 - The Beta and horizontal dispersion (Eta) functions, coordinate-wise

We use the **dipole approximation**: we consider the beam as a succession of tiny dipoles, each having its own direction and curvature. The advantage is that the resulting synchrotron radiation can be relatively easily described analytically, and a fast photon generation algorithm can be implemented. At the same time, this approach isn't suitable for coherent sources such as undulators, where interference modifies the SR pattern substantially.

Starting from the first point in the region, our code calculates successive ones with the user-defined step length, until the region boundary is reached. For each new point, we calculate the following properties:

2.2.1 Location, curvature and beam direction

We first determine the Lorentz force:

$$\vec{F}_L = q \vec{v} \times \vec{B}$$

Then we calculate the curvature $\bar{\rho}$ using the practical formula [34]⁴:

$$\bar{\rho} [m] = \frac{E [GeV]}{0.3 \cdot B_{\perp} [T]} \cdot \|\vec{F}_L\|$$

where

$$B_{\perp} = \sqrt{|B|^2 - B_{\parallel}^2} = \sqrt{|B|^2 - \|\vec{v}\| \|\vec{B}\|}$$

Then the offset dL to the next point's position can be calculated as

$$\vec{dL} = dL \cdot \vec{v}$$

and the next point's direction \vec{v}^* is obtained by rotating \vec{v} around

$$\vec{O} = \vec{v} \times \vec{\rho}$$

axis by

$$d\alpha \sim \frac{dL}{|\vec{\rho}|}$$

⁴ $\|\vec{F}_L\|$ stands for a unit length vector parallel to \vec{F}_L , while $|\vec{F}_L|$ is the length of \vec{F}_L




Figure 2.2: Notations for trajectory step calculation (in this figure we assumed a positive charge particle)

2.2.2 Beam width and divergence

We describe non-ideal beams probabilistically: when generating a photon at a given pre-calculated trajectory point, I choose the starting location with an offset in the plane perpendicular to the (local) beam direction. This offset follows a bivariate Gaussian distribution, with mean 0 and st.dev. σ_x and σ_y corresponding to the beam horizontal/vertical size, obtained as:

$$\epsilon_x = \frac{\epsilon}{1 + \text{coupling}}$$

$$\epsilon_y = \epsilon_x \cdot \text{coupling}$$

$$\sigma_x = \sqrt{\epsilon_x \cdot \beta_y + \left(\eta \cdot \frac{\Delta E}{E}\right)^2}$$

$$\sigma_y = \epsilon_y \beta_y$$

$$\sigma'_x = \sqrt{\frac{\epsilon_x}{\beta_y} + \left(\frac{\eta'}{\frac{\Delta E}{E}}\right)^2}$$

$$\sigma'_y = \sqrt{\frac{\epsilon_y}{\beta_y}}$$

where $\frac{\Delta E}{E}$ is the energy spread, and σ'_x and σ'_y are the st.dev. of the photon's directional divergence due to the beam width, added to the natural divergence of an ideal beam (see later).

At this point, we have our sources in the form of magnetic regions, consisting of a large number of pre-calculated trajectory points. In the dipole approximation each trajectory point can serve as short dipole source, individual from the others. In the real world, these dipoles generate a number of physical photons every second, the exact number depending on the magnetic field, the beam energy and the current. Our simulations are not time-dependent, so the flux from such a dipole source is considered constant. During our Monte Carlo simulations, we will generate virtual photons, distributed equally between these trajectory points.

When generating such a virtual photon, we need to determine the following properties:

- Energy
- Polarization
- Flux (the number of physical photons per second ⁵ it represents)
- Power (the SR power that this flux carries)
- Starting location and direction, as described earlier

These parameters aren't independent of each other: for example, as the spectrum of SR is well defined; choosing a different photon energy will change the flux the virtual photon represents. To determine the exact relation between these properties, we need to recall the analytic description of SR.

2.3 Synchrotron radiation properties

2.3.1 Spectrum

Synchrotron radiation consists of photons of different energy: its wide spectrum is one of the main advantages of synchrotron light sources. The spectrum of a dipole can be described analytically. If γ is the relativistic factor (depending on the beam energy), ρ is the radius of curvature (depending on the magnetic field strength), then for wavelength λ , we will have the following number of photons per second [35]:

$$N(\lambda, t) = \frac{3^{3/2}}{4\pi} \frac{e^2 \gamma^4}{h \rho^2} \left(\frac{\lambda_c}{\lambda} \right)^2 \int_y^\infty K_{5/3}(\eta) d\eta \quad (2.1)$$

per unit $d\lambda$, where $K_{5/3}$ is a modified Bessel function. Such spectrums are plotted in Fig.2.3a.

The expression includes the energy dependence through parameter $y = \frac{\lambda_c}{\lambda}$ (equal to E/E_c since the energy and the wavelength can be converted to each other through equations $E = hf$ and $c = \lambda f$). Because of this, it is convenient to express the spectrum on a relative energy scale, in units of critical energy E_c ⁶. If we plot this relative spectrum on log-log scale (Fig.2.3b), we can notice that its low-energy tail is straight ($N \sim 1.3 * x^{1/3}$), and it has a high-energy cutoff before $E = 10 * E_c$.

2.3.1.1 Fluxwise/powerwise generation

Even if we know the distribution of the SR energy, virtual photons can represent any number of physical photons, so when we begin our Monte Carlo simulation, we still have the freedom to

⁵Usually for Monte Carlo simulations each virtual particle represents a given number of physical ones. In our case, as the photon flux is constant, each generated virtual photon corresponds to a certain flux, not an absolute number of photons

⁶The critical energy is defined so that half the power of the SR is carried by photons with energy below it




Figure 2.3: Synchrotron radiation spectrum, on absolute (2.3a) and relative (2.3b) scales
Source: [36] (for 2.3a) and [37] (for 2.3b)

sample the spectrum arbitrarily. We could, for example, generate virtual photons distributed evenly in a certain energy range, or we could oversample energies with high photon flux (fluxwise generation), or concentrate on high energy photons that carry most of the energy (powerwise generation). Choosing the correct strategy depends on the task we need to solve.

Traditionally, SR calculations for engineering tasks concentrate on the SR power: when designing chamber walls, absorbers and shielding, the thermal load can cause permanent damage, therefore detailed simulations and an adequate cooling design is necessary.

If we are only interested in the heat load, it is therefore recommended to generate photons powerwise, omitting low-energy photons. For example, limiting the generated photon energy to $0.1 E_c \dots 20 E_c$ contains more than 95% of the total power, whereas this range contains less than half of the photons (speeding up a possible heat load simulation by a factor of two). Moreover, as we will see later in this section, the high-energy part of SR is more collimated, and scatters with a lower probability, therefore simulations converge relatively fast.

In accelerator design, however, low-energy photons are very important. The reason is that while these photons carry less heat, even photons with energy as low as 4 eV are able⁷ to increase the dynamic gas load through photon stimulated desorption, as described in chapter 3.

As low energy photons are distributed over a larger angle, and they scatter more easily, they are able to reach farther parts of an accelerator, where traditional heat load calculations show very little SR presence.

Considering the above, Synrad+ allows the user to choose between fluxwise and powerwise generation. If ran for a sufficiently long time, both methods will yield the same result, however – as statistical accuracy increases with the number of traced virtual photons – for heat load calculations, powerwise generation will converge faster, whereas fluxwise generation will concentrate on low-energy photons calculating the absorbed flux faster.

For both strategies, we generate photons using numerical inversion: at the beginning of a sim-

⁷4 eV is an average photoelectric work function value for common metals




Figure 2.4: Cumulative distribution functions for synchrotron radiation

ulation, we sample the energy range between $1 \times 10^{-10} E_c$ and $100 E_c$ on 5000 logarithmically increasing intervals. For each interval we calculate the number of photons and the power it contains⁸. Instead of the analytic expression 2.1 containing the integral $G_0(x) = \int_y^\infty K_{5/3}$, we use a fast numerical approximation method[38]. Summing these interval fluxes and powers, we construct one fluxwise and one powerwise cumulative distribution function for the whole energy range, shown in Fig.2.4.

Once these cumulative distributions are stored in memory, for each generated photon, we reverse look up uniformly generated random numbers and scale the interpolated (relative) photon energy with the critical energy of the actual trajectory point.⁹

Given the sampling at the beginning, the energy range $[10^{-10}..10^2] E_c$ presented above is the widest for which simulation is possible in Synrad. The user can narrow the generated photon range for speedup, though. In that case random numbers are generated only between the C.D.F. values of the lower and upper energy limits.

2.3.2 Representation

As mentioned earlier, all our SR simulations are steady-state, so each traced virtual photon represents a given SR flux. To assign a given physical flux and power to our virtual photons, we calculate the following way:

- The total SR flux for an electron beam of a full revolution in practical units is $8.084 \times 10^{17} * E[\text{GeV}] * I[\text{mA}]$ photons/sec [40]
or $4.131 \times 10^{14} * \gamma * I[\text{mA}]$ for a general particle, γ being the relativistic factor
- Each of our short dipole sources at the pre-calculated trajectory points represents a fraction of the full revolution:

⁸The power is the number of photons in the interval multiplied by the mean interval energy

⁹It is worth mentioning that one author has analytically approximated[39] the inverse of these C.D.F. functions by splitting the energy range to three intervals, then expressing the inverse functions with a combination of regular and Chebyshev polynomials with approximately 30 coefficients in each interval. That method allows direct generation of the spectrum.

revolution_ratio = $dL / (2\rho\pi)$

with dL the trajectory step length and ρ the local radius of curvature

- We choose a starting point uniformly among the precalculated trajectory points. We would have a “full scan” when we have chosen each starting point exactly once. As we choose starting points randomly, we can express a full scan in statistical terms as having generated the same number of test photons as the number of trajectory points. At this point, the number of scans will be 1. In general, we can introduce the factor

$$\#scans = \frac{\#generated\ test\ photons}{\#trajectory\ points}$$

- Having all the quantities above, in case of fluxwise generation, one test photon will represent

$$K_{real/virtual} = revolution_ratio * \gamma * 4.131 \times 10^{14} * I[mA] / \#scans \quad (2.2)$$

photons per second.

Remarks:

- The power each test particle represents is the flux calculated above, multiplied by the average photon energy for the whole SR spectrum¹⁰
- Consequently, if we generate test photons powerwise, the represented flux, calculated above, is divided by the average photon energy for the whole energy range
- As mentioned earlier, the user can limit the photon generation range to a narrower range than $[10^{-10}..10^2]E_c$. In this case two new factors are calculated, which describe what ratio of the full SR flux and power spectrum is contained in the limited energy range. The reverse lookup process for finding the photon energy and the representation above will be corrected by these factors. For simplicity these correction factors are omitted in this description.

2.3.3 Vertical distribution and polarization

Synchrotron radiation from an accelerated particle has the highest intensity towards the direction of the beam. The exact angular distribution depends on the particle's energy, and is described by Schwinger's equation[41]. A qualitative representation is seen in Fig.2.5.

Exactly 7/8 of the SR from a dipole is polarized parallel to the orbit plane[44], and 1/8 orthogonal to it. Because of this strong dominance of the parallel polarization, it is desirable to treat the two components separately. The ratio between the parallel and orthogonal planes depends on the vertical angle ψ between the observation point and the orbit plane.

The exact dependence can be described analytically (eq. 13. of [35]). For a given vertical angle ψ and for a photon wavelength λ , the two components can be expressed:

¹⁰ $E_{avg} = 0.308E_c$ for the range $[10^{-10}..10^2]E_c$




Figure 2.5: Illustration of Synchrotron Radiation in the laboratory frame.
Source: fig.3. of [42], based on [43]

At $\psi = 0$

$$F\left(\frac{\lambda_c}{2\lambda}, 0\right) = F_{\parallel}(0) = K_{2/3}^2\left(\frac{\lambda_c}{2\lambda}\right)$$

And generally at angle ψ

$$F_{\parallel}(\psi) = (1 + \gamma^2 \psi^2) K_{2/3}^2 \left[\frac{\lambda_c}{2\lambda} (1 + \gamma^2 \psi^2)^{3/2} \right]$$

$$F_{\perp}(\psi) = \gamma^2 \psi^2 (1 + \gamma^2 \psi^2) K_{1/3}^2 \left[\frac{\lambda_c}{2\lambda} (1 + \gamma^2 \psi^2)^{3/2} \right]$$

where $K_{1/3}$ and $K_{2/3}$ are modified Bessel functions. The solution of this analytic expression can be visualized (see Fig.2.6) on a relative scale where the vertical angle is expressed in units of γ , and with a different X axis range for each energy. Finally, we often use the expression *degree of polarization* which is defined as $P_{\text{lin}} = (F_{\parallel} - F_{\perp}) / (F_{\parallel} + F_{\perp})$.

We can notice that in the plane of the orbit the polarization is purely linear.

As the modified Bessel functions can be approximated, one approach to generate photons according to this angular distribution would be to – similar to the spectrum generation – construct the cumulative distribution function for vertical angles between 0 and $\pi/2$, as the function is symmetric (the first version of Synrad in the 1990s used that method). However, this is computationally expensive: we need to integrate a rather complex analytic expression for every generated photon.

I chose to numerically pre-calculate the integrals. However, several challenges need to be tackled to sample the distribution for a large range of energies.

In Fig.2.6 the curves for the different energies are similar, but not the same: I did not find a convenient yet precise scaling method to derive all distributions from a single stored one, therefore I decided to sample them for all energies.

We can see that the sampling must be performed on an angular range that depends on the energy (corresponding to the fact that SR is more collimated for high energies). I have chosen

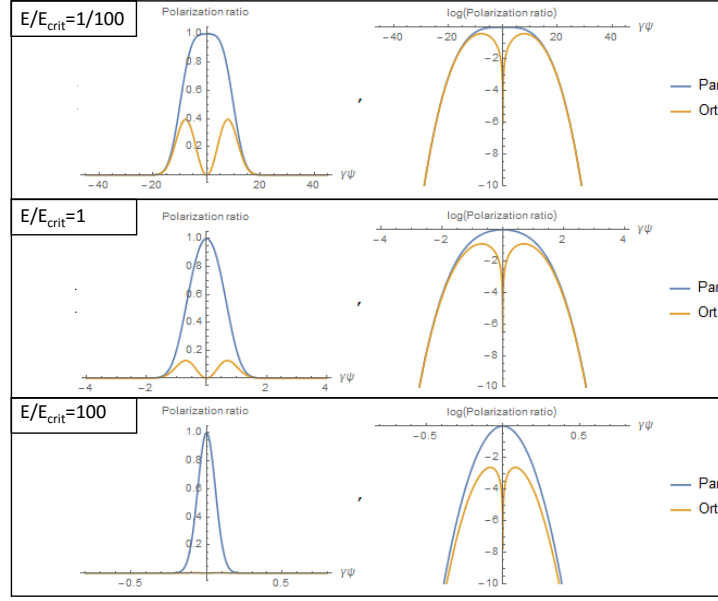


Figure 2.6: Polarization components for different lambda ratios. Each plot has an X scale that depends on the relative photon energy:

$$\gamma\psi = [-4/(E/E_c)^{0.35} .. +4/(E/E_c)^{0.35}]$$

Left: linear scale, Right: logarithmic scale

to include all the angles where the radiated flux is larger than 10^{-10} th of the peak. I have found that setting this sampling range to $\gamma\psi = [-4/(E/E_c)^{0.35} .. +4/(E/E_c)^{0.35}]$ is adequate for the range $E/E_c = [1E-10 .. 100]$, as shown on the surface plot 2.7.

Taking advantage of the symmetry of the distribution, I have sampled one half of these surfaces (individually for the two polarization components and also for their sum) the following way:

- Energy: Logarithmically in the range $E/E_c = [1E-10 .. 100]$, along 121 values, with two consecutive values having the ratio of $10^{0.1} = 1.259$
- $\gamma\psi$: Linearly along 200 values from 0 to $4/(E/E_c)^{0.35}$, with $\Delta = 0.005 * 4/(E/E_c)^{0.35}$
- Post-processing: integrating and normalizing each line (corresponding to given E/E_c values), to get a cumulative distribution function

And stored the resulting three tables of 121x201 values as input files, loaded at each startup of Synrad.

2.3.4 Horizontal distribution

Although information regarding the vertical distribution of SR from a dipole is abundant in literature, and experimental measurement is straightforward (for example, using a photosensitive film after a collimator), less information is available on the horizontal distribution: to measure it, we would require a sufficiently short dipole source, where the beam rotation is

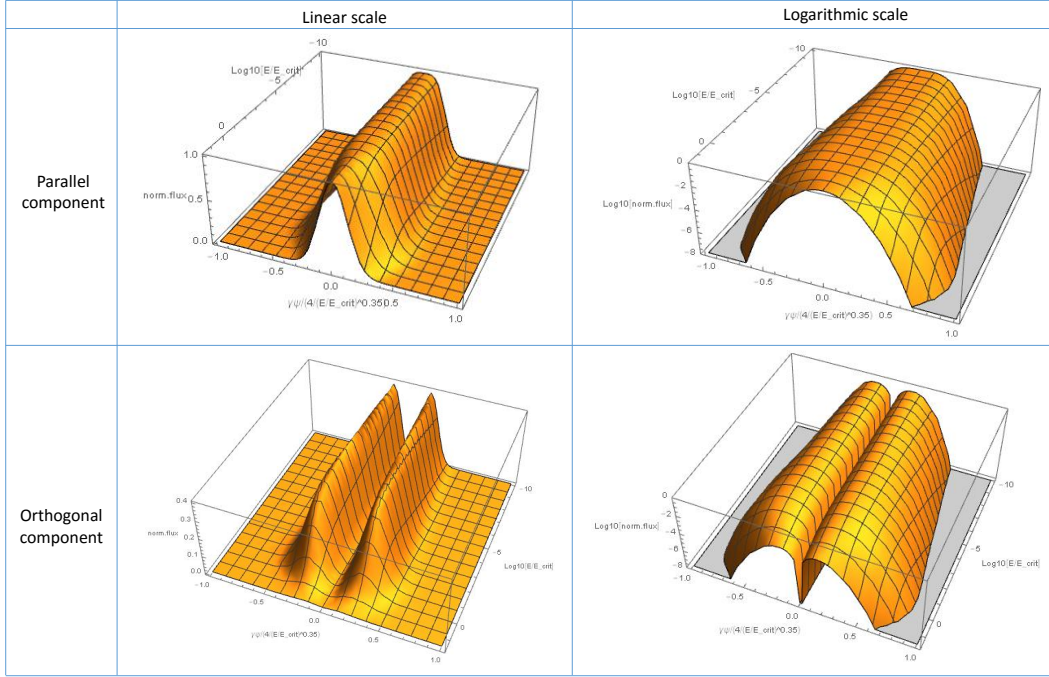


Figure 2.7: Flux distribution for the parallel (top) and orthogonal (bottom) polarization components, as a function of the relative photon energy and the vertical angle. Please note that the angle range is normalized for each energy between $\gamma\psi = [-4/(E/E_c)^{0.35} .. +4/(E/E_c)^{0.35}]$

small compared to the horizontal distribution. As Synrad treats the beam as a succession of such tiny dipoles, we need the angular description of such short sources. We refer to [45] which derives the angular distribution from the far field expression of a point charge. The analytic solution presented there for the field strength of the two polarization components is:

$$f_{\parallel} = (1 + \gamma^2 \theta^2)^{-3} \cdot (1 - \gamma^2 \theta^2 + 2\gamma^2 \theta^2 \sin^2 \phi)$$

$$f_{\perp} = -2(1 + \gamma^2 \theta^2)^{-3} \cdot \gamma^2 \theta^2 \sin \phi \cos \phi$$

and the intensity (flux) is the square of these. The expressions use the angles ϕ (elevation angle of the observation point relative to the orbit plane) and θ (angle between the local beam direction and the observation point), as shown on Fig.2.8, expressed in terms of ψ (vertical) and χ (horizontal) angles.

Similar to the vertical distribution, instead of evaluating the analytic expressions containing 11 trigonometric and 15 power expressions for every generated photon, I can achieve a significant speedup by pre-calculating and sampling the distribution.

Luckily for our case, the divergence of SR scales with $1/\gamma$, therefore the 2-dimensional horizontal/vertical distribution can be stored on a relative scale where the angles are expressed in units of $\gamma\psi$ and $\gamma\chi$.




Figure 2.8: Angular notations used in the expression for the horizontal SR distribution from a short dipole

The distributions are symmetric both in horizontal and vertical directions (in the analytic expressions angles are arguments of even trigonometric functions or are raised to squares), therefore it is sufficient to evaluate only the first (positive) quarter of the full range. As shown in Fig.2.9, the parallel component is peaked at $(\psi, \chi) = (0, 0)$, whereas the orthogonal exhibits a peak with an offset, creating four orthogonal lobes around the beam direction in the SR pattern.

In Fig.2.9 (right) the sum of the two components is plotted on a logarithmic scale, revealing the "dip" on the horizontal axis, which is also visible in Fig.2.14 (bottom right).

The problem with sampling these distributions is that they extend to a very large angular range




Figure 2.9: Horizontal/vertical distributions of SR components originating from a short dipole for $\gamma = 100$, normalized to the peak flux at $(\psi, \chi) = (0, 0)$




Figure 2.10: Logarithmic sampling of the directional distribution of SR, with sampling focusing on the narrow peak

(in theory the full hemisphere around the beam direction), but their peak is contained in a very small solid angle. Omitting larger angles where the distribution has small values is not an option: for any generated vertical angle we need Synrad+ to be able to generate a valid horizontal distribution, so every "line" (fixed $\psi\gamma$ value) of the digitized 2D map needs to have precise values in order to construct a cumulative distribution function.

In practice, we can limit the sampled range, if we decide that – as before – the lowest photon energy that we generate (which will have the largest angular divergence) will be $E_{\text{lowest}} = 10^{-10} E_c$

With the vertical angle limited to $4 * E/E_c^{0.35}$, and $E/E_c = 1 \times 10^{-10}$ in this worst case, we get the largest angle, in units of gamma, as $\psi_{\text{max}} * \gamma \sim 12500$. Since as seen in Fig.2.10, the vertical and horizontal distribution has more or less the same range, for simplicity, we will – for both angles – sample between 0 and $10000/\gamma$. This sampling will need to concentrate on the very limited peak near $(\psi, \chi) = 0, 0$, but at the same time contain values for all the range, therefore a logarithmic sampling is convenient (illustrated in Fig.2.10).

I have stored this angular distribution in three tables (parallel and orthogonal polarization and their sum), as follows:

- Vertical/horizontal angles, in units of $1/\gamma$: from 0 to 10000, sampled in 176 logarithmically increasing intervals (so that the ratio of two consecutive intervals is 1.0965)
- For each ψ vertical angle, the distribution is integrated on the range of χ values then normalized, constructing therefore 176 cumulative distribution functions. (At the integration care must be taken that the integration interval increases logarithmically)

2.4 Photon generation process

Having the energy, vertical and horizontal distributions in the memory, the exact photon generation algorithm goes along the following steps:

2.4.1 Source point choice

During a simulation, as described earlier, SR sources are represented as magnetic regions, each storing the ideal beam trajectory as a series of short dipoles. Among these trajectory points, one is chosen randomly.¹¹

If the user wishes to approximate the beam as ideal, we can use the position, direction, curvature vectors and the critical energy stored in memory during the trajectory calculation. If, however, the beam's emittance is also taken into account, the pre-calculated position is offset according to two Gaussian distributions describing the beam width. Depending on the type of the SR source, a position offset can change the magnetic field (in case of a quadrupole, for example), in which case the magnetic field vector is recalculated and the curvature and critical energy are updated.

2.4.2 Choosing photon energy

Depending whether the user generates virtual photons flux- or powerwise, one of the two stored cumulative distribution functions describing the SR spectrum is used for a reverse interpolation. The resulting relative energy is converted to absolute by multiplying with the local critical energy.

2.4.3 Choosing vertical angle

The photon's "natural" vertical angle is generated the following way:

Knowing the relative photon energy (E/E_c) from the previous step, in the preloaded vertical angle / energy distribution from section 2.3.3 the two lines digitized for the energies just below and above the searched energy are looked up by binary search.

From the values of the cumulative distribution functions stored in these two lines (for a slightly lower and slightly higher energy), an interpolated distribution is calculated for the searched energy.

Then a reverse interpolation is performed in this new distribution constructed on-the-fly, to assign a vertical angle to a pseudo-random number. If the beam's emittance isn't considered to be zero, this natural vertical angle is offset by a σ'_y width Gaussian-distributed random number, to account for the divergence.

2.4.4 Choosing horizontal angle

Similar to the previous step, the generated vertical angle's lower and higher neighbors are looked up in the angular distribution table described in section 2.3.4, an interpolated C.D.F is

¹¹I have found that it is slower to interpolate the photon starting positions and beam / magnetic field properties between the trajectory points, than to reduce the step length, and use more of these pre-calculated locations as photon starting points. To give an approximate order of magnitude, it is suitable to model a simple dipole with 10.000 points, and not to have more than a total of 1.000.000 trajectory points for all regions of the model.




Figure 2.11: Looking up a random number rnd in an angular C.D.F., linear interpolation will result in a step p.d.f. function




Figure 2.12: Quadratic interpolation in the C.D.F. results in a continuous p.d.f.

constructed, then a pseudo-random number is reverse interpolated to find the corresponding horizontal angle.

If the beam isn't ideal, this natural horizontal angle is then also offset with a σ'_x width Gaussian-distributed random number.

2.4.5 Reverse interpolating in a cumulative distribution function

During the photon generation process, we apply the numerical inversion method three times: we look up a random number in a cumulative distribution function for the photon energy, then for the vertical and the horizontal angles. I have experienced that for numerical inversion, linear interpolation is not the most efficient method: if we assume that the C.D.F. is linear between the sampled points, then it will correspond to a step function in the probability distribution function, illustrated in Fig.2.11.

This will result in discontinuity in the generated distributions, which we could compensate by increasing the number of sampled points. This, however, increases both the memory requirement and the steps required for the binary lookup. For larger tables, especially the vertical/horizontal angle distribution, the sample size can be significantly reduced (without a noticeable change in the generated distributions) by using quadratic interpolation, which takes into account adjacent points to make the first derivative of the C.D.F. continuous, corresponding to a continuous p.d.f. This is illustrated in Fig.2.12.

The forward quadratic interpolation (looking up a y value at a given x location) is a well-established problem: for three data points $(x_1, y_1), (x_2, y_2), (x_3, y_3)$, the interpolated $y = P(x)$ value and the $y' = P'(x)$ derivative can be expressed by the Lagrange interpolating polynomials:

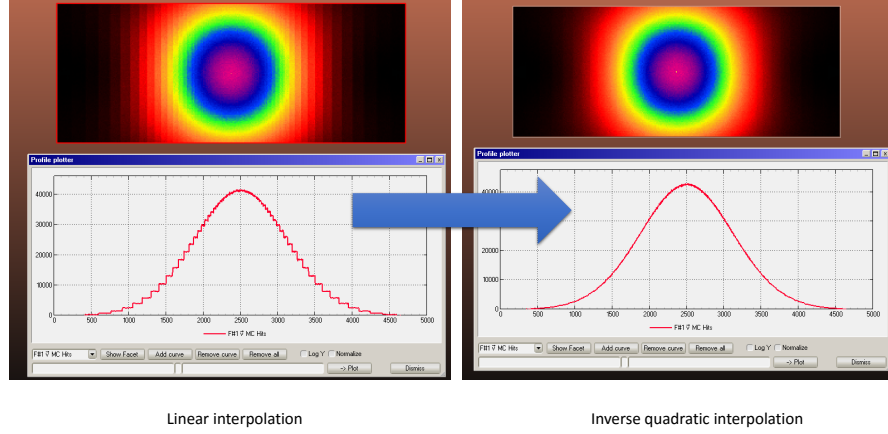


Figure 2.13: Comparison of the linear (left) and inverse quadratic (right) interpolation methods, for the same number of sampled points

$$\begin{aligned}
 P(x) &= \frac{(x-x_2)(x-x_3)}{(x_1-x_2)(x_1-x_3)}y_1 + \frac{(x-x_1)(x-x_3)}{(x_2-x_1)(x_2-x_3)}y_2 + \frac{(x-x_1)(x-x_2)}{(x_3-x_1)(x_3-x_2)}y_3 \\
 P'(x) &= \frac{2x-x_2-x_3}{(x_1-x_2)(x_1-x_3)}y_1 + \frac{2x-x_1-x_3}{(x_2-x_1)(x_2-x_3)}y_2 + \frac{2x-x_1-x_2}{(x_3-x_1)(x_3-x_2)}y_3
 \end{aligned} \tag{2.3}$$

We, however, do the inverse: we look for the x value corresponding to a y value. Therefore we solve the above equation for x , which will result in two roots¹², one of which – in our interest – will be between x_1 and x_2 . The improvement in the generated distribution - with the same number of sampled points - is illustrated in Fig.2.13.

Compared to generating photons by evaluating the analytic expressions, this sampling and reverse interpolation method results in a speedup of more than 10x (see Fig.2.14).

2.4.6 Validity

The photon generation algorithm described earlier (without my optimizations) has been used in multiple MC simulators including GEANT4[46], MAD-X[47], and its validity is discussed in many places of literature, of which a good summary is published in chapter 1.3 of [39]. Briefly:

- In reality, the spectrum is not continuous as we assume but discrete. However, for accelerator applications the mode number will be too high to make this noticeable ($n_c = 1\text{E}13$ for 10 GeV electrons)
- The spectrum is valid as long as the photon energy is small compared to the particle energy, which is safe to assume for electron beams in modern (GeV) accelerators

¹²I solved eq. 2.3 for x with the symbolic solver Mathematica. The resulting expressions giving the two roots are too long to state here.




Figure 2.14: angular distributions for ~ 22 million fixed energy photons originating from a short dipole.

Left: generated by evaluating the analytic expressions (generated angles limited for speedup)

Right: generated by reverse lookup in sampled distributions

- It assumes a homogenous magnetic field for each trajectory point, omitting edge effects of short dipoles
- It does not include coherence effects of short bunches with many particles

As for my implementation, it can be tested against calculated data for real machines. Based on [48], I can analytically calculate the spectra for a bending magnet and a wiggler of the Brookhaven National Laboratory's VUV Storage Ring, which – operating from 1981 - was one of the first 2nd generation light sources in the world. The referenced article gives a practical version of eq.2.1, expressed as a function of the photon energy E instead of the wavelength:

$$\text{Flux}(E)[\text{photons/sec}/0.1\% \text{ bandwidth}] = N \cdot \alpha [\text{mrad}] \cdot 2.457 \text{E}13 \cdot E_{\text{beam}} [\text{GeV}] I [\text{A}] G_1(E/E_c)$$

where E_c is the critical energy, α is the horizontal bending angle, N is the number of wiggler half-periods (1 in case of a bending magnet), E_{beam} is the beam energy, I the beam current and

$$G_1(y) = y \int_y^{\infty} K_{5/3}(y') dy'$$

The article then integrates the SR light on a 5 mrad horizontal opening and publishes the resulting plots. Comparing my results to the analytically calculated examples test Synrad's beam trajectory calculation, photon energy and angular divergence calculation and ray tracing modules at the same time.

Table 2.2: Machine parameters of a dipole and a wiggler of the BNL VUV storage rings, as published in [48].

	VUV bending magnet	U13U wiggler
Beam	0.745 GeV, 800mA	
Magnetic field [T]	1.28	0.81 (max)
Radius of curvature [m]	1.91	3.06 (min)
Critical energy [eV]	472.3	298.8 (max)
Period length [mm]		100
2N (Half-periods)		45

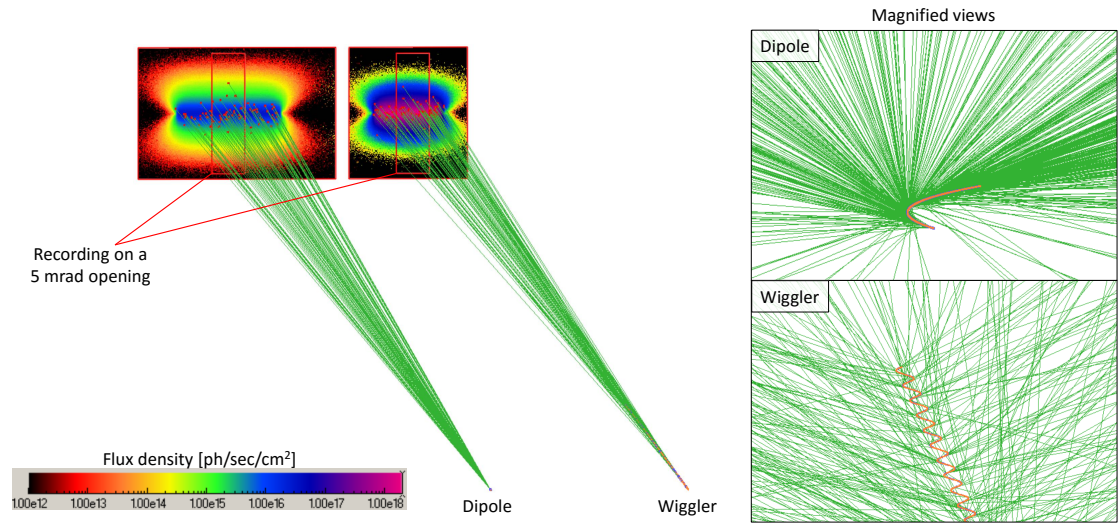


Figure 2.15: Synrad+ simulation of the VUV bending magnet and the U13U wiggler

Parameters for the bending magnet designated as VUVBM and the wiggler referred to as U13U are summarized in Table2.2.

The magnetic regions in my simulation are described by 3000 points for the dipoles and 2200 points for the wiggler, and SR patterns are registered on two textures of 300x200 and 180x200 cells. The spectrum is vertically integrated on a 5 mrad wide opening. The simulation (Fig.2.15) took 1 minute on a standard desktop PC, during which 60 million test photons were generated.

Results, plotted in Fig.2.16 are compared to the analytic formula (also plotted in Fig.1. of the original article [48]), and show a very good match. ¹³

¹³The simulated flux results for the wiggler are slightly lower than the analytic calculations. It turns out that when using a collimator to filter for 5 mrad angle, in case of the 2.25 m long wiggler, the 5 mrad opening is not uniform along all its length: the beginning of the wiggler sees the collimator at a slightly smaller angle than its end. Moving the wiggler back- or forwards (so for example its center, instead of its beginning) is aligned with the dipole would change the results slightly.




Figure 2.16: SR spectrum of the VUVBM bending magnet and the U13U wiggler of the VUV Storage Ring of Brookhaven National Laboratory, as calculated by Synrad's spectrum plotter tool and as calculated analytically

2.5 Photon tracing

Once virtual photons are generated from a source, they are traced until they are absorbed or until they leave the modeled region. Therefore, at every collision with a surface, reflectivity and photon scattering must be taken into account. Both photon reflectivity and scattering from metallic surfaces are well-established subjects. Many approaches exist to simulate them, from solving the Maxwell equations to different levels of approximations, which in turn are faster to calculate. As our Monte Carlo codes aim for possibly thousands or even millions of simulated hits per second on standard PCs, and usually a geometry consists of thousands of facets, I chose methods which are acceptably accurate on a macroscopic scale, at the same time don't slow down photon tracing excessively.

2.5.1 Photon reflection

Synrad+ uses the same raytracing engine as Molflow+. However, whereas in vacuum simulations reflection is usually executed simply (Lambertian reflection, optionally with a sticking factor), in case of photons, the model, overviewed in Fig.2.17 is more complex. Its steps are detailed as follows:

Upon collision with a surface, a virtual photon can be absorbed, transmitted or reflected. In Synrad+ it is possible to define surfaces where these probabilities are user-defined constants (for example, 70% of the incident photons are absorbed, the rest reflected). I consider these cases to be straightforward to implement therefore I don't detail them here.

For a real metal vacuum chamber surface, the reflection probability depends on the material itself, but also on the wavelength (energy) and incident angle of the photon. The idea is that for each material, we store a table of reflection probabilities for a range of photon energies and incident angles, and we interpolate the exact value at each collision.

The model to construct such material tables is described in [49], and can be summarized as

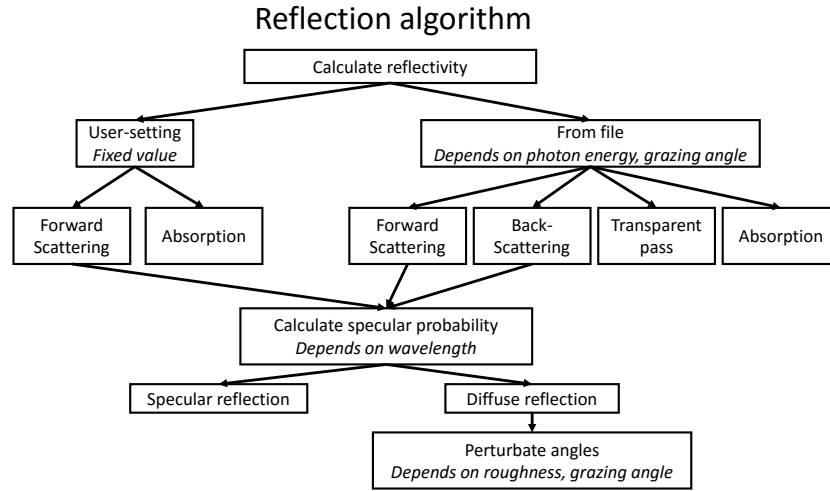


Figure 2.17: Reflection algorithm main steps in Synrad+

follows:

The refractive index of a metallic surface, a complex number, can be calculated from the atomic scattering factors as¹⁴:

$$n = \frac{r_e}{2\pi} \lambda^2 n' f(0) \quad (2.4)$$

Where r_e is the classical electron radius, λ is the wavelength, and n' is the number of atoms per unit volume. The complex atomic scattering factor can be expressed as

$$f(0) = f_1 + i f_2$$

Where the imaginary part f_2 is derived from the atomic photoabsorption cross section, and the real part f_1 can be obtained from the imaginary part through the Kramers-Kronig dispersion relation. I obtained these factors from the Henke database [50] which contains these scattering factors for most elements of the periodic table for energies between 30 eV and 30 000 eV in a downloadable format.

As for energies below 30 eV (important for accelerator applications), I extend our refractive index database for common vacuum materials from [51]. It is worth mentioning that at the moment of writing, there is a significant effort to organize reflection data for many elements originating from various scientific papers, and present them – with references – in an interactive and downloadable format on a website [52].

With the refractive index at hand (example in Fig.2.18), we can derive the reflectivity values using the Fresnel formulas:

¹⁴See Eq.1 of chapter 1.7 of [49]




Figure 2.18: Real and imaginary components of copper refractive index

Left: Calculated from eq.2.4

Right: As presented on [52], containing data for low energies as well

$$R = |r|^2 = \left| \frac{k_{iz} - k_{tz}}{k_{iz} + k_{tz}} \right|^2$$

Where $k_{iz} = \frac{2\pi}{\lambda} \cos \theta$ and $k_{tz} = \frac{2\pi}{\lambda} \sqrt{n^2 - \sin^2 \theta}$

Using the method above, I built a library of reflectivity tables for common materials, such as the one plotted in Fig.2.19 for copper, and stored it in Synrad+. I also let the user add his own materials through CSV-formatted input files.

These input files – apart from the reflectivity - also allow to specify a probability for photon backscattering and transparent pass, useful in two special cases.

2.5.1.1 Sawtooth surfaces

When Synrad+ was used for a photon flux calculation problem in the LHC, the whole, 27 km long vacuum chamber was imported, consisting of 200.000 facets and 2500 magnetic regions (each LHC dipole appeared as two different sources in the two beampipes). The geometry, containing only the vacuum-related parts, is shown for illustration in Fig.2.20.

The memory usage at this point was approximately 2 GB. The LHC beam screen has a sawtooth profile (of 40 μm height and 500 μm pitch, as in Fig.2.21) on a 14 mm-high strip centered on the plane of the orbit, where practically all primary photons hit: since, as explained before, photons with near-orthogonal incidence have low scattering probability, this geometry significantly reduces scattering of SR photons down the beamline[53].




Figure 2.19: Copper reflectivity tables for different grazing angles (left) from [50] and Synrad-simulated reflections (right) of SR from a dipole on a copper surface at different grazing angles




Figure 2.20: The vacuum system of the LHC ring (left) and the ATLAS detector (right), in a Synrad-importable geometry format. Horizontal distances between the beam pipes increased for visibility.

Images: Jan Sopousek (Brno University, CERN)




Figure 2.21: Microscopic surface view of the LHC beam screen, measurements by Nicolaas KOS (CERN)




Figure 2.22: Sawtooth surface model and reflectivity measurements in Synrad+. Geometry: Jan Sopousek

In theory, we could describe these surfaces without approximations, i.e. using one polygon for each sawtooth period. However, this would prevent us from simulating such large structures. In our case, modeling each saw-tooth indentation would increase the number of facets to 108 million (LHC length divided by sawtooth period), which is impossible to keep in memory, and even if we could, ray tracing would be extremely slow.

Therefore, to treat this problem, the need came to describe such periodic surfaces with a single facet, and to include the forward- and backscattering through input files.

The forward- and backscattering probabilities can be determined by simulating a simple reflectivity measurement (see Fig.2.22). 20 magnetic regions, representing quasi-straight dipoles of the same parameters, emit photons of 10 eV with 20 different grazing angles increasing in logarithmic steps.

Each region is separated from the others by opaque facets, so the emitted photons can either be absorbed on the periodic sawtooth surface, or get reflected forward or backwards. The sawtooth surface is assumed to be made of copper, applying the corresponding reflectivity table.

The forward- and backscattered photons are summed by textures on the target hemispheres, and the ratios are summarized in Fig.2.23.

Having determined these values, we can repeat the experiment for different photon energies. With these data, we can finally define the sawtooth beam screen as a special surface, described by a single facet, where the input file contains not only single-value reflection probabilities, but – for each energy and grazing angle – there is a probability for absorption, forward- and backscattering.

We can see in Fig.2.24 that replacing the periodic structure by such a facet, the reflection pattern is nearly identical in the real and the approximated case.




Figure 2.23: Simulated forward- and backscattering probabilities of 10 eV photons from a copper sawtooth surface




Figure 2.24: Reflected SR pattern from a sawtooth surface (bottom) and from its approximation as a single facet (top), with different grazing angles, increasing to the right




Figure 2.25: Mass attenuation coefficient of Aluminum, downloaded from NIST database[54], and the data points (in red) show the digitized data points.

One notable difference is that the real surface allows backscattering via two different geometrical paths, which is visible through a bifurcation of the scattering pattern for low angles. Also, because the teeth of the surface are tilted compared to the general wall direction, the reflected patterns have some offset.

2.5.1.2 Transparency

Once our input files can contain several reflection components for each grazing angle and photon energy, it is straightforward to extend the absorption, forward- and backscattering probabilities with a fourth one: transparency. This allows modeling of vacuum windows made of Beryllium, for example. Such a window is partially transparent to synchrotron radiation, but seals the vacuum chamber from the atmosphere.

In the example below, we perform a simple radiation protection calculation, where we calculate how much SR can get out from an aluminum vacuum chamber of 1 cm thickness.

First we obtain and digitize the aluminum mass attenuation coefficient for X-ray wavelengths from the NIST material database (Fig.2.25). This graph is adjusted to include the Compton effect and pair production for higher energies.

To get the ratio I/I_0 of the incident radiation that traverses the aluminum, we use the equation

$$I = I_0 \exp\left(-\frac{\mu}{\rho} x\right)$$

where $\rho = 2.69 \text{ g cm}^{-3}$ is the aluminum density, and

$$x = \frac{1 \text{ cm}}{\sin \theta}$$

is the path the X-rays travel in the aluminum wall (θ is the grazing angle).




Figure 2.26: Transmission probabilities of X-rays through 1 cm of aluminum

Table 2.3: Radiation protection example simulation parameters

Beam	10 GeV, electrons
Dipole magnetic field	15.6 T
Radius of curvature / critical energy	2.138m / 1MeV
Photon generation range	10eV .. 10 MeV
Photon flux / power	1E18 ph/s, 52.5kW
Virtual photons generated	3 Million
Simulation time (avg. desktop PC)	30 seconds

With this information, we can fill all cells of the input file, each describing the ratio of transmitted light for a given energy and grazing angle. The resulting transmission probability file is plotted in Fig.2.26.

Our simulation (Fig.2.27) will measure the radiation that originates from a dipole and can escape from a vacuum chamber. The chamber has a diameter that is reduced by a taper after the dipole. Three targets are placed to measure the radiation outside the chamber: (1) placed just after the dipole, (2) after the taper and (3) after the end flange that closes the chamber.

Since aluminum is almost perfectly opaque below energies of 10 keV, I have chosen beam and dipole parameters to include high energies that fall in the semi-transparent region. Simulation parameters are summarized in Table2.3.

Results show that 29% of the dipole's flux and 62.4% of the power can escape the vacuum chamber. Since aluminum is more transparent at higher photon energies, the power ratio is – as expected – higher. This is also visible on the spectrums recorded in the targets (in Fig.2.28).




Figure 2.27: Synchrotron radiation photons (green lines) originating from a 10 GeV beam (orange) bent by a 15.6 T dipole escaping from a 1 cm thick aluminum chamber. External targets are marked with red outlines; photon absorption locations are marked with red dots.




Figure 2.28: Transmitted spectrum results on the three targets of the radiation protection simulation

2.5.2 Rough surface scattering

Many approaches exist for simulating photon reflection from rough surfaces. In general, these can be classed into:

- Geometric optics methods, which either describe the actual surface with very high detail, or generate it according to a statistical distribution, then perform geometric ray tracing treating the surface points as locally smooth. Usually these methods require storing (or generating on-demand) a large number of polygons for a single surface, which isn't suitable for our simulations which already use a large number of facets for the description of the accelerator geometry.
- Physical optics methods, based on electrodynamics, describing diffuse scattering based on the Maxwell equations. The theory of diffuse scattering of EM waves from random rough surfaces is described in detail in [55] and [56]. These methods are generally too expensive computationally to include in a Monte Carlo simulation where thousands of collisions are expected to be simulated every second.

My code uses therefore a mixed method that is based on a model presented in [57], applied in the synchrotron radiation simulator Synrad3D, developed at Cornell University. That model uses the scalar Kirchhoff theory [58],[59].

The statistical description of the surfaces assumes Gaussian distribution for the surface height variations (of rms σ) and also for the transverse distribution (with autocorrelation coefficient T). Although the code is closed-source, the mathematics behind it is published. I decided to implement it as Synrad3D is an etalon in the electron-cloud community since several years to estimate the seeding position of photoelectrons [60]. I have, however, applied approximations and optimizations to the model to bring it to a speed that is acceptable for our simulations.

The theory, as described in [61] is as follows:

Upon colliding with a surface, a photon can be absorbed or reflected. The reflection probability of the surface is determined the same way as with smooth surfaces, using the database for materials, as described earlier in this section. If a photon is reflected, there is a certain probability that it does it specularly. In this case, to get the new direction, we mirror the incident direction to the reflecting surface. This specular probability is described by the Debye-Waller factor:

$$P_{spec}(y, \lambda) = \exp \left[- \left(\frac{4\pi\sigma y}{\lambda} \right)^2 \right]$$

In this expression $y = \cos(\theta_{in})$, with θ_{in} the angle between the incident direction and the surface normal, λ is the photon's wavelength and σ is the RMS surface roughness.

If – based on the probability above - a reflection isn't specular, we apply the diffuse scattering model, where the probability of a certain combination of the incident, outgoing and out-of-plane angles are described by the equations:

$$\frac{dF_{diff}}{d\Omega} = \frac{J(y, \lambda, x, \phi)}{J_0(y, \lambda)}$$

$$J(y, \lambda, x, \phi) = \frac{(1 + xy)^2}{(x + y)^4} (1 - a \cdot \cos\phi)^2 g e^{-g} \sum_{m=1}^{\infty} X_m$$

$$g = \left(\frac{2\pi\sigma(x + y)}{\lambda} \right)^2$$

$$x_m = \frac{gm}{m!m} e^{-gq/m}$$

$$q = \frac{(2 - x^2 - y^2 - 2h\cos\phi)\tau^2}{4(x + y)^2}$$

$$\tau = \frac{T}{\sigma}$$

$$h = \sqrt{1 - (x^2 + y^2) + x^2 y^2}$$

$$a = \frac{h}{1 + xy}$$

In these expressions, $x = \cos(\theta_{out})$ is the cosine of the reflected direction's angle with the surface normal, and ϕ is the "out of plane" angle, i.e. the angle of the reflected direction with the plane defined by the incident and surface normal vectors.

As mentioned in [61], these expressions simplify in the $g(\sigma, \lambda) \gg 1$ limit corresponding to technical vacuum chambers and to high energy photons, for which typically $\sigma \gg \lambda$. In this limit, the distribution can be described without the infinite sum:

$$\frac{dF_{diff}}{d\Omega} \simeq \frac{J_a(y, x, \phi)}{J_{a0}(y)}$$

where

$$J_a(y, x, \phi) = \frac{(1 + xy)^2}{(x + y)^4} (1 - a \cdot \cos\phi)^2 e^{-q}$$

$$J_{a0}(y) = \int_0^1 dx \int_{-\pi}^{\pi} d\phi J_a(y, x, \phi)$$

As a compromise between the speed and accuracy, my code uses this high energy approximation (it would be numerically impractical to calculate an infinite sum at every photon reflection with an acceptable speed). However, even this approximation uses a formula that consists of too many mathematical operations to evaluate efficiently (four power operations, two multiplications, three sums and one trigonometric function - and the integral for all





Figure 2.29: Angle notation for the reflection model in Synrad+

Figure 2.30: Diffuse reflection directional distributions for 3 different θ_{in} incident angles at $\tau = 30$

angles). By plotting the distributions, however, it is possible to notice their characteristics that give way to simpler, analytic approximations.

We are interested in the scattering angle $d\theta = \theta_{out} - \theta_{in}$ and the out-of-plane angle ϕ , shown in Fig.2.29.

Plotting the diffuse part of the scattered distributions on the $\theta_{out}\phi$ plane (Fig.2.30), we can notice that the distributions are Gaussian-like, with a different width in the two angular directions, and the peak is at the specular reflection point ($\theta_{out} = \theta_{in}, \phi = 0$).

I have therefore approximated these distributions as bivariate Gaussians, and made an attempt to describe the standard deviations for both angles analytically, as a function of the roughness ratio τ and the incident angle θ_{in} .

To determine the standard deviations, for each θ_{in} incident angle, I have used a root finding algorithm to find the offsets $d\theta$ and ϕ at which the value of the distribution is $\exp(-1/2) \sim 60.6\%$ of the peak value (Gaussian distributions have this value at their 1σ offset).

Results show that the width of the $d\theta$ distribution is the same for all incident angles, and depends only on the surface roughness, described in our case by the parameter τ . We get the best fit if we choose a Gaussian with $\sigma_\theta = C/\tau$, which has the physical meaning that the more the surface is rough, the wider our angular distribution becomes for the scattered photons.

As for the out-of-plane angle ϕ , the case is more difficult, as the width of the distribution




Figure 2.31: Standard deviation of the reflected SR light's out-of-plane angle distribution, as a function of the incident angle θ_{in}

Table 2.4: Fit parameters for the reflection angles

Scattering angle θ_{out}	Gaussian, mean= θ_{in} $\sigma_{\theta} = \frac{2.9267}{\tau}$
Out-of-plane angle ϕ	Gaussian, mean=0 $\sigma_{\phi} = \frac{2.80657 * \theta^{-1.00238} - 1.00293 * \theta^{1.2266}}{\tau}$

depends on the incident angle and the surface roughness. The dependence on the latter can be quickly found to be similar to the scattering angle, $\sigma_{\phi} \sim 1/\tau$.

Executing the root-finding algorithm for 100 distinct θ_{in} values shows (Fig.2.31) that the width of the ϕ distribution shows a near-hyperbolic dependence on the incident angle, suggesting to try fitting with an analytic function similar to $\sigma = A * \sigma^{-1} * \tau^{-1}$. Using the best fit for parameter A , a second fitting run, introducing a correction factor added to the hyperbolic term has resulted in a very accurate fit, as seen on Fig.2.31. The final solution is summarized in Table2.4.

Comparing the original distribution with our Gaussian approximations shows practically perfect match for most incident angles. The exception is the when the incident photons are almost parallel to the surface ($1.55 < \theta_{in} < \pi/2$) where the shape of the out-of-plane's distribution is somewhat different (Fig.2.32).

This discrepancy could be solved by fitting a higher order expression or by fitting different functions for different domains of the incident angles. As these practices would cost computation time for all incident angles, I have decided to not implement them for the moment.

2.5.3 Validation

As results of the model I approximate are published in [61], a practical benchmarking of my approximation can be done by simulating a reflectivity measurement in Synrad+. I set up five identical dipole sources, each surrounded by a long and narrow tube. The highly collimated SR light exiting the tubes was then projected to surfaces tilted at different angles.

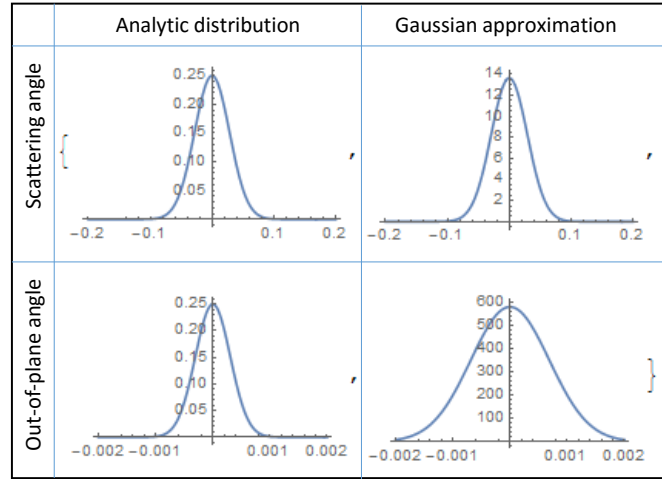


Figure 2.32: Angular distributions of the analytic expressions, as published in [61] and their Gaussian approximations in Synrad, for a near-parallel incidence $\theta_{in} = 1.56$ and $\tau = 100$




Figure 2.33: Specular reflectivity test in Synrad+

2.5.3.1 Specular reflectivity

At this test, the five reflective surfaces were placed so that the grazing angle of the incident photons were 0.5, 1, 2, 5 and 10 degrees, respectively. I calculated the specular part of the reflected light by placing and carefully aligning five very small targets in the specular direction. That way the diffusely reflected light, which is distributed over a larger solid angle misses the targets, so measuring the spectrum on the targets gives us information of the specular part¹⁵. Simulation parameters are summarized in Table 2.5 and an overview is shown in Fig. 2.33.

Having the spectrum of the incident flux and the specularly reflected photons at each target, for every photon energy, we can calculate the specular reflection ratio. This allows comparison with published data from Fig. 2 of [61]. As shown in Fig. 2.34, the simulated ratios match reasonably well with the Debye-Waller factor, except for some imprecisions at very low energies

¹⁵Strictly speaking, the small targets in the specular direction also collect the peak of the diffuse part, however my measurements show that its flux is negligible compared to the specular flux.

Table 2.5: Synrad+ simulation parameters for the specular reflectivity test

Beam	1 GeV, electron, 1 mA
Dipole	0.3T, 1 mm length
Collimators	20 cm length, 10 μm diameter
Surface roughness	$\sigma=100$ nm, $\tau=100$
SR generation	1..250 eV, with $E_c=199.4$ eV
Generated photons	400 million
Photons reaching targets	1 million




Figure 2.34: Ratio of specularly reflected SR at different grazing angles (in degrees, see legend) of a Synrad+ simulation, compared to the analytic Debye-Waller factor

(where the incident flux is low) and high energies (where the reflected flux is low).

2.5.3.2 Scattering and out-of-plane angles

In this test I measure the angular distribution of the diffusely reflected pattern. Compared to the previous setup, in this case only a fixed energy of the SR spectrum is generated, and the targets are larger, each containing a texture of 1042x201 cells. The reflective surfaces are tilted at 1, 2, 5, 10 and 20 degrees, and the targets are placed exactly 5 cm from the reflection point, so that each texture cell location on their surface can be converted to $d\theta$, ϕ angles by trigonometric formulas (the targets are centered around the specular reflection point).

The test parameters are summarized in Table 2.6, and the angular distributions are calculated (Fig.2.36) from the texture values, for comparison with data published in the original article (Figures 5 and 6 of [61]).

Comparison shows the aforementioned discrepancy for low grazing angles, most visible for the out-of-plane distribution for 1 and 2 degrees (my code approximates the distribution to be wider than the analytic solution). For other incident angles the difference is negligible.




Figure 2.35: Diffuse reflectivity test in Synrad+




Figure 2.36: Comparison of scattering angle (left) and out-of-plane angle (right) distributions on a rough surface of $\sigma = 100$ nm and $\tau = 100$, for 6 keV photons, simulated by Synrad+ and as published in Figures 5 and 6 of [61]

Table 2.6: Synrad+ simulation parameters for the diffuse reflectivity test

Beam	3 GeV, electron, 1 mA
Dipole	1 T, 1 mm length
Collimators	20 cm length, 10 μ m diameter
Surface roughness	$\sigma=100$ nm, $\tau=100$
SR generation	fixed at 6000 eV, with $E_c=5983$ eV
Generated photons	225 million
Photons reaching targets	2.5 million

Depending on the frequency of Synrad+ usage for very low grazing angles, it might be necessary in the future to perform the analytic approximation of the out-of-plane angle by fitting several functions on different domains. That would yield a better match for low grazing angles, at the expense of making the code somewhat slower and more complex.

2.5.4 Low-flux mode

When presenting my ultra-high vacuum simulations in chapter 1, I mentioned that large pressure differences cause a problem since for every test particle in the low pressure region, we might need up to thousands of traced particles in high pressure regions.

We face a similar problem during SR simulations: as described later in chapter 3, in some cases we need to simulate 7 orders of magnitude of flux density. In such a case, for a single Monte Carlo hit in the low flux regions, 10 million test photons are traced in the high-flux parts.

A way to mitigate this problem is to dynamically change the represented photon flux of a test photon as we trace it in our system. When a test photon is generated, it represents a certain number of real photon flux, the exact number expressed in eq.2.2. When this test photon hits a reflective surface, there is a certain probability that it will be reflected, otherwise it will be absorbed locally. Our random generators assure that the ratio of reflected photons will match the surface reflectivity.

If this reflectivity is small, for example 10%, then the chance of - let's say - 5 consecutive reflections is very low ($0.1^5 = 0.00001$), meaning that to register a single hit on areas reachable only by 5 or more reflections, 100.000 test photons are required on average.

We can therefore change the algorithm and – as shown in Fig.2.37 – do ray-tracing in such a way that the photon is always reflected, but after the reflection the real flux it represents is reduced. If R is the surface reflectivity, and F_{in} is the photon flux it represents at incidence,

- the photon will be reflected, and the outgoing photon will represent $F_{out} = R * F_{in}$ flux of real photons
- and $(1 - R) * F_{in}$ absorbed flux will be registered at the reflection point (in form of textures, etc.)

In this case, however, we must assure that the ray-tracing eventually stops. I let the user

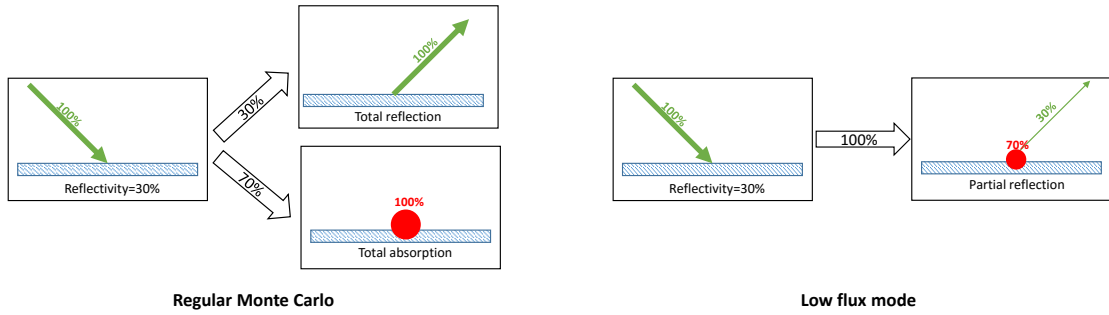


Figure 2.37: Synrad+ low flux mode overview

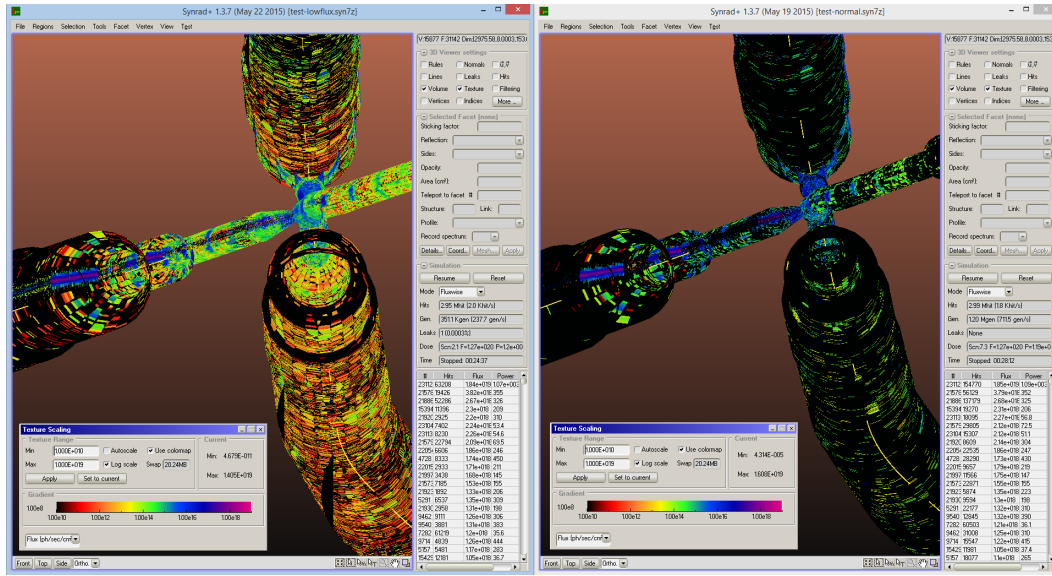


Figure 2.38: Flux density of the SuperKEKB interaction region, after 3 million Monte Carlo hits, simulated with low-flux mode enabled (left) and disabled (right), shown on the same color scale. While flux density values are exactly the same for high-flux areas (blue/violet colors), low-flux areas (red, yellow) have significantly more statistics with the low-flux mode enabled.

define a cutoff ratio. If the ratio of the test particle's actual representation (reduced through multiple reflections) and the original representation falls below this cutoff ratio, we eliminate the photon.

As a thumb rule, if the user is interested in N orders of magnitude of flux density, then the optimal cutoff ratio is 10^{-N} .

This low flux mode was implemented for the simulation of low-flux areas of the SuperKEKB interaction region (see section 4.3). We were interested in 9 orders of magnitude of flux density, calculations often requiring overnight simulations. In Fig.2.38 the result of this low-flux mode (with cutoff ratio of $1\text{E-}7$) is compared to the standard algorithm after 3 million traced hits.

3 Photon stimulated desorption

In the previous chapters, vacuum and synchrotron radiation (SR) simulations were discussed. The link between the two fields is photon stimulated desorption (PSD), a phenomenon where SR photons impinging on the accelerator wall desorb gas molecules. According to [62], the majority of the gas load in an electron storage ring originates from PSD, and as such, this effect significantly influences the cost, design and performance of an accelerator.

The process is closely related to electron stimulated desorption (ESD), gas load caused by photoelectrons: there is evidence[63] that when a SR photon hits the chamber wall, a photoelectron is first extracted, and it is the extraction and recapture process of this electron that desorbs a gas molecule.

We quantify the process statistically by the molecular yield η :

$$\eta = \frac{\text{molecule desorption rate [molecules/sec]}}{\text{SR impingement rate [photons/sec]}}$$

The yield η depends on:

- The type of desorbed gas
- The material and treatment history of the surface
- The energy of the impinging SR photons

For a typical vacuum surface, H_2 is the dominant gas type, followed by CO , CO_2 and CH_4 . The yield increases with the photon energy: the dependence of η on the spectrum's critical energy E_c can be approximated by a power function:

$$\eta \sim E_c^\beta$$

Experimental data shows β values between 0.74 and 1.12 [64].

When operating a machine, the yield is not constant in time: after an initially high value, as molecules leave the surface (and subsequently the whole system through pumping), the surface cleans, and the yield decreases. This surface conditioning, in accelerator operation terms, is called scrubbing: when (re)starting an accelerator, its current is ramped up slowly for the very reason of leaving time for the pressure due to PSD to decrease.

The conditioning is quantified by expressing the η yield as a function of the integrated number of absorbed photons, the dose (D). On a physical level, it is the surface dose (absorbed photons per cm^2) that determines the yield, however in literature the measurements are expressed almost exclusively in linear dose (absorbed photons per m). I'll discuss the reason for this and a possible conversion between the two units later in this chapter.

3.1 PSD data in literature

PSD measurements are usually carried out at synchrotron light sources, using the SR fan originating from a bending magnet. Most often the magnet is isolated from the experiment area by a radiation protection wall, and therefore the synchrotron light source is relatively far. The originally wide SR fan is thus collimated so that it can travel through a long pipe under vacuum, resulting in a narrow, well delimited, practically straight light by the time it arrives to the experiment area.

This is then projected either to a perpendicular target[65], or to the wall of a long pipe which is slightly tilted compared to the SR direction, so that the incidence area spreads all along the pipe's wall, such as in Fig.3.5.

By the time such an experiment is set up, out of the four parameters determining η (photon energy, material, gas type and dose), the energy and the material are already fixed. The experiment therefore lasts for several days, during which the dose accumulates, and one or two calibrated residual gas analyzers (RGAs) measure the partial pressure rise of several gas species, most often H_2 , CO , CO_2 and CH_4 , and sometimes H_2O , N_2 , O_2 , C_2O , etc.

Yield results (deduced from the pressure rise) are then usually plotted on a log-log plot (like in Fig.3.11) where the PSD yield η is expressed as a function of the accumulated linear dose. These curves typically start with a constant, high yield, until the *cutoff point* of about 1×10^{20} ph/m, corresponding to the time where the cleaning of the surfaces starts. The second part - after the cutoff - is a straight line on log-log the plot, described by the power law

$$\eta \sim D^\alpha$$

where α , the slope is typically between -0.6 and -1. At this point it's worth highlighting that the actual desorption rate is the product of the incident flux and the molecular yield. The extreme case of $\alpha = -1$ would mean that across different parts of an accelerator, the outgassing rate would be the same regardless of the impinging SR flux: high flux areas – though receiving more photons – would in turn condition faster and the two terms of the product would cancel each other out.

The desorption curves are published for several common materials, such as copper, steel and aluminum[62][66]. As we'll see later that the treatment of a surface can largely alter the molecular yield, the history of each sample is carefully documented in these articles.

3.1.1 Pressure calculation methods

When doing vacuum calculations with the SR flux already available, a physicist can decide to use a PSD conversion curve in literature, if his accelerator's parameters (photon energy, surface material and treatment type) are close to the published experiments. Otherwise, a dedicated experiment is necessary. In this chapter I'll cover both methods. For large machines, often a rough pressure estimation is sufficient in the initial design phase: by making a list of the accelerator parts (such as bends, straight sections, etc.), each list element can have its own photon flux (calculated analytically, often by geometrically tracing the SR fan on the blueprints), molecular yield, conductance and pumping speed. As presented in the analytic calculations section in chapter 1, such a list – typically realized in Excel - would then calculate the outgassing for each element (for a given integrated dose), then output a single pressure value for each part. Though sufficient for a first approximation, the method doesn't take into account:

- Local surface features (absorber masks, pumping ports)
- Photon scattering to adjacent parts
- Pressure variations within elements

A more detailed approach, still for large structures, is implemented by the Monte Carlo code Synrad3D[57], where the geometry is defined as a succession of cross-sections. The chamber wall is then interpolated longitudinally between these cross-sections, and photon scattering is taken into account. Although the code doesn't do the PSD conversion nor vacuum simulations, its photon tracking output can be used for a more sophisticated outgassing map definition.

Doing PSD conversions for multipart geometry features (like crotch absorbers, grids) is not possible with the above methods: the geometry would need to be split into too many small parts (each receiving a different flux thus conditioning at a different rate).

My codes therefore fit in the palette of SR simulation tools by:

- Being able to simulate SR in small, detailed structures (grids, pumping ports, crotch absorbers), using the interface with CAD programs
- Keeping account of the received flux on all surface points (in form of textures, with user-defined resolution)
- Performing the PSD conversion automatically for each texture cell
- Staying within the same software package, performing a vacuum simulation by desorbing gas according to the converted values.




Figure 3.1: The original experiment schematics from Fig.1. of [69] and its simulation in SynRad+

3.1.2 Conversion between linear and surface dose

In theory, programming the PSD conversion is straightforward: the user defines the conversion plot either by an equation or by digitizing a plot published in literature. Then, based on the machine time, the surface flux for each texture cell is converted to a surface dose, which will allow interpolating the actual molecular yield. Having that yield, the outgassing rate for that segment of the wall is determined, and repeating this for every cell in the system, we define the outgassing map, to be respected during the subsequent vacuum simulation.

The problem with using data from literature is that the published photon dose is linear (expressed in photons/meter). It is determined by dividing the total flux passing through the collimator by the length of the irradiated area, a calculation method which is preferred by authors because its exactness. However, it is clear that on an elementary level it's the surface dose that determines the cleanness of a surface point, and for that reason Synrad+ uses this latter for conversion (it isn't possible to talk about linear dose on a jaw element of a crotch absorber, for example).

During the experiments, photons get reflected from the primary incidence location, many getting absorbed on the end flange of a tilted tube, or some on the opposite side. In practice, the measured PSD yield is the average of these differently conditioned surfaces. (Some authors have modified their experimental setup to be able to measure the forward-scattered photon quantity[67], and others estimated that in case of stainless steel with electrodeposited copper coating, photons reflected to the opposite side amount to about 25% of the primary incidence[68]).

As a result, before using a curve from literature, we have to recreate the experiment in Synrad+ using the published geometric dimensions, and calculate the actual surface covered by the photons. An example for the process is illustrated in Fig.3.1, where Groebner et al.'s experiment[69] is modeled. The 1×10^{15} photons/sec flux of the original experiment spreads over a 3.6 m long impingement length on the 11 mrad tilted pipe wall, resulting in a linear flux of 2.7×10^{14} photons/m. Synrad+ simulations show an average flux of 4×10^{12} photons/sec/cm² on the approx. 1 cm high, 3.6 m long slice of the tube where the majority of the SR is absorbed.

Comparing the two quantities, it is deduced that 70 photons/m linear dose of the published

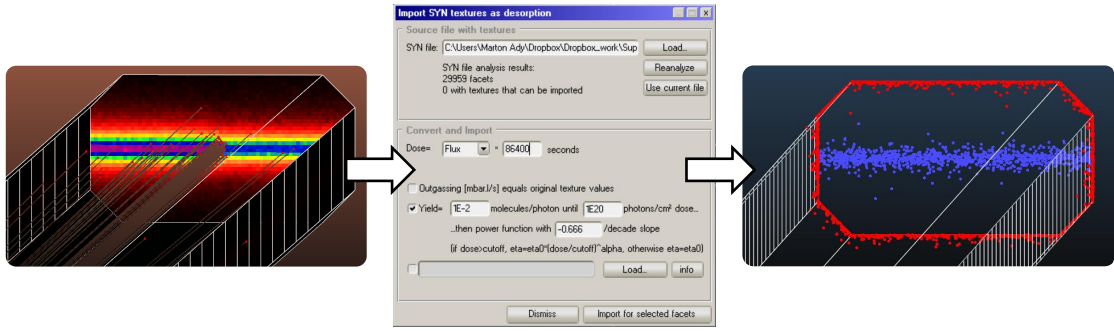


Figure 3.2: The PSD conversion process. Textures from a Synrad simulation are converted to an outgassing map. On the right side, we can see Monte Carlo test particles desorbing (blue dots) following the converted flux data (left)

data corresponds to 1 photon/cm^2 surface dose simulated by Synrad+ on the irradiated parts. This means that the conversion is in practice a shift of the X axis on the yield plot.

This extra step must be done before using any literature data in Synrad+, although we will see that the conversion factor will have values in the same order across experiments. For example, repeating the analysis for a NEG PSD experiment at ESRF[70], we get a factor of $52 \text{ photons/m} = 1 \text{ photon/cm}^2$, close to the previous value. The process becomes more complicated if we take into account photon scattering – one has then to decide what to count towards the irradiated area.

Using Synrad+, an extensive comparative study calculating the conversion factor for real (LEP, NSLS, APS) accelerator geometries has been carried out by J. Carter at APS [71] for aluminum surfaces. His conclusion was that depending on the actual geometry, the PSD yield curves can vary by up to an order of magnitude across different experiments – for the same surface material.

Once we perform the analysis above for the published experiment in question, the resulting PSD yield curve can be used as input for Molflow+, which – as illustrated in Fig.3.2 - extracts flux data from a Synrad+ simulation and creates the outgassing map for the selected facets.

3.2 Experiments at KEK¹

In accelerator design, we often use elements of which the PSD data is not yet available in literature. In CERN's vacuum group, thin films coatings of NEG and amorphous carbon have been the subject of interest for many years, because of their efficiency in reducing PSD and also photoelectron (PE) yield.

Since these coatings are candidates for the FCC machines (both FCC-ee and FCC-hh), with their design underway, a partnership between CERN and KEK has allowed us to carry out two experimental campaigns in 2014 and 2015. The goal was to quantify PSD and PE yield from these surfaces, and to compare the results with uncoated reference samples.

We have designed a vacuum chamber, shown on Fig.3.3 to be irradiated with synchrotron light

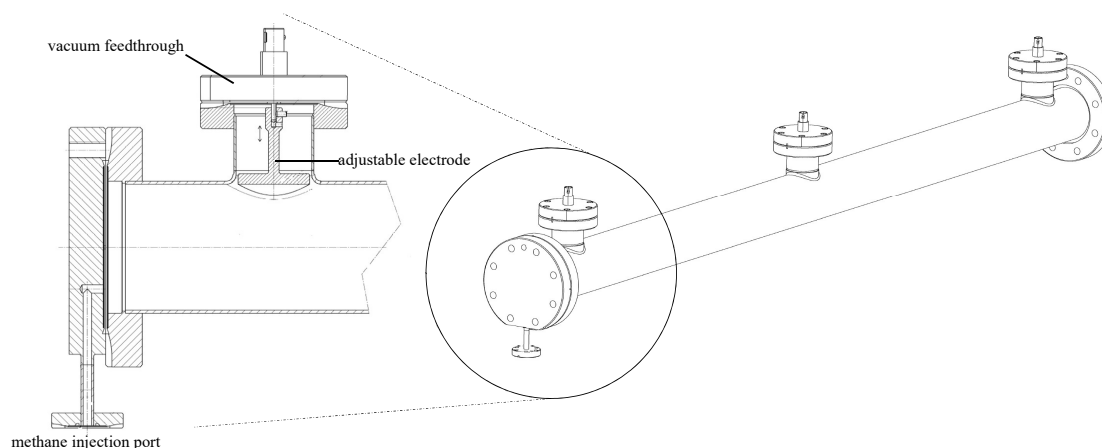


Figure 3.3: Vacuum chambers used

in similar conditions to those of the FCC-hh (see table 3.1). CERN has manufactured six of these, then prepared and coated them differently. The chambers were then installed in one of the experimental hutches of KEK's Photon Factory (PF).

3.2.1 Samples

The 1200 mm long, 61 mm internal diameter vacuum chambers were manufactured of 316L type stainless steel, with three 35 mm ports opened by mechanical extrusion and subsequent welding. All ports use standard ConFlat flanges, one rotatable (machine side) and four fixed.

The three ports are plugged with vacuum feedthroughs with a pin on the vacuum side - these pins hold stainless steel disc-shaped electrodes. A socket/side screw fitting solution allows precise height and orientation adjustment to ensure that the bent discs follow the tube profile, while a 1 mm isolation gap is left allowing them to be kept at high voltage. The tube is closed from one side by a blank flange with a 4 mm injection tube drilled from its side: this allows gas injection during the experiment, for calibration of the vacuum gauges. The other side is connected to the front end of the beamline.

All chamber components were cleaned following standard CERN chemical degreasing procedure [72, 73] while small parts such as screws were cleaned with ultrasound in an ethanol solution. Following this basic treatment, each sample was conditioned and coated differ-

¹Special chapter acknowledgements

This collaboration between CERN and KEK was a real teamwork of many people. First and foremost, I'd like to thank prof. Yasunori TANIMOTO who prepared the experiment on KEK's side, and once I arrived, supervised both me and the whole experiment. As this thesis will probably get published before our article describing this experiment, it is very important that he is to be credited equally for all results.

I'd also like to thank T. Honda, P. Chiggiato and R. Kersevan for making the collaboration possible.

I had indispensable help in the design and manufacturing of the samples – at a very tight schedule – from Luca Dassa, Jean-Marc Malzacker and Raphael Fernandez-Gomez from the mechanical design team of CERN. The cleaning and coating of the samples couldn't have been done without Marc Thiebert, Antonios Sapountzis, Benjamin Holliger, Pedro Costa Pinto, Mauro Taborrelli and Wilhelmus Vollenberg from the CERN vacuum group.

Finally, I thank Anthony Harrison and Fabrice Valceschini for their help in transports and logistics.




Figure 3.4: Amorphous carbon (left) and NEG (right) coatings on our pickup antennas standing on the coating support.

Photos: Wilhelmus Vollenberg, CERN

Table 3.1: SR parameters for our setup and the FCC-hh[75]

	Photon Factory BL-21	FCC-hh (16-20 T)
Stored beam	2.5 GeV electron	50 TeV proton
Beam current	450 mA	510 mA
SR critical energy	4.0 keV	4.3-5.4 keV
SR power density	20.1 W/m (10 mrad incidence)	28-44 W/m (per aperture)

ently. Table 3.2 summarizes the treatment and coating types. It is important to mention that stainless steel and NEG-coated chambers' PSD yield was measured previously [70, 74] and NEG's PSD mitigation capability is already known, but our goal was to systematically test different coatings in the same setup with same conditions, starting from uncoated (reference) chambers, moving towards coated ones, so that an exact comparison is possible.

As part of the pre-treatment, samples 2, 3 and 6 were vacuum fired at 950°C for 2 hours to liberate the H₂ gas stored in the bulk.

3.2.2 Experimental setup

We performed our experiment in PF's BL-21 hutch, because of its SR parameters with similar parameters to the FCC-hh. We tilted the pipe so that the SR hits the wall of the chamber with a 10 mrad incidence angle, similar to the beam pipe of an accelerator. The length of the sample was chosen to use all the available space in the hutch.

The SR light has properties introduced in Table 3.1, and its calculated energy spectrum is shown in Fig.3.6. The SR fan is shaped by a water-cooled copper collimator of 10 mm x 5 mm orifice, then it passes a known vacuum conductance (7.05 l/s for nitrogen at room temperature), until it finally hits the wall of the tube over a length of 1 m. At full current, the SR

Chapter 3. Photon stimulated desorption

Table 3.2: Treatment types and experiment durations for the six vacuum chambers

Sample	Treatment and coating	Vacuum firing	NEG activation	Experiment duration
1	Stainless steel			10 days
2	Stainless steel	O		7 days
3	Amorphous carbon coating	O		14 days
4	TiZrV NEG coating			6 days
5	TiZrV NEG coating		O	6 days
6	TiZrV NEG coating	O	O	7, then 12 days

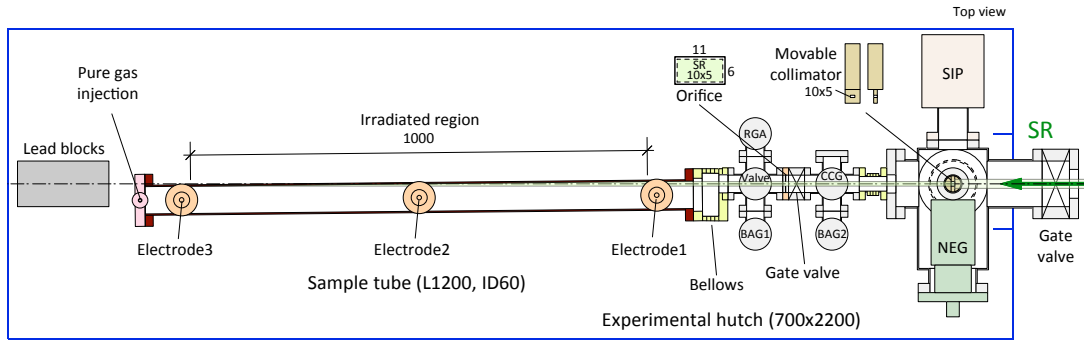


Figure 3.5: Experimental setup overview. Image: Y. Tanimoto

flux arriving on the sample is 5.5×10^{16} photons/sec and the power is 20.1 W. Depending on the coating type, the majority of these photons are absorbed either at the primary incidence location or on the end flange after one reflection. In case of sample 6, the nominal beam current of 450 mA was reduced to 400 mA to test a hybrid machine mode.

The system is pumped by a NEG pump of 1500 l/s and a noble-gas-capable sputter ion pump of 55 l/s, both located upstream. The vacuum sample is connected to the machine through a bellows to handle the 10 mrad rotation, and is aligned based on SR radiation patterns on photosensitive films fixed on the exterior. The alignment – as done by naked eye – has a precision of about 1 mm. (Because of the very small tilting angle, this 1 mm error can lead to a 10 cm shift in the longitudinal position of the SR incidence.) When absorbed, the SR photons desorb gas molecules from the surface, increasing the pressure in the chamber. This desorbing gas can be quantified by the pressure difference on the two sides of the known conductance, indicated by the Bayard-Alpert gauges BAG1 and BAG2 in Fig.3.5. A residual gas analyzer logs the ion current for atomic masses 1 to 100. Before our experiment, this RGA was calibrated at the National Metrology Institute of Japan by Dr. Hajime Yoshida (using the method presented in[76]) for CH₄, CO, CO₂ and H₂, therefore the ion currents of these gases can be converted to partial pressures, giving us an indication on the composition of the desorbed gas.

With the vacuum chamber grounded, the three electrodes are biased to 84 V, the highest available voltage of the used power supply, and the currents of collected electrons are measured by picoammeters. Near the electrodes, three solenoids of 40 turns are wound around the vacuum pipes (see Fig.3.10), capable of transmitting currents up to 20 Amps (temporarily, to avoid




Figure 3.6: Synchrotron Radiation spectrum of the dipole of the BL-21 hutch, as calculated by Prof. Tsuchiya in KEK, and as simulated by SynRad+




Figure 3.7: Assembled experiment before and after bakeout




Figure 3.8: PSD desorption yield calculation method, with p_s and p_r the sample- and ring-side pressures, C the known conductance, S the pumping speed and F_{photon} the incident SR flux

overheating). This allows to create weak magnetic fields inside the tube, to study the effect of a known electron cloud mitigation method by external solenoids (heavily implemented on KEK-B's positron ring). The injection port of the end flange is connected to a CH_4 bottle allowing injection through a variable leak valve and a SAES MC1-905FV gas purifier which removes trace-level water molecules from the methane.

3.2.3 Experiment

All samples were baked out at 120°C for 20 hours, and on samples 5 and 6 the NEG coating was activated at 250°C during 4 hours. Then, with the lead-plated experimental hutch closed, the ambient temperature stabilized at 31°C , and the bias voltage was applied to the electrodes. Opening a shutter, the samples were subjected to the SR for a duration between 6 and 11 days per sample, as described in Table 3.2, reaching an accumulated photon dose in the order of 1×10^{22} photons/m. The pressures of the BA gauges, the currents on the pico-ammmeters, the sample's wall temperatures near the electrodes and the beam current were logged every second. The RGA ion currents were logged separately every 53 seconds (time of a full scan), and we have merged the two data files using linear interpolation, based on timestamps.

We were interested in the following data:

3.2.3.1 Total photon stimulated desorption yield, as a function of the total absorbed dose

It is calculated from the pressure difference measured between BAG1 and BAG2, the photon flux derived from the measured beam current and the known conductance, as shown in Fig.3.8. and explained in eq.3.1.

$$\eta = \frac{Q_{psd}}{F_{ph} \cdot kT} = \frac{C(\Delta p_s - \Delta p_r)}{F_{ph} \cdot kT} \quad (3.1)$$

The background pressures are measured prior to opening the shutter and deducted from the gauge pressures during yield calculation.

The absorbed photon dose is calculated by integrating the beam current over time, and therefore estimating the total number of photons that enter our sample.




Figure 3.9: Photoelectron trajectories under 84 V pickup voltage
Simulations with Opera by Thibaut Richard [CERN]

3.2.3.2 Partial (CH_4 , CO , CO_2 and H_2) molecular yields

They are calculated from the converted ion currents of the RGA. Since we had only one calibrated RGA installed, we could not measure the partial yields with the orifice method used for the total yield. Instead, we assumed that the pumping speed S , deduced from the total pressure difference to be 21.4 l/s for the mixture of all gases, scaled with the molecular speed (i.e. the inverse of the square root of the molecular mass m). Using that speed (different for each type of gas), and using the equations

$$\frac{1}{S_{\text{eff}}} = \frac{1}{S} + \frac{1}{C(m, T)}$$

$$Q_{\text{PSD}} = \Delta p_{\text{sample}} S_{\text{eff}}$$

where S_{eff} is the local pumping speed at the sample side, we could convert the partial pressures indicated by the RGA to partial outgassings.

3.2.3.3 Photoelectron current, as a function of the electrode bias voltage

To get this data, we have performed a bias test for each sample where we varied the bias voltage simultaneously for the three electrodes in discrete steps between 0 and 84 V, and recorded the electrode currents measured by the pico-ammmeters. The originally straight electron trajectory is bent by the bias voltage, as illustrated in Fig.3.9.

3.2.3.4 Photoelectron current as a function of the absorbed photon dose

This can be derived from the output of the pico-ammmeters, which was logged all the time. The goal is to see if the PE yield decreases over time.




Figure 3.10: External coil on the sample wall allowing to create a magnetic field

3.2.3.5 Photoelectron current as a function of the external magnetic field

For each sample, we have used our power supplies to circulate different currents in 1 A steps in both directions. Although the extrusions of the test chambers around the electrodes didn't allow us winding the coil evenly, based on our calculations, we could create a fairly uniform magnetic field of approximately 30 Gauss at the maximum current of 20 A. This test helped us to confirm that electron cloud effects can be mitigated in an accelerator by an external magnetic field.

3.2.3.6 Wall pumping of methane

We have investigated a phenomenon that was observed during a PSD measurement[77]: although NEG doesn't pump methane, when SR was incident on the wall of a NEG-coated vacuum chamber, the reduction of the partial pressure of methane was observed. It is believed that this wall pumping of methane is due to its photo stimulated cracking to other hydrocarbons[78], which in turn can be pumped by NEG.

We tested this hypothesis by injecting a constant flux of purified CH_4 to our system, and measured its partial pressure with the SR shutter open and closed.

3.2.4 Results

3.2.4.1 Total and partial PSD yields

Figure 3.11 shows the PSD molecular yields of the six samples as a function of the absorbed photon dose. Since the RGA supplies information about the gas composition, the total yields, measured by the BA gauges were corrected for the different gas sensitivities. This (corrected) total yield is summarized for all samples on Fig.3.13. Please note that this figure contains a second run for the activated, vacuum fired NEG sample (no. 6). This was performed

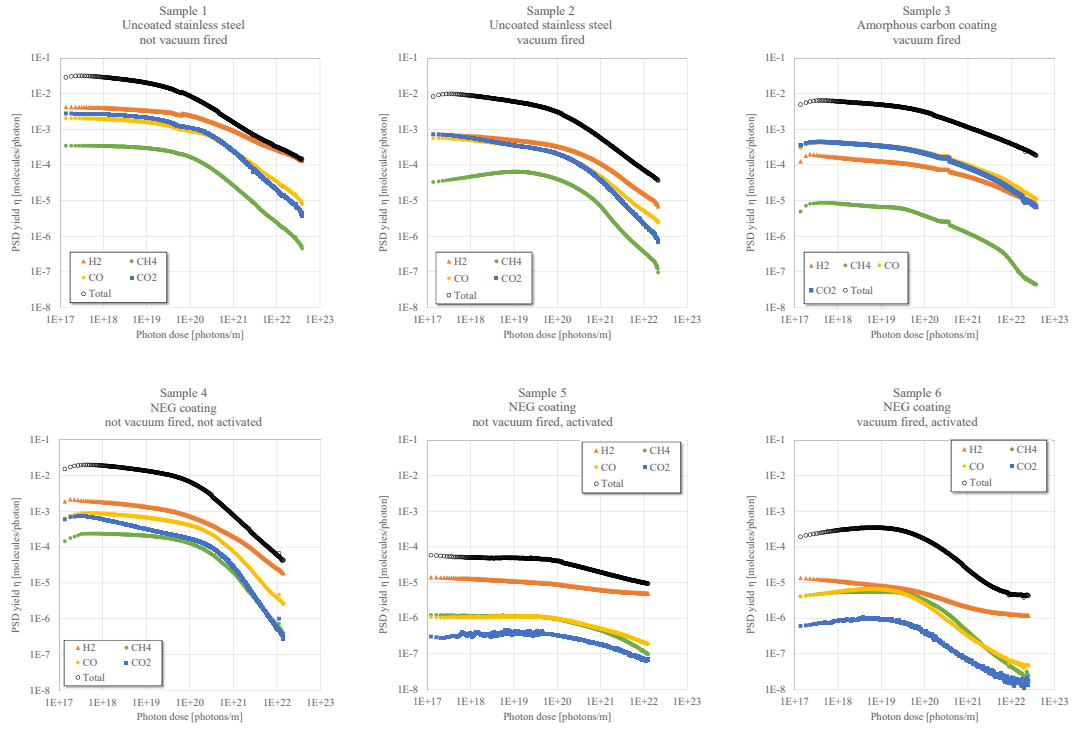


Figure 3.11: Total and partial molecule yields for each sample, as a function of the absorbed photon dose

five months later, to investigate if the higher than expected initial PSD yields were due to a contaminant that disappeared during conditioning. During these months the sample was kept under vacuum. To reset the SR conditioning of the primary incidence location, the tube was briefly exposed to air and rotated by 45 degrees around its longitudinal axis compared to the first run.

The following observations can be made:

- NEG coating, when activated, was verified to have the lowest PSD yield: it reduced the molecular yield by approximately two orders of magnitude.
- For uncoated samples, vacuum firing can reduce the H_2 outgassing (see sample 1 vs. sample 2). For activated NEG samples, where the coating itself is pumping, this effect - though observable at higher doses - is less significant (see samples 5 vs 6).
- For sample 6, the second run's PSD yield was proven to be approximately half an order of magnitude lower, indicating that a brief exposure to air doesn't reset an irradiated NEG sample to its original state. As the gas composition of this second run was practically the same as the first (absolute values scaling with the total yield), only the second run is plotted in Fig.3.11.
- In case of activated NEG coatings, the pressure drop comes from two effects: lower intrinsic PSD yield reduces the amount of gas desorbed, but also these desorbed molecules

have a certain probability to be repumped by the NEG surface before reaching the RGA or the BAG gauges. We could observe this second effect directly: after turning off SR irradiation, despite the ring side pumping, the sample side reached lower pressure than the ring side. Distinguishing the two effects is difficult, and similar to [70] we attempt to do it by modeling the desorption and pumping process inside the tube with Monte Carlo simulations: we use Molflow+ to create a 3D model of the vacuum parts of our system, as shown in Fig.3.12. We desorb molecules at locations marked red (primary incidence location along the tube wall and the end flange where photons are reflected). Then we apply different sticking factors to the NEG tube, and see how the pressure difference recorded by the facets representing the BAG gauges (on the two T-intersections visible in the back) scales compared to the no-sticking reference case. Results (plotted on the right of the figure) show that even a low sticking probability of 10^{-3} decreases the PSD-induced pressure rise by an order of magnitude, while a sticking of 1 would let only those molecules escape the tube that fly parallel to its axis, and would therefore drastically reduce the measured pressure rise (by 6 orders of magnitude). We have, however, decided not to correct for this effect in figures 3.11 and 3.13, and - similar to [70] - calculate the molecular yield directly from the perceived pressure rise. We argue for this choice as in case of CH_4 , there is no NEG pumping[79], yet we see the decrease of the partial pressure, clearly showing that (at least for that gas) NEG - when activated - has a lower intrinsic desorption. As for CO and CO_2 , the one-week duration of the experiments asserts that the NEG surface eventually saturates and its sticking decreases, yet we see that the PSD yield curves decrease steadily, i.e. the PSD yield vs. accumulated photon dose curve has constant slope, showing that the pumping properties of NEG are probably altering only the low-dose measurements. Finally, we publish uncorrected results as the main purpose of our measurements is to facilitate engineering calculations for future accelerator design, in which case distinguishing lower intrinsic yield from NEG pumping is not necessary as the SR-induced pressure rise (result of the two effects together) is used.

- Considering the above, to compare with previously published results² in [74], we take into account our non-activated sample (no. 4) for photon doses below 1×10^{21} ph/m and activated samples (no. 5 and 6) above. Results are summarized in Table3.3. The conditioning slopes match very well for samples 4 and 5, although the absolute values differ. Our activated, vacuum-fired sample 6 conditioned significantly faster and although molecular yields are almost the same at our comparison point of 1×10^{22} ph/m, at other doses the match is not good. This is contrary to our expectations as the vacuum chamber in [70] was vacuum-fired and therefore we expected its yields to be closer to our sample no. 6. This discrepancy of the yield across different experiments shows that executing dedicated PSD measurements are warranted where really precise calculations are necessary.

3.2.4.2 Photoelectron yield

The photoelectron current - normalized by the beam current - is shown as a function of the absorbed dose on Fig.3.14 (We label the ring side, middle and absorber end-side electrodes as

²In that experiment the NEG sample was activated at 1×10^{21} ph/m dose.




Figure 3.12: Left: Monte Carlo simulation of pressure distribution in the system for a NEG-coated tube with sticking 0.007; Right: Calculated pressure rise on the two BAG gauges for a molecular yield of 0.001, as a function of the NEG sticking factor

Table 3.3: Comparison of our experimental results with those published at [70]

	[70]	Sample 4	Sample 5	Sample 6 (1st run)	Sample 6 (2nd run)
Yield at 10^{21} ph/m [mol/ph]	2×10^{-3}	6×10^{-4}			
Slope below 10^{21} ph/m	-1	-1.1			
Yield at 10^{22} ph/m [mol/ph]	3×10^{-6}		8×10^{-6}	4×10^{-6}	2×10^{-6}
Slope above 10^{21} ph/m	-0.38		-0.38	-1.14	-0.77

no. 1, 2 and 3).

We can conclude the following:

- Amorphous carbon coating has the lowest PE yield, moreover this yield is constant
- Similarly, activated NEG samples have constant, low yields
- The stainless steel and the non-activated NEG samples, however, have a significant conditioning of their PE yield, up to approximately 1×10^{22} ph/m dose.
- During the second run of sample 6 we have had an approximately 20% higher photoelectron yield. We have performed a bias test (ramping up the electrode voltage and measuring the PE current response) on each sample. Although these were performed at slightly different doses (usually near the end of the measurement), we can see in Fig.3.14 that the measurement times (marked by crosses) are at high doses where conditioning is already slow, therefore test results across samples remain comparable. Figure 3.15 shows the individual response of each electrode for all samples. As higher electrode voltages can attract electrons up to a higher energy, and also the action radius (the volume from which electrons are attracted) is extended, the result is a current response curve that




Figure 3.13: Total molecular yields as a function of the absorbed photon dose, for all samples from the opening of the shutter to the end of the experiment

scales with approximately the square root of the bias voltage. It has the same shape over different samples, and differs only in absolute value (see summary in Fig.3.16). Up to our maximum available voltage of 84 V we did not observe any saturation effect.

3.2.4.3 External magnetic field

We have performed the PE mitigation by magnetic field test for samples 1, 3 and 4 (reference stainless steel, amorphous carbon and non-activated NEG). For each sample we have performed the test twice: one with the 84 V electrode bias and one without. The PE current response of each electrode is shown in Fig.3.17.

We can see that applying an external magnetic field can reduce the PE current, but the response of the samples - especially in the unbiased case - is asymmetric. To interpret the results, Fig.3.18 helps to visualize what is happening inside the vacuum chamber: the solenoid's magnetic field (parallel to the tube's axis) steers the photoelectrons through the Lorentz force, making them fly on a curved path. In an accelerator, if this magnetic field is high, the radius of the path can be small enough that the electrons get absorbed on the wall again before they could get into center of the beam pipe where they would interact with the beam.

The asymmetry in our case rises from the geometry of our experiment: the incidence location of the SR (and thus the origin of the photoelectrons) has an angle of 90 degrees with the electrode. Depending on the direction of the current and the resulting magnetic field, those electrons that aren't directed back to the wall can be steered either away from the electrode, or towards it. Therefore, especially in the unbiased case where there is no additional attraction of electrons by the electrode, we can see that there is a "good" direction where electrons are steered away from the electrode thus the PE current is reduced, and a bad one where the PE current is more or less constant, regardless of the magnetic field.




Figure 3.14: PE current per sample as a function of the photon dose. Cross positions mark the point where the bias tests were performed

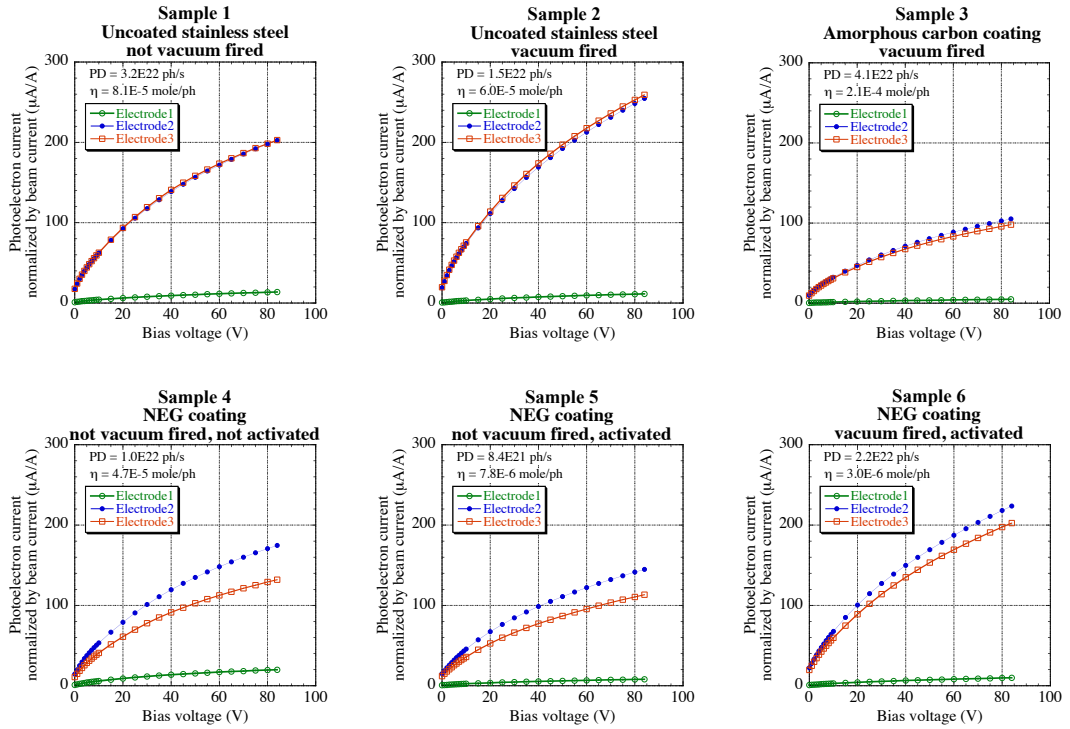


Figure 3.15: Photoelectron current response to bias voltage

Chapter 3. Photon stimulated desorption




Figure 3.16: Comparison of photoelectron current response on the middle electrode to bias voltage for all samples

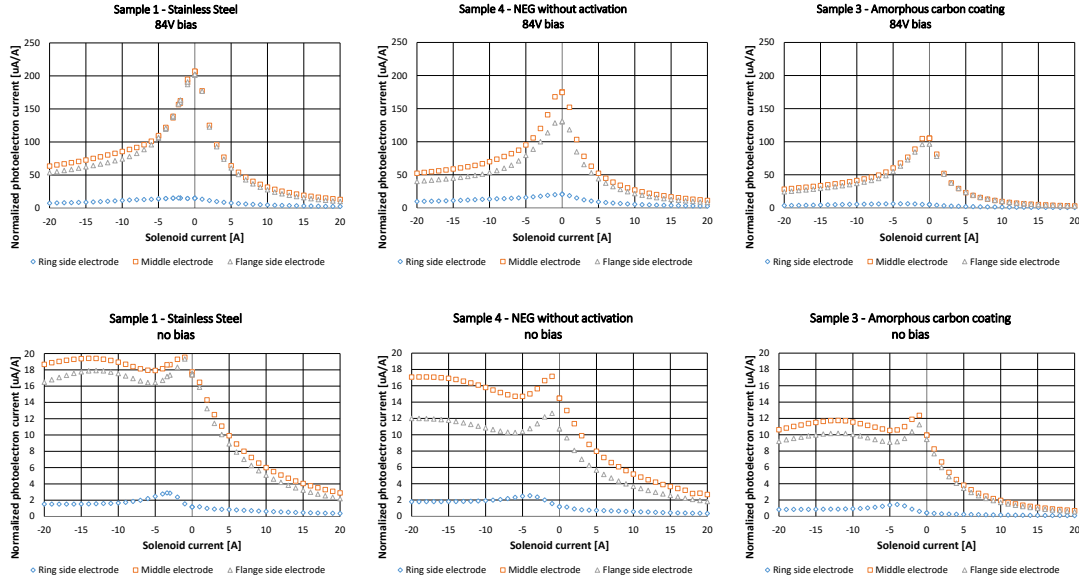


Figure 3.17: Photoelectron currents as a function of the current in the coils for biased and unbiased cases. Please note that the Y scale is different for the two cases, since the peak photoelectron current was approximately 10 times higher when the electrode was biased.




Figure 3.18: Photoelectron trajectories under external magnetic field, for explanation of the solenoid test

In case of a real accelerator, where we want to reduce the overall PE yield (and not the one measured on an electrode), the current direction doesn't have importance.

3.2.4.4 Methane pumping

To reproduce the effect mentioned in section 3.2.3.6, at the end of the test run for each NEG-coated tube (samples 4, 5 and 6) we have injected a constant flux of methane through the purifier. This flux was set by a variable leak valve to increase the methane partial pressure (indicated by the mass 15 current on the RGA) to approximately 3 orders of magnitude above the value without injection (this large difference was necessary to clearly distinguish the pressure rise due to the injected flux from the one due to PSD). Then we have opened the SR shutter and checked if the RGA current changed. Figure 3.19 shows the measured data for sample 5. (Samples 4 and 6 behaved similarly thus they are not plotted).

For H_2 , CO and CO_2 we can see a pressure rise coincident with the shutter opening duration, due to PSD. In case of CH_4 , the pressure either increased slightly (sample 4), or decreased without any change in the rate (sample 5 and 6). Therefore - contrary to our hypothesis - we didn't find any evidence of SR induced wall pumping of methane on NEG coatings.

3.2.4.5 Error of the results

Of the eleven PSD articles this thesis references, ten don't discuss the error of the results [62, 64, 65, 66, 68, 69, 70, 74, 77, 80], while one analyzes the sensitivity of the gauges and assumes the reproducibility to be within 20% [67]. The reason for this is probably that PSD molecule yields are converted from pressures measured by gauges which are calibrated thus considered to be precise. Nevertheless, I attempt to give an error range for our results, taking into account the following:

- Based on my colleagues' experience at CERN's vacuum group, a Bayard-Alpert gauge, unless carefully kept at constant temperature, will have an error of around 15%. Since PSD yields are converted from the gauge values as per eq.3.1, this 15% might be a good error estimation for the molecular yields. As the calculated values vary across several orders of magnitude, this error is relatively low.

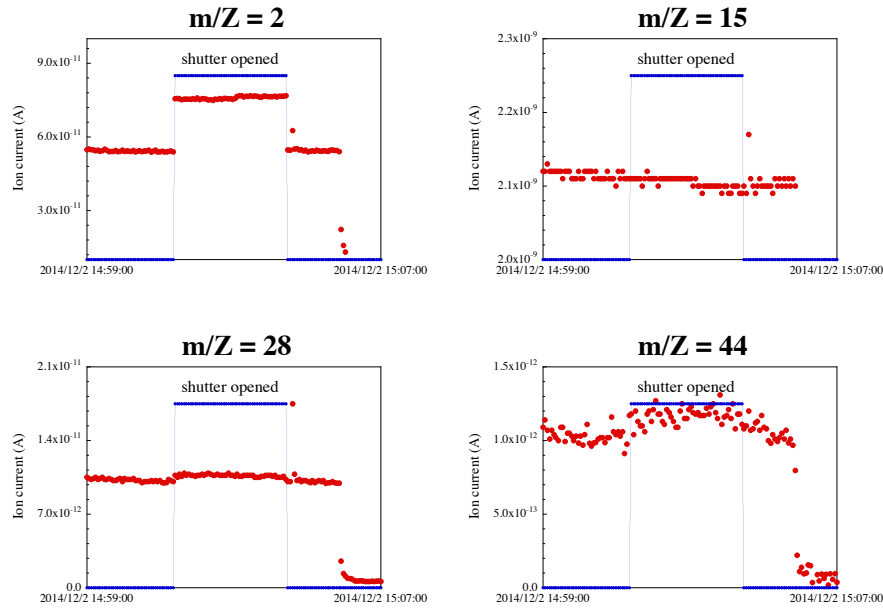


Figure 3.19: RGA currents at shutter opening (Sample 5: NEG, not vacuum fired, activated)

- The pico-ammeters used for the photoelectron yield measurements were calibrated for the zero current at every startup. Therefore I estimate their error to be those decimals on its display that fluctuated during constant experimental conditions. As such, it is in the order of 0.1%.
- The beam current value, used for calculation of the integrated dose, was sent to the experiment area through a measurement point of varying voltage, sampled by our recorder. Expressed in mA, it was also constant up to four digits (including the decimals), giving a negligible error of <0.02% for the nominal current of 450 mA. Assuming that our time measurements (synchronized from internet NTP servers) were precise, this error is a good assumption for the integrated dose.
- Finally, for our bias tests, the power supply voltages could be digitally regulated up to two decimals, therefore I estimate the error to be <0.01% for the 84 V maximal bias.

3.2.5 Experiment summary

Six different vacuum chambers were irradiated with SR similar to FCC-hh at KEK, and the photodesorption yield and the photoelectron currents were measured. Quantitative results allow exact comparison of different coatings, and because of the SR parameters and the sample geometry, they can be used for future machine design. We have confirmed previous results that activated NEG coating reduces the photodesorption significantly and that amorphous carbon coating reduces the photoelectron yield. We have also confirmed that the photoelectron yield can be (further) reduced by an external magnetic field. We have tested a hypothesis of methane pumping by NEG under SR, but we have found no evidence of the phenomenon. It




Figure 3.20: Modified chambers allowing liquid N₂ cooling

should be mentioned that in a real machine, where not only the SR fan is present but also the particle beam, a noticeable methane pumping is visible and attributed to ionization by the latter (presented in [81], numerical value in [82]).

3.3 Experiments at liquid N₂ temperature

One year after the first measurement campaign the experiment was repeated at KEK, this time cooling a NEG coated chamber with liquid N₂ – as close to the temperature of the FCC-hh beam screen as possible. Since this change required a different chamber, a modified setup and a modified measurement protocol, it is discussed in this separate section.

3.3.1 Samples

As shown in Fig.3.20, I have "recycled" the two last tubes from the initial test campaign and modified them to allow cooling the full body with circulated liquid nitrogen.

The modification consisted of cutting all original flanges, and retrofitting an external shell of stainless steel. This shell, with a 1 cm larger diameter than the original tube, is made of two halves, and kept in place by spacer half-rings and electron beam welding.

An entrance/exit tube and the alternated upper/bottom positions of the spacer half-rings ensure that a liquid can be inserted between the original tube and this shell, and its flow uses all the surface, not just the shortest path between the entrance and exit. The flow pattern is very well visible in Fig.3.22 through the ice buildup.

Apart from the modification of the original chambers, two fixed end-flanges were also manufactured to close the tube from one side. The flange bodies are also cooled: they have a hollow




Figure 3.21: Tube's wall deviation from the straight line
Metrology by Cyril HAERINCK, CERN

structure with an entry and exit tube: serially connected with the main tube, liquid N₂ can cool them from the inside.

It is important to mention that during the cutting, welding and extrusion procedures, the tube suffered a minor deformation (probably due to thermal shock). Laser metrology shows (in Fig.3.21) that the wall's longitudinal line has a maximum deviation of 0.12 mm from a perfectly straight line. Although higher than the design tolerance of 0.1 mm, and much higher than the room temperature samples of the previous campaign (max. deviation 0.02 mm), we decided to carry on with the experiment.

Since the tube's structure is relatively complex, the usual room temperature leak detection was extended: I have assembled the tube and the flanges (without the pickup antennas), and cooled down the whole assembly by liquid N₂ (see Fig.3.22).

It turned out that the more than 200 K of temperature change causes considerable mechanical strains in the tube: during the cool down process it visibly arched (standing on the two side flanges, the middle flanged raised by 3 mm), but fortunately regained its straight shape when it cooled down entirely³.

After passing the room temperature and the cooled leak tests, the samples were degreased then vacuum fired the same way as the samples in the previous campaign.

³To those repeating such experiments, it might be interesting that several times leaks manifested at cool state between the special end flange and the tube (despite being entirely leak tight at room temperature), suggesting us that the CF flanges don't tolerate thermal shock, probably due to the different expansion of steel and the copper gasket. After warmup, changing gaskets and a repeated cool down, I could achieve complete leak-tightness. Later, at KEK the leaks manifested again, finally fixed by a vacuum glue seal.




Figure 3.22: Leak detection on the assembly at liquid N₂ temperature

One sample was coated with NEG, one with amorphous carbon. As explained later, we could only test the NEG-coated sample at KEK.

3.3.2 Experimental setup

Although the location and the goal of the experiment was the same as the previous year, the liquid N₂ cooling required technical modifications.

Since we were targeting to collect data for several hours, we had to deal with the problems of keeping the sample at liquid N₂ temperature for a prolonged period. To increase thermal insulation and accordingly reduce the amount of liquid N₂ used, the KEK team designed an insulating vacuum jacket around the tube, as sketched⁴ up in Fig.3.23. This box was intended to be pumped to rough vacuum, significantly reducing the conductive heat transfer from the outside to the sample, and at the same time preventing ice buildup from the condensed air humidity.

During operation, it turned out that the rotating CF flanges of the sample aren't leak-tight on the air side, so this rough vacuum couldn't be achieved. Still, the box turned out to be practical: bakeout efficiency was better in a closed volume, and during the experiment, we filled the

⁴The reader might notice a slight difference between the sketch in Fig.3.23 and the final box in the photos of Fig.3.24: although the jacket was designed to be tubular, due to manufacturing constraints, it was bent from a planar steel sheet and thus had an octagonal shape in the end




Figure 3.23: Vacuum insulating jacket, designed by Yasunori Tanimoto (KEK)



Figure 3.24: Assembly steps of the liquid N₂ cooled sample with the thermal insulating jacket

space between the jacket and the sample with Mylar foil, which helped to achieve a very good insulation – moreover the box collected all the condensed water dripping from the sample keeping the setup clean. The assembly process is illustrated in Fig.3.24.

From a physics standpoint, the experiment had essentially the same parameters as the previous year: the photon flux of 5.5×10^{16} photons/m (at the design current of 450 mA) reached a 1 m-long section of the 10 mrad tilted tube. PE yield was measured by the same picoammeters on the biased pickup antennas. Two changes are important:

- This time, as opposed to last year, we had a second calibrated RGA, which we placed on the ring side (opposite to BAG2 in Fig.3.5)
- Due to the liquid N₂ cooling system requiring constant supervision, the experiment length was reduced to 17 hours, corresponding to a maximum accumulated dose of 2×10^{21} photons/m.

Most practical difficulties came from the cooling. Liquid N₂ was circulated by the generated pressure of the dewar, connected to the special flange entrance through insulated tubing.

3.3. Experiments at liquid N₂ temperature

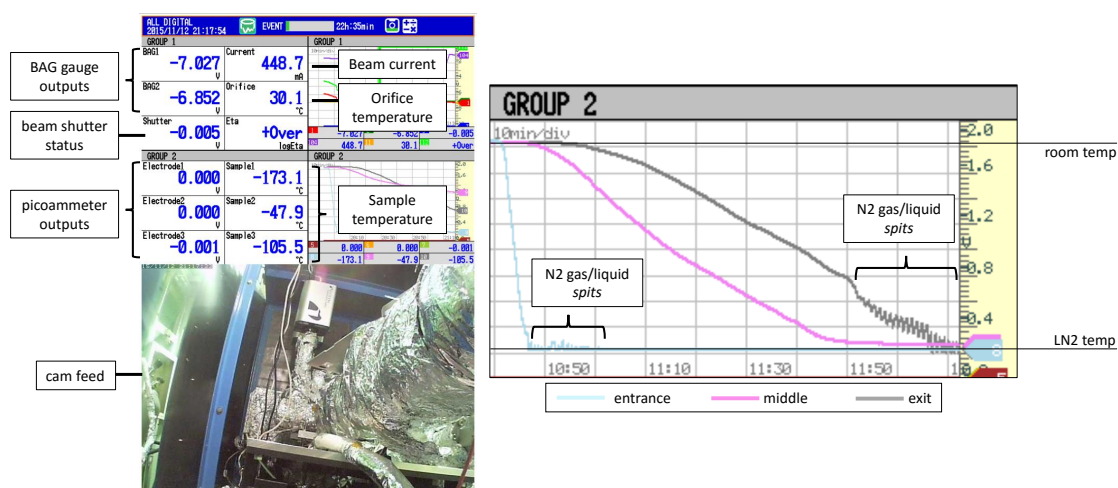


Figure 3.25: Control screen of several parameters of the closed experimental hutch (left) and cooldown process of the sample (right)

The exit of the flange was connected in sequence to the sample's cooling input port, and the coolant was finally evacuated by connecting the sample's exit port to the PF's gas extraction line.

The cooling of the sample was done gradually to reduce mechanical stress: temperature was monitored at three locations on the sample, taking about 90 minutes to bring the whole tube body from room to liquid N₂ temperature. In Fig.3.25 the control screen is shown, where it is visible that near its boiling point, the circulated nitrogen changed phase periodically, resulting in "spits" of gas/liquid mixture, before reaching a stable liquid flow.

During the experiment, the delivery rate of nitrogen had to be regulated carefully through the dewar's pressure builder valve – the flow had to be sufficient to keep the full sample at liquid N₂ temperature, at the same time, a too powerful cooling could have frozen the whole PF extraction line.

Data collection – similar to the previous campaign – was performed by an electronic recorder logging all parameters in Fig.3.25 every second, and data of the two RGAs was collected on two computers, subsequently merging their data with the recorder's through interpolation using data timestamps.

Approximately once an hour we closed the SR shutter and waited around 30 seconds until the pressure stabilized. During post-processing, I determined the background pressure (changing in time) by interpolating between these manual background measurements.

Data collection was performed in two runs, with a 3-day pause at 1×10^{21} ph/m integrated dose, visible as a discontinuity on the results.




Figure 3.26: Pressure data from BAG and RGAs on the sample and ring sides

3.3.3 Results

3.3.3.1 PSD yields

BAG gauges have different sensitivities for different gases. Since this time we had RGAs opposite to both gauges, I could correct their measured (N_2 -equivalent) pressure to real pressure by taking into account the gas composition at every moment.

Since a very cold, previously activated NEG-coated tube is a powerful pump on its own, the pressure in the sample side was very low (below 1×10^{-6} Pa), in fact lower than the pressure on the ring side. This means that we cannot apply eq.3.1 to determine the molecular yield as we would obtain a negative number.

Before deciding the correct post-processing strategy, it is advisable to look at the RGA signals in Fig.3.26.

We can see that the sample side pressure (RGA1) is dominated by H_2 , whereas the ring side (RGA2) has high CO and CO_2 pressures. The shape (conditioning) of the three gases are very similar on the two sides, though. From this, we can conclude that

- PSD in the sample is dominated by H_2 , and as the H_2 pressure is higher on the sample side, this gas flows predominantly from the sample towards the ring
- In exchange, CO and CO_2 pressures are higher on the ring side, so they must contribute to the pressure rise on the sample side (flowing from ring to sample through the conductance between the two locations)

I investigated this ring to sample "backflow" by analytic calculations (the conductance and the ring pumping speeds are known), and I found that any outgassing on the ring side causing Δp_{ring} pressure rise will also cause a $\Delta p_{sample} = 0.02 \cdot \Delta p_{ring}$ pressure rise on the sample side.

In Fig.3.26 I have plotted this "extrinsic" pressure rise labelled as *due to ring* (on the plot, approx. two orders of magnitude lower than the ring pressure), and treated it as a background pressure that I deducted from the total Δp_{sample} for PSD yield calculations.

3.3. Experiments at liquid N₂ temperature




Figure 3.27: Total and partial molecular yields of the liquid N₂ sample, and comparison with previous samples.




Figure 3.28: Lack of PE yield conditioning (left) bias test results (right)

I argue for the validity of this correction because with it the final PSD yield plot of the cooled NEG sample, as shown in Fig.3.27, will have a straight conditioning slope, which is what we expect. Without the correction, the calculated molecular yield would have increased above 1×10^{21} ph/m dose, which we can claim that isn't likely in view of previous experiments.

3.3.3.2 Photoelectron yield

Results show that there isn't a clear conditioning trend of PE yield as the dose accumulates. Nevertheless, the flange-side electrode (no. 3) shows a slight yield decrease while the other two remain constant. The exact reason is not known, one possible explanation is that the power of the SR beam warmed the chamber's inner side at the incidence location, slightly reducing the reflectivity and thus less photons reached the end flange area. An other plausible scenario is that we began the measurements right after the whole sample reached liquid nitrogen temperature - but at that point uncooled extremities of the tube remained warmer, and subsequent cooling of those parts caused a slight thermal deformation of the sample.

Similar to the previous campaign, a bias test was performed at 1.7×10^{21} ph/m dose, with the results in Fig.3.28.




Figure 3.29: Pressure history during the whole experiment, showing the effect of the warmup

3.3.4 Discussion

Comparison with previous samples shows that – as expected – the liquid N₂ cooled sample has the lowest yield among all measured samples. Its PSD is dominated by H₂, and the CH₄ yield is very low, around an order of magnitude lower than the other gases.

In fact, we can prove that the cool temperature on its own reduces the yield: when finishing our experiment, we closed the cooling circuit with the SR shutter still open, measuring the yield while the chamber warmed up. We couldn't determine the exact temperature of the chamber's inner side with the external thermocouples, so just as a qualitative illustration, Fig.3.29 shows the pressure rise on both sample and ring sides – confirming the PSD reduction effect of the temperature alone.

As for photoelectrons, the cooled sample had the highest yield of all samples.

As said at the beginning of this section, in addition to the NEG-coated sample, a second, amorphous carbon coated tube was planned to be measured at cryogenic temperatures. However, two consecutive coating attempts have resulted in the peel-off of the carbon layer (see Fig.3.30).

As the used technology is exactly the same as a year earlier when the coating was successful, CERN's vacuum group has decided to cut the failed chamber for additional analysis. Until the exact cause of the peel-off is determined, a substitute chamber is in production, scheduled to be tested at KEK in May 2016 – the results will be published therefore outside of this thesis.



Figure 3.30: Amorphous carbon peel-off locations after the second coating attempt
Photo: Antonios Sapountzis, CERN

4 Case studies

In the previous chapters I have shown how Molflow+ simulates ultra-high vacuum systems and how Synrad+ is able to calculate synchrotron radiation properties. Although some examples of the codes' application were given, they mainly served to demonstrate that basic calculations match literature.

The goal of developing the codes was evidently to apply them to solve real accelerator problems. Being available as a free download, during the development period they have been applied on several occasions both in science and in industry for case studies. In this chapter, I show three projects where I was actively involved in using my own codes to deliver practical results for different systems. I have chosen one problem from CERN, one from Max IV and one from KEK.

4.1 Simulation of SR power for the LHeC interaction region quadrupoles¹

4.1.1 Introduction

The LHeC[84], which is a project to combine the intense hadron beams of the LHC and the possible Future Circular Hadron Collider (FCC), is in the design stage and various interaction region optics have been considered. The power emitted by the synchrotron radiation in the quadrupoles have been previously calculated by analytic and numerical methods, and published as a CERN technical note[85]. In those calculations an approximation was made: the beam is considered to pass through the center of the quadrupoles, as opposed to the real case where it has a 1 mrad angle.

Repeating the calculations with Synrad+ served two goals: I can benchmark our code against already existing results, and also I can do a simulation which is more accurate by using the precise geometry – as opposed to the approximated one in [85].

¹Copublished with R. Kersevan as a CERN technical note[83]




Figure 4.1: Quadrupole configuration of the LHeC High Acceptance Machine configuration

Table 4.1: LHeC beam parameters

I [mA]	100
E [GeV]	60
ϵ_x [nm]	5
ϵ_y [nm]	2.5

4.1.2 The IR region of the LHeC

In the case of the LHeC, two configurations are proposed[86]: High Luminosity and High Acceptance. In this section I will simulate the optics for the latter. This consists of a pair of focusing and defocusing quadrupoles sketched in Figure 4.1 which - in the approximated geometry - are positioned in a way that the beam passes through their center. If we had an ideal (zero-emittance) beam, we wouldn't have SR as the magnetic field in the very center of the quads is zero. However, a real beam has a width which is represented by a bivariate Gaussian distribution with sigma values calculated from the beam properties.

In order to set up a simulation, we need the following parameters:

- Quadrupole geometries (position, orientation, length)
- Quadrupole strength (magnetic gradient)
- Beam position and direction at the entrance point
- Beam properties (current, energy, emittance, coupling, beta functions)

The beam properties are obtained from the design report[86] and summarized in Table 4.1. The beta values along the quadrupoles were approximated by digitizing Figure 1 of this design report and are summarized in Table 4.2.

4.1.3 Magnetic regions

In Synrad+, the two quadrupoles are defined as *magnetic regions*. Some parameters, such as the beam properties and region positions can be entered within the GUI, while others through external files. With these information, Synrad+ can calculate the beam trajectory. The interface can be seen (for an other case study) in Fig.4.23.

4.1. Simulation of SR power for the LHeC interaction region quadrupoles

Table 4.2: LHeC beta function values near the simulated quadrupoles

Z Position [cm]	β_x [cm]	β_y [cm]
-230	5000	74000
-180	5500	68000
-130	7000	49000
-100	8500	38000
-50	10000	25000
0	9500	19000

4.1.4 Geometry

In a general case, setting up the geometry consists of exporting the CAD drawing in STL format, then opening it in Synrad+, performing a *collapse* where adjacent, coplanar triangles are merged into polygons, then defining material properties (roughness, reflectivity) for each surface.

In this case study, the process is greatly simplified, since we are not interested in the trajectories and absorption locations of the photons but only in the emitted SR power by Q1 and Q2. For this, we will simply set up two facets at $Z=-110\text{cm}$ and $Z=1\text{cm}$ that will serve as targets to absorb the emitted photons from each quadrupole.

For post-processing purposes, we create a texture on these surfaces that will help us visualize the SR flux and power distribution, and add spectrum counters to make Synrad+ store the SR energy spectrum in memory.

4.1.5 Results (approximate geometry)

In this section I show the results of the case which assumes that the beam passes perfectly at the center of the quadrupoles. This is approximate since as discussed in section 4.1.6, the actual beam has a small angle with the geometrical centerline and thus an offset resulting in much higher SR power. We will, however, solve the simplified case first to compare our results with previous ones.

4.1.5.1 Geometry overview, flux and power distribution

Synrad+ simulation helps us to see an overview of the geometry, the beam and the photon trajectories as shown in Fig.4.2. In the same figure, the textures on the two target facets help us visualize the flux and power distribution. We can see that the photon flux is more distributed vertically, originating from the β_y function being larger than β_x .

4.1.5.2 SR spectrum

The spectrum of the radiation is shown on Fig.4.3, showing that Q2 generates slightly more flux and power than Q1. This is mainly because the beam is larger in the Q2 region, passing through stronger magnetic fields because of its larger offset from the quad center.




Figure 4.2: Synrad+ simulation of the approximate solution and resulting flux and power densities of Q2. Note that the orthogonal projection compresses the longitudinal distance - the beam offset thus becomes enlarged.




Figure 4.3: Flux and power spectra of Q1 and Q2 in the approximate solution

4.1. Simulation of SR power for the LHeC interaction region quadrupoles

Table 4.3: Generated SR power and its convergence (approximate solution)

Generated photons	P_{Q1} [W]	P_{Q2} [W]	Time [s]
10	1862.8	3639.7	<0.01
10^2	5061.2	5008.5	<0.01
10^3	4077.4	5038.4	<0.01
10^4	4120.6	5237.8	0.01
10^5	4262.7	5190.3	0.08
10^6	4262.9	5211.3	0.65
10^7	4258.2	5199.8	5.50

Table 4.4: Comparison of generated power. Synrad+ error is approximated by eq.4.1

Element	Power [W]		
	Analytic[85]	Geant4[85]	Synrad+
Q1	4208.3	4231.8±93	4253.56±53
Q2	5131.9	5173.7±92	5190.3±66
Total	9340.2	9405.5±130	9453.0±120

4.1.5.3 Generated power

The ultimate goal of this section is to calculate the SR power generated by the quadrupoles. Table 4.3 shows the result after different number of generated Monte Carlo photons. For reference, the calculation time on a desktop PC with an Intel i7-3770 3.4GHz processor is shown. It can be concluded that the simulation has converged after about one second, corresponding to 1 million generated photons.

As discussed in section 1.2.1.9 the error of the results scales with $\sqrt{1/N}$ where N is the number of generated virtual photons. Although the exact number depends on the mesh size and the geometry, empirical tests with Synrad+ show that the 95% confidence interval for calculated quantity X is well approximated by

$$\Delta X_{95\%} = \frac{4}{\sqrt{N}} X \quad (4.1)$$

Table 4.4 shows that our results are within the error range of the previous analytic and numerical calculations.

4.1.6 Precise solution

According to the design report, the beam has a 1 mrad angle with the Z axis at the interaction point, resulting in a small transverse offset from the center of the quadrupoles (that way the quads will bend the beam). Because of this, a much larger magnetic field will act on the beam and the SR power will increase drastically. Analytic calculations omitted this offset for simplicity, but Synrad+ allows us to model it accurately.




Figure 4.4: Quadrupoles and beam trajectory of the precise solution

Table 4.5: Generated power (precise solution)

Element	Power [kW]
Q1	148.79 ± 0.9
Q2	64.21 ± 0.4
Total	213.0 ± 1.3

To find the beam trajectory, we take the 1 mrad angle at the interaction point as the fixed parameter, and backtrace the beam through the quadrupoles Q1 and Q2. The results (beam entry/exit points and directions at the quadrupole boundaries) are shown in Fig.4.4.

We use the same targets and visualization textures as in the previous case.

4.1.6.1 Results

Using the mesh set up for the approximate case, we can visualize the SR from Q2 on Fig.4.5.

Figure 4.6 shows that due to the higher magnetic field resulting in a higher critical energy (831 keV on average instead of 221 keV in the approximate case), the energy of the generated photons is higher.

The generated power is shown in Table 4.5. We can see that as a result of the stronger magnetic field, the generated power is an order of magnitude higher than that obtained with the simplified model.

4.1.7 Conclusion

In both cases, setting up and running the simulation was fast and required only a stand-alone program and a text editor. Since our results in the approximate case matched with previous calculations, and it was relatively easy to tweak our model for the more precise case, Synrad+ is a viable alternative to other calculation methods of SR radiation in both simple cases and complex geometries.

4.1. Simulation of SR power for the LHeC interaction region quadrupoles

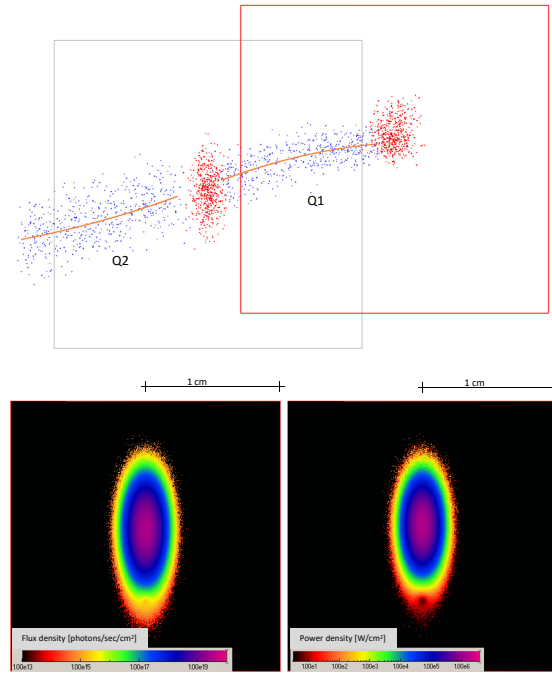


Figure 4.5: Synrad+ simulation and flux and power densities from Q2 (precise solution). Blue dots show the photon generation locations while red dots are where photons get absorbed on the targets.




Figure 4.6: Normalized Flux and Power spectrum of the precise solution

4.2 The MAX IV crotch absorber problem²

4.2.1 Introduction

MAX IV is a new synchrotron light source facility, currently being commissioned in Lund, Sweden. The facility comprises of a 3.4 GeV linac, 1.5 GeV and 3 GeV storage rings and a short pulse facility. The 3 GeV ring (528 m in circumference) is an ultra-low emittance ring (bare lattice emittance of 0.33 nmrad) based on multi-bend achromat (MBA) magnet concept.

Due to the compact magnet design and small magnet aperture of 25 mm, the majority of the vacuum chambers are designed as circular tubes of 22 mm inside diameter and 1 mm wall thickness. Their specific conductance is only 1.3 l m/s for N₂ at room temperature. To reduce pressure, all the conductance limited chambers have been coated with non-evaporable getter (NEG) film of TiZrV. The coating was deposited by magnetron sputtering method that was developed at CERN and is used in other synchrotron light facilities around the world.

Each of the 20 vacuum achromats of the 3 GeV ring is equipped with three sputter ion pumps and one crotch absorber in the place where the electron and photon beams split. The synchrotron radiation fan between the two beams impinges on the crotch absorber. To cope with large outgassing at that position, one of the ion pumps is located directly below the absorber. The upstream keyhole profiled structure and the downstream beam chamber are NEG-coated starting from 6 cm from the ion pump chamber. The wall of that pump chamber is also coated.

In 2013, when the decision to do the NEG coating was approved, I was asked to try the new Monte Carlo codes to estimate the pressure profile near the crotch absorber, and in particular try to estimate the saturation speed of the getter. Working with Marek Grabski from the MAX IV team and Roberto Kersevan from CERN, a full analysis of the region was performed, which I believe is a representative example of coupled SR and UHV simulations, including some hard-to-simulate phenomena, like the NEG saturation requiring iterative simulations.

4.2.2 Synchrotron radiation

The first calculation step is determining the absorbed SR power and flux on the chamber wall, which serves as base to calculate the thermal load and the dynamic outgassing. To achieve this, a CAD model describing the region around the absorber is imported to Synrad+. It consists of 8600 facets (see Fig.4.7), the majority of them describing the grids of the pumping ports next to the absorber.

The dipole magnet (DIPm) 451 mm upstream of the modeled section has a 0.5 T average strength and 1.5° bending angle (see Fig.4.8). In Synrad+, it is described as a magnetic region.

The SR simulations were carried out using Synrad's usual ray tracing algorithm: virtual photons were generated from the pre-calculated trajectory points, traced until their absorption location through several reflections. Surface reflectivity data was interpolated (depending on the incident angle and energy) from the Henke database[50] for copper.

²Based on the IPAC '14 proceeding [87]

4.2. The MAX IV crotch absorber problem




Figure 4.7: Top view of the vacuum chamber with the MAX IV crotch absorber




Figure 4.8: DIPm 1.5° magnet field strength (in red) generating the SR flux in Synrad+




Figure 4.9: The dipole magnet and SR radiation schematics in Synrad+

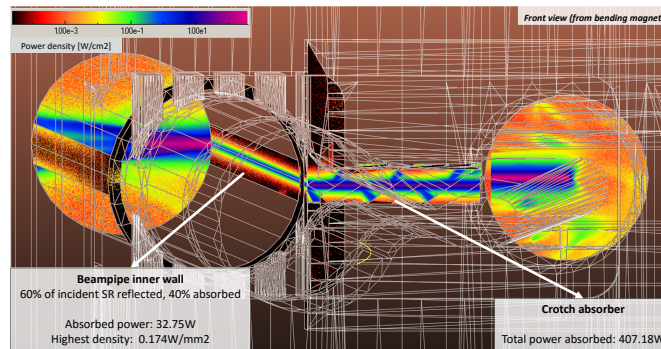


Figure 4.10: Power distribution on the walls




Figure 4.11: Power density on the lower jaw of the absorber

5 million test photons were generated, representing the real flux of 4.2×10^{18} photons/sec and power of 700.7 W generated by the magnet at the nominal current of 500 mA.

The absorber jaws receive 407 W of the SR power, with the highest power density of 6.1 W/mm^2 , a value which is below the thermal limit for copper. A further 33 W is deposited on the outer wall of the downstream beam pipe, distributed on a meter-long distance thus causing negligible heating. The remaining power leaves the modeled region through the light and beam ports.

As for the photon flux, 50% is absorbed by the crotch absorber, approximately 40% leaves the region through the beam and SR extraction ports, and the rest is absorbed at various locations, typically after multiple reflections. The highest SR flux density, corresponding to 6.8×10^{17} photons/sec/cm² is also observed on the absorber's jaws.

4.2.3 Vacuum simulations

Since file formats of Synrad+ and Molflow+ are compatible, the previously saved SR simulation can be opened directly for vacuum simulations. All SR data (magnetic regions, flux, material roughness) are removed so the vacuum chamber's pure geometry remains, thus before starting the simulation the gas desorption and the pumping have to be added.

Four different vacuum simulations are performed – one for CH₄, H₂, CO and CO₂. Taking

advantage of the linear behavior of ultra-high vacuum systems, the partial pressures are summed at the conclusion.

To set up desorption, the flux data from the saved SR simulation is used, along with the four molecule yield conversion files for each gas, taken from literature[66] for copper. In chapter 3, I have described how I recreated Groebner's experiment in Synrad+ to allow a conversion of linear and surface doses, so here we will only use the result: 70 photons/m linear dose equals 1 photon/cm² surface dose simulated by Synrad+.

To convert the incident flux to (integrated) dose, the conditioning time is required for which the approximate value of 200 hours (100 Ah dose) was chosen.

Performing the conversion, we obtain that the total outgassing from the four gases is 5×10^{-7} mbar l/s in the modeled region.

As for the pumping, the sputter ion pump under the absorber has a nominal pumping speed of 75 l/s for H₂ in the 1×10^{-9} mbar regime. (For heavier gases it was adjusted based on catalogue data.)

A sticking factor of 0.1 is assigned to the three ports of the vacuum chamber (one upstream and two downstream). This simulates that at those locations the gas is in equilibrium with the rest of the machine, except for an approximate 10% of the particles that leave the modeled region and are pumped elsewhere.³

The pumping properties of the NEG-coated sections are different for each gas, as described below.


4.2.3.1 Methane

NEG does not pump methane[79], therefore its sticking is disabled in the simulation. This means that simulation for this gas is straightforward, and all the gas from the absorber will be extracted by the sputter ion pump and the unmodeled parts of the system. The pressure distribution is shown in Fig.4.12.

4.2.3.2 Hydrogen

Hydrogen is pumped with a sticking factor of about 0.008 [79], and NEG saturation is not significant.

³Parameter analysis for sticking factors 0.1, 0.3 and 0.5 have been performed, the resulting pressure profiles along the model's centerline (with the absorber in the middle) are shown below. Increasing the 10% value results only in local pressure drops, but due to the low conductance the pressure peaks are barely affected:






Figure 4.12: CH₄ simulation at 100Ah dose

4.2.3.3 CO and CO₂

Estimating pressure profile is more complex for these two gases. According to [79], as shown in Fig.4.13, initially they are pumped by NEG with a sticking factor of approx. 0.7, until they saturate the coating at an approximate molecule dose of 1×10^{15} 1/cm², at which point the sticking is already reduced to 0.002.

Therefore the simulation was performed in an iterative way: first the NEG was assumed unsaturated, but the NEG area was sliced longitudinally allowing individual adjustment of each slice. The conditioning process was modelled in increasing time steps corresponding to an elapsed time of 10, 100 ... 100.000, 720.000 seconds. At each time step, the simulation was stopped, then the number of molecules pumped by every NEG element up to that moment was determined.

If saturation occurred, based on interpolation from data in Fig.4.13 a new, lower sticking factor was assigned locally for the next time slice, and the simulation resumed. In a total of 6 steps the 200 h conditioning time was reached, by which approximately 12 cm of the NEG film closest to the absorber reached saturation condition.

In Fig.4.15 it can be seen that at the border of the saturated region, the first (unsaturated) NEG section pumps the most intensively, therefore reducing pressure in a steep drop.

The final pressure profile is shown in Fig.4.16. It can be seen that since the dynamic outgassing originates almost exclusively from the absorber surface, it introduces a locally elevated pressure in the 1×10^{-9} mbar region, which corresponds to a beam-gas scattering lifetime that is acceptable for the operation of MAX-IV.




Figure 4.13: Figure 1.a from [79]: NEG pumping speed of CO as a function of its surface coverage




Figure 4.14: Pump locations defined for CO and CO₂ for the Molflow simulations.
 Left: End of modeled regions (sticking 0.1) and ion pump
 Right: NEG locations (The central vertical cylinder is also NEG-coated, not shown for visibility)

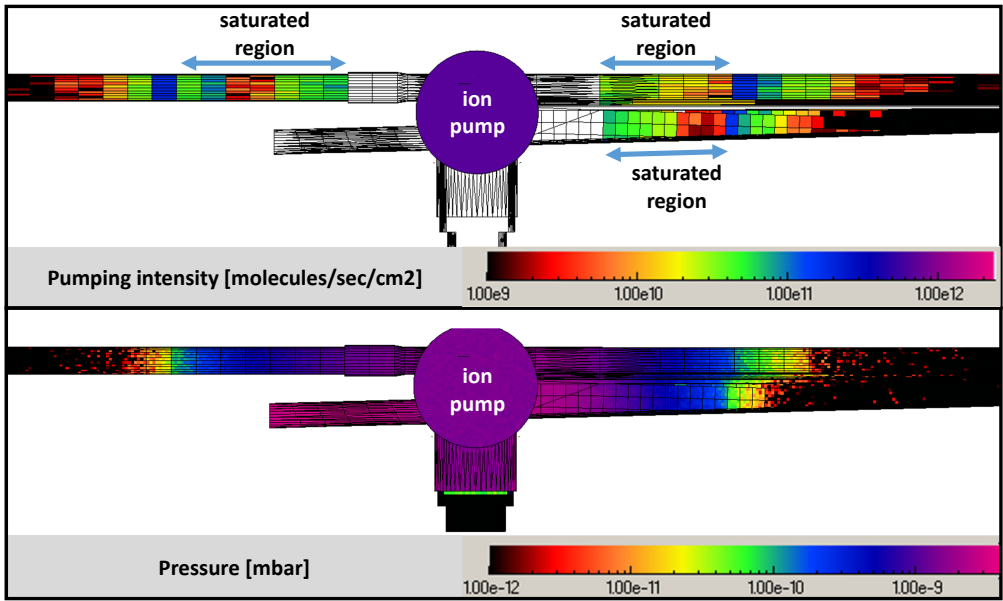


Figure 4.15: NEG saturation and its effect on pumping speed and pressure profile




Figure 4.16: Partial gas pressures around the absorber

4.3 The SuperKEKB interaction region⁴

4.3.1 The Belle experiment⁵

Between 1998 and 2010 the Belle experiment, located at KEK (High Energy Accelerator Research Organization), has analyzed the characteristics of B- and anti-B-mesons pairs and confirmed the effect of CP- violation as predicted by the theory of M. Kobayashi and T. Maskawa, who both were awarded the physics Nobel prize in 2008.

The experiment was installed on KEKB, a 3 km circumference electron-positron collider. An upgrade project, SuperKEKB[89] is currently taking place, aiming to increase the instantaneous luminosity from $2.11 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ to $8 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$. The upgrade requires changing several parts, including the detector itself, the DAQ and trigger system, and – where we are concerned – the vacuum system.

The electron and positron rings of the KEK cross each other at the interaction point (IP). Although a low pressure in that region is critical for the operation of the machine, its calculation is a difficult task: it requires knowledge of the exact SR flux, contributing to the dynamic outgassing of several materials, such as copper, aluminum and tantalum. Vacuum-wise, several types of pumps are deployed with varying efficiency across gas types, moreover because of the narrow design of the new Belle 2 detector, they are placed more than 3 m downstream so their pumping performance is conductance-limited, requiring additional calculations.

The Belle 2 Design Report[90] estimates the IP pressure as $6 \times 10^{-5} \text{ Pa}$, by assuming a general photodesorption rate of 1×10^{-5} molecules/photon for the scrubbed machine. As this is a rough calculation, an exact simulation was considered important before the actual commissioning of the machine. During my stay at KEK for the PSD experiments presented in chapter 3, I was supplied with blueprints and optics parameters of the interaction region by Prof. Ken-Ichi Kanazawa, and asked to utilize my codes to perform a full simulation. The process, described below, goes through the following milestones:

1. **Vacuum chamber modeling:** based on the drawings a 3D model, suitable for Synrad+ and MolFlow+ is created
2. **Optics definition:** based on beam surveys the 8 bending magnets and 10 quadrupoles near the IP are defined in Synrad, and the trajectory is calculated
3. **Surface definition:** the materials, with their reflectivity and roughness properties are defined

⁴Special chapter acknowledgements

This case study has been carried out as a cooperation with the Argonne National Laboratory: **Jason Carter**, a young engineer from the Advanced Photon Source (APS) has been visiting CERN's vacuum group for three months in 2015. As this case study requires coupled simulation of SR and UHV, the idea was that the methods we apply here can be reused for his vacuum calculations of the APS upgrade (APS-U) project[88]. In practice, he has put to use his excellent 3D modelling skills, acquired material data for the PSD conversion and did the simulation runs himself, while I was developing the codes in parallel to suit our needs and supplied the vacuum and optics properties of the modelled region. As the work described in this chapter – apart from an oral presentation to the KEK team – wasn't published elsewhere, I'd like to attribute credits to his contribution, without which this case study couldn't have been completed.

⁵Introduction from Desy website: <http://belle2.desy.de/e103206/>




Figure 4.17: CAD drawing of the IR left side showing the magnet positions
From Ken-Ichi Kanazawa (KEK)




Figure 4.18: Technical drawings of individual components serving as base of our 3D model
From Ken-Ichi Kanazawa (KEK)

4. **SR simulations:** photons are generated and traced until their absorption location
5. **PSD conversion** (for each gas type): results from the previous point are imported to Molflow, converting absorbed flux density to a dynamic outgassing map
6. **Pump definition** (for each gas type): pumping locations and speeds are input
7. **Vacuum simulations** (for each gas type): regular UHV simulations and post-processing of results

4.3.2 Vacuum chamber modeling

We have received AutoCAD drawings (Fig.4.17) of the region ranging longitudinally from -14 m to +14 m around the interaction point. Its extremities, delimited by gate valves, are sufficiently far from the IP to be considered to be in equilibrium with the rest of the ring, so simulating this reduced part is a good approximation.

Based on these assembly drawings and the technical details for individual elements (such as in Fig.4.18), a 3D model of the vacuum chamber has been created in Solidworks. As for almost every Molflow+ simulation, this model is *inside out*, meaning that based on the actual walls, the vacuum volume is modeled as a solid body.⁶

The geometry, imported from Solidworks and collapsed in Molflow+ (Fig.4.19) consists of

⁶As my codes look for collisions with surfaces, it is unnecessary to define each wall with two boundaries, instead we only need the boundary of the vacuum region

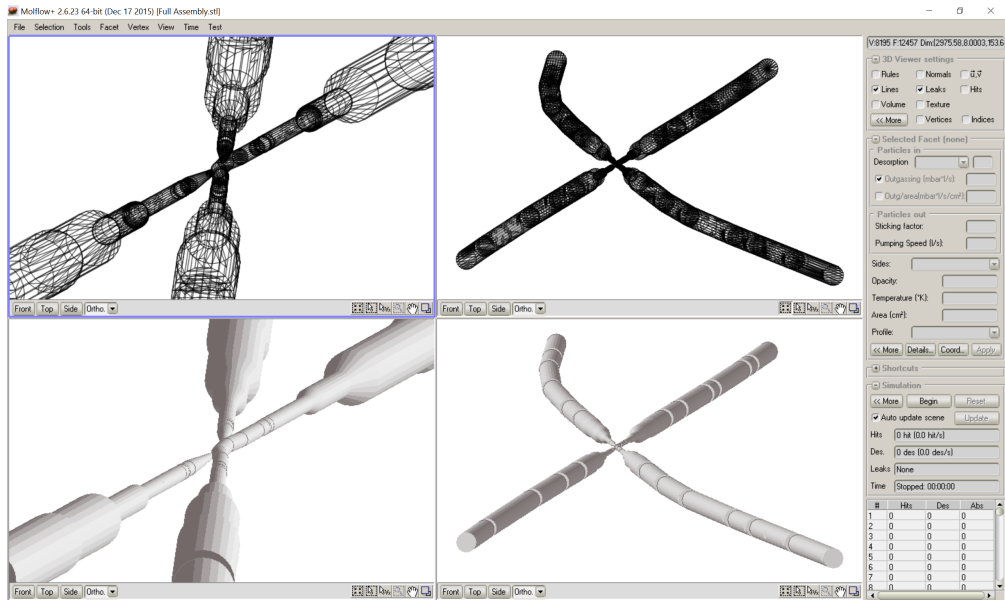


Figure 4.19: Molflow+ model of the interaction point (left) and the full modeled region (right)




Figure 4.20: KEK magnet naming convention

12.500 facets. This relatively low number is achieved through simplifications of the pumping ports, as presented later.

4.3.3 The interaction region optics

Within the modeled, approximately 28 m long region, we have included 8 dipoles and 10 quadrupoles, each having a 5 to 7 character identifier. For better understanding, the naming convention is explained in Fig.4.20.

Each optical element will be defined in Synrad+ as a magnetic region. Having both the vacuum chamber model and these regions, very careful alignment of the two is necessary.

Our starting point was four beam survey files⁷ from KEK: they contained the location-dependent beam parameters (beta and dispersion functions) and the magnetic elements.

All entries in the beam survey were parametrized in absolute beam length coordinates. Each

⁷Received from Prof. Yukiyoishi Ohnishi and Prof. Katsonobu Oide, KEK




Figure 4.21: Step function approximation of quadrupoles with edge gradients

entry had one value:

- Drifts: the drift length
- Bending magnets: the bending angle (proportional to $K_0 L$)
- Quadrupoles: integrated gradient ($K_1 L$)

In Synrad+, regions have to be defined in Cartesian coordinates, and by the magnetic field strength instead of the bending angle. To perform the conversion, we have considered the interaction point as the origin of our coordinate system. Based on machine parameters (and the blueprints), we already knew that the positron and electron lines cross each other with an angle of 83 mrad. We have chosen the starting directions as 50 mrad (relative to the Z axis) for the electron and 33 mrad for the positron beams, and for each element, using trigonometric functions based on the element length and direction, calculated the Cartesian coordinates. The direction was updated when we reached a bending magnet.

One issue was to treat the quadrupoles. They were represented in the survey files as a succession of 1 cm long elements, each having their own $K_1 L$ value. Plotting the values revealed that this solution was chosen to represent the varying fields at the edges. As Synrad+ doesn't support gradient quadrupoles yet, we have approximated these with constant quads, so that the integrated $K_1 L$ for the whole element is identical to the one in the survey file (see Fig.4.21).

The β and η (dispersion) functions (Fig.4.22) were saved in a special .BXY file format (one for the electron and one for the positron line). Synrad+ can load these files when defining the regions, and interpolate the local value when a photon is generated. As illustrated in Fig.4.23, this external file contains the β_x , β_y , η_x , η_y and $\Delta E/E$ (energy spread) functions for a list of locations referred to by beam coordinates.

We have made one more assumption. In reality, the interaction region is not entirely horizontal. Two bending magnets (*BC1LP* and *BC1RP*) and two quadrupoles with an offset (*QC2LP* and *QC2RP*) bend the beam vertically. To avoid introducing a third, vertical dimension in our simulations (which would have made the 3D modelling and most coordinate calculations more difficult), we have chosen to model the beam trajectory in the horizontal plane⁸. To

⁸Of course, our vacuum chamber model and simulation still had three dimensions, but the beam and the chamber centerline was confined to the horizontal plane




Figure 4.22: Beta functions of SuperKEKB in the modeled region

Table 4.6: SuperKEKB beam parameters

	Energy	Current	Emittance	H/V coupling	Energy spread
Electron beam	7 GeV	2620 mA	324 nm	35%	0.0649%
Positron beam	4 GeV	3600 mA	448 nm	40%	0.0814%

have the equivalent SR fans on the chamber wall, we have placed these magnets so that their starting point is bent, and their last trajectory point has a horizontal tangent, as shown in Fig.4.24.

The actual definition of the regions happened in Synrad's graphical interface, requiring the starting point and direction, and the constant (see table 4.6) and location-dependent beam parameters, and the boundary of the region. Trajectory points are then calculated automatically, allowing verification of our calculations by comparing the end point coordinates with the next element's starting location in the beam orbit files.

The 18 modeled regions had a total of 166.000 pre-calculated trajectory points. They are shown, placed in the vacuum chamber, in Fig.4.25.

4.3.4 Synchrotron radiation simulation

As we are performing a coupled SR/vacuum simulation, where low energy photons are just as important (they contribute to the dynamic gas load), we have chosen fluxwise photon generation (as explained in section 2.3.1.1). The approximate generation speed was 12.500 photons/second. The process, assuming immediate absorption, is illustrated in Fig.4.26 for a single dipole, where it is visible that the beam has a certain width (the source points are distributed around the ideal orbit).

The total flux generated by the 18 elements is 1.5×10^{20} photons/sec, with a power of 12 400 W. It is worth noting that (as shown in Table 4.7) practically all flux and power originates from the bending magnets (including the two bending quadrupoles): the 8 straight quadrupoles together generate less than 0.1% of the total flux and power.

Chapter 4. Case studies

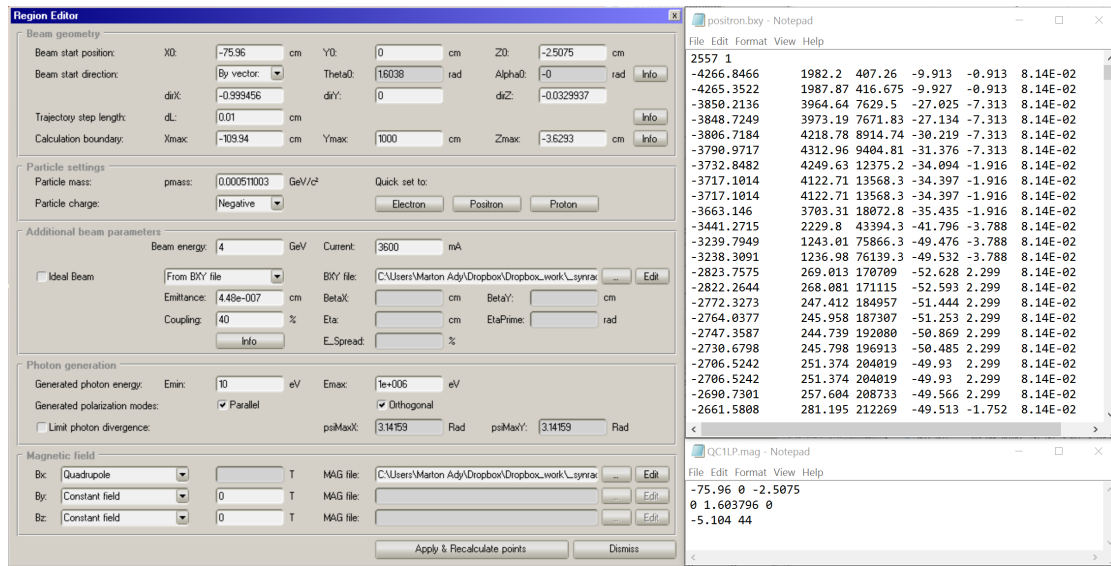


Figure 4.23: Region definition in Synrad for the QC1LP quadrupole of the positron line (left). The beta functions are referenced through an external BXY file (top right), and the quadrupole parameters (starting point, centerline direction, strength and length) through a MAG file (bottom right)




Figure 4.24: Horizontal approximation of vertical bends for the *BC1LP*, *BC1RP* dipoles and *QC2LP* and *QC2RP* bending quadrupoles

4.3. The SuperKEKB interaction region




Figure 4.25: Interaction region optics, as modeled in Synrad. Pre-calculated trajectory points in orange for each magnetic region.




Figure 4.26: Photon generation in Synrad+ from the BC1LP1 dipole, with generated source points in blue, color-coded photon lines representing the carried power in green, and absorption points in red

Table 4.7: Flux and power of the SuperKEKB interaction region optics

Element	Flux [ph/s]	Power [W]	Critical energy [eV]
BLC1RP	4.2E+19	4460	1764
BLC1LP1	3.8E+19	4960	2193
BLCWRP	2.1E+19	1280	946
BLC1LP2	1.3E+19	543	616
BLC2RE	1.2E+19	1380	2012
BLC1LE	1.1E+19	856	1247
QC2RP	3.2E+18	175	
QC2LP	2.8E+18	133	
BC1LP	8.4E+17	22	351
BC1RP	3.8E+17	5.8	181
QC1LP	1.7E+17	2.1	
QC1RP	1.6E+17	2.4	
QC1LE	1.3E+17	2.8	
QC2LE	4.8E+16	4.5E-01	
QC2RE	4.6E+16	4.7E-01	
QLC3LE	4.3E+14	1.2E-03	
QLC1RP	1.5E+13	3.2E-05	
QLC2LE	2.2E+12	4.4E-06	

Once the optics is set up, we have to define the reflective properties of surfaces. Different parts of the chamber are made of different materials: aluminum, copper, copper-coated tantalum and – at the direct vicinity of the interaction point - tantalum. We have the reflection probability, dependent on the incident angle and the photon energy, as input files for these materials, so we can apply them directly to the corresponding surfaces (as marked for aluminum and copper in Fig.4.38).

As for the surface roughness, we did not have metrology data available, so we have approximated surface roughness with parameters of a typical technical vacuum chamber surface of RMS roughness $\sigma=400$ nm and autocorrelation length $T=10000$ nm. We will see in the conclusion that the pressure profiles aren't very sensitive to a moderate variation of the roughness value.

This RMS roughness value of 400 nm was chosen – as shown in Fig.4.27. so that it sufficiently smooths the flux density map to remove artifacts that originate from the polygonal approximation of originally curved surfaces.

A typical overnight simulation run with scattering and low flux mode enabled has generated 450 million virtual photons, each bouncing 9 times on average before absorption (4 billion total hits).

The resulting flux density map for the whole modeled region (Fig.4.28) was later converted to dynamic outgassing.

It is useful to analyze the effect of photon scattering on the flux and power map. As shown in

4.3. The SuperKEKB interaction region




Figure 4.27: Roughness value choice: with $\sigma=10$ nm facet boundaries and planar surfaces create artifacts on the flux map. At a higher value of $\sigma=400$ nm these effects disappear through the smoothing effect of rough (diffuse) reflection.




Figure 4.28: Flux density for the SuperKEKB interaction region assuming $\sigma=400$ nm surface roughness

Chapter 4. Case studies

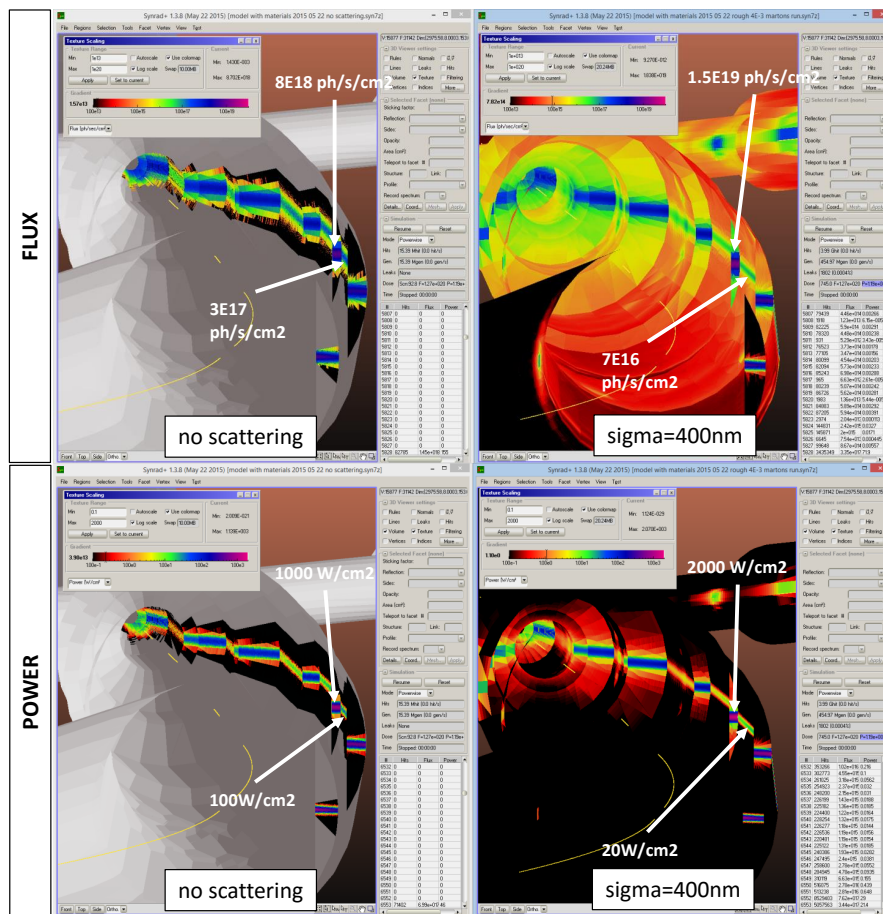


Figure 4.29: Effect of photon scattering (right) versus the instant absorption assumption (left) on the peak flux (top) and power (bottom) peaks

Fig.4.29, the flux and power peaks can be observed on the absorber masks, followed by the direct incidence locations.

Enabling material-interpolated reflectivity and rough surface scattering changes the ratio of the peaks compared to the immediate absorption case: the direct incidence locations will receive lower flux and power, whereas the absorber masks will receive higher.

The former phenomenon can be explained by the high reflectivity of copper and aluminum for low grazing angles: instead of total absorption, a significant ratio of the incident SR flux is reflected, first to the opposite side of the chamber, then to other locations through rough surface scattering.

The absorbers in turn receive approximately double flux and power density: this is due to reflected photons travelling downstream either almost parallel to the chamber wall, or reflected from the opposite side, until being collected by the first wall protrusion – which happen to be the absorbers in our case.

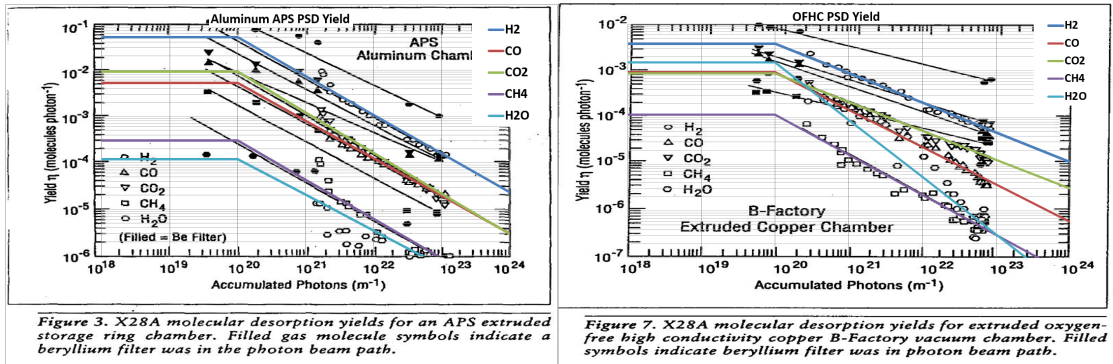


Figure 4.30: Original figures from[80] showing the molecule yields for Aluminum (left) and Copper (right), with overlay lines showing the digitized values assuming a cutoff at 1×10^{20} ph/m

4.3.5 PSD conversion

To move from SR simulations to vacuum calculations, we have to import Synrad's flux density results to Molflow+ while performing a conversion to dynamic outgassing. For each texture cell, based on the machine operating time given by the user, Molflow+ calculates the dose, then interpolates a molecular yield, which – multiplied by the flux density – will give us an outgassing density.

The most crucial parameter for this conversion is the molecular yield vs. beam dose graph.

In our case, we have copper, aluminum and tantalum materials. For the first two, PSD measurements have been performed already thus data is available[80] from the literature.

We will simulate four gases: H_2 , CO, CO_2 and CH_4 . The molecule yield for each of them can be digitized, as shown in Fig.4.30. Post-processing requires converting the X scale from linear dose [ph/m] to surface dose [ph/cm²] used in Synrad. As discussed in chapter 3, the conversion factor - determined by recreating the original PSD experiment – tells us that 1 ph/cm² surface dose equals 70 ph/m linear dose.

Although the data from [80] contains the yield for each individual gas type, we will not use it directly. In fact, in-situ PSD measurements are available from the electron (HER) and positron (LER) rings. Data supplied by prof. Kanazawa from KEK, plotted in Fig.4.31 shows similar initial yields as the one in literature, but a faster conditioning. Since PSD measurements collected directly in the machine to be simulated is certainly the most precise data available, we will use it.

However, the received KEK data contains only the total desorption yield, and not the gas composition. What we do is that we use the data from [80] to estimate the gas composition: we will assume that the ratio of the desorbing gases is the same in the KEK ring as in Foerster's experiments, therefore we will scale the individual molecule yields with the ratio of the literature data and the KEK data. That way we will have at least an approximate yield for each gas and we can be certain that the total yield matches actual measured data in the machine.

Finally, the tantalum chamber is coated with 10 μm of copper. Given the average energy of the

Chapter 4. Case studies

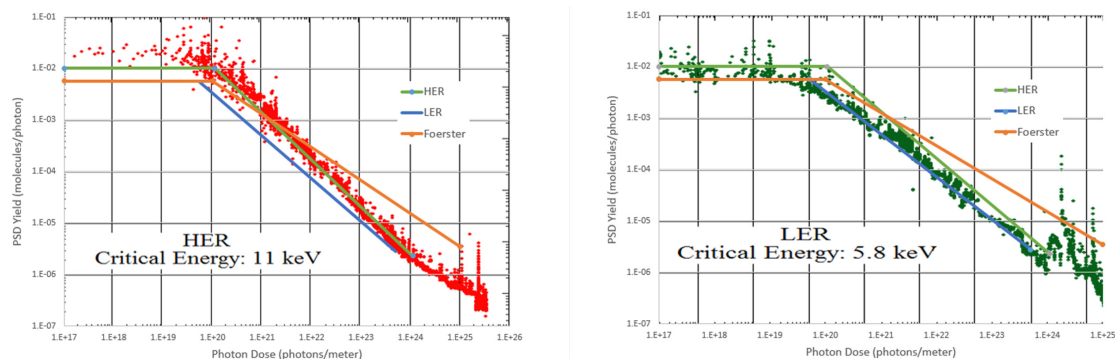


Figure 4.31: Total copper PSD yield measurements for the HER and LER rings, and comparison with the data from Foerster's experiment




Figure 4.32: PSD outgassing density as a function of incoming SR flux density. The cutoff point depends on the machine operation time.

photons and the grazing angle of incidence, most photons would not penetrate the tantalum layer - we will hence estimate that it has a copper surface's PSD yield.

Having the conversion plots, we can demonstrate the importance of Synrad's low flux mode, which facilitates calculating the flux density on areas that receive only very few photons. In fact, because of the steep conditioning curves, those areas that receive high flux will in turn condition faster, decreasing the local molecule yield. As emphasized previously in chapter 3, in an extreme case, with a conditioning slope of -1 on a log-log PSD plot, a ten times higher flux density would decrease the local molecule yield by a factor of ten. Since the outgassing is the product of the flux and the molecule yield, that high-flux area would outgas with the same rate as the low-flux areas. (In practice, conditioning slopes are usually below -0.75). The point is that a seemingly extreme difference in flux density is "compressed" to a much smaller difference in outgassing. In the example in Fig.4.32, having executed the multiplication of yield and flux, we plot the outgassing density as a function of the incident photon flux density. We can see that 7 orders of magnitude flux density difference converts to less than 3 orders of magnitude outgassing difference.




Figure 4.33: Vacuum system overview on the left side of the interaction region




Figure 4.34: Pumping performances for different gas types of a CapaciTorr D1000 (left) and a NEX Torr D1000-10 pump (right)

Source: SAES getters

Therefore, to approximate the outgassing map with 0.1% accuracy for our vacuum simulations, Synrad+ has to run with the low flux mode, with a cutoff ratio of at least $1\text{E-}7$.

4.3.6 The vacuum system

The interaction region's vacuum system (overview in Fig.4.33) uses four kinds of pumps:

- NEG strips in antechambers
- CapaciTorr D1000 NEG cartridge pumps
- Nextorr D1000-10 NEG cartridge pumps
- IP400 ion pump

The pumping speeds are readily available from the manufacturer's catalogues (Fig.4.34), and depend on the type of the gas and the pumped gas quantity. In our simulations we have assumed that during normal machine operation, NEG is kept unsaturated, taking a pumping speed value from the flat part of the curves. In fact, the simulations aren't very sensitive to these theoretical pumping speeds, as they are significantly reduced by the conductance limitation of the pumping ports.

The cartridge and the ion pumps have four different pumping ports. All types consist of an antechamber where the pump is placed, separated from the beam by a grid, which minimizes




Figure 4.35: Pump locations in our Molflow+ model (left) and the pumping port they represent (right)




Figure 4.36: Vacuum simulation to determine the pumping speed reduction of the pumping port grids (top) and the resulting pressure profile along the port's centerline (bottom)

impedance loss but also reduces the effective pumping speed. To determine the reduced, effective speeds, we have created 3D models of each type. The locations and a representative model is shown in Fig.4.35.

To calculate the effective pumping speed on the opposite side of the grid, we do a vacuum simulation where we introduce a gas load Q in the modeled section of the main vacuum chamber, and pump it at a facet representing the orifice of the pump, with the pumping speed from the manufacturer. We sample the pressure along the centerline of the port with a transparent facet containing a profile.

According to vacuum theory, at every point in the system,

$$p_{\text{local}} = \frac{Q}{S_{\text{local}}}$$

where S_{local} is the effective pumping speed of the pump across the grid. In Fig.4.36. we can see

Table 4.8: Effective pumping speeds (Port type A and B)

Gas	NexTorr/CapaciTorr S_{pump} [l/s]	Port Type A S_{eff} [l/s]	Port Type B S_{eff} [l/s]
H ₂	1000	400	370
CO	500	128	119
CO ₂	500	110	98
CH ₄	20	19	18

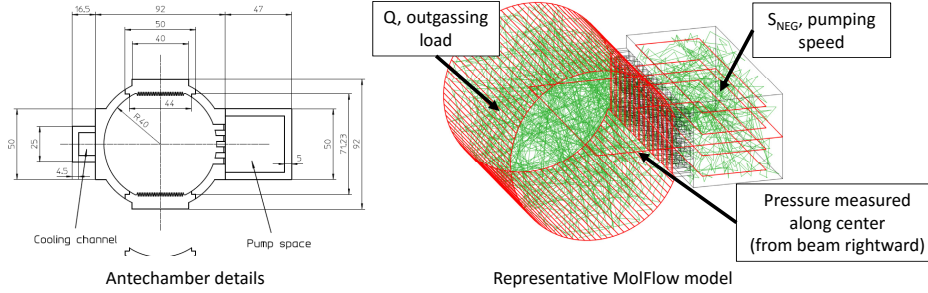


Figure 4.37: NEG pumping ports (left) and their Moflow+ simulation to determine the effective pumping speed (right)

that the pressure is the lowest at the pump's location, where S_{local} is the highest (and equal to the theoretical pumping speed). In the beam chamber, however, we see an increased pressure and thus a decreased effective pumping speed.

Repeating the simulation for all port types and all gases, we get the effective pumping speeds in table 4.8.

As for the NEG strips, the process is similar. Four facets representing the NEG strips are placed in the antechamber, and a transparent facet is measuring the pressure profile in the center. As shown in table 4.9, the pumping speed reduction due to the grid is very significant for CO and CO₂ where the sticking factor is high. On the other hand, this shows that the simulation isn't very sensitive to our assumption that the NEG surfaces are unsaturated.

The main purpose of simulating pumping ports individually is that once we determine their effective pumping speed in the main beam chamber, we can replace their complicated geometry with a limited number of facets – having the equivalent pumping speed. As shown in red in

Table 4.9: Effective pumping speeds (NEG ports)

Gas	Sticking factor	4x10 cm length S_{pump} [l/s]	4x10 cm length S_{eff} [l/s]
H ₂	0.008	85	60
CO	0.7	1977	54
CO ₂	0.7	1577	43
CH ₄	0	-	-




Figure 4.38: Material definitions in our Molflow+ model

Fig.4.35, these pumps are represented as facets of the chamber wall.

4.3.7 Vacuum simulations

At this point we have defined the pumps (represented by wall facets) and the gas source (dynamical outgassing converted from Synrad's flux maps). As UHV systems are linear, we will do four individual simulations for the four gas types, and sum the results at the end.

Each simulation begins with importing the same flux map – calculated for the $\sigma=400$ nm case – from Synrad. Then we convert it to dynamical outgassing with the conversion plot corresponding the given gas and material type loaded from an external file. We then apply the pumping speeds for the given gas on all facets corresponding to pumping locations. We assume that the boundaries of the simulated region are in equilibrium with the rest of the ring, so we represent them by simple circular facets without sticking.

During a typical overnight simulation, we can simulate 200 million hits, corresponding to around 200.000 virtual particles traced from desorption to the pumping location.

Results are visualized as textures on the wall, and sampled through transparent facets running in the centerline of the beam pipe along the whole system. Each of these facets will generate a longitudinal pressure profile, which we will concatenate outside of Molflow – obtaining two continuous pressure curves (one for the electron and one for the positron lines) meeting in the center.

The series of four simulations (one for each gas) were repeated twice: one for 100 Ah and one for 1000 Ah dose to see the effect of conditioning. As shown in Figures 4.40 and 4.41, the ten times longer conditioning results in a pressure drop of approximately half an order of magnitude.

4.3. The SuperKEKB interaction region




Figure 4.39: Pressure profiles around the interaction point for CO₂ (left) and H₂ (right) gases at 1000 Ah dose

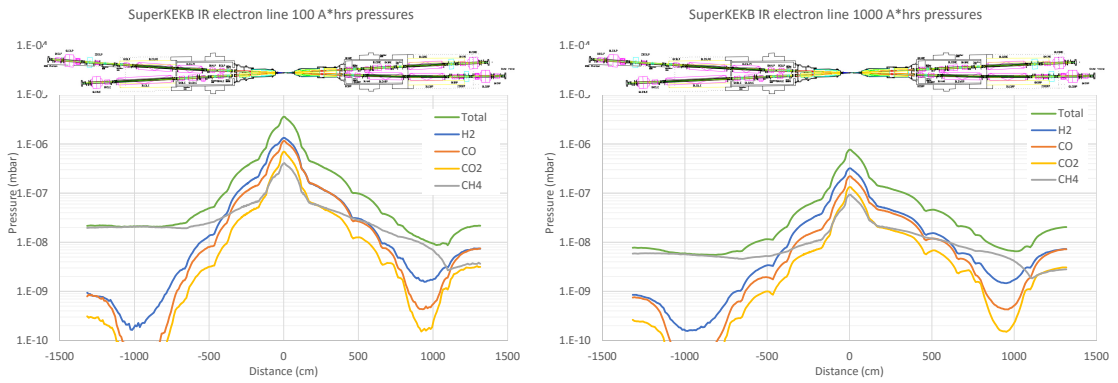


Figure 4.40: Total and partial pressures of the SuperKEKB interaction region (electron line) at 100 Ah (left) and 1000 Ah (right) machine times




Figure 4.41: SuperKEKB interaction region total pressure summary at different conditioning times

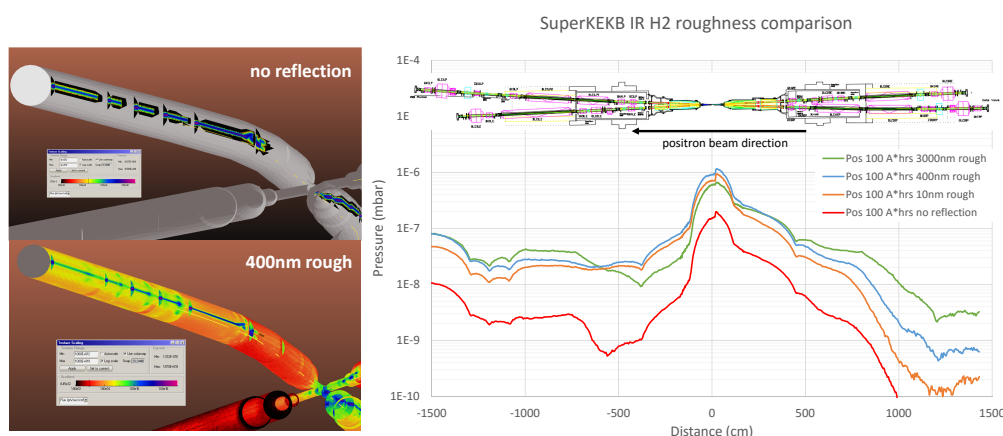


Figure 4.42: H₂ pressure profiles near the interaction point at 100 Ah conditioning comparing immediate absorption with three different rough surface scattering cases

4.3.8 Effects of roughness

Before discussing the pressure results, we need to mention a certain point about their validity: previously we have made an assumption for the surface roughness which is not backed by actual metrology results from the machine. Therefore, it is important to discuss the sensitivity of our results to that assumption. We have run four Synrad+ simulations: one assuming immediate absorption, and three other simulations with material reflection, with increasing roughness values (RMS roughness of 10 nm, 400 nm and 3000 nm).

To speed up the analysis, we did the conversion to dynamic outgassing and the subsequent vacuum simulation only for H₂. The pressure profiles in the positron line for 100 Ah conditioning are plotted in Fig.4.42.

The pressure is the lowest in the immediate absorption case. This confirms what I have emphasized in the PSD conversion subsection: a very high flux density will in turn cause very fast local conditioning, therefore the outgassing (product of the molecule yield and the flux) won't be extremely large. Since in this scenario only the direct incidence locations act as gas sources, the total outgassing – and the resulting pressure – will be low.

As soon as we turn on material reflection and rough surface scattering, photons reach much higher area of the chamber wall. If we distribute the total photon flux (which is fixed across the cases) to a larger surface, the resulting conditioning will be slower (higher molecule yield compared to the immediate absorption case), whereas the total flux remains the same. The resulting outgassing (product of molecule yield and flux) will therefore increase.

That explains the trend of pressure increasing with roughness in Fig.4.42, and without actual metrology data of the SuperKEKB chambers, we have chosen to calculate with the middle case of the 400 nm roughness. This certainly affects our results, but fortunately rather at the low pressure areas (where direct incidence is negligible) – in high-pressure locations, and at the peak located at the IR point, roughness values vary the pressure within about one-fourth order of magnitude, which is an acceptable incertitude.

4.3.9 Discussion of results

To put in context our final results in Fig.4.41, I quote from the Belle 2 Design Report[90]:

"The target vacuum pressure in the ring is on the order of 1×10^{-7} Pa on average" (chapter 2.6.4) and "Assuming a sufficiently scrubbed value of the photo-desorption coefficient of 1×10^{-5} , the pressure at the IP and at the pump is estimated as 6×10^{-5} and 6×10^{-6} Pa, respectively, for the design current" (chapter 3, p64)

Given that 1 mbar = 100 Pa, our calculated pressure peak in the 1000 Ah case of 8×10^{-7} mbar is very close to the design value, and we also see that the pressure at the first pump is approximately one order of magnitude lower. Consequently, our coupled simulations seem to confirm that the KEK team's "rough" estimations were in fact very precise.

Summary and future plans

Summary

We have seen the steps of implementing a vacuum simulator, which - based on the TPMC method - can simulate steady-state and time-dependent vacuum systems. We have also seen how to generate and trace test photons for synchrotron radiation simulations - covering the difficulties in treating complex analytic distributions with a sufficient speed and finding among the various rough surface scattering models one suitable for our needs. The fields of SR and UHV were connected through photon stimulated desorption: a synchrotron simulation determines the flux map in the vacuum system, which in turn can be converted to dynamic outgassing, allowing us to proceed with a vacuum simulation. Upcoming machines require precise PSD data, therefore samples were manufactured at CERN and tested at KEK with conditions similar to the FCC-hh. The validity of the codes was benchmarked against textbook cases and compared to experimental data. The programs were used to solve three accelerator problems, two requiring coupled SR and vacuum simulations.

Future plans

Molflow+ and Synrad+ are available to the public as a free download. The website⁹ has a forum where apart from signaling bugs and requesting help, new features can be suggested. Implementation of many features (among others: radioactive decay, wall sojourn time, ANSYS export, moving parts) were driven by user feedback, and in some cases even changes in the physics engine were due to suggestions of fellow users. Based on the four years I've spent developing the software, I learned that changing the physics engine, or implementing a certain physical effect is the smaller part of the work. Subsequently changing the user interface, updating the file formats, uploading the new version, and at the same time ensuring compatibility with previous versions, then synchronizing relevant changes between Molflow and Synrad, extending the documentation and eliminating bugs that manifest after the modifications is a much longer undertaking. It is by no wonder that most "competing" codes used by the physics community are command-line based, and input/output happens through text files rather than through an interactive graphical interface. Nevertheless, I believe that the added value of our programs are the straightforward, 3-dimensional, easy-to-learn interface, and the success of this approach is confirmed by the fact that according to website analytics, Molflow+ is downloaded all over the world, and used - apart from the accelerator applications - from

⁹<http://cern.ch/test-molflow>

the semiconductor industry to fusion physics. Having said that, I see two ways of further developing the codes.

The popular approach

If we aim to keep the wide user base, the development needs to go in a direction that would ensure that Molflow+ is as compatible as possible with CAD programs, and works out of the box. This would mean implementing an advanced geometry editor, capable of building (rather than importing and modifying) geometries. Part of the project is already done: Molflow+ 2.6 can move, edit, rotate, scale, align, etc. facets. These functions would be extended with smart selections (automatically recognizing surface features), splitting and merging planes, reading popular CAD file formats and having similar zoom/rotate controls as popular 3D modeling software.

This approach would please many engineers within CERN's vacuum group and outside, at the same time its scientific value is questionable. Therefore the other way to go is to concentrate on the physics:

The scientific approach

In this case development would focus on the basic algorithm behind the interface, and features would be implemented quicker, but at the expense of user-friendliness: input would be executed through parameter files, and some versions would stay internal. Currently interesting features that we'd include are:

- **Iterative simulations** In section 4.2.3.3, I have run into a problem where the pumping speed depends on the pressure, but in turn the pressure depends on the pumping speed. Working with NEG surfaces, we face this problem frequently, and the solution is to solve the problem by advancing the system state in small time steps (during which sticking factors are assumed constant), then update the surface properties based on the calculated quantity of pumped molecules. This can be done manually, as I did in case of the Max IV crotch absorber analysis, but then we have to use side calculations and the process is somewhat cumbersome. Automatizing it would have the immediate benefit of being able to include NEG saturation in our simulations. In theory, implementing it in Molflow+ is straightforward: by digitizing a saturation plot similar to 4.13, surface properties could be recalculated at user-defined time steps based on the number of pumped molecules, then simulation could advance to the next step. Difficulty is in the details, such as choosing the correct step length, when to consider that simulation has converged for each step, and once again, most of the work would come from updating to user interface to merge already existing parameters: so far we include a "short" time-dependent mode, simulating all molecule hits thus suitable up to a few seconds, this would need to be reconciled with a "long" time-dependent mode that models saturation processes taking hours or days.
- **GPU computing** GPUs are very efficient in applying the same operation on a very large set of data. In principle, the ray tracing of Molflow+ is a problem that is suitable to

solve on the GPU: we have a large number of test particles that all behave the same way. In 2007 J.L.Pons at ESRF has created a test version of Molflow that ran on the GPU. His conclusion was that due to the frequent individual changes of test particles (thermalization, teleport, etc.) the advantage of the GPU is negligible: if we have 3 or more cores, Molflow runs faster on the CPU. As computer hardware advances at an astonishing speed, with 9 years passed since those tests, the situation should be reexamined: a GPU implementation, preferably through a cross-platform language such as OpenCL, should be attempted - it is possible that a significant speedup could be achieved by a good code.

Currently development of the codes and maintenance of the website is a one-man project. At the moment of writing, although I know that I'll have the chance to continue my work in CERN's vacuum group for two more years, my future tasks are not defined yet. The amount of work I can still invest in the development of Molflow+ will therefore depend on my CERN tasks, and deciding which of the two development directions we'll go is a question of the coming months.

A High-pressure experiments

A.1 Introduction

As the whole thesis is devoted to high vacuum conditions, in this appendix chapter I discuss a special case of higher pressures – the shock tube problem. A shock tube in this sense is a volume under high vacuum that is suddenly exposed to a higher – in our case atmospheric – pressure, which enters the tube with a very high speed. Until the whole volume is equalized, multiple pressure regimes are present in the tube.

In the atmospheric regime UHV theory doesn't hold because gas dynamics isn't governed by wall-particle collisions, but by the much more frequent intermolecular collisions. The gas dynamics can be macroscopically described as compressible flow, classed into different regimes (laminar and turbulent, based on the Reynolds number).

Several assumptions of the UHV regime aren't valid to compressible flows. For example, the conductance between two points will depend not only on the geometry, but also on the pressure. Because of the interaction between gas molecules, the gas can't be simulated as a set of independently moving particles. Usually finite element solvers handle these high-pressure systems by dividing the gas volume to a large number of mesh elements, and model interactions between the cells based on fluid dynamics equations.

The most complicated case is the mixed flow, where the system to be simulated has multiple pressure regimes present at the same time. Although some early trials – mainly in aerospace engineering [91, 92] - treat the problem through hybrid solvers, even private correspondence with some authors confirmed that these codes are extremely intensive computationally and can treat only very specific problems (for example: space vehicle reentry) with major simplifications.

Some commercial simulation software, like Comsol can handle mixed viscous/free molecular flow, but with the condition that the user defines the regime boundaries manually which stay fixed during the simulation.

On the other hand, the problem of sudden opening of a membrane separating two regions of high and low pressure is well known in continuous fluid dynamics: it is the vacuum shock tube problem[95]. This problem is solved analytically starting from energy and momentum

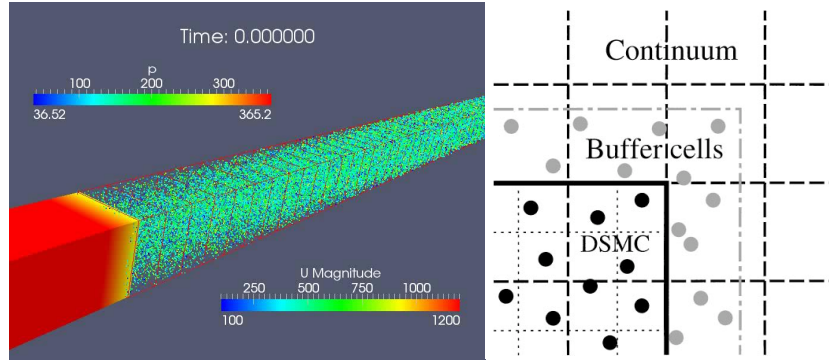


Figure A.1: Visualization of a hybrid fluid dynamics – DSMC solver[93] computing the boundary between two different pressure regimes.
Images from [94]

conservation. The most startling feature in this treatment is the shock wave, travelling at constant speed

$$u_{\max} = 2/(\gamma - 1)c_0$$

into vacuum, but carrying vanishing density. Here, c_0 is the local speed of sound, while γ is the specific heat ratio. For air, $\gamma=1.4$ and $c_0=331$ m/s, so the shock wave speed is 1655 m/s, or Mach=5. The particular case of constricted flow, i.e., where an orifice of diameter d separates a cylindrical tube of diameter D at initial pressure p_1 from a region at (higher) pressure p_4 , is developed in [96], in relation to the sudden opening of a leak at the extremity of a vacuum beam line. Similar to the shock tube problem, the propagation of a pressure wave inside vacuum results in a shock wave travelling at speed Mach=6.7, while the contact surface between gas and vacuum follows at a constant speed of ~ 1800 m/s.

Several experimental studies have been performed on the issue of expansion of air into vacuum through an orifice. Parameters like initial base pressure [96] and initial pressure ratio across the orifice, ratio between the orifice size and the tube diameter [96, 97], gas species[98] were widely explored. Results vary considerably, with propagation speed of the pressure front between 50 m/s and 1100 m/s.

A.2 HIE-ISOLDE and SPIRAL2 safety systems

I faced[99] the problem of mixed regime gas dynamics through a collaboration with the HIE-ISOLDE experiment of CERN and SPIRAL2 of GANIL.

The Isotope mass Separator On-Line facility (ISOLDE) is a source of low-energy beams and radioactive nuclides at CERN. A high-energy proton beam is directed into special thick targets, yielding a large variety of atomic fragments. Different devices are used to ionize, extract and separate nuclei according to their mass, forming a low-energy beam that is delivered to various

experimental stations. This beam can be further accelerated to 3 MeV/nucleon. The post acceleration of radioactive beams has opened new fields of research, allowing the study of nuclear reactions with light and medium-mass radioactive projectiles. Presently an upgrade of the machine, HIE-ISOLDE, is underway that will improve the experimental capabilities of ISOLDE in many aspects. From spring 2016 radioactive beams of 5.5 MeV/nucleon will be available.¹.

In order to cope with space limitations of existing infrastructure, the cryomodules of the HIE-ISOLDE superconducting linac feature a common insulation and beam vacuum, imposing the severe cleanliness standard of RF cavities to the whole cryostat. Protection of the linac vacuum against air-inrush from the three experimental stations through the HEBT (High Energy Beam Transport) lines relies on fast valves, triggered by fast cold cathode gauges.

The protection of beam vacuum against an accidental vacuum hazard in an experimental station, crucial in the beam lines of synchrotron light sources, is present also in superconducting linacs and even more so in those with common beam and insulation vacuum in the cavity cryomodules. Indeed, vacuum rupture results in a propagating air front, transporting air and dust towards the clean cavities, causing loss of beam time and requiring a lengthy reconditioning or even cleaning of the cavities. Usually, protection is obtained by a fast acting valve along the beam line, together with air propagation delaying devices, so called Acoustic Delay Lines (ADL), first proposed in [100].

In HIE-Isolde, the compact layout of the beam transfer line system does not permit the insertion of a delay device. A fast valve, triggered by fast Penning gauges, reacts within 15 ms to the pressure signal rise; this time includes the initial pressure rise, the delay induced by the Penning gauge, cables, electronic cards, up to the movement onset of the valve and until complete leak-tightness is obtained. However, the overall distance between the nearest experimental station and the last cryomodule is shorter than 6 m: a shock wave travelling at 1000 m/s will reach the cryomodule before the fast valve has completely closed.

The problem is similar in case[101] of SPIRAL2, which is a radioactive ion beam accelerator at GANIL dedicated to fundamental and applied research in nuclear physics. This machine aims principally at producing high intensities of exotic nuclei issued from the interaction between the ion beam and the targets. Different kinds of beams will be delivered by a superconducting Linac (for example: Deuteron, E=40 MeV), and targets could be stable or radioactive. During interaction, a large quantity of contamination will be produced. So, if the vessel containing radioactive targets is broken, an air inrush could transport the contamination from the targets to the LINAC whose operation can't be maintained unless the system is clean. For this reason, a fast closing valve system between the targets and the accelerator is installed. The gauges to detect the inrush, the fast valves and the electronics are classified as *Equipment Important for Safety* and are used to prove to the Nuclear Safety Authority of France that the LINAC is correctly protected.

Although – as presented above – analytical and experimental data is available on the shock tube problem, because of the often contradicting numerical results and the common (HIE-Isolde and SPIRAL2) interest, we have decided to carry out our own measurement campaign

¹Introduction from the ISOLDE website, <http://home.cern/about/experiments/isolde>, retrieved 2016-02-25

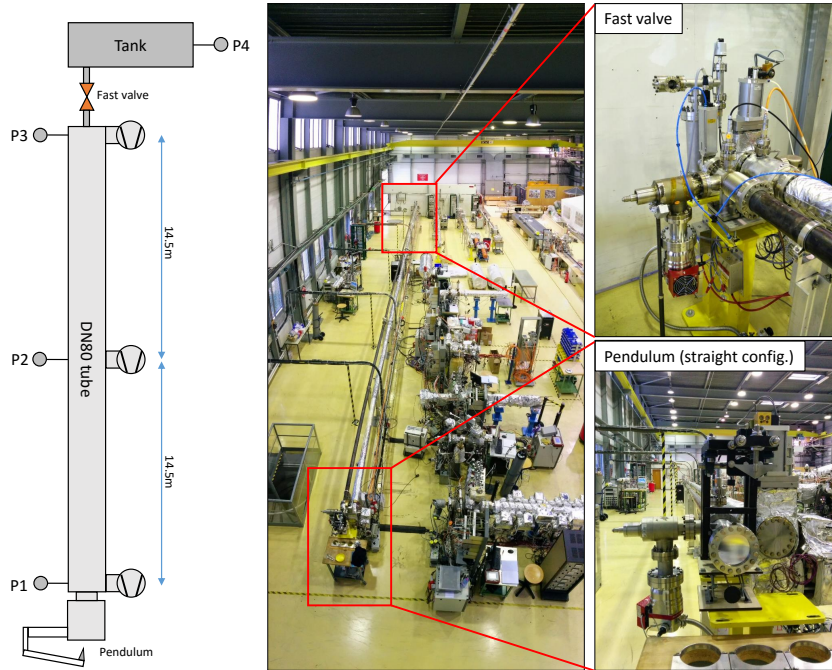


Figure A.2: Experimental setup of our high pressure experiments at CERN

to determine the propagation speed and at the same time test the safety system. This chapter thus describes our test station and the results we obtained.

A.3 Experimental setup

The experimental bench at CERN is presented in Fig.A.2. It consists of a 29 m long, 80 mm diameter tube, with intercalated 100 mm T-sections on which pumps and instrumentation are installed. The system is evacuated by 3 turbomolecular pumping groups of nominal pumping speed 100 l/s. Pressure is measured by commercial fast-reacting IKR070 Penning gauges placed next to the pendulum, the middle, and before and after the fast-closing valve.

At one extremity, the tube is closed by a thin vacuum window of a diameter 100 mm, obtained by clamping a pure aluminum foil of 100 μm between two copper gaskets. A pendulum armed with a sharp tip and a shock-recorder accelerometer is attached to the flange. Two configurations were tested: in the first case, the pendulum is parallel to the tube, in the second, it is rotated by 90°. At the other extremity, a fast valve (VAT, DN40) separates the main tube from a 3-liter volume vacuum chamber. Vacuum in the secondary chamber is also monitored by a Penning gauge; this serves to measure the gas that passes the valve before its full closure in future test runs.

The fast gauge signals are recuperated from the analog output modules of the VAT fast valve controller and sent to an 8-channel high frequency (19.8 kHz) 24bit A/D converter and amplifier (QuantumX MX840A by HBM). The amplifier sends data to a PC via an Ethernet bus. Data processing software is provided by the amplifier's supplier.




Figure A.3: Puncture holes of the pendulum on the aluminium vacuum window

A.4 Tests and results

Before starting a new test, the tube is pumped down for several hours until a base pressure of 1×10^{-7} mbar is reached. The gate valves isolating the turbomolecular pumps from the tube are then closed and pressure is stabilized for some minutes. At time t_0 , the window is broken.

For all tests, the window turned out to be rigid enough to keep its form after the rupture, therefore each sample was photographed, measured and catalogued after the tests. Some examples are shown in Fig.A.3. The pendulum knife's adjustable position allowed to experiment with different hole sizes. When keeping the position constant, the holes were almost identical over consecutive tests.

There are several ways to determine t_0 . With the gauges already installed, the most obvious is to check the pressure signal of P1, closest to the vacuum window. Its voltage output being between 0 and 10 V (with 2 V corresponding to an order of magnitude pressure difference between 1×10^{-8} and 1×10^{-3} mbar), we have chosen t_0 as the moment when its signal reached 0.4 V above the background level.

Since we suspect that there is a delay of gauge 1's signal (partly physical since it isn't precisely at the break location, and partly electronic – the signal has to go through a 30 m long cabling and through several controllers), a more precise way to catch the break moment is to install accelerometers on the system.

Using a special two-component glue that doesn't dampen vibrations, we have fixed an accelerometer on the pendulum and one on the fast-closing valve. In Fig.A.4 we can see the pressure signal of gauge 1 along with the accelerometer data.

Appendix A. High-pressure experiments




Figure A.4: Spectrum analyser signal. Top: pressure on P1; Bottom: accelerometers (identical time scale)

Table A.1: Events after foil break

Event	Foil break	Pressure rises at P1	4V threshold reached	Valve moves	Valve closes
Time (s)	5.1746	5.1785	5.1805	5.1817	5.1830

According to table A.1, there is a 4 ms delay between the pendulum breaking the window and the first noticeable change on the gauge. The safety system – actuated by this pressure rise measured by the first gauge – takes a further 3 ms to start closing the valve. The shock of the valve’s gate closure drives the accelerometer over range 1 ms later. The total 8 ms required for the valve to close is well below the manufacturer’s guarantee of 15 ms therefore the safety system works correctly.

As for the gas propagation, Fig.A.5 shows the four pressure signals of a typical test run. Since we know the gauge locations, and also the time when the pressure wave arrives (defined by the moments when the signal rises 0.4 V above the background), the propagation speed can be calculated.

Table A.2 shows the propagation speeds on different sections, and Fig.A.6 shows the data on a time vs. distance plot.

A.5 Additional test

The braking effect of the fast-closing valve’s reduced diameter (which is in fact a simple acoustic delay device) urged us to change the configuration of the pendulum. With the accelerometers in place, we repeated the experiment with the pendulum rotated by 90 degrees relative to the tube direction. As such, the direct line of sight, allowing straight flight of air molecules through the system is broken and all gas is forced to take a 90 degree turn.




Figure A.5: Penning gauge signals after breaking the vacuum window at $t = 98s$

Table A.2: Propagation speeds for different sections across nine experiment runs

Section speed (m/s)	gauge 1 -> gauge 2	gauge 1 -> gauge 3	gauge 1 -> gauge 4	gauge 2 -> gauge 3
1	868	n/a	719	n/a
2	895	803	678	729
3	873	803	650	744
4	879	810	639	751
5	843	782	629	729
6	843	790	723	743
7	852	763	581	690
8	829	652	467	537
9	829	707	575	617
Average	857	764	629	693
St. dev.(% of avg.)	2.42%	6.46%	11.39%	9.76%




Figure A.6: Pressure wave arrival moments as a function of gauge position, data plotted over 9 test runs

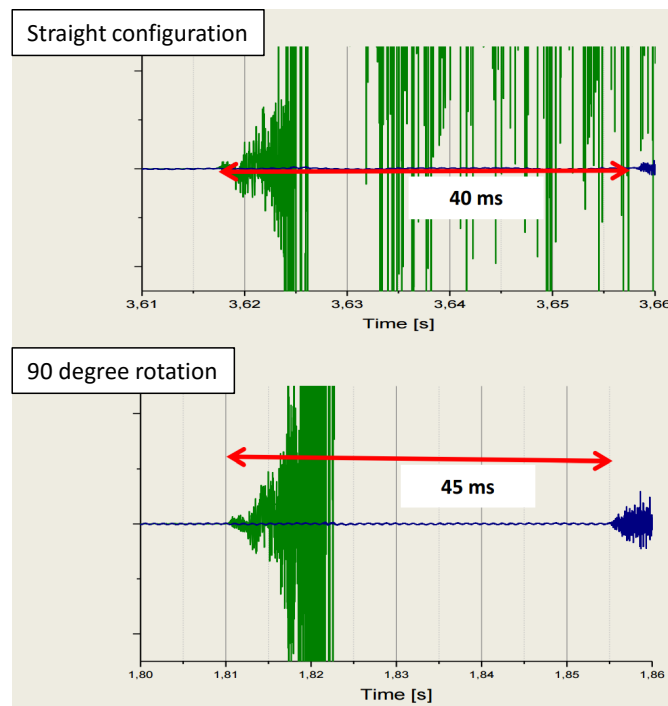


Figure A.7: Accelerometer signals at the foil and at the end of the line, showing the delay between the two locations. Source: [101]

Results (shown in Fig.A.7) show that the 90 degree turn introduces a 5 ms (or 12.5%) delay to the total flight time from entrance to opposite end.

A.6 Conclusion

- Reproducibility of the experiments is good (st.dev. 11% for the whole tube and a very low 2.4% for the first 14.5 meters)
- The air spreads initially with 857 m/s (Mach 2.5) then slows down by about 10% in the second section. This is visualized in Fig.A.6 where the delay compared to the P1->P2 propagation speed is visible for all tests
- Obstacles or turns further delay the inrush: the approximately 1 m distance between gauge 3 and gauge 4 cause a 1 ms delay, corresponding to a significantly lower propagation speed of 90 m/s
- We do not see evidence of an initial, low-density shockwave travelling at Mach 5 (predicted by [95]), only the subsequent high-pressure density wave at approximately half the speed

B Miscellaneous simulations

This last chapter gives the reader an outlook on some interesting usage areas of Molflow+ outside of regular vacuum simulations. During the years of development, we faced various problems that we could solve with relatively small modifications to the original code. These examples are to be treated as demonstrations of what the codes could do – their scientific accuracy depends on many parameters (in the examples below: real surface emissivity, light source properties, pressure in space, etc.) that weren't fully investigated. Unless noted otherwise, results are qualitative.

B.1 Thermal radiation simulations - McCryoT

In the introduction of chapter 1, I already mentioned the similarity between radiative heat transfer and ultra-high vacuum.

In heat transfer, the emissivity of a material quantifies its ability to emit and to absorb radiation. The emissivity of an ideal black body is 1, and that of a perfectly reflective surface is 0.

Using this quantity, the emitted radiation – which is analogous to outgassing – is described by the Stefan-Boltzmann law: the power P radiated from an object is

$$P = \sigma A \epsilon T^4$$

where $\sigma = 5.67 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ is the Stefan-Boltzmann constant, A is the area of the object, ϵ is the emissivity and T is the body temperature.

The ratio of the radiation absorbed and received by a body (analogous to sticking factor in vacuum) also equals to ϵ .

Since heat radiation obeys the cosine law, our vacuum ray-tracing engine can be reused with small changes for radiative heat transfer calculations.

I have written a program that we named *McCryoT*, which is essentially a mod of Molflow. The interface allows the user to define a temperature and an emissivity value for each facet.

Appendix B. Miscellaneous simulations

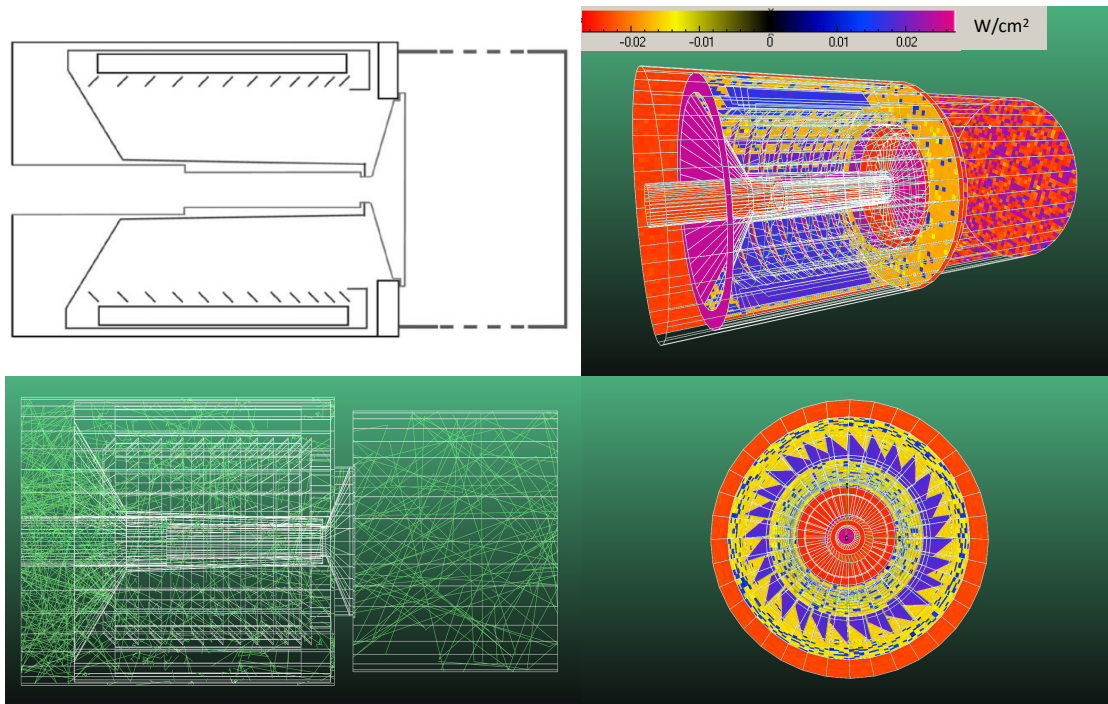


Figure B.1: Published[102] model of the ITER cryopump and its McCryoT simulation. Texture colors show the heat balance of elements (yellow and red parts lose while blue and violet parts gain power)

From these two parameters, "outgassing" or the radiated power is then calculated based on the Stefan-Boltzmann law, and the absorbed power is determined by the emissivity playing the role of the sticking factor.

Since non-isothermal systems can have large temperature differences (for example, a beam screen can have 10 K and 300 K parts), I argue that Monte-Carlo test particles shouldn't represent the power they carry – that way the fourth power scaling would confine virtually all test particles to hot parts. Instead, my emitted test particles scale with the area and emissivity of a surface, and remember the temperature of the facet they originate from. That way, upon absorption, the power they carry is determined by multiplying their number with σT^4 .

Each facet (with non-zero emissivity) is absorbing and emitting at the same time, and the user interface shows both the incident and radiated power, along with the power balance for each facet. Simulations are steady-state, so temperatures are considered constant.

Profiles and textures can count either only the emitted or absorbed power density, or the balance, therefore I modified the color scale to be able to visualize negative values as well.

As an example, Fig.B.1 shows the radiative heat balance of an ITER pre-production cryopump, based on [102].

We chose this problem as the article presenting the cryopump also recognizes the possibility to modify the author's MC code Provac3D to perform radiative heat transfer calculations. As

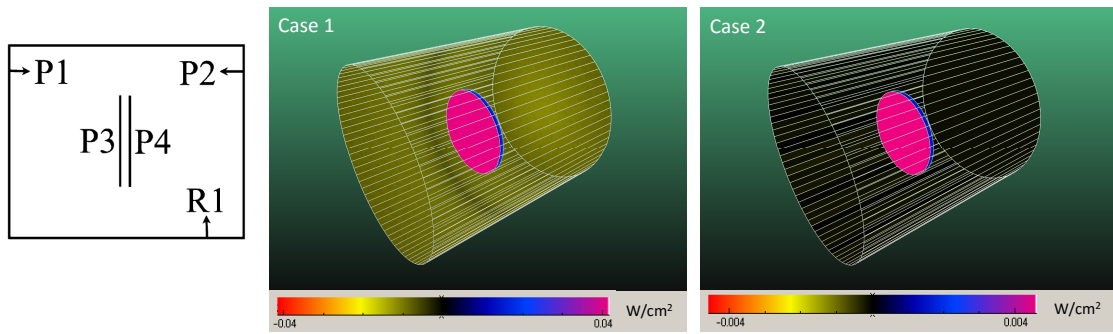


Figure B.2: simulated model and notations from [102] (left) and the two simulated cases in McCryoT (right)

Table B.1: Comparison of published[102] heat exchange results by Provac3D with McCryoT

	Case 1		Case 2	
	Provac3D	McCryoT	Provac3D	McCryoT
P1	-78.0	-77.8	-9.2	-9.2
P2	-78.1	-77.8	-9.1	-9.0
P3 right	18.7	18.4	8.5	8.5
P3 left	192.3	191.9	21.5	21.5
P4 right	192.3	191.9	21.6	21.5
P4 left	18.7	18.3	8.5	8.5
R1	-266.0	-265.1	-41.9	-41.8

that article contains a simple model to compare the view factor and the MC methods, we can actually benchmark our code.

The model we create in McCryoT is a cylinder at 300 K temperature of 178.4 cm length and 154.6 cm diameter, with two discs in the center kept at 4.5 K, having half the cylinder's diameter. Radiated power from the components is indexed as shown in Fig.B.2.

Two cases are simulated: in case 1, all components have 0.9 as emissivity, in case 2, the central discs' emissivity is reduced to 0.1.

Table B.1 shows the heat balance of the components, as published in [102] and as calculated by McCryoT. Results match within 1.8%.

B.2 Lighting calculations

On rough surfaces visible light scatters following the cosine law – just like molecules in vacuum. Therefore – without any modifications – we can use Molflow+ for approximate illumination calculations.

In this simulation, I used an existing model of a room. During the day, light comes from the window. At night, it comes from a lamp on the ceiling.

Appendix B. Miscellaneous simulations

Table B.2: Transmission probabilities between the inlet and extraction port of the Gaede pump, as a function of the rotor size and speed

Rotor speed [Hz]	Rotor diameter [cm]				
	15.6	16.3	17.1	18	18.9
-1500	0.013	0.011	0.008	0.006	0.004
-1000	0.017	0.014	0.011	0.008	0.005
-500	0.031	0.026	0.020	0.015	0.009
0	0.113	0.099	0.082	0.061	0.040
500	0.300	0.298	0.283	0.260	0.212
1000	0.407	0.406	0.400	0.384	0.345
1500	0.451	0.451	0.447	0.437	0.408

The light sources are emitting with cosine distribution. Walls have a sticking factor of 0.65, shiny surfaces such as the table have 0.35. Lighting is visualized in Fig.B.3 as a grayscale texture counting the density of absorbed photons.

B.3 Moving parts

Molflow allows the user to define some facets as moving. Motion can be fixed (in a certain direction) or calculated depending on the facet location (in case of rotation around an axis). At every reflection the local facet's velocity vector is added to the diffuse reflected direction.

B.3.1 Gaede pump

W. Gaede built one of the first turbomolecular pumps in 1905. It consists of a rotor spinning with edges very close to the stator, and the drag force of the moving rotor surface orients the molecules from the inlet pipe to the outlet.

Because of its simplicity, we can easily model it in Molflow. I designed the stator to have a 20 cm internal diameter, then checked the pumping efficiency as a function of the rotor size and speed.

In Fig.B.5, the model is shown. The facets highlighted in red that are parts of the rotor are defined as moving parts. The rotation axis is defined in Molflow+ by two vertices going through the rotor's centerline, and the rotation speed is varied between 0 and 1500 Hz in both directions.

Both the inlet and outlet have perfect sticking, and the pumping efficiency is measured by the transmission probability, that is the ratio of molecules leaving through the outlet and all molecules.

Looking at the results in table B.2, we can see that there is a good and bad spin direction, i.e. one that drives molecules towards the outlet and one that pushes them back to the inlet.

The transmission probability rises abruptly compared to the stationary case with a small spin

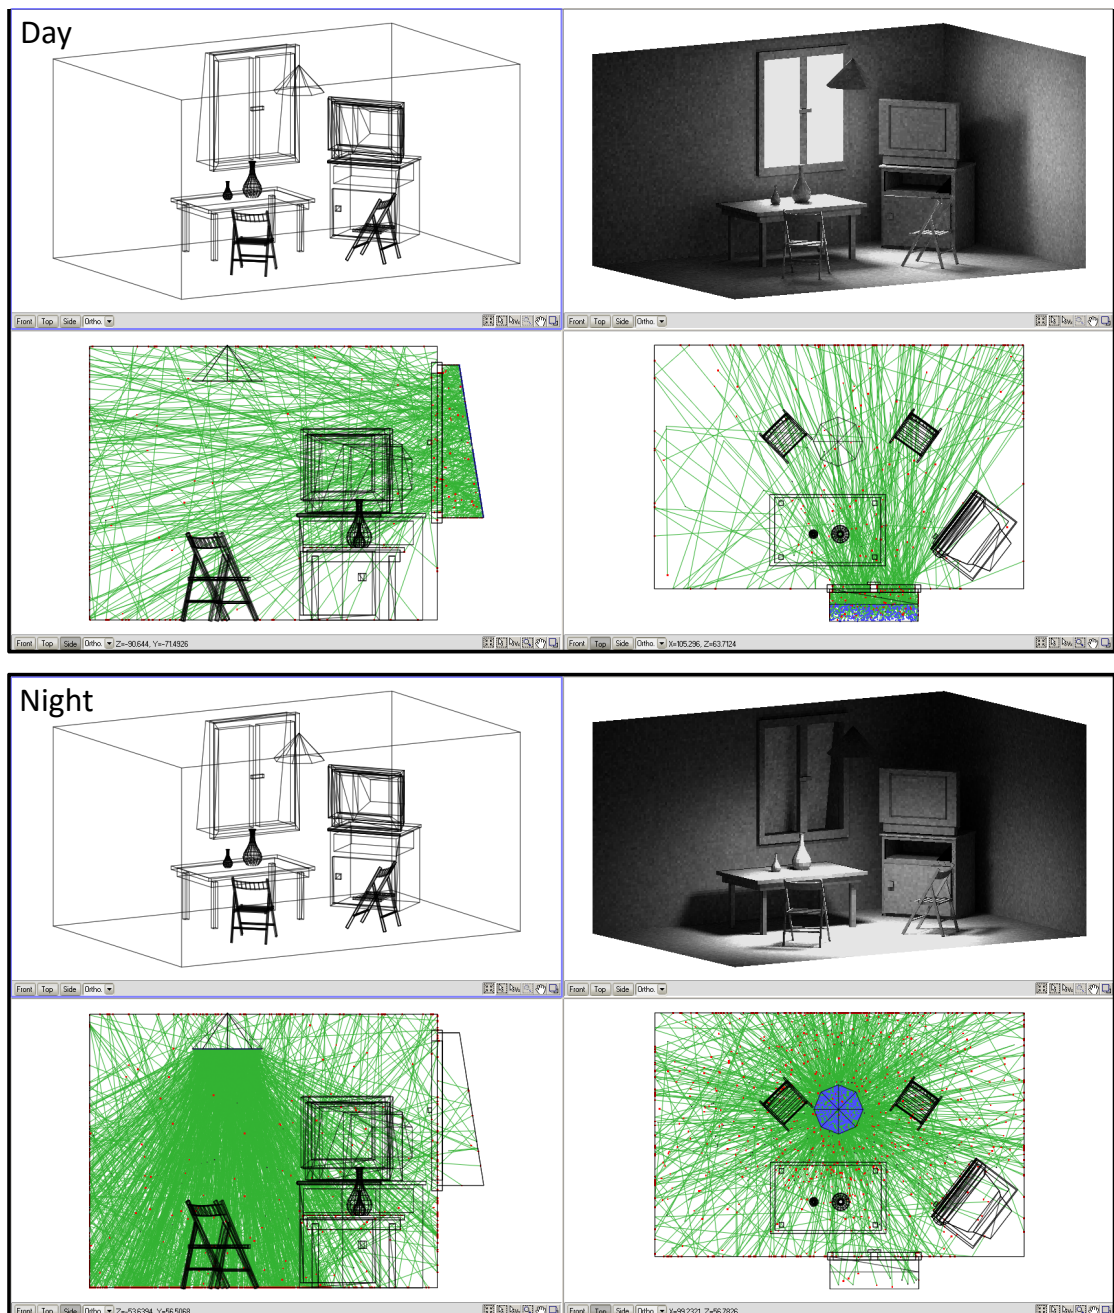


Figure B.3: Illumination in a simple room during day (top) and night (bottom)
Geometry: R. Kersevan / J.L. Pons




Figure B.4: W. Goede's rotary molecular pump, image from [103]




Figure B.5: Pressure distribution and direction vectors for the Gaede pump model with 18 cm rotor size at 1500 Hz rotation speed. One can notice particles "stuck" at the outlet causing a locally elevated pressure.

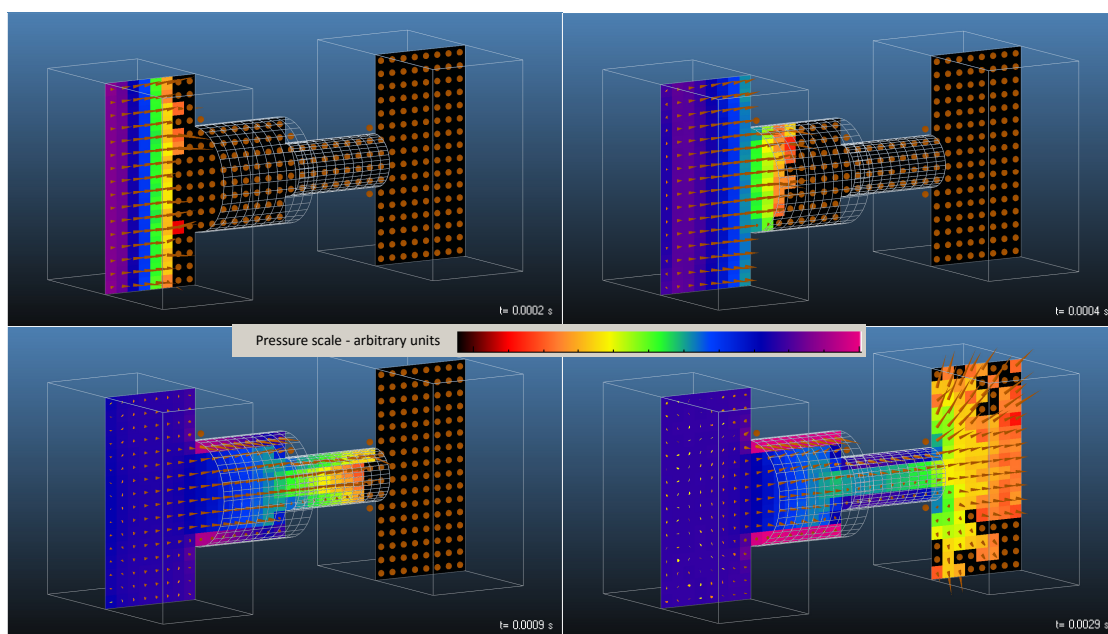


Figure B.6: Pressure profile at four different moments after gas injection from the left through a spinning rotor in the middle

speed (0 vs. 500 Hz) but reaches saturation around 1500 Hz.

Decreasing the rotor diameter also increases the efficiency because the conductance becomes higher. This effect is important in the stationary case (driven entirely by the conductance: the smallest diameter rotor has 180% higher transmission probability than the largest), but at high spin speeds it becomes less significant (only 10% improvement at 1500 Hz).

An interesting phenomenon to notice in Fig.B.5 is the particles "stuck" under the outlet causing a local increase in the pressure (the rotor drives them towards the corner).

B.3.2 Time-dependent problems

Combined with the time-dependent mode, simulations with moving parts can create informative demonstrations of molecule dynamics. In Fig.B.6 a system state is shown at four different moments.

At $t=0$, gas is desorbed in the left reservoir. The gas then passes through a spinning tube before arriving to the reservoir at the opposite end with sticking walls. Frames in Fig.B.6 depict four characteristic states of the system:

1. **$t=0.2$ ms** At this point the gas is confined to the source reservoir following the pulsed desorption. Consequently, all direction vectors are either zero or point towards the expansion direction of the gas
2. **$t=0.4$ ms** The first particles enter the spinning tube. However, their dominant direction




Figure B.7: Simulation of a space shuttle moving at 10 km/s speed in space at 1×10^{-6} mbar

is still towards the expansion.

3. **t=0.9 ms** At this point the fastest gas molecules have almost crossed the spinning tube. Accordingly, those molecules that are close to the tube have a tangential direction. Pressure near the wall also increased because the rotor augments the molecule speeds. The effect is stronger near the rotor entrance, where molecules already had several collisions with the tube wall, and almost negligible at the gas expansion front.
4. **t=2.9 ms** Quasi-steady state reached. Molecules exiting the tube have their directions perturbed by the spin (at the top they come towards the viewer, at the bottom they move away). In the tube, the pressure is higher near the wall along all the length.

B.3.3 Space shuttle

100 km above the Earth's surface, the pressure drops to around 1×10^{-2} Pa [104]. At even higher altitudes, we can approximate space as a medium under high vacuum, which we can simulate with Molflow.

I have downloaded a reference space shuttle model from NASA's 3D model database[105], and inserted it in a bounding box.

Each side of the box has a small outgassing and a small sticking, with values chosen to create a near-equilibrium state inside of 1×10^{-6} mbar. The space shuttle's surface is defined as a moving part with 10 km/s velocity towards the shuttle's nose. All parts are considered to be at 3K.

Figure B.7 shows the pressure distribution on the shuttle's surface using a logarithmic color scale. As the stationary pressure of 1×10^{-6} mbar corresponds to the blue color, we can see that the pressure deviates from the uniform case: there is a higher load on the nose and wing entries and zero pressure on the shielded parts (tail).

This simulation is adequate for a flight at high altitudes. When the shuttle reenters the atmosphere (below the Kármán line at 100 km, at a speed of 7.8 km/s), the gas is compressed and the temperature increases, therefore DSMC simulations are more adequate.

Bibliography

- [1] J Kouptsidis and AG Mathewson. Reduction of the photoelectron induced gas desorption in the petra vacuum system by in situ argon flow discharge cleaning. Technical report, Deutsches Elektronen-Synchrotron (DESY), 1976.
- [2] R. Kersevan, P. Chiggiato, *Examples of Design of Vacuum Systems for Accelerators, JUAS 2013 proceedings*, 7 January 2013 to 14 March 2013.
- [3] Oerlikon Leybold Vacuum. Fundamentals of vacuum technology, 2007.
- [4] Pfeiffer - The vacuum technology Book, Vol 1-5.
- [5] O Bruning et al. *LHC design report*. European Organization for Nuclear Research, 2001.
- [6] R. Kersevan. Analytical and numerical tools for vacuum systems. *CAS CERN Accelerator School, Vacuum in Accelerators: Platja D'Aro, Spain, 16-24 May 2006*, (3):285, 2007.
- [7] G. Bregliozi, G. Lanza, Simulation of the LHC vacuum system with the Vasco Code in presence of electron clouds, CERN.
- [8] M.Sullivan, VACCALC: A Program to find Beam pipe Pressure Profiles, Nov. 3, 2011, JLAB, CASA, MEIC .
- [9] M. Knudsen. Die Gesetze der Molekularströmung und der inneren Reibungsströmung der Gase durch Rohren. *Annalen der Physik*, 333(1):75–130, 1909.
- [10] P. Clausing. The flow of highly rarefied gases through tubes of arbitrary length. *Journal of Vacuum Science & Technology*, 8(5):636–646, 1971.
- [11] J. Lettry C. Pasquino, P. Chiggiato. Vacuum simulation of the LINAC4 H- source. May 2012.
- [12] R. Kersevan and J-L. Pons. Introduction to MOLFLOW+: New graphical processing unit-based Monte Carlo code for simulating molecular flows and for calculating angular coefficients in the compute unified device architecture environment. *Journal of Vacuum Science & Technology A*, 27(4):1017–1023, 2009.
- [13] X. Luo, T. Giegerich, and C. Day. Transient gas flow studied by a test particle Monte Carlo approach with ProVac3D. *AIP Conference Proceedings*, 1501(1):857–863, 2012.

- [14] A.G. Popeko, O.N. Malyshev, R.N. Sagaidak, and A.V. Yeremin. Monte carlo simulation of ion trajectories in the kinematic recoil separator “vassilissa”. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 126(1–4):294 – 296, 1997. International Conference on Electromagnetic Isotope Separators and Techniques Related to Their Applications.
- [15] Y. Suetsugu. Application of the monte carlo method to pressure calculation. *Journal of Vacuum Science & Technology A*, 14(1):245–250, 1996.
- [16] A. Pace and A. Poncet. Monte carlo simulations of molecular gas flow: some applications in accelerator vacuum technology using a versatile personal computer program. *Vacuum*, 41(7):1910 – 1913, 1990.
- [17] C Garion. Monte Carlo method implemented in a finite element code with application to dynamic vacuum in particle accelerators. *Vacuum*, 84(1):274–276, 2009.
- [18] Graeme Austin Bird. Molecular gas dynamics. *NASA STI/Recon Technical Report A*, 76:40225, 1976.
- [19] Y. Kusumoto. Reflection rules preserving molecular flow symmetry in an arbitrarily shaped pipe. *Journal of Vacuum Science & Technology A*, 25(2):401–408, 2007.
- [20] J. Greenwood. The correct and incorrect generation of a cosine distribution of scattered particles for monte-carlo modelling of vacuum systems. *Vacuum*, 67(2):217–222, 2002.
- [21] M. Matsumoto and T. Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 8(1):3–30, 1998.
- [22] J.C. Maxwell. Illustrations of the dynamical theory of gases. *The Kinetic Theory Of Gases. Series: History of Modern Physical Sciences, ISBN: 978-1-86094-347-8., vol. 1, pp. 148-171, 1:148–171*, 2003.
- [23] W. Hörmann and J. Leydold. Continuous random variate generation by fast numerical inversion. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 13(4):347–362, 2003.
- [24] D.J. Rader. *Measurements of thermal accommodation coefficients*.
- [25] VL Kovalev, AN Yakunchikov, Deborah A Levin, Ingrid J Wysong, and Alejandro L Garcia. Accommodation coefficients for hydrogen molecules on graphite surface. In *AIP Conference Proceedings-American Institute of Physics*, volume 1333, page 481, 2011.
- [26] H. Matsumoto and K. Kanamori. Monte carlo simulation of rarefied gas flow induced by wall temperature gradient. In *28th International Symposium on Rarefied Gas Dynamics 2012*, volume 1501, pages 661–666. AIP Publishing, 2012.
- [27] Low-Pressure effusion of gases, Hope College, Krueger Laboratory for Biophysics.
- [28] C Benvenuti. Molecular surface pumping: the getter pumps. *CERN-Reports*, pages 43–50, 1999.

-
- [29] M. Maietta: A study of the propagation of neutral isotopes at Isotopes Separator On Line Device (ISOLDE): an analysis by Monte Carlo simulation and spectroscopic methods; Universita degli Studi di Napoli "Federico II".
- [30] S.B. Damelin and W. Miller Jr. *The mathematics of signal processing*. Number 48. Cambridge University Press, 2012.
- [31] F. Sharipov. Non-isothermal gas flow through rectangular microchannels. *Journal of Micromechanics and Microengineering*, 9(4):394, 1999.
- [32] R Assmann, R Bingham, T Bohl, C Bracco, A Butterworth, A Caldwell, S Chattopadhyay, S Cipiccia, E Feldbaumer, RA Fonseca, et al. Proton-driven plasma wakefield acceleration: a path to the future of high-energy particle physics. *Plasma Physics and Controlled Fusion*, 56(8):084013, 2014.
- [33] R. Kersevan. Synrad, a monte carlo synchrotron radiation ray-tracing program. In *Particle Accelerator Conference, 1993., Proceedings of the 1993*, pages 3848–3850. IEEE, 1993.
- [34] J Rossbach and Peter Schmueser. Basic course on accelerator optics. In *Conf. Proc.*, volume 9209071, pages 17–88, 1992.
- [35] GK Green. Spectra and optics of synchrotron radiation. Technical report, Brookhaven National Lab., Upton, NY (USA), 1976.
- [36] J.D. Jackson. *Classical electrodynamics*. Wiley, 1999.
- [37] K.T. McDonald: Notes on Synchrotron Radiation, Princeton U., Feb 10 1997.
- [38] H H Umstätter. Efficient computation of synchrotron radiation spectrum. Technical Report CERN-PS-SM-81-13, CERN, Geneva, 1981.
- [39] H Burkhardt. Monte Carlo generation of the energy spectrum of synchrotron radiation. Technical Report CERN-OPEN-2007-018. CLIC-Note-709. EUROTEV-Report-2007-018, CERN, Geneva, Jun 2007.
- [40] A. Hofmann. *The Physics of Synchrotron Radiation*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 2004.
- [41] J. Schwinger. On the classical radiation of accelerated electrons. *Physical Review*, 75(12):1912, 1949.
- [42] S Mobilio and A Balerna. Introduction to the main properties of synchrotron radiation. In *Conference Proceedings-Italian Physical Society*, volume 82, pages 1–24. Editrice Compositori; 1999, 2003.
- [43] Do H Tomboulia and PL Hartman. Spectral and angular distribution of ultraviolet radiation from the 300-mev cornell synchrotron. *Physical Review*, 102(6):1423, 1956.
- [44] A.A. Sokolov and I.M. Ternov. Synchrotron radiation. *Akademia Nauk SSSR, Moskovskoie Obshchestvo Ispytatelei prirody. Seksia Fiziki. Sinkhrotron Radiation, Nauka Eds., Moscow, 1966 (Russian title: Sinkhrotronnoie izluchenie)*, 228 pp., 1, 1966.

Bibliography

- [45] R. Coisson. Angular-spectral distribution and polarization of synchrotron radiation from a "short" magnet. *Phys. Rev. A*, 20:524–528, Aug 1979.
- [46] S. Agostinelli, J. Allison, K al Amako, J Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G Barrand, et al. Geant4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.
- [47] Madx website. <http://mad.web.cern.ch/mad/>. Accessed: 2016-02-28.
- [48] Steven L.H. and Jill M.W. Flux and brightness calculations for various synchrotron radiation sources. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 319(1–3):25 – 31, 1992.
- [49] The X-Ray Data Booklet, Lawrence Berkeley National Laboratory, Oct. 2009.
- [50] B.L. Henke, E.M. Gullikson, and J.C. Davis. X-ray interactions: Photoabsorption, scattering, transmission, and reflection at $e = 50$ -30,000 eV, $z = 1$ -92. *Atomic Data and Nuclear Data Tables*, 54(2):181 – 342, 1993.
- [51] H-J Hagemann, W Gudat, and C Kunz. Optical constants from the far infrared to the x-ray region: Mg, Al, Cu, Ag, Au, Bi, C, and Al₂O₃. *JOSA*, 65(6):742–744, 1975.
- [52] Refractive Index Database - Mikhail Polyanskiy. <http://refractiveindex.info/>. Accessed: 2016-02-28.
- [53] O. Groebner, The Impact of SR in the LHC, VLHC Workshop, 18-20 Sept 2000.
- [54] NIST Database. <http://physics.nist.gov/PhysRefData/XrayMassCoef/ElemTab/z13.html>. Accessed: 2016-02-29.
- [55] P. Beckmann and A. Spizzichino. The scattering of electromagnetic waves from rough surfaces. *Norwood, MA, Artech House, Inc., 1987, 511 p.*, 1987.
- [56] J. A. Ogilvy and Harold M. Merklinger. Theory of wave scattering from random rough surfaces. *The Journal of the Acoustical Society of America*, 90(6):3382–3382, 1991.
- [57] G. Dugan and D. Sagan. Synrad3d photon propagation and scattering simulations. 2013.
- [58] H. Hogrefe and C. Kunz. Soft x-ray scattering from rough surfaces: experimental and theoretical analysis. *Applied optics*, 26(14):2851–2859, 1987.
- [59] RJ Noll and P Glenn. Mirror surface autocovariance functions and their associated visible scattering. *Applied optics*, 21(10):1824–1838, 1982.
- [60] G Dugan, M Billing, K Butler, J Chu, J Crittenden, M Forster, D Kreinick, R Meller, M Palmer, G Ramirez, et al. Observations and Predictions at CesrTA, and Outlook for ILC. *arXiv preprint arXiv:1310.0261*, 2013.
- [61] GF Dugan, KG Sonnad, R Cimino, T Ishibashi, and F Schäfers. Measurements of x-ray scattering from accelerator vacuum chamber surfaces, and comparison with an analytical model. *Physical Review Special Topics-Accelerators and Beams*, 18(4):040704, 2015.

- [62] CL Foerster, HJ Halama, and G Korn. Photodesorption from copper, beryllium, and thin films. *Journal of Vacuum Science & Technology A*, 10(4):2077–2081, 1992.
- [63] M Bernardini and L Malter. Vacuum problems of electron and positron storage rings. *Journal of Vacuum Science & Technology*, 2(3):130–141, 1965.
- [64] J Gómez-Goñi, O Gröbner, and AG Mathewson. Comparison of photodesorption yields using synchrotron radiation of low critical energies for stainless steel, copper, and electrodeposited copper surfaces. *Journal of Vacuum Science & Technology A*, 12(4):1714–1718, 1994.
- [65] Y. Tanimoto, K. Kanazawa, PSD Yield Measurement from a Gold-coated Surface, The 54th Annual Symposium of the Vacuum Society of Japan & the 33rd Annual Symposium of the Surface Science Society of Japan, 26 November 2013.
- [66] O Grobner, AG Mathewson, and PC Marin. Extended study of photon stimulated gas desorption from ofhc copper by 3.75 kev critical energy photons. In *EPAC 92: Third European Particle Accelerator Conference: Berlin, 24-28 March, 1992*, page 132. Editions Frontières, 1992.
- [67] O.B. Malyshev, V.V. Anashin, R.V. Dostovalov, N.V. Fedorov, A.A. Krasnov, I.R. Collins, and V.L. Ruzinov. Method and setup for photodesorption measurements for a nonevaporable-getter-coated vacuum chamber. *Journal of Vacuum Science & Technology A*, 23(3):570–576, 2005.
- [68] V.V. Anashin, O.B. Malyshev, V.N. Osipov, I.L. Maslennikov, and W.C. Turner. Investigation of synchrotron radiation-induced photodesorption in cryosorbing quasiclosed geometry. *Journal of Vacuum Science & Technology A*, 12(5):2917–2921, 1994.
- [69] Oswald Gröbner, AG Mathewson, and PC Marin. Gas desorption from an oxygen free high conductivity copper vacuum chamber by synchrotron radiation photons. *Journal of Vacuum Science & Technology A*, 12(3):846–853, 1994.
- [70] P. Chiggiato and R. Kersevan. Synchrotron radiation-induced desorption from a neg-coated vacuum chamber. *Vacuum*, 60(1):67–72, 2001.
- [71] J. Carter, APS, private communications, to be published at IPAC’16, Busan.
- [72] M Andritschky. Damage of oxide layers on an Al-alloy by electron bombardment. (CERN-LEP-VA-88-53):11 p, Oct 1988.
- [73] R J Reid. Cleaning for vacuum service. (OPEN-2000-276), 1999.
- [74] V.V Anashin, I.R Collins, R.V Dostovalov, N.V Fedorov, A.A Krasnov, O.B Malyshev, and V.L Ruzinov. Comparative study of photodesorption from tizrv coated and uncoated stainless steel vacuum chambers. *Vacuum*, 75(2):155 – 159, 2004.
- [75] R. Kersevan. Synchrotron radiation vacuum concepts. Future Circular Collider Study Kickoff Meeting. Feb 2014.

Bibliography

- [76] H. Yoshida, K. Arai, H. Akimichi, and M. Hirata. Two-stage flow-dividing system for the calibration of vacuum gauges. *Journal of Vacuum Science & Technology A*, 26(1):128–132, 2008.
- [77] V.V. Anashin, O.B. Malyshev, I.R. Collins, and O. Gröbner. Photon-stimulated desorption and the effect of cracking of condensed molecules in a cryogenic vacuum system. *Vacuum*, 60(1):15–24, 2001.
- [78] M.A. Gondal, Zain H. Yamani, A. Dastgeer, M.A. Ali, and A. Arfaj. Photo-conversion of methane into higher hydrocarbons using 355 nm laser radiation. *Spectroscopy Letters*, 36(4):313–326, 2003.
- [79] P. Chiggiato and P. Costa-Pinto. Ti-Zr-V non-evaporable getter films: from development to large scale production for the Large Hadron Collider. *Thin Solid Films*, 515(CERN-TS-2006-001. 2):382–388. 13 p, Dec 2005.
- [80] CL Foerster. Synchrotron light source measurements of photon stimulated desorption from vacuum chamber materials. *Synchrotron Radiation News*, 11(5):20–26, 1998.
- [81] R. Kersevan, presented at ALCATEL Workshop, Gen. 2003.
- [82] A. Rossi and N. Hilleret. Residual gas density estimations in the lhc experimental interaction regions. In *Prepared for 2nd national conference of the german vacuum society, berlin, germany*, pages 13–26, 2003.
- [83] M. Ady and R. Kersevan. Monte-Carlo simulation of Synchrotron Radiation power for the LHeC interaction region quadrupoles using SynRad+. Dec 2013.
- [84] M Klein, L Rinolfi, S Chattopadhyay, Oliver Sim Brüning, TPR Linnecar, R Tomás, JM Jowett, R Garoby, Saleh Sultansoy, A Vivoli, et al. Prospects for a large hadron electron collider (lhec) at the lhc. Technical report, LHeC-Note-2008-003-ACC, 2008.
- [85] N. Bernard. Analytic method to calculate the power produced by synchrotron radiation in a quadrupole magnet. Technical report, LHeC-Note-2010-002-ACC, 2010.
- [86] L. Thompson, R.B. Appleby, B.J. Holzer, M. Fitterer, P. Kostka, M. Klein, and NR Bernard. Interaction region design for a ring-ring lhec. *IPAC11*, 2011.
- [87] M. Ady, R. Kersevan, and M. Grabski. Monte carlo simulations of synchrotron radiation and vacuum performance of the max iv light source. 2014.
- [88] E.S. Reich et al. Ultimate upgrade for us synchrotron. *Nature*, 501(7466):148–149, 2013.
- [89] S Hashimoto, M Hazumi, J Haba, JW Flanagan, and Y Ohnishi. Letter of intent for kek super b factory. Technical report, High Energy Accelerator Research Organization, 2004.
- [90] T. Abe, I. Adachi, K. Adamczyk, S. Ahn, H. Aihara, K. Akai, M. Aloï, L. Andricek, K. Aoki, Y. Arai, et al. Belle ii technical design report. *arXiv preprint arXiv:1011.0352*, 2010.
- [91] KA Stephani, DB Goldstein, and PL Varghese. Parametric study of hypersonic boundary layer flow over discrete surface roughness using a hybrid dsmc/navier-stokes solver. In *41st AIAA Fluid Dynamics Conference and Exhibit*, page 3250, 2011.

-
- [92] T.E. Schwartzentruber and I.D. Boyd. Detailed analysis of a hybrid cfd-dsmc method for hypersonic non-equilibrium flows. In *38th AIAA Thermophysics Conference, Toronto, Ontario Canada*, pages 6–9, 2005.
- [93] S. Pantazis and H Rusche. A hybrid continuum-particle solver for unsteady rarefied gas flows. *Vacuum*, 109:275–283, 2014.
- [94] S. Pantazis, Hybrid continuum-particle simulations for vacuum flows; School on Vacuum Gas Dynamics via Kinetic Theory; Volos, 07-11 May 2012.
- [95] HP Greenspan and DS Butler. On the expansion of a gas into vacuum. *Journal of Fluid Mechanics*, 13(01):101–119, 1962.
- [96] T. Takiya, F. Higashino, Y. Terada, and A. Komura. Pressure wave propagation by gas expansion in a high vacuum tube. *Journal of Vacuum Science & Technology A*, 17(4):2059–2063, 1999.
- [97] W Peatman and EW Weiner. Measurements of shock wave flight times in long pipes. *Journal of Vacuum Science & Technology*, 17(5):1237–1240, 1980.
- [98] H Betz, P Hofbauer, and A Heuberger. Measurements on the efficiency of acoustic delay lines in view of beam lines for synchrotron radiation. *Journal of Vacuum Science & Technology*, 16(3):924–926, 1979.
- [99] M. Ady, D. Ziemianski, G. Vandoni, R. Kersevan, and M. Hermann. Leak propagation dynamics for the HIE-ISOLDE superconducting linac. Technical report, 2014.
- [100] R.Jean and J.Rauss, *Le Vide* 111 (1964) p.123.
- [101] M. Faye, R. Levallois, M. Ady - Measurement of atmospheric air inrush speed into a high vacuum vessel, Proceedings of the 2nd European Conference on Non-equilibrium Gas Flows – NEGF15; December 9-11, 2015 – Eindhoven.
- [102] X. Luo, V. Hauer, and C. Day. Monte Carlo calculation of the thermal radiation heat load of the ITER pre-production cryopump. *Fusion Engineering and Design*, 87(5):603–607, 2012.
- [103] S.P. Khromov. Great Soviet Encyclopedia (a translation of the third edition). *Macmillan London, Collier Macmillan Publishers*, 17:129, 1978.
- [104] T. Squire. US Standard Atmosphere, 1976. *Thermal Protection Systems Expert and Material Properties Database*, 2000.
- [105] NASA 3D model database. <http://nasa3d.arc.nasa.gov/>. Accessed: 2016-03-02.



MÁRTON ADY

BORN: BUDAPEST, 13 NOVEMBER 1986

NATIONALITY: HUNGARIAN

Education

Apr 2013 to May 2016	Ecole Polytechnique Federale de Lausanne <ul style="list-style-type: none">▪ Monte Carlo simulations of Synchrotron Radiation and Ultra High Vacuum	PhD in physics
Sep 2010 to Jan 2013	Budapest University of Technology and Economics <ul style="list-style-type: none">▪ Thesis subject: nonlinear effects in optical cables	MSc in physics
Sep 2007 to June 2009	Ecole Centrale Paris <ul style="list-style-type: none">▪ Double-degree programme with French Government <i>Eiffel</i> scholarship	Ingénieur Généralist
Sep 2005 to June 2007	Budapest University of Technology and Economics	Engineering Physics

Experience

CERN Geneva, Switzerland TE dept., Vacuum and Surface Coatings	PhD student Apr 2013 to Apr 2016
KEK Tsukuba, Japan Vacuum group, Photon Factory	Guest researcher Oct to Nov 2014 Oct to Nov 2015
CERN Geneva, Switzerland TE dept., Vacuum and Surface Coatings	Technical Student Feb to Dec 2012
Budapest University of Technology and Economics	Evening lecturer Autumn semester 2011
CERN Geneva, Switzerland LHCb online team	Summer Student June to Sep 2011
TOTAL E&P Indonesia Balikpapan, Indonesia Drillings department	Intern Oct 2009 to Mar 2010
Ginza Megane , Tokyo, Japan	Intern Apr to June 2010
Hotel New Shiobara , Nasushiobara, Japan	Intern July to Aug 2008