Machine Protection for FLASH and the European XFEL

Lars Fröhlich
Deutsches Elektronen-Synchrotron
Notkestraße 85
22607 Hamburg, Germany
lars.froehlich@desy.de

German title:
Maschinenschutz für FLASH und den Europäischen XFEL

Published under DESY report numbers:
DESY-THESIS-2009-012
TESLA-FEL 2009-03

Printed by the DESY copy center.
This document was typeset using LaTeX and KOMA-Script.

Universität Hamburg
Department Physik

Machine Protection for FLASH and the European XFEL

Dissertation
zur Erlangung des Doktorgrades
des Department Physik
der Universität Hamburg

vorgelegt von
Lars Fröhlich
aus Brake (Unterweser)

Hamburg
2009

Gutachter der Dissertation:	Prof. Dr. Jörg Roßbach Prof. Dr. Rüdiger Schmidt Prof. Dr. Leonid Rivkin
Gutachter der Disputation:	Prof. Dr. Jörg Roßbach Priv.-Doz. Dr. Bernhard Schmidt
Datum der Disputation:	29. Mai 2009
Vorsitzender des Prüfungsausschusses:	Dr. Georg Steinbrück
Vorsitzender des Promotionsausschusses:	Prof. Dr. Robert Klanner
Dekan der Fakultät für Mathematik, Informatik und Naturwissenschaften:	Prof. Dr. Heinrich H. Graener

Abstract

The Free-Electron Laser in Hamburg (FLASH) and the future European X-Ray Free-Electron Laser (XFEL) are sources of brilliant extreme-ultraviolet and X-ray radiation pulses. Both facilities are based on superconducting linear accelerators (linacs) that can produce and transport electron beams of high average power. With up to 90 kW or up to 600 kW of power, respectively, these beams hold a serious potential to damage accelerator components. This thesis discusses several passive and active machine protection measures needed to ensure safe operation.

At FLASH, dark current from the rf gun electron source has activated several accelerator components to unacceptable radiation levels. Its transport through the linac is investigated with detailed tracking simulations using a parallelized and enhanced version of the tracking code Astra; possible remedies are evaluated.

Beam losses can lead to the demagnetization of permanent magnet insertion devices. A number of beam loss scenarios typical for FLASH are investigated with shower simulations. A shielding setup is designed and its efficiency is evaluated. For the design parameters of FLASH, it is concluded that the average relative beam loss in the undulators must be controlled to a level of about 10^{-8}.

FLASH is equipped with an active machine protection system (MPS) comprising more than 80 photomultiplier-based beam loss monitors and several subsystems. The maximum response time to beam losses is less than 4 µs. Setup procedures and calibration algorithms for MPS subsystems and components are introduced and operational problems are addressed.

Finally, an architecture for a fully programmable machine protection system for the XFEL is presented. Several options for the topology of this system are reviewed, with the result that an availability goal of at least 0.999 for the MPS is achievable with moderate hardware requirements.

Kurzfassung

Der Freie-Elektronen-Laser in Hamburg (FLASH) und der geplante Europäische Röntgenlaser (XFEL) sind Quellen brillanter Strahlungspulse im extrem-ultravioletten und Röntgen-Spektralbereich. Beide Einrichtungen basieren auf supraleitenden Linearbeschleunigern (Linacs), die Elektronenstrahlen mit hoher mittlerer Leistung erzeugen und transportieren können. Mit Leistungen von 90 kW bzw. 600 kW haben diese Strahlen ein hohes Gefährdungspotenzial für verschiedene Beschleunigerkomponenten. Diese Arbeit behandelt passive und aktive Maßnahmen des Maschinenschutzes, die zur Gewähr eines sicheren Beschleunigerbetriebs notwendig sind.

Dunkelstrom aus der Elektronenquelle hat bei FLASH zur Aktivierung mehrerer Komponenten auf unakzeptable Strahlungsniveaus geführt. Der Transport des Dunkelstroms durch den Linac wird anhand von Simulationen mit einer parallelisierten und erweiterten Version des Programmes Astra untersucht. Mögliche Gegenmaßnahmen werden diskutiert und bewertet.

Strahlverluste können zur Entmagnetisierung von Undulatoren führen, die aus Permanentmagneten aufgebaut sind. Mehrere für FLASH typische Strahlverlustszenarien werden anhand von Schauersimulationen untersucht. Eine Abschirmung wird entworfen und ihre Effizienz wird untersucht. Um den Betrieb des Beschleunigers bei vollem Strahlstrom zu erlauben, muss der mittlere relative Strahlverlust in den Undulatoren auf ein Niveau von etwa 10^{-8} begrenzt werden.

FLASH ist mit einem aktiven Maschinenschutzsystem (MPS) ausgerüstet, das verschiedene Subsysteme und mehr als 80 auf Photomultipliern basierende Verlustmonitore umfasst. Die maximale Reaktionszeit dieses Systems beträgt weniger als 4 µs. Prozeduren und Algorithmen zur Einrichtung und Kalibrierung von MPS-Subsystemen und Komponenten werden vorgestellt, und mit dem Betrieb zusammenhängende Probleme werden behandelt.

Schließlich wird die Architektur eines neuen, durchgängig programmierbaren Maschinenschutzsystems für den XFEL vorgestellt. Mehrere Optionen für die Topologie dieses Systems werden untersucht. Eine Verfügbarkeit von mindestens 0.999 für dieses System ist mit moderaten Hardware-Anforderungen erreichbar.

We have found a strange footprint on the shores of the unknown. We have devised profound theories, one after another, to account for its origins. At last, we have succeeded in reconstructing the creature that made the footprint. And lo! It is our own.

Sir Arthur Eddington, *Space, Time, and Gravitation*

Contents

List of Tables

List of Figures

1 Introduction

Since the advent of the first betatron in 1941 [Ker40, Ker41], particle accelerators have undergone a tremendous development. From few MeV in those early years, the available beam energy will have increased by more than six orders of magnitude with the commissioning of CERN's Large Hadron Collider (LHC) at 7 TeV. Large accelerators primarily serve as colliders for nuclear and elementary particle physics, and as sources of synchrotron radiation. These areas of application have both driven an evolution towards higher beam currents—to increase luminosity in the first case, or the intensity of emitted light in the latter. Table 1.1 illustrates this trend by comparing the stored beam energy in several historic, working and projected circular machines. For linear accelerators (linacs), the average beam power listed in Tab. 1.2 is a more meaningful quantity.

As beams become more powerful, they also become more dangerous to the accelerator transporting them. It has been calculated that the energy stored in the LHC beams will be sufficient to melt one ton of copper [Cho07], and modern electron linacs exceed the power of welding machines by many orders of magnitude. Even before causing mechanical damage, beam losses produce strong radiation fields that can lead to permanent or intermittent failure of electronic components. Over longer periods of time, materials deteriorate in the radiation field—plastics get brittle, optical glasses darken—, and some substances become radioactive themselves, which complicates or prevents necessary maintenance work.

The domain of *machine protection* comprises all measures that protect an accelerator and its infrastructure from the beam. Collimators and lead bricks can be counted among the oldest tools used for this purpose. Today, active electronic systems can provide increased safety by monitoring critical devices, detecting beam losses, and taking the necessary steps to avoid damage—for example, by dumping the beam. Such an active system is usually referred to as *machine protection system* (MPS).

This thesis has been written during the commissioning of the MPS for FLASH, the Free-Electron Laser at Hamburg, and highlights some of the studies needed to facilitate a continuous, safe operation of the machine. After a brief motivation for the use of single-pass free-electron lasers, the layout of the FLASH accelerator and its basic parameters are discussed. The introduction concludes with an outlook on the design of the European X-Ray Free-Electron Laser whose construction is about to start at the time of writing.

In chapter 2, the relevant physical processes for the interaction of an electron beam with matter are outlined. Chapter 3 applies these considerations to the discussion of various beam loss detection techniques. At FLASH, two major problems

Accelerator	Start	Stored beam energy	
AdA	1962	400 µJ†	[Ber63]
Spear	1972	90 J†	[Wie77]
Doris	1974	870 J†	[Koc79]
LEP	1981	56 kJ†	[Bro01]
Tevatron	1987	1.3 MJ (p^+), 140 kJ (p^-)	[Zha03]
Hera	1992	2 MJ (p^+), 29 kJ ($e^\pm$)	[Hof01]
PEP-II	1999	41 kJ (LER), 70 kJ (HER)	[See02]
SNS (accumulator ring)	2006	5 kJ, 24 kJ*	[Hen05]
LHC	2008	362 MJ*†	[Sch06a]

Table 1.1: Stored beam energy of selected storage rings (∗ planned, † per beam)

have become the focus of machine protection measures: First, the loss of dark currents, and second, the damage of undulator magnets. These topics are covered in chapters 4 and 5. Especially the requirements for undulator protection have had a profound impact on the design of the machine protection system described in chapter 6 and on some of the calibration procedures discussed in chapter 7. Finally, chapter 8 presents the concept of an MPS for the European X-Ray Free-Electron Laser.

Before proceeding to the introduction to free-electron lasers, some general remarks on machine protection systems have to be made.

1.1 Tasks and requirements for a machine protection system

To reiterate, the task of a machine protection system is to protect the accelerator and its infrastructure from the beam. This sets it apart from both personnel protection systems and technical interlock systems; the former prevent radiation exposure or physical injury to people, and the latter prevent damage of equipment due to malfunctions. While interfaces between the three system types are common, they have quite different purposes and should therefore not be mixed.

There are three main requirements a machine protection system has to fulfill. In the order of their importance, they can be formulated as follows:

1. Protect accelerator components and devices in the vicinity of the accelerator from direct and indirect damage by the beam
2. Minimize downtime and unneccessary interference with the operation of the accelerator
3. Limit activation of accelerator components to preserve their maintainability

While the first and third points are obvious, the requirement to minimize downtime deserves some elaboration. Beam time at accelerators is expensive and usually

Accelerator	Start	Max. average beam power	
SLC	1989	150 kW	[Ros91]
TTF	1994	8 kW	[Sch08]
Jefferson Lab FEL	2003	1.5 MW	[Nei06]
FLASH	2004	4 kW, 86 kW*	[Fro06a]
SNS (linac)	2006	310 kW, 1.4 MW*	[Hen05]
LCLS	2008	1.7 kW*	[Art02]
European XFEL	2014	600 kW*	[Agh07]
ILC	2020	10.8 MW*†	[ILC07]
Cornell ERL	2020	500 MW*	[Hof08]

Table 1.2: Maximum average beam power of selected linear accelerators (∗ planned, † per beam)

in high demand. Therefore, an MPS should only impair the operation if it is necessary. The system must have a high availability to avoid downtime due to failure of MPS components. Thresholds should not be too tight. And finally, since a lot of time can be lost due to operating errors if the MPS interface is inefficient or complicated, the system must present the reason for any intervention in a clear and concise way.

1.2 Accelerators as light sources

Synchrotron light from particle accelerators is a powerful tool in solid state physics and material science due to its high intensity and broad spectrum, ranging from the infrared to hard X-rays. At first merely a by-product of electron accelerators used for elementary particle research, the unique properties of this radiation have attracted a huge community of users and given rise to a growing demand for higher brilliance of the synchrotron light.

In the last decades, more and more accelerator facilities have either been converted to synchrotron light sources (e.g. Tantalus[1], Doris[2]) or have been constructed solely for that purpose (e.g. ESRF[3], SLS[4], Diamond[5]). These installations of the so-called second and third generation are storage rings that use wigglers and undulators to improve the quality of the generated radiation.

It has become increasingly difficult to improve the peak brilliance of these light sources. The beam energy as a main factor is typically limited by the radius of the

1 Operated by the Midwest Universities Research Association near Stoughton, WI, USA, "Tantalus was the first storage ring operated exclusively for the production of synchrotron radiation, although it was not designed for that purpose." [Lyn97]

2 The "Doppel-Ring-Speicher" at DESY, built in 1974 primarily as an electron-positron collider, was converted into a dedicated synchrotron radiation source in 1993.

3 European Synchrotron Radiation Facility, Grenoble, France

4 Swiss Light Source, Paul-Scherrer-Institut (PSI), Villigen, Switzerland

5 Diamond Synchrotron Light Source, Rutherford Appleton Laboratory, Didcot, UK

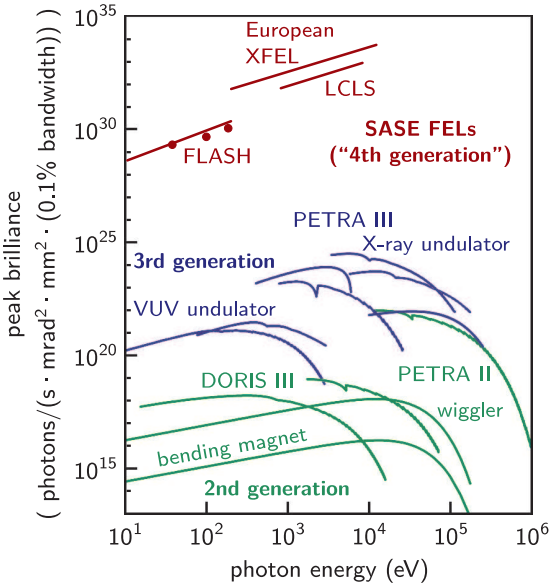

Figure 1.1: Peak brilliance of various light sources. SASE FELs produce radiation with a brilliance several orders of magnitude higher than that of conventional synchrotron light sources.

accelerator because of the scaling of synchrotron radiation power with E^4/r^2. The transverse emittance in storage rings is an equilibrium property that is fixed by the optical lattice; very small emittances can only be reached with large installations and high technical effort, as proposed e.g. for the Petra III light source [Bal08a]. The bunch length is also an equilibrium property; although longitudinal bunch compression to the picosecond regime has been reported [Fei04], this as yet remains a special setup that is incompatible with standard light source operation.

These limitations of second- and third-generation facilities have spurred the development of *Free-Electron Lasers* (FELs) driven by linear accelerators. Linacs give substantially more freedom in choosing time structure and shape of the beam and of the synchrotron light pulses. Equipped with modern injectors, they can produce bright electron beams with extremely low emittances. Furthermore, they provide a simple way of increasing the charge density in the particle bunches by longitudinal bunch compression. While the last two factors themselves help to increase the peak brilliance of a light source, they are also the key to the high gain FEL process that dramatically improves both radiation power and coherence level. As Fig. 1.1 illustrates, FELs exceed classical synchrotron light sources by several orders of magnitude in brilliance.

1.3 Free-electron lasers

The principle of the free-electron laser was first discussed in 1971 when J. M. J. Madey suggested the "stimulated emission of bremsstrahlung in a periodic magnetic field" [Mad71] as a mechanism to generate radiation of high power in an undulator structure. In 1976, the applicability of the theory was demonstrated by a group from Stanford University when the first FEL amplified infrared light of

10.6 µm wavelength [Eli76].

The setup of this first experiment is archetypal for virtually all low-gain FELs: An electron beam traverses a short undulator in which it overlaps with a photon beam. The photon beam is captured in an optical cavity and may initially be supplied by an external source or by the spontaneous undulator radiation. During the passage of electron bunches through the undulator, the light intensity grows by stimulated emission of photons from the relativistic electrons. The optical cavity is needed because each electron bunch amplifies the light only little, typically in the range of few percent—hence the term *low-gain.* This principle works well in the infrared, visible or soft ultraviolet range, but fails for shorter wavelengths due to the lack of adequate mirrors to set up an optical cavity. Hence, for hard ultraviolet and X-ray radiation, a high FEL gain must be obtained within a single pass through the undulator.

1.3.1 High-gain free-electron lasers

The key to free-electron lasers in this spectral range is the *microbunching* process that arranges the electrons into thin longitudinal slices spaced at the radiation wavelength. Radiation emitted from these slices can then interfere constructively and enable a *high-gain*, single-pass FEL. A perfectly microbunched ensemble of N particles would emit coherent radiation with an intensity proportional to N^2, as opposed to spontaneous undulator radiation scaling linearly with N.

Microbunching develops as a consequence of the interaction between the electron and the photon beam. All electrons perform transverse oscillations while passing through the alternating magnetic field of the undulator. The oscillation amplitude of electrons that lose energy to the radiation field grows, and consequently the average longitudinal velocity of these particles decreases. This velocity modulation gradually transforms into a density modulation. To allow the development of this microbunched structure, undulators for high-gain FELs generally have to be much longer than those for low-gain FELs.

In addition, high-gain FELs critically depend on bright electron beams. This can be seen directly from the simplified equations for the case of a monoenergetic electron beam with an energy matching the resonance condition of the undulator [Wie03a]. In this case, the power P of the electromagnetic wave grows exponentially with the longitudinal position z in the undulator,

$$P(z) \propto \exp\left(\frac{z}{L_\mathrm{G}}\right) \quad \text{with} \quad L_\mathrm{G} \propto \gamma \left(\frac{1}{n_\mathrm{e}}\right)^{1/3}.$$

L_G is called the power gain length, γ denotes the relativistic Lorentz factor, and n_e is the number of electrons per unit volume. As this electron density is influenced by transverse and longitudinal properties of the bunch, the gain length is often parametrized as

$$L_\mathrm{G} \propto \gamma \left(\frac{\sigma_\mathrm{t}^2}{I_0}\right)^{1/3},$$

where σ_t is the transverse rms size of the beam, and I_0 is the peak current of the electron bunch. Both a high peak current and a small transverse beam size are needed to render the generation of laser pulses feasible within an undulator of limited length. A detailed introduction to FEL theory can be found in [Doh08].

1.3.2 Self-Amplified Spontaneous Emission (SASE)

As set out above, the FEL process can only be started in the presence of a seeding radiation field. While suitable seed lasers are available for infrared and optical wavelengths, there are no such sources for the hard ultraviolet or X-ray regime. While other seeding schemes have been proposed, the simplest and currently most widely used is the principle of *Self-Amplified Spontaneous Emission* (SASE). In simple terms, it states that the spontaneous undulator radiation created by the electron bunch itself can start the microbunching process provided that the undulator is long enough.

First discussed by A. M. Kondratenko and E. L. Saldin in 1981 [Kon81] and independently derived by R. Bonifacio et al. in 1984 [Bon84], the applicability of the SASE principle was demonstrated almost two decades later in some prototype machines. In the year 2000, the Low Energy Undulator Test Line LEUTL at Argonne National Laboratory showed SASE lasing at 530 nm wavelength [Mil00], which could later be improved to 385 nm. In the same year it was followed by the TESLA Test Facility (TTF) at DESY which produced the first hard ultraviolet radiation with a peak intensity at 109 nm [And00]. Tunability down to 80 nm was demonstrated few weeks later. After a major upgrade, the accelerator achieved first lasing at 32 nm in 2005 and became a user facility [Ayv06]. It was renamed to FLASH in April 2006 and pushed the FEL wavelength record down to 13.5 nm [Ack07] and to 6.5 nm in 2007. Another energy upgrade for FLASH is planned in late 2009.

The demonstration of the first X-ray FEL is expected in the second half of 2009 from the Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory, with a design wavelength of 150 pm [Art02]. At DESY, the European X-Ray Free-Electron Laser (XFEL) project aims at FEL wavelengths down to 100 pm several years later [Agh07].

1.4 FLASH—Free-Electron Laser in Hamburg

At the time of writing, the *Free-Electron Laser in Hamburg*—known by its acronym FLASH or by its former designation TESLA Test Facility (TTF)—is the world's only light source that can produce ultrashort flashes of extreme ultraviolet light with wavelengths as short as 6.5–32 nm and with an energy content of up to 100 µJ per pulse. This capability has attracted a huge community of experimentalists from various disciplines and the facility is therefore expected to continue working for several years even when dedicated X-ray FELs have started routine operation at much shorter wavelengths.

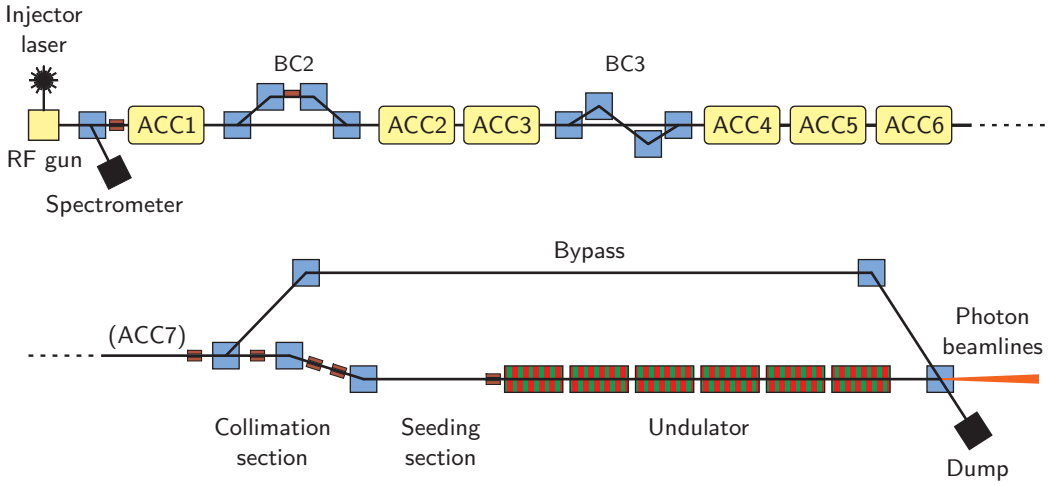

Figure 1.2: Sketch of the FLASH beamline. The elements shown include main dipole magnets (blue), acceleration/rf structures (yellow), collimators (brown), and FEL undulators (green/red).

The TTF was conceived mainly as a testbed for the proposed linear collider TESLA (TeV Energy Superconducting Linear Accelerator) and its superconducting cavity technology. Its goal was to produce an electron beam of high average power and of considerably lower emittance and shorter bunch length than attainable in a storage ring. Work on the superconducting 1.8 K niobium cavities for the accelerator modules started in 1992, and construction of the 100 m long linac commenced about two years later. In the final stage, TTF comprised an rf photocathode electron gun, one cryogenic acceleration module with eight nine-cell cavities fed with an rf frequency of 1.3 GHz, a bunch compressor, and a 15 m long undulator that allowed to demonstrate the SASE FEL process. The accelerator was operated until 2003, when a major upgrade of the machine extended its length to 260 m.

Commissioning of the second phase of the test facility, named TTF2 or Vacuum-Ultraviolet Free-Electron Laser, began in early 2004. The main change was the installation of additional acceleration modules and a longer undulator to allow lasing at shorter wavelengths. In 2006, the machine was officially turned into a user facility and renamed to FLASH.

Figure 1.2 shows the layout of the linac in the end of 2008: its main components are an rf photocathode gun, six acceleration modules similar to the type used in TTF, two bunch compressors, and an undulator composed of six 5 m long segments. In front of the main undulator, space is reserved for a high harmonic gain (HHG) laser seeding option. A bypass line allows to send the beam to the dump without traversing the sensitive undulator structures.

Besides providing photon beams for experiments, FLASH also serves as a pilot machine for the *European X-ray FEL* (XFEL), a 17.5 GeV linac-driven free-electron

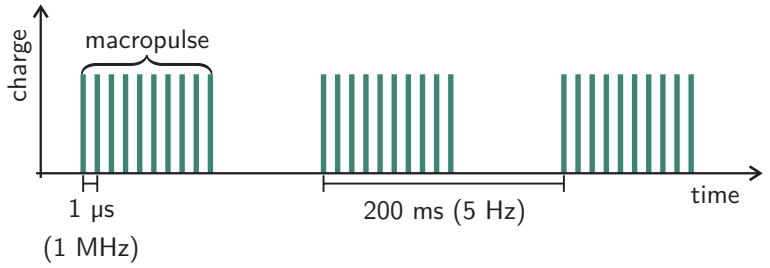

Figure 1.3: Time structure of the FLASH beam. A macropulse or bunch train contains up to 800 individual bunches at a minimum spacing of 1 µs. The 5 Hz repetition rate of the rf system defines the time between successive macropulses. The design provides for up to 7200 bunches per train with a bunch frequency of 9 MHz and a repetition rate of 10 Hz.

laser being built in Hamburg [Wei04]. Notwithstanding its dedication as a user facility, FLASH is also a major testing ground for various novel types of electron beam diagnostics.

1.4.1 Beam production

The FLASH beamline starts with a normal conducting rf gun. Its core part is a 1.5-cell copper cavity for the standard rf frequency of 1.3 GHz. At the backplane of the half cell, an exchangeable plug coated with Cs_2Te serves as a high quantum efficiency photocathode. The photoemission is stimulated by short laser pulses from a widely configurable injector laser system. The most important parameters controlled by this system are

- the number of bunches per macropulse,
- the bunch frequency, and
- the bunch charge.

Typical bunch charges are in the range of 0.5–1 nC, but higher and lower values are possible. The bunch frequency can be set to 1 MHz or to a number of subharmonics (cf. section 6.6) which means that the bunch spacing is 1 µs or a multiple thereof. Because the usable rf pulse length is limited to 800 µs, up to 800 bunches can be produced in quick succession. They form a *macropulse* or *bunch train* as illustrated in Fig. 1.3. The repetition rate of these macropulses is limited to 5 Hz by the rf systems, particularly by the klystron tubes. Future modifications should allow FLASH to be operated with its design parameters of 9 MHz bunch freqency and 10 Hz repetition rate.

A small copper collimator of 1 cm length with a circular inner aperture of 8 mm diameter is located about 120 cm behind the photocathode. This *gun collimator*

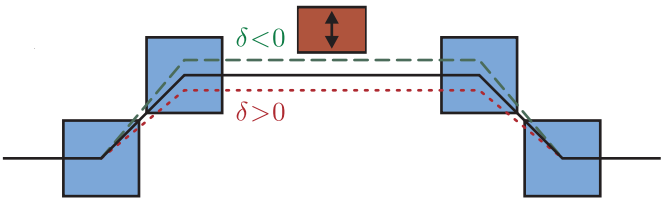

Figure 1.4: Schematic of a 4-bend bunch compressor. Longitudinal compression is possible because high energy particles take a shorter path through the magnetic chicane than low energy particles. At bunch compressor BC2, a movable copper collimator can be inserted to intercept low energy electrons.

prevents a big part of the dark current generated by the rf gun from entering the linac. Its influence is discussed in detail in section 4.5.9.

If the phase of the rf wave relative to the injector laser pulse is correct, the bunches are accelerated away from the cathode and reach an energy of about 5 MeV. After this, they drift into the superconducting cavities of cryomodule ACC1 where they are brought to the relativistic energy of 127 MeV.

1.4.2 Bunch compression

Behind ACC1, the longitudinal charge density of the electron bunch is increased by a magnetic chicane, the bunch compressor BC2. Four identical bending magnets introduce a section of closed dispersion to the linac. Because the deflection of charged particles in dipole fields is energy dependent, electrons with higher energy have a shorter trajectory through the bunch compressor than those with lower energy (Fig. 1.4).

The influence of beam line elements on individual particle coordinates is usually described in a coordinate system that moves with the electron bunch along the trajectory of a reference particle of momentum p_0. In this frame of reference, the momentum p of a particle is expressed as a relative momentum deviation

$$\delta = \frac{p - p_0}{p_0},$$

and the longitudinal distance to the reference particle is denoted by s. The influence of a beam line element on the longitudinal coordinate of a passing particle can be written as

$$\Delta s = \alpha_c \, \delta + O(\delta^2).$$

α_c is called the *momentum compaction* of the beam line element (see e.g. [Wie03]). At standard settings, bunch compressor BC2 has a momentum compaction of $\alpha_c = 180.6\,\text{mm}$ [Flo03a].

In order to decrease the bunch length in the chicane, a longitudinal energy chirp must be imprinted along the bunch; particles in the tail should have a higher energy than those in the head. This kind of chirp is achieved by accelerating the electron

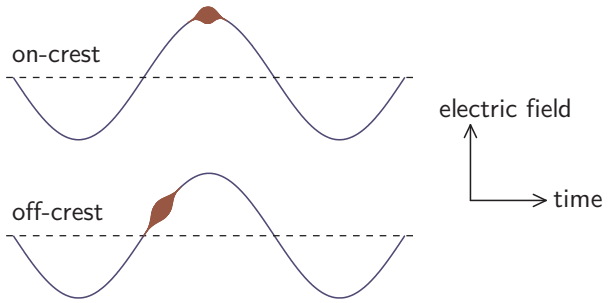

Figure 1.5: On- and off-crest acceleration. If the phase of the rf wave with respect to the bunch arrival is zero, the bunch is accelerated *on-crest*, i.e. on the crest of the rf wave. If the bunch arrives on a slope of the rf wave *(off-crest)*, a longitudinally correlated energy spread is imprinted along the bunch.

bunch not on the crest, but on the rising slope of the rf wave. As illustrated in Fig. 1.5, *off-crest* operation increases the energy spread of the particle ensemble while lowering the total energy transfer from the rf wave as compared to *on-crest* acceleration. In typical FEL runs, ACC1 is operated about 8–10° off-crest.

BC2 is also equipped with a movable collimator, the *BC2 collimator* or *BC2 scraper*. It is a copper block of 35 cm length that can be moved horizontally into the vacuum chamber of bunch compressor BC2 as shown in Fig. 1.4. It can thus be adjusted to intercept a variable part of the low-energy electrons. From the horizontal dispersion between the inner dipoles of [Cas03]

$$D_x = \frac{\mathrm{d}x}{\mathrm{d}\delta} = \frac{\mathrm{d}x}{\mathrm{d}E/E} \approx 34\,\mathrm{cm}$$

and the fixed beam energy of $E \approx 127\,\mathrm{MeV}$, the energy calibration for this collimator can be calculated to

$$\frac{\mathrm{d}E}{\mathrm{d}x} = \frac{E}{D_x} \approx 374\,\mathrm{keV/mm}.$$

It is mainly used to remove gun dark current and low energy tails from bunches accelerated off-crest.

In the following two acceleration modules, ACC2 and ACC3, the electron energy is typically increased to 380–450 MeV at an off-crest phase of 20–30°. Afterwards, further longitudinal compression takes place in bunch compressor BC3 with $\alpha_\mathrm{c} = 49\,\mathrm{mm}$ [TTF02].

1.4.3 Collimation

The acceleration modules ACC4–6 bring the beam to the desired end energy. About 1 GeV can be reached by on-crest acceleration in all modules. The maximum practical beam energy for an FEL run is about 950 MeV because of the big off-crest phase in ACC2–3.

Behind the acceleration modules, the electron bunches pass the main collimation system. It consists of four copper blocks of 50 cm length with four bore holes each. To the beam, these holes represent circular apertures of 4, 6, 12, and 34 mm diameter. Towards the entrance and exit of the collimator, the bore diameter increases slightly to minimize wakefield effects. With a motorized table, any of the apertures can be selected by bringing the respective hole in line with the beam axis. The first two of these copper blocks serve as transverse collimators by intercepting electrons with too large betatron amplitudes. The second two are located in a dispersive section between two dipole magnets and can therefore provide energy collimation. In the standard setup, the transverse collimators are used with 4 mm diameter, the first energy collimator with 6 mm, and the second one with 4 mm. This results in an energy acceptance of about ±3 % [Bal03].

1.4.4 Undulators and dump

A few meters upstream of the first FEL undulator, a copper block of 14 cm length protects the transition to a smaller beam pipe diameter. It absorbs most of the radiation energy released by electrons with excessive transverse offsets. A detailed discussion of this *undulator absorber* is found in chapter 5.

After passing the absorber, the beam enters the main undulator section consisting of six 4.5 m long undulator segments. The undulators are planar permanent magnet devices with a fixed gap of 12 mm and a magnetic field amplitude of 486 mT [Pfl03a, Pfl04]. Behind this section, the electron beam is separated from the photon beam by a strong vertically deflecting dipole. The electron beam is disposed of in a beam dump able to withstand an average power of more than 100 kW [Mas99].

1.5 XFEL—European X-Ray Free-Electron Laser

The next big accelerator to be built (at least partly) on the DESY site is the *European X-Ray Free-Electron Laser* (XFEL). The superconducting linear accelerator will bring an electron beam to the energy of 17.5 GeV and use it to generate extremely brilliant pulses of spatially coherent X-rays in an array of undulators using the SASE process. In many respects, the design of the XFEL is based on experience from FLASH.

Figure 1.6 shows a simplified schematic of the XFEL beamline. The accelerator will be built in an approximately 3.4 km long underground tunnel accomodating the following major sections:

- photoinjector(s)
- beam conditioning sections
- main linear accelerator (linac)
- beam distribution system

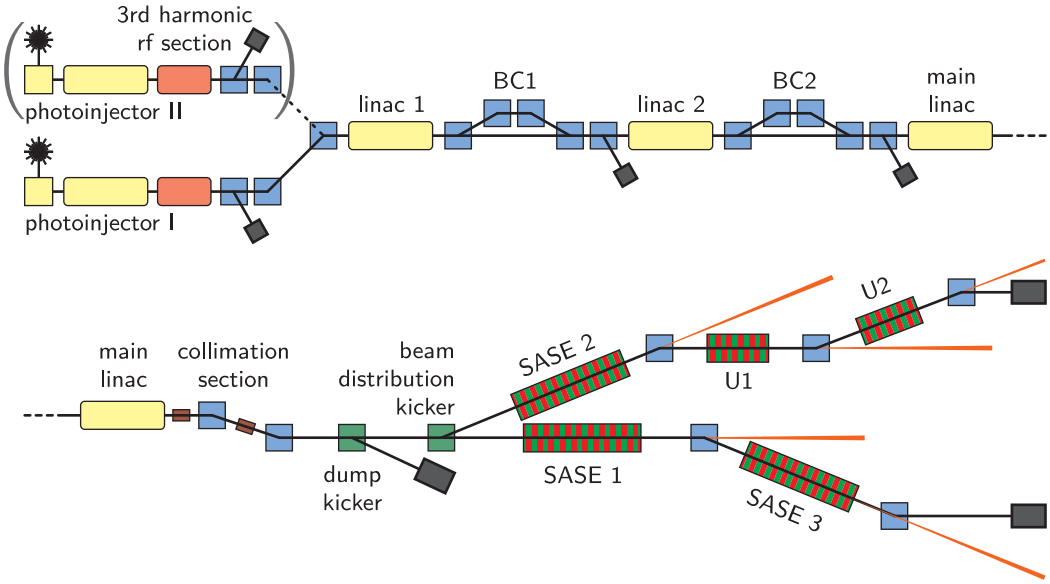

Figure 1.6: Overview of the XFEL beamline. Yellow boxes indicate 1.3 GHz rf sections, red boxes 3.9 GHz rf sections. The sections designated *SASE 1–3* are undulator lines long enough to support the SASE FEL process, *U1–2* are comparatively short undulator lines for the generation of spontaneous radiation.

- undulators
- photon beamlines
- experimental stations

The laser-driven photoinjector consists mainly of a normal-conducting rf gun and a superconducting acceleration module for the rf frequency of 1.3 GHz, producing electron bunches with an energy of 120 MeV. An optional second photoinjector would facilitate maintenance work or studies without interruption to the accelerator operation. The designated usable rf pulse length for the XFEL is 600 µs; at a maximum bunch frequency of 5 MHz, this allows bunch trains of up to 3000 bunches.

After further acceleration, the bunches are longitudinally compressed in the two bunch compressors BC1 and BC2 at energies of 500 MeV and 2 GeV. A number of cavities operated at 3.9 GHz (the third harmonic of the main rf frequency) allow to linearize the longitudinal phase space and to obtain a more homogeneous compression.

In the main linac, the beam is brought to the final energy of 17.5 GeV; an upgrade to 20 GeV is already foreseen in the original design. Considering the pulse repetition rate of 10 Hz and the nominal bunch charge of 1 nC, this results in a

maximum beam power of

$$20\,\frac{\mathrm{GeV}}{e}\cdot 1\,\mathrm{nC}\cdot 600\,\mu\mathrm{s}\cdot 5\,\mathrm{MHz}\cdot 10\,\mathrm{Hz} = 600\,\mathrm{kW}.$$

Behind the last acceleration module, transverse and energy collimation take place in a manner similar to FLASH. The next active component is a fast kicker that can send selected bunches from the macropulse into a beam dump.

The remaining bunch train is distributed between two (or possibly more) undulator lines by a comparatively slow beam distribution kicker with a rise time of the order of 10 µs. Each of the undulator lines ends in an electron beam dump. A peculiarity of the accelerator design is that each of the three main beam dumps can withstand only half of the nominal beam power, i.e. 300 kW. This limitation has been introduced to avoid the necessity for more complicated systems like liquid dumps. Hence, the protection of the beam dumps is an additional task for the machine protection system.

Further details of the XFEL facility and its technical systems can be found in the technical design report [Agh07].

2 Hazards

The goal of machine protection at an accelerator is to prevent any kind of damage by the beam, or, in a wider sense, by particles accelerated or stored within the machine. This includes damage caused by direct interaction of the particles with matter as well as effects in the generated radiation field. In the following sections, the basic physical mechanisms relevant for electron and positron accelerators will be discussed. Effects for heavy particles like muons or protons are explicitly excluded.

2.1 Interaction of electrons with matter

Electrons passing through matter generally experience a deflection from their original direction, and a loss of kinetic energy. Several processes contribute to both effects, most importantly

1. bremsstrahlung
2. inelastic scattering with atomic electrons
3. elastic scattering with nuclei.

All of these phenomena are caused by Coulomb interaction of the projectile with the atoms of the target material. In particular, bremsstrahlung is the emission of a photon in the wake of Coulomb scattering—or, in the classical picture, the emission of radiation by a decelerated charge.

Inelastic scattering requires that some of the electron's kinetic energy is converted into internal energy of the target atom, usually causing electronic excitation or ionization. Although even nuclear excitation may be caused, the cross section for this process is extremely small and it can be safely neglected.

In an *elastic* scattering event with a nucleus, on the other hand, the projectile loses only a tiny fraction of its kinetic energy because of the huge mass difference between both collision partners. Even if the absolute value of the particle momenta remains virtually unchanged, elastic Coulomb scattering in the lattice of the target material can cause significant deflection from the incident direction of the electrons. The classical Rutherford formula states that the probability of deflection by an angle θ in a single scattering event is proportional to $\sin^{-4}(\theta/2)$. The treatment of multiple Coulomb scattering is rather complex and hardly needed for the evaluation of hazards in the operation of an accelerator. A description of the expected angular distribution is found in [Lyn91].

The deposition of energy in matter by the first two processes, bremsstrahlung and inelastic scattering, is of much greater practical importance. Figure 2.1 shows the contributions to the energy loss of an electron traveling through various materials. Inelastic scattering is most important at low energies while emission of bremsstrahlung dominates the high energy region. Together with a brief discussion of these processes, some practically usable formulas for the quantitative calculation of energy losses will be given.

2.1.1 Inelastic scattering

The primary process by which *heavy* particles lose energy in matter is inelastic scattering with atomic electrons. Depending on the amount of energy transferred, it leads to electronic excitation of the atom or to ionization. For very light charged particles such as electrons and positrons, inelastic scattering is most important at low energies where radiative processes are still weak.

A quantum-mechanical theory of the deceleration of charged particles in matter has been derived by Bethe and Bloch in the 1930s [Bet30, Blo33]. Since then, several empirical corrections have been added to account for various effects. An accurate evaluation of the energy loss is rather difficult especially for the case of electrons and positrons. For practical purposes, a fit like the following from [Sel84] is often used:

$$-\frac{1}{\rho}\frac{\mathrm{d}E}{\mathrm{d}x} = \frac{1}{\beta^2}\frac{Z}{A}B(T)\cdot 0.153536\,\frac{\mathrm{MeV\,cm^2}}{\mathrm{g}} \tag{2.1}$$

This equation gives the so-called *mass collision stopping power*, i.e. the mean energy loss per track length, normalized to the density ρ of the medium. The kinetic energy of the particle is denoted by T, β is its velocity in units of c, Z the number of electrons in the atom or molecule, A the atomic or molecular weight in units of g/mol. The *stopping number* $B(T)$ is defined piecewise:

$$B(T) = \begin{cases} B_0(T) + b_0 - b_4\left(\dfrac{p}{mc}\right)^2 & \text{for} \quad T \leqslant T_0 \\ B_0(T) + 1 - 2\ln\dfrac{p}{mc} + b_1 - b_2\left(1 - \dfrac{2\ln\frac{p}{mc}}{b_3}\right)^k & \text{for} \quad T_0 < T < T_1 \\ B_0(T) + 1 - 2\ln\dfrac{p}{mc} + b_1 & \text{for} \quad T \geqslant T_1 \end{cases}$$

Here, b_0 to b_4, k, T_0 and T_1 are material constants. Tab. 2.1 shows an excerpt of the comprehensive list of values found in [Sel84]. With p denoting the particle momentum and m its rest mass,

$$\frac{p}{mc} = \left(\frac{1}{\beta^2} - 1\right)^{-\frac{1}{2}}$$

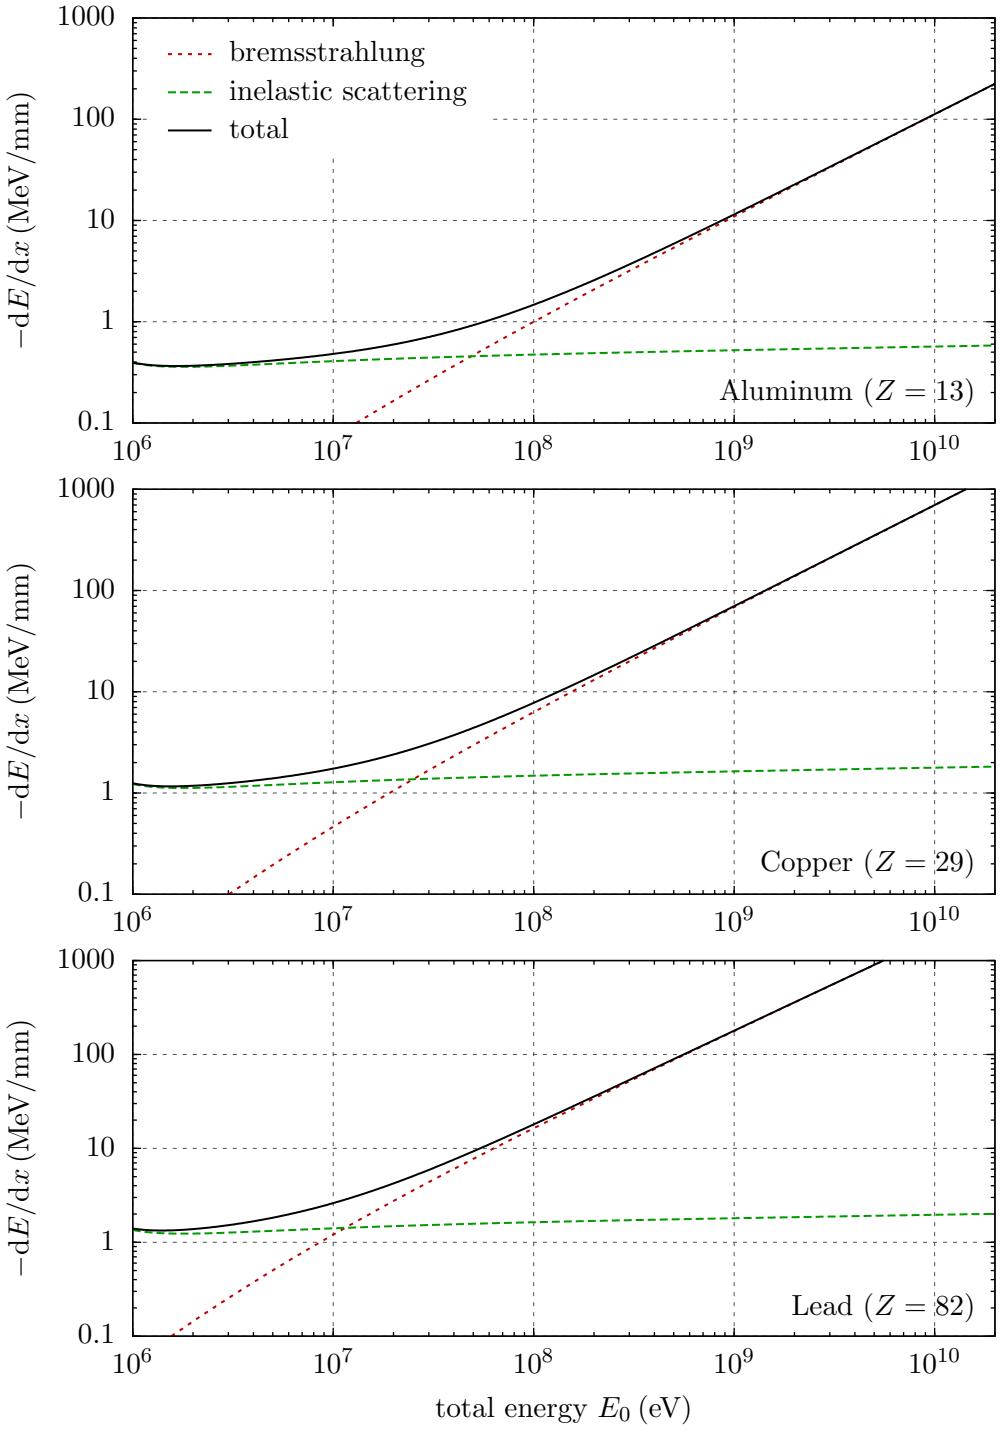

Figure 2.1: Energy loss by electrons in aluminum, copper, and lead as a function of total electron energy.

Material	T_0	T_1	b_0	b_1	b_2	b_3	b_4	k
Aluminum	0.403	525.6	16.0643	19.3038	4.4172	13.8739	0.05464	3.6345
Iron	0.211	726.5	14.9763	18.2673	4.4116	14.5205	0.12065	2.9632
Copper	0.191	971.4	14.7392	18.1581	4.5136	15.1013	0.08993	2.9044
Niobium	0.414	847.7	14.2221	18.2362	5.1677	14.8289	0.06153	3.0930

Table 2.1: Material constants for calculating the stopping power for electrons and positrons according to (2.1). T_0 and T_1 are given in MeV, the other constants are unitless. Data from [Sel84] for pure materials under standard conditions. For non-standard densities, minor deviations have to be expected.

may be substituted. Finally, $B_0(T)$ is a function of only the kinetic energy of the electron or positron:

$$B_0(T) = B_0(\tau mc^2) = \ln \frac{\tau^2(\tau+2)}{2} + B_0'(\tau)$$

$$B_0'(\tau) = \begin{cases} \dfrac{1 + \frac{\tau^2}{8} - (2\tau+1)\ln 2}{(\tau+1)^2} & \text{for electrons} \\ 2\ln 2 - \dfrac{\beta^2}{12}\left(23 + \dfrac{14}{\tau+2} + \dfrac{10}{(\tau+2)^2} + \dfrac{4}{(\tau+2)^3}\right) & \text{for positrons} \end{cases}$$

As charged particles pass through matter, their electric field polarizes the surrounding atoms. A consequence of this is the so-called *density effect*: The polarization effectively shields distant electrons from the field of the particle, and their contribution to the total energy loss is reduced. Equation (2.1) and the material constants in Tab. 2.1 describe this behavior correctly for standard densities of the materials; for non-standard conditions, results may deviate. However, for solids, density variations are comparatively small, so that the influence on the stopping power can be neglected for simple calculations.

Figure 2.1 shows that the energy loss by inelastic scattering has a minimum at low electron energies and grows only slowly with increasing electron energy E_0. It can be inferred from the high-energy limit of (2.1) that $\mathrm{d}E/\mathrm{d}x$ scales logarithmically with E_0. The dependence on the atomic number Z is roughly linear. On a final note, the calculation of stopping powers for electrons is often facilitated by use of the *estar* online database [Ber05].

2.1.2 Bremsstrahlung

In the classical picture, it can be stated that when an electron collides with an atom, both collision partners emit radiation—*bremsstrahlung*—because they are subject to a certain amount of acceleration. Since the mass of the electron is at least three orders of magnitude lower than that of the nucleus, the acceleration of the atom is tiny and its contribution to the radiation field can be neglected. From

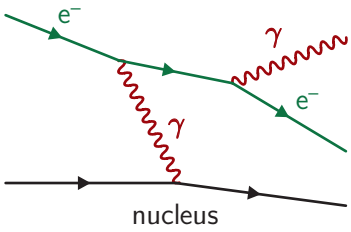

Figure 2.2: Feynman diagram for the emission of a bremsstrahlung photon by an electron scattered at an atomic nucleus

the same argument, it can be understood that energy loss by bremsstrahlung does not play a major role for heavy particles at energies below a few hundred GeV.

To obtain quantitative statements on radiation spectrum and energy loss, the cross section for the emission of a bremsstrahlung photon (as illustrated in Fig. 2.2) has to be calculated. Most analytical formulas found in the literature are approximations that are valid in only a restricted energy range. For relativistic electrons and positrons of at least 5 MeV, [Leo94] gives the differential bremsstrahlung cross section as

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\nu}(\nu, E_0) = \frac{4Z(Z+1)r_e^2\alpha}{\nu}\left\{\left(1+\frac{(E_0-h\nu)^2}{E_0^2}\right)\left(\frac{\phi_1(\xi)}{4}-\frac{\ln Z}{3}-f(Z)\right)\right.$$
$$\left.-\frac{2}{3}\frac{E_0-h\nu}{E_0}\left(\frac{\phi_2(\xi)}{4}-\frac{\ln Z}{3}-f(Z)\right)\right\}.$$

$\alpha \approx 1/137$ is the fine structure constant, r_e denotes the classical electron radius. The function $f(Z)$, the so-called *Coulomb correction,* is given by

$$f(Z) = \alpha^2 Z^2\left(\frac{1}{1+\alpha^2Z^2}+0.20206-0.0369\alpha^2Z^2+0.0083\alpha^4Z^4-0.002\alpha^6Z^6\right) \tag{2.2}$$

for elements up to uranium [Dav54].

$\phi_1(\xi)$ and $\phi_2(\xi)$ describe the *screening effect* which deserves a brief discussion. Although the atomic electrons scarcely contribute as collision partners in the bremsstrahlung emission process, they have a major influence because they screen the projectile from the electric field of the nucleus. For an atom with atomic number Z, this effect is quantified by a *screening parameter*

$$\xi(\nu, E_0) = \frac{100 m_e c^2 h\nu}{(E_0-h\nu)E_0 Z^{\frac{1}{3}}}$$

that depends on the frequency ν of the emitted photon and on the total energy E_0 of the electron before the collision [Koc59]. There is no screening for $\xi \gg 1$, complete screening takes place for $\xi \approx 0$. As Fig. 2.3 shows, complete screening applies for most photon energies except for very high ones. In the limit of $h\nu = E_0 - m_0c^2$, where the electron loses all of its kinetic energy in a single photon emission, the screening parameter converges against a constant value:

$$\lim_{\nu\to\nu_{\max}} \xi(\nu, E_0) = \frac{100}{Z^{\frac{1}{3}}}\left(1-\frac{m_e c^2}{E_0}\right), \qquad \nu_{\max} = \frac{E_0 - m_e c^2}{h}$$

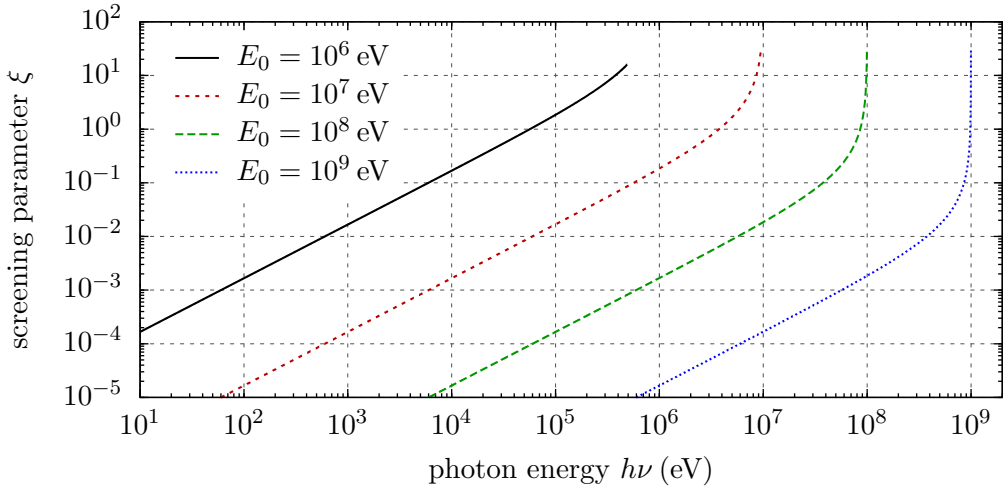

Figure 2.3: Bremsstrahlung screening parameter ξ for various electron energies in copper. No screening of the electric field of the nucleus is effective for $\xi \gg 1$, complete screening takes place for $\xi \approx 0$.

To conclude the discussion of the cross section formula, [Leo94] gives the following approximations for the screening functions:

$$\begin{aligned}\phi_1(\xi) &= 20.863 - 2\ln\bigl(1 + (0.55846\xi)^2\bigr) \\ &\quad - 4\bigl(1 - 0.6\exp(-0.9\xi) - 0.4\exp(-1.5\xi)\bigr)\end{aligned} \tag{2.3}$$

$$\phi_2(\xi) = \phi_1(\xi) - \frac{2}{3 + 19.5\xi + 18\xi^2}. \tag{2.4}$$

The approximations are valid for boron and heavier elements.

Figure 2.4 shows the differential cross section for various electron energies. The displayed quantity is actually $\mathrm{d}\sigma/\mathrm{d}\nu$ multiplied with the frequency ν of the emitted photon. Considering that $\mathrm{d}\sigma/\mathrm{d}\nu$ is proportional to the emission probability of a photon of energy $h\nu$, the plotted cross section is proportional to the bremsstrahlung power spectrum. The maximum energy of bremsstrahlung photons is the kinetic energy of the electron. The spectrum is almost flat on a logarithmic frequency scale; the cross section only decays to zero in the upper two orders of magnitude, starting at about $\nu_{\max}/100$. The low frequency part of the spectrum is also independent of the electron energy.

The energy loss can now be obtained by integrating over the product of photon energy and differential cross section for the allowed range of frequencies:

$$\frac{\mathrm{d}E}{\mathrm{d}x}(E_0) = -\frac{\rho N_a}{A^*}\int_0^{\nu_{\max}} h\nu\,\frac{\mathrm{d}\sigma}{\mathrm{d}\nu}(\nu, E_0)\,\mathrm{d}\nu \tag{2.5}$$

The quantity $\rho N_a/A^*$ is the number of atoms per volume, calculated from density ρ and atomic mass A of the material. While the cross section scales mainly with

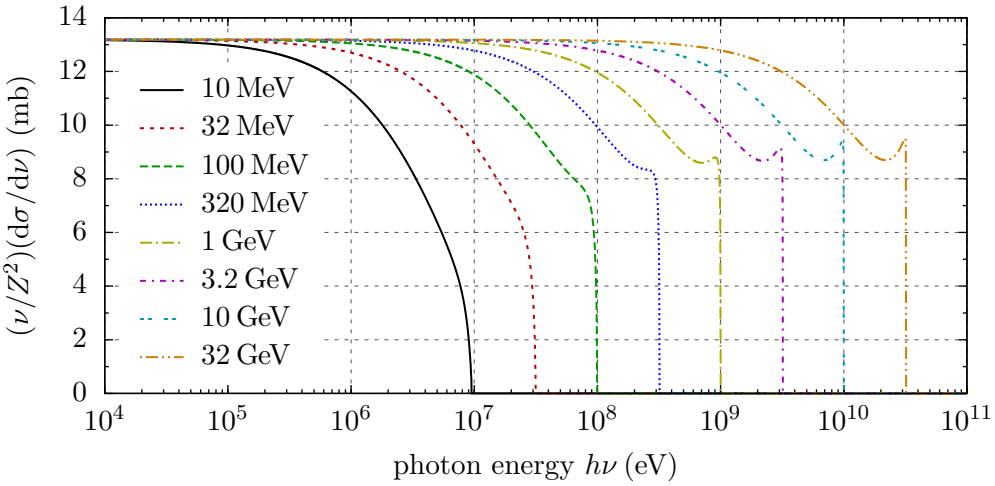

Figure 2.4: Cross section for bremsstrahlung emission vs. photon energy in millibarn, for various total electron energies. The plotted quantity is the differential bremsstrahlung cross section $\mathrm{d}\sigma/\mathrm{d}\nu$ multiplied with the frequency ν of the emitted photon (and normalized to $Z = 1$). It is proportional to the power spectrum.

$1/\nu$, the screening parameter ξ is also ν-dependent and complicates an analytical evaluation of (2.5). Analytic solutions have been derived only for the extreme cases of complete or no shielding and can be found in [Leo94]. Usually, the integration is performed numerically.

Because bremsstrahlung is so ubiquitous at electron accelerators, it is often desirable to make simple shielding calculations without using the cumbersome relations just described. For this purpose, two quantities find wide application:

Critical energy: As shown in Fig. 2.1, the energy loss by bremsstrahlung is almost proportional to the total electron energy. This explains why radiative losses soon exceed those by inelastic scattering which scale only logarithmically with E_0. The point where both losses are equal is usually referred to as the *critical energy* of the material; it is about 26 MeV for copper, 11 MeV for lead. A widely quoted formula for the critical energy is given by [Ber64]:

$$E_c \approx \frac{800\,\mathrm{MeV}}{Z + 1.2}$$

Radiation length: Radiative losses scale like $\mathrm{d}E/\mathrm{d}x \approx E \cdot const$ as long as the particle is well above the critical energy. Hence, it is possible to write

$$E(x) \approx E_0 \exp\left(-\frac{x}{L_{\mathrm{rad}}}\right).$$

Material	L_{rad} (cm)	X_0 (g/cm^2)
Aluminum	8.90	24.01
Titanium	3.56	16.17
Iron	1.76	13.84
Copper	1.43	12.86
Tungsten	0.35	6.76
Lead	0.56	6.37

Table 2.2: Radiation length of various materials; data from [PDG06].

The quantity L_{rad} is a material constant called the *radiation length*. It specifies the distance after which the energy of an ultrarelativistic electron has decreased to $1/e$ of its initial value. It should be noted that in the literature, a value $X_0 = L_{\mathrm{rad}} \cdot \rho$ including the density of the material is often referred to as the radiation length. Table 2.2 shows L_{rad} and X_0 for some materials. A convenient method of calculating X_0 to a precision of few percent is given by [PDG06]:

$$X_0 = \frac{A}{Z(Z+1)\ln(287Z^{-0.5})} \cdot 716.4 \,\frac{\mathrm{g}}{\mathrm{cm}^2}$$

2.2 Interaction of photons with matter

It has already been stated that the energy loss by high energy electrons is dominated by the emission of radiation. In a sufficiently thick target, a photon beam develops and co-propagates with the electrons. Of course, these photons can also interact with the material. The most important processes are

1. photoelectric effect
2. Compton scattering
3. pair production
4. photonuclear reactions.

The first two effects lead to ionization of the material—in Compton scattering, the incident photon transfers a part of its energy to an atomic electron, in the photoelectric effect it is absorbed completely—while pair production creates a positron and an electron from a γ-ray of sufficient energy. Compared to these three effects, photonuclear reactions are extremely rare. Their main importance lies in the creation of free neutrons.

To understand the importance of these effects in various energy ranges, it is helpful to calculate and compare their total cross sections. For aluminum, copper, and lead, the results are shown in Fig. 2.5.

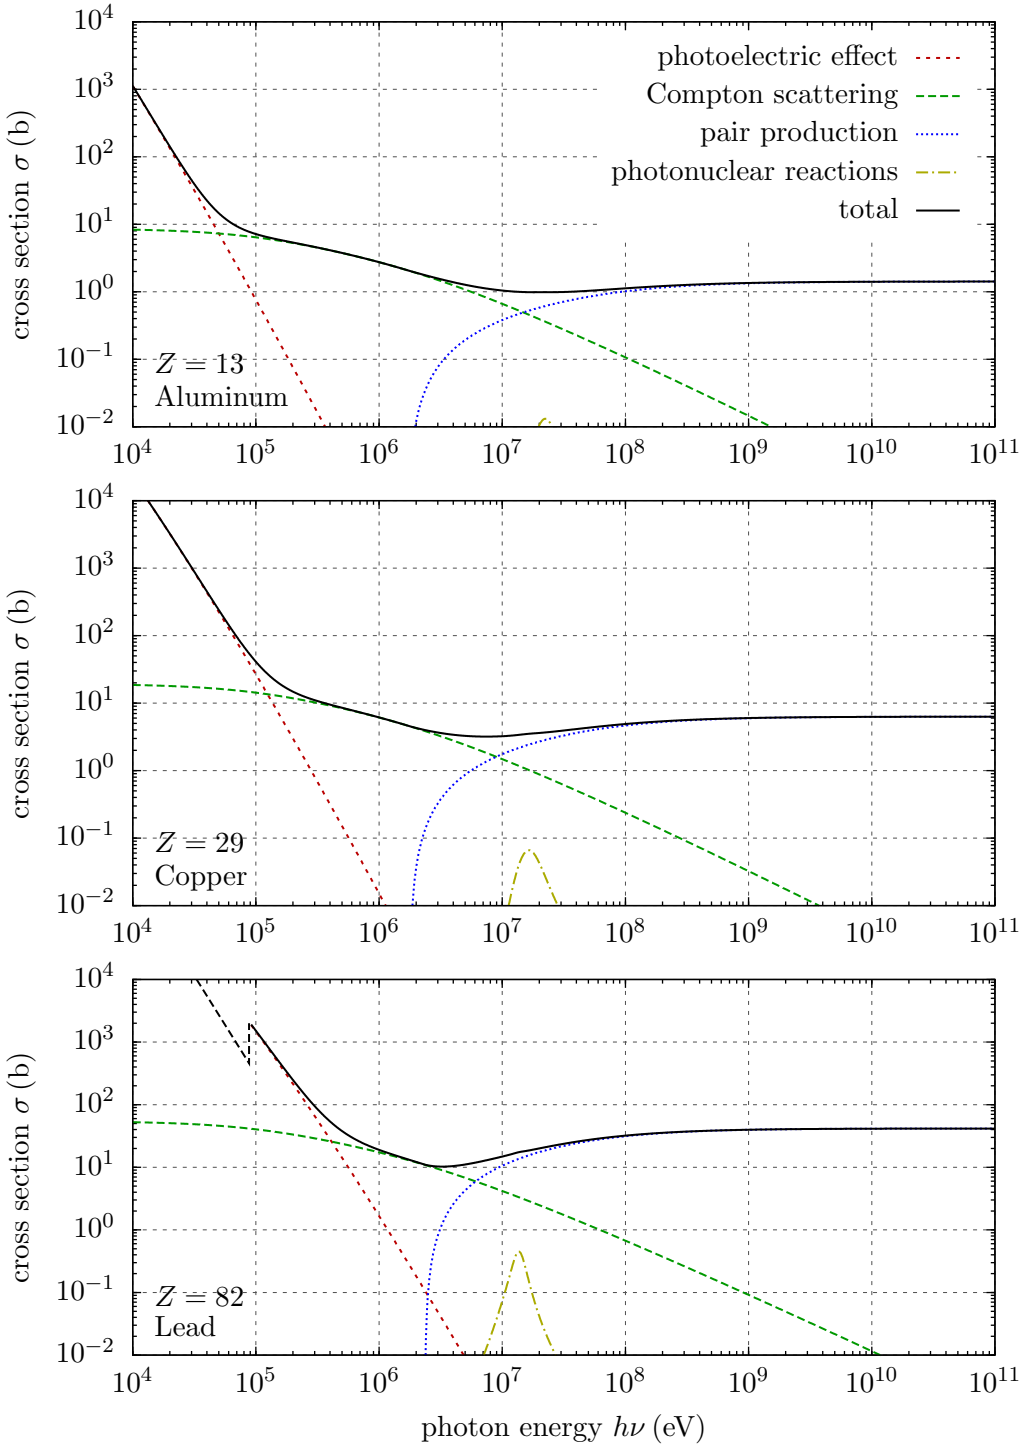

Figure 2.5: Total cross sections for photonic interactions in aluminum, copper, and lead. In the last plot, the K edge of lead is visible at 88 keV. For lower energies, the photoelectric cross section formula (2.6) is not valid anymore.

2.2.1 Photoelectric effect

Low-energy photons have a high probability of being absorbed by an atomic electron. If the photon energy is higher than the binding energy E_b, the electron is released with a kinetic energy of $h\nu - E_b$ while the nucleus receives a minor recoil momentum.

The binding energy depends on the atomic orbital the electron is occupying. Therefore, the photoelectric absorption spectrum has several sharp edges that mark the energies at which the electrons of a certain orbital can contribute to the total absorption. The electrons with the highest binding energy are those from the *1s* shell, often called K electrons in spectroscopy. For photons above this binding energy—above the *K edge*—, all atomic electrons can contribute to photoelectric absorption and the total cross section per atom can be written in closed form as follows [Leo94]:

$$\sigma(\nu) = \frac{8}{3}\pi r_e^2 \cdot \frac{2^7\pi}{\alpha^3 Z^2}\left(\frac{\nu_k}{\nu}\right)^4 \cdot \frac{\exp\bigl(-4\xi\arctan(1/\xi)\bigr)}{1-\exp\bigl(-2\pi\xi\bigr)} \tag{2.6}$$

with

$$\nu_k = \frac{(Z-0.03)^2 m_e c^2 \alpha^2}{2h}, \qquad \xi = \sqrt{\frac{\nu_k}{\nu-\nu_k}}.$$

The K edge is found at about 7 keV for iron, 9 keV for copper, and 88 keV for lead.

2.2.2 Compton scattering

The Compton process describes the elastic scattering of a photon at a free electron. In matter, bound electrons can be regarded as quasi-free as long as their binding energy is small with respect to the energy $h\nu$ of the incident photon. The outgoing photon is deflected by an angle θ and has a lower energy of $h\nu'$. This energy difference is most compactly described in terms of the wavelength of the photons by the well-known formula

$$\lambda' - \lambda = \frac{h}{m_e c}\left(1-\cos\theta\right).$$

From this relation, it can be shown that the maximum energy loss allowed by the kinematics of the process is

$$h\nu - h\nu' = \frac{2h\nu}{\frac{m_e c^2}{h\nu}+2}$$

which corresponds to backscattering into the direction of incidence. In order to classify the importance of the Compton effect in various energy ranges, the total cross section can be calculated from [Leo94]

$$\sigma(\epsilon) = 2\pi r_e^2\left\{\frac{1+\epsilon}{\epsilon^2}\left(\frac{2(1+\epsilon)}{1+2\epsilon}-\frac{\ln(1+2\epsilon)}{\epsilon}\right)+\frac{\ln(1+2\epsilon)}{2\epsilon}-\frac{1+3\epsilon}{(1+2\epsilon)^2}\right\}$$

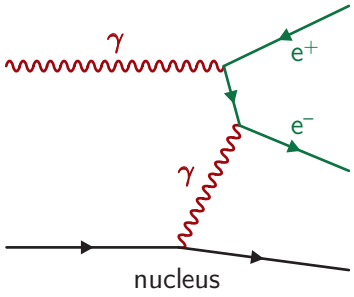

Figure 2.6: Feynman diagram for electron-positron pair production

with $\epsilon = h\nu/(m_e c^2)$. From the plot in Fig. 2.5, it can be seen that Compton scattering is most important at moderate photon energies. High energy photons have a much greater probability to take part in pair production events.

2.2.3 Pair production

Pair production is the creation of a particle-antiparticle pair from a photon. The presence of another charge, usually an atomic nucleus, is required to guarantee momentum conservation (Fig. 2.6). The energy of the incident photon must be at least as big as the sum of the rest masses of the products, i.e. 1.022 MeV for an electron-positron pair, or 211 MeV for a muon-antimuon pair. For practical purposes, the discussion can be restrained to e^+/e^- events which are far more probable than the production of heavier particles. However, it should be remarked that especially muons have to be considered for personnel radiation protection because of their long range.

As indicated by the respective Feynman diagrams, the pair production process is closely related to bremsstrahlung emission. The formula for the differential cross section (from [Leo94]) is therefore very similar:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_\pm}(\nu, E_+, E_-) = \frac{4Z(Z+1)r_e^2\alpha}{(h\nu)^3}\left\{({E_+}^2 + {E_-}^2)\left(\frac{\phi_1(\xi)}{4} - \frac{\ln Z}{3} - f(Z)\right) + \frac{2}{3}E_+E_-\left(\frac{\phi_2(\xi)}{4} - \frac{\ln Z}{3} - f(Z)\right)\right\} \tag{2.7}$$

The screening functions ϕ_1, ϕ_2 and the Coulomb correction f are the same as for bremsstrahlung ((2.2), (2.3), and (2.4)). The screening parameter has to be redefined in terms of the total energy of the outgoing positron (E_+) and electron (E_-):

$$\xi(\nu, E_+, E_-) = \frac{100 m_e c^2 h\nu}{E_+E_-Z^{\frac{1}{3}}}$$

If the minor kinetic energy transfer to the nucleus is neglected, the electron and the positron share the energy of the photon: $h\nu \approx E_+ + E_-$. With this assumption, it is possible to obtain the total cross section for pair production from numerical integration of (2.7).

Isotope	$E_{\min}$ (MeV)	E_0 (MeV)	$E_{\max}$ (MeV)	Γ_{fwhm} (MeV)	$\hat{\sigma}$ (mb)	$\int \sigma$ (mb·MeV)
^{12}C	18.72	∼23	37.4	∼5	∼7	46.8
^{27}Al	13.03	∼22	36.7	∼11	∼14	167
$^{\text{nat}}$Cu	9.91	∼17	27.8	∼7	∼70	635
^{93}Nb	8.81	16.59	24.3	5.05	200	1331
^{208}Pb	7.38	13.46	26.4	3.90	491	2646
$^{\text{nat}}$Pb	6.73		22.0			4100

Table 2.3: Parameters for neutron production by giant dipole resonance, from [Swa79]. $E_{\min}$—threshold energy, E_0—resonance energy, $E_{\max}$—upper integration limit, Γ_{fwhm}—full resonance width at half maximum, $\hat{\sigma}$—maximum resonance cross section, $\int \sigma$—integrated cross section. Values for E_0, Γ_{fwhm}, $\hat{\sigma}$ are the result of Lorentzian fits to the resonance curve unless marked as approximate. In this case, only visual approximations are given because "the giant dipole resonance is fragmented into considerable structure for nuclei lighter than ^{60}Ni" [Ber75].

Figure 2.5 shows that pair production dominates the attenuation of photon beams in matter above an energy of about 10 MeV. Since the cross section scales approximately with Z^2, heavy elements like lead are the most favorable choice for shielding against high energy X-ray beams. In the energy range dominated by pair production, there is also a useful approximation for the mean free path of the photons:

$$L_{\text{pair}} \approx \frac{9}{7} L_{\text{rad}}$$

The typical path length a photon can travel in matter until it is consumed in a pair production event is about 30 % higher than the radiation length of the material.

2.2.4 Photonuclear reactions

Although the attenuation of a photon beam is completely dominated by interactions with the electrons of the target material, also reactions with the atomic nuclei are quite common above a threshold energy of few MeV. The most important one is known as *giant dipole resonance*, and usually described as an induced collective oscillation of the nuclear protons against the neutrons. This excited state typically relaxes by emitting one or more so-called *photoneutrons.* The resonance energy for these (γ,n) reactions lies between about 13 and 26 MeV for most elements. It decreases slowly with the mass of the nucleus as $A^{-1/3}$ [May02], although there are big deviations from this rule especially for light nuclei. The typical width of the resonances lies in the range of 3–6 MeV and the total cross section reaches several hundred millibarns for heavy elements. An excerpt from the comprehensive data collection given in [Swa79] is shown in Tab. 2.3.

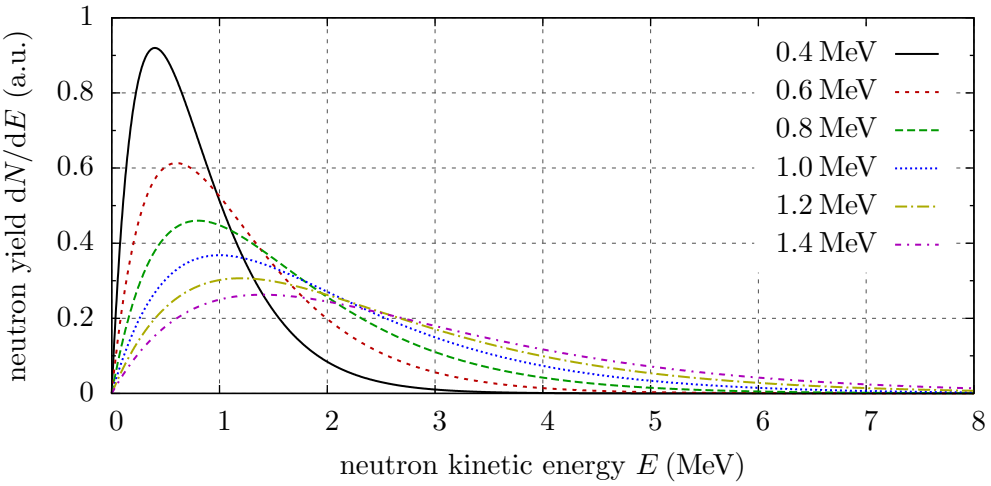

Figure 2.7: Typical shapes of neutron evaporation spectra for various (nuclear) temperatures. The curves are normalized to unit area.

The major part of the neutrons from the giant dipole resonance follows a Maxwellian energy distribution that is usually referred to as *evaporation spectrum:*

$$\frac{\mathrm{d}N}{\mathrm{d}E}(E) = \frac{E}{T^2}\exp\left(-\frac{E}{T}\right)$$

E denotes the kinetic energy of the neutrons, N the number of neutrons produced, and the energy T is an abstract parameter called the *(nuclear) temperature* of the spectrum. From the properties of the Maxwellian distribution follows that T is the most probable and $2T$ the average neutron energy. Figure 2.7 shows neutron evaporation spectra for the typical range of temperatures, 0.4 to 1.4 MeV. Deviations from these shapes have to be expected at very low energies due to photofission (for heavy nuclei) and in the high-energy tail due to direct neutron emission [Swa79].

Several other types of photonuclear reactions contribute to the production of neutrons at higher energies, most notably the *quasi-deuteron effect* and, above photon energies of ∼140 MeV, *photopion production*. Because the cross sections for these processes are at least one order of magnitude below the giant dipole resonance, they can be neglected in the current discussion. In personnel radiation protection, however, the high energy neutrons from these processes play an important role and need to be considered. Details on this can be found in [Swa79].

2.3 Electromagnetic showers

At sufficiently high energies, the energy loss of electrons is dominated by bremsstrahlung, and the main interaction of photons with matter is the production of

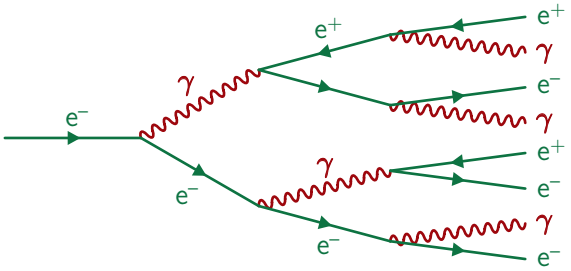

Figure 2.8: Illustration of an electromagnetic cascade created by alternating pair production and emission of bremsstrahlung.

electron-positron pairs. Combined, these two effects lead to the phenomenon of an *electromagnetic cascade* or shower. As illustrated in Fig. 2.8, bremsstrahlung photons induce pair production, and the newly created electrons and positrons will again generate bremsstrahlung when they interact with the nuclei of the material. These new photons can produce additional e^+/e^- pairs, and therefore the number of particles involved in the cascade rises exponentially until the energies are low enough to favor different processes. Hence, the effect of an electromagnetic cascade is the dispersal of transported energy from few high-energy particles to many low-energy particles. These low-energy particles are mainly responsible for the energy transfer to the material.

The characteristic length scale for an analysis of electromagnetic showers is the radiation length of the material—on average, electrons have lost all but $1/e$ of their initial energy after a path length of L_{rad}, and the mean free path of photons for pair production is $\frac{9}{7}L_{\mathrm{rad}}$. This is exploited in analytical shower theory to derive some basic properties of the cascade. Because results derived from these analytical methods are generally only coarse approximations or only valid for very simple boundary conditions, computer based Monte Carlo simulations are now widely used for the analysis of cascades. Well-established codes are for instance EGS [Nel90], Fluka [Fas03, Fer05], and Geant4 [Ago03, All06].

2.4 Radioactivation of materials

Accelerator components exposed to the electron beam or to electromagnetic showers may become radioactive. This is problematic because it can prevent access to parts of the accelerator even when the machine is turned off. Therefore, it may complicate or inhibit necessary maintenance work.

Residual radioactivity is mainly caused by the photonuclear reactions discussed before, namely the giant dipole resonance and the quasi-deuteron effect. The loss of neutrons (or protons, deuterons, α-particles) frequently leaves unstable nuclei that subsequently decay to stable end products. For heavy nuclei, also photospallation may occur.

The neutron fluences released by an electron accelerator are usually not high enough to generate serious activation by neutron capture. This means that induced radioactivity is localized in the core region of the electromagnetic showers,

i.e. in beamline components like beam pipes, collimators, flanges, magnet yokes, etc. Naturally, the characteristics of the generated radioactivity depend on the type of isotopes produced. Therefore, it is found that especially stainless steel is susceptible to activation due to its nickel, molybdenum, and cobalt contents. Also tungsten and gold fall into this category. Moderate amounts of radioactivity build up in copper, iron, or carbon steel; other typical materials like aluminum, lead, or plastics are not particularly prone to activation.

2.5 Radiation field effects

Through the discussed mechanisms, beam losses in an electron accelerator can release extreme amounts of ionizing radiation including considerable neutron flux. This radiation field will interact with materials and components in the vicinity of the machine and can cause a variety of undesired effects:

Degradation of material properties: By removing electrons from chemical bonds, ionizing radiation can cause changes in the chemical composition of materials. The displacement of atoms in crystal lattices is another common effect, often caused by knock-on neutrons. As a result, various properties of the affected substance change as well. Typical examples are cable insulations or plastic parts turning brittle, glasses or plexiglasses becoming intransparent, and glues losing their adhesive power. Effects of this kind progress with the amount of dose deposited in the material and are therefore known as *total ionizing dose* or TID effects. The term subsumes processes depending on neutron fluence as well as those depending on electromagnetic dose.

Influence on electronics: Electronic components are susceptible to TID radiation damage to a varying degree. As an example, the current amplification of most bipolar transistors decreases with accumulated neutron fluence. In addition to TID degradation, electronic circuits may be affected by *single event effects* or SEEs, which are defined as "individual events which occur when a single incident ionizing particle deposits enough energy to cause an effect in a device" [NAS96]. Hence, the probability for such an event to occur is governed by dose rate rather than integrated dose. SEEs may have transient or permanent consequences for the operability of a circuit. There are several sub-categories of SEEs:

A *single event upset* (SEU) is a transient change of state, e.g. a bit-flip in a memory cell. It can be caused if an ionizing particle modifies the stored charge in a semiconductor device. Another example is a temporary current surge in an analog circuit. SEUs may cause a device to malfunction, but are non-destructive. The device can resume normal operation after it has been reset to a well-defined state.

A *single hard error* (SHE) permanently modifies the operation of a device, like a memory cell that is forever stuck to a logic "1".

Single event burnout (SEB) and *single event gate rupture* (SEGR) are special conditions in which a power MOSFET is destroyed by high breakdown currents caused by local ionization. These events often occur in switching power supplies, which practically forbids their use in high-radiation areas.

A comprehensive discussion of radiation-induced failure modes and corresponding design guidelines is found in [NAS96] and [MIL83].

3 Beam loss detection

A machine protection system usually has a multitude of different inputs to detect the failure of components or other hazardous conditions like, e.g., excessive orbit deviations of the beam. This kind of precautionary monitoring is important and in fact sufficient for many accelerators. Machines with high average beam power or with high stored energies additionally require the direct detection of beam losses. Especially at superconducting linacs, beam loss detection is one of the key functions of the MPS.

Because reliable information on beam losses is also of high diagnostic value for the operation of an accelerator, a multitude of different techniques has been developed for their measurement. This chapter provides a brief overview of this field of diagnostics. Of course, it cannot claim completeness.

3.1 Current monitoring

One of the simplest and oldest diagnostic methods is the measurement of the beam current in a storage ring. In the presence of beam losses, the stored current gradually decreases until the ring is refilled with new particles. The time when the current has reached $1/e$ of its initial value is known as the *lifetime* of the filling. Obviously, a lifetime measurement is a relatively slow procedure and gives no information on the position of the losses.

In a linear accelerator, current measurements can be used much more fruitfully because the current difference between two points in the machine provides a simple and direct measure of beam loss. Typical devices for such measurements are current transformers as depicted in Fig. 3.1, sometimes called *toroids* for the type of construction. Because this kind of setup can measure only alternating currents, it is also known as *AC current transformer* (ACCT). In contrast to this, a DCCT is a more sophisticated setup that also allows to measure DC currents as caused by a coasting beam. Detailed discussions on these types of current transformers and on lifetime measurements in storage rings are found in [CAR04].

Passive rf cavities can be used as an alternative to current transformers. When the beam passes such a cavity, it excites a number of electromagnetic modes.

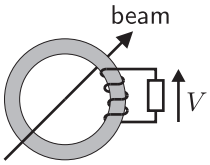

Figure 3.1: Schematic of a toroid (AC current transformer). For bunched beams, output voltage is a measure of beam current.

Because the amplitude of the monopole mode depends only on the parameters of the cavity and on the charge passing through it, it can be used to measure the current.

The mechanism for obtaining the current difference between two measurement points and the implementation in the framework of an MPS vary strongly according to the specific needs of individual accelerators. To give a few examples, the beam loss accounting system of the CEBAF electron accelerator measures beam currents with pillbox cavities in the injector and in front of each beam dump. An analog comparator circuit causes a beam stop if the difference between the injected and the dumped beam current exceeds a predefined threshold [Urs95]. At FLASH, signals from toroids near the electron gun and near the dump are digitized and compared in an FPGA-based system that implements several programmable thresholds (cf. section 6.5). For the BNL ERL test facility, it has been proposed to link two DCCTs in a special current-cancelling mode to reach a resolution of 10^{-7} for the measurement of lost beam current [Cam05].

3.2 Detection of electromagnetic showers

Instead of monitoring the beam current, beam losses can also be detected by the electromagnetic shower they create. Several of the processes discussed in the last chapter have been exploited to build *beam loss monitors (BLMs);* a few of the more popular designs are discussed in the following.

A suitable BLM can detect small beam losses that would not be noticed by even the best available current monitoring system. However, the electromagnetic showers are quite localized, and therefore the positioning of the monitors is critical. So, unless a BLM system provides very good coverage of the accelerator, it is always in danger of missing potentially hazardous beam losses; this is a strong argument for a combination with one of the current monitoring techniques discussed above.

It should be noted that the term *beam loss monitor* is slightly misleading; in fact, BLMs do not measure the amount of beam loss, but rather some different quantity like electromagnetic dose. In the best case, the signal of a BLM is proportional to the amount of beam loss at a specific location under well-defined conditions.

In the following, the sensitivity S of a given BLM will be defined as the charge Q generated at the output of the device per dose D deposited in the active detector volume:

$$S := \frac{Q}{D}$$

3.2.1 Ionization chambers

One of the oldest device classes used for beam loss detection is the ionization chamber. It can be described as two conducting surfaces held at high voltage against each other and separated by a gas. Two typical geometries are shown in Fig. 3.2.

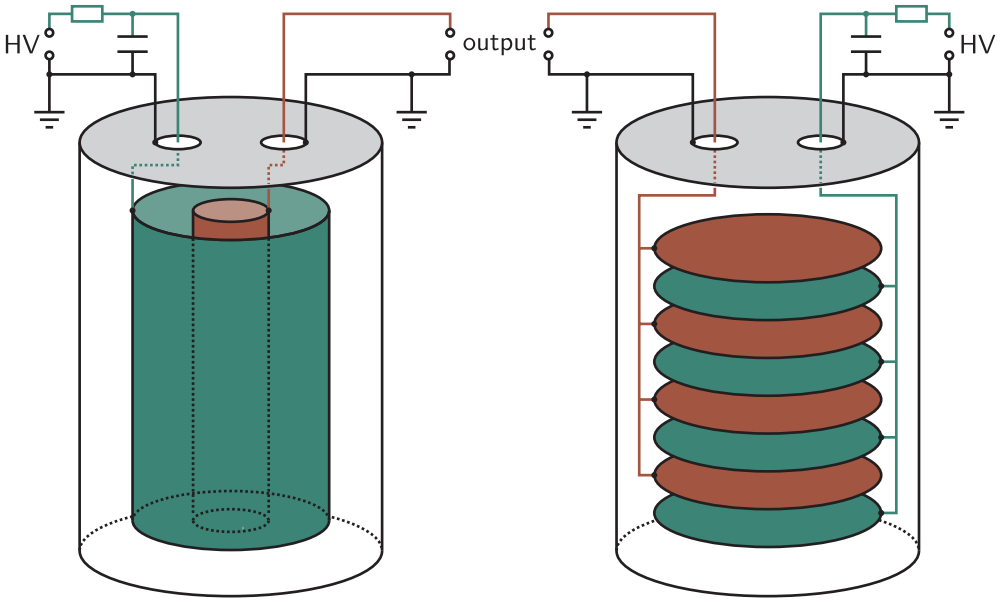

Figure 3.2: Schematic of an axial and a parallel-plate ionization chamber. In the shown configurations, the enclosure is grounded and the output does not carry high voltage.

If ionizing radiation penetrates the gas container, electron-ion pairs are created. The electrons and ions then drift towards anode and cathode and create a measurable current. Usually, the high voltage is adjusted to a point where all the primary ionization charges in the gas volume are collected by the electrodes. Since this behavior is stable over a relatively wide voltage range, the calibration of the chamber does not change with small variations of the voltage. In fact, the calibration is almost completely determined by the volume V of the enclosed gas and by the gas type:

$$S \approx V\rho \cdot \frac{e}{E_{\text{ion}}} \tag{3.1}$$

ρ denotes the density of the gas and E_{ion} is the mean energy deposition needed to create an electron-ion pair. Since this is an average value for all kinds of primary particles and processes, E_{ion} is considerably higher than the well-known ionization potential of the gas. A table of values is found in [Leo94]; useful examples for E_{ion} are 35 eV for nitrogen and 26 eV for argon.

As mentioned before, this simple relation is only valid in a confined high voltage range. At lower operating voltages, the sensitivity is reduced by recombination of charge carriers. At higher voltages, on the other hand, the primary ionization electrons are accelerated so strongly that they produce secondary ionization charges along their path. At extreme voltages, electrical breakdown occurs.

Another important operational parameter for an ionization chamber is its response time, i.e. the time the ionization charges need to reach the electrodes. Drift

velocities for electrons in gases at atmospheric pressure are of the order of 50 km/s [Jea79] so that an electrode separation in the centimeter range gives a response time of few microseconds or below. For the ions, collection times can be as high as several hundred microseconds, or even milliseconds in some chamber designs. Depending on the geometry of the electrodes, these collection times may change when the polarity of the high voltage is reversed. The sum of electron and ion currents is measured at the output of the ionization chamber. In the configurations depicted in Fig. 3.2, the enclosure is grounded and the output does not carry high voltage.

Ionization chambers are used as beam loss monitors in many laboratories; two recent examples are a 110 cm^3 argon-filled chamber with cylindrical electrodes used at the Spallation Neutron Source [Wit02], and a 1.5 l nitrogen-filled parallel-plate design used at the Large Hadron Collider [Sto07].

Sensitivity example: The gas in the LHC chamber is kept at the pressure of 110 kPa, so the density of the N_2 filling is approximately 1.38 g/l. Inserting these parameters into (3.1) gives

$$S = 1.5\,\mathrm{l} \cdot 1.38\,\frac{\mathrm{g}}{\mathrm{l}} \cdot \frac{e}{35\,\mathrm{eV}} = 59\,\frac{\mu\mathrm{C}}{\mathrm{Gy}}.$$

3.2.2 Long ionization chambers

The *long ionization chamber* is a beam loss monitor whose origins can be traced back to the Stanford Linear Accelerator Center in the 1960s. Instead of monitoring beam losses with a huge number of discrete ionization chambers, W. Panofsky proposed to measure ionization currents in a long gas-filled coaxial cable with a diameter of few centimeters which has a high voltage applied between the inner and outer conductors [Pan63].

Information about the longitudinal position of the beam loss can be derived from the arrival time of the ionization current pulse at the front end of the cable. Following the designations given in Fig. 3.3 and assuming the loss of a bunch at position B, Δt can be defined as the time between the passage of the bunch and the

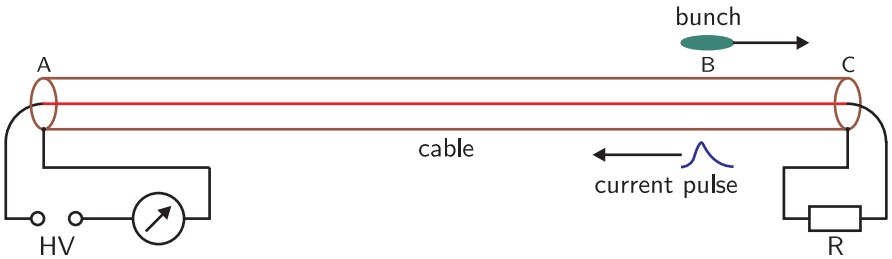

Figure 3.3: Schematic of a long ionization chamber. From the delay of the ionization current pulse, the longitudinal position of the beam loss can be identified.

arrival of the current pulse at A. Δt is proportional to the distance A–B; under the assumption that the beam is ultrarelativistic and that the signal propagation speed in the coaxial cable is about 90 % of the speed of light in vacuum, the distance A–B can simply be calculated as

$$s_{\mathrm{AB}} \approx \frac{c}{1.9}\Delta t. \tag{3.2}$$

To avoid the necessity of an external bunch trigger, also the time difference between the direct current pulse and its reflection at the far end of the cable can be measured; (3.2) has to be adapted accordingly.

The longitudinal resolution reached by long ionization chambers is limited by dispersive broadening of the pulses in the coaxial line, by the electron collection time (typically few hundred nanoseconds), and by the sampling rate of the measurement system. For the first cables at SLAC, several kilometers long, a resolution of few ten meters has been reported [Fis67]. With more recent designs, resolutions of 2 m or below have been demonstrated [Pat92].

The positional information becomes ambiguous if the cable is too long with respect to the bunch spacing. This problem arises if a second bunch reaches point A before all current pulses caused by the first bunch have reached it. The maximum cable length for an unambiguous determination of the loss location is defined by (3.2); hence, for a bunch frequency of 1 MHz, the maximum allowable length is about 150 m. In the case of CW beams, no position information can be obtained at all. As a consequence, some facilities use segmented coaxial cables with a separate current readout at each segment, e.g. the ELBE radiation source [Sch02b].

Sensitivity example: As for a discrete ionization chamber, the sensitivity of a gas-filled cable can be estimated from the geometry and the gas type with (3.1). Assuming that a shower penetrates an argon-filled cable of 2 cm diameter on a length of 1 m,

$$S = 310\,\mathrm{cm}^3 \cdot 1.8\,\frac{\mathrm{g}}{\mathrm{l}} \cdot \frac{e}{26\,\mathrm{eV}} \approx 21\,\frac{\mu\mathrm{C}}{\mathrm{Gy}}.$$

is obtained.

3.2.3 PIN diodes

A PIN diode consists of three semiconductor layers as shown in Fig. 3.4. An undoped, intrinsically conducting layer (i) is sandwiched between a p-doped and an n-doped layer. If the diode is reverse biased, movable charge carriers are depleted from the undoped layer and the component behaves like an ionization chamber. Electron-hole pairs created by ionizing radiation are separated by the bias voltage and create a measurable current.

Compared with a gas-filled ionization chamber, the mean energy deposition required to create a pair of charge carriers in a semiconductor is considerably lower. For silicon, a mean energy of 3.6 eV is required to lift an electron to the conduction

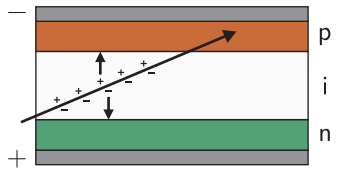

Figure 3.4: Schematic of a PIN diode. An undoped semiconductor layer (i) is sandwiched between p- and n-doped layers. If the diode is reverse biased as shown, the i layer is depleted of movable charges, and the diode acts like an ionization chamber.

band [Wit00]. In addition, the density is about three orders of magnitude higher than that of gases at atmospheric pressure. This high specific sensitivity is counterbalanced by the small active volumes of commercially available PIN-diodes. The highest volumes are found in planar PIN-photodiode geometries which are available up to surface areas of 1–2 cm^2 with depletion depths of the order of 100 µm.

Except for extreme voltages, the drift velocity of electrons and holes in the semiconductor scales linearly with the applied electric field as $v = \mu E$. Assuming a mobility μ of the charge carriers of 1350 $\mathrm{cm}^2/(\mathrm{Vs})$ for electrons, 480 $\mathrm{cm}^2/(\mathrm{Vs})$ for holes [Leo94], and a typical bias voltage of 24 V across a semiconductor thickness of 300 µm, drift velocities of 11 km/s for electrons and 4 km/s for holes are obtained. Consequently, the charge collection time for the PIN diode is lower than 100 ns. Depending on the geometry and on the applied voltage, response times of few nanoseconds are possible.

One of the most prominent examples of PIN diodes being used for beam loss monitoring is the HERA collider. Here, pairs of back-to-back mounted diodes were read out by a coincidence detection circuit. A counter was incremented each time that a simultaneous current pulse from both diodes was observed. This counting scheme had the advantage of being almost blind to photons from the strong synchrotron radiation background while retaining a relatively high sensitivity for charged particles [Wit94, Wit00]. Recently, PIN diodes have been successfully employed in a current readout configuration at SPring-8 [Shi08].

Sensitivity example: A typical silicon PIN photodiode with 1 cm^2 surface and a depletion depth of 100 µm has a sensitivity of about

$$S = 10\,\mathrm{mm}^3 \cdot 2.3\,\frac{\mathrm{g}}{\mathrm{cm}^3} \cdot \frac{e}{3.6\,\mathrm{eV}} \approx 6\,\frac{\mu\mathrm{C}}{\mathrm{Gy}}.$$

3.2.4 Secondary emission monitors (SEMs)

A simple beam loss monitor with low sensitivity is obtained by applying a static high voltage between two or more typically metallic surfaces under vacuum (Fig. 3.5). The dominant mechanism for charge extraction from the cathode(s) is the emission of secondary electrons stimulated by the passage of a charged particle. Because of the absence of a retarding medium, these secondaries can reach the anode(s) within few nanoseconds, making the secondary emission monitor a very fast BLM.

The secondary emission yield, i.e. the average number of secondaries produced by one primary particle, varies strongly with the material and with the energy of

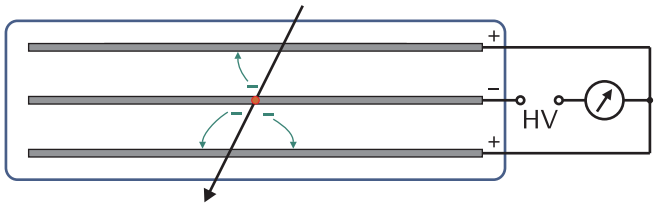

Figure 3.5: Schematic of a secondary emission monitor (SEM).

the incident particle. Values of few percent up to about 20 are typical. A carefully designed secondary emission monitor can be extremely radiation hard; its response is also very linear over a wide range of radiation intensities due to the absence of saturation effects. At extreme radiation levels, the linearity is usually limited by ionization of residual gas.

Sensitivity example: About 300 SEMs are used in high radiation areas at the LHC. The monitor consists of three titanium discs of 8.9 cm diameter held at a high voltage of 1.5 kV. Its geometry is similar to the one shown in Fig. 3.5, and the device is oriented in a way that the discs are largely perpendicular to the expected particle trajectories. Due to this arrangement, it is expected that the currents caused by *delta electrons*—fast knock-on electrons with energies sufficient to overcome the high potential difference between the electrodes—from the first anode and from the cathode cancel out. [Kra07]

Assuming an average emission yield of 5 % for the remaining low-energy secondaries, the sensitivity of the SEM can be roughly estimated as

$$S = 62\,\mathrm{cm}^2 \cdot 0.05 \cdot \frac{e}{2\,\mathrm{MeV\,cm}^2/\mathrm{g}} \approx 2\,\frac{\mathrm{nC}}{\mathrm{Gy}}.$$

For this estimate, it has also been assumed that the primaries are minimum ionizing particles which have a mass collision stopping power of about $2\,\mathrm{MeV\,cm}^2/\mathrm{g}$ (cf. section 2.1.1).

3.2.5 Photomultipliers

Photomultiplier tubes (PMTs) are detectors for extremely low light intensities. Figure 3.6 illustrates the typical design of a head-on PMT. Inside an evacuated glass tube, a number of electrodes are mounted: a photocathode, several dynodes, and an anode. A high voltage power supply and a resistor chain provide voltage drops of few 10 to few 100 V between the electrodes.

When an incident photon strikes the cathode, it can release an electron via the photoelectric effect. This particle is electrostatically accelerated towards the first dynode where it creates several more electrons by secondary emission. The secondaries are accelerated towards the next dynode and the process repeats until the charges are picked up by the anode. In this avalanche-like process, the number of electrons is multiplied by an average factor G, the so-called *gain* of the PMT.

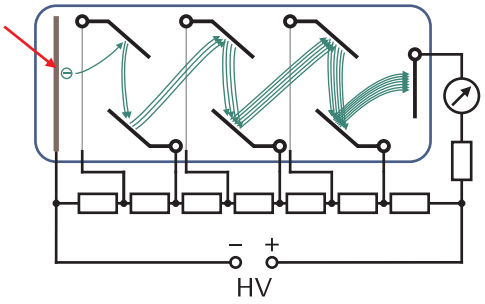

Figure 3.6: Schematic of a photomultiplier tube (PMT). The electrons released from a photocathode by incident photons are multiplied in several dynode stages by secondary emission.

To obtain a high gain with a moderate number n of dynode stages—for conventional tubes, n is in the range of 8–14, and G can reach 10^6 or more—, materials with high secondary emission coefficients δ are used. Typical examples are $SbCs_3$, GaP, or CuBe. Because the coefficient δ is a function of the incident electron energy, the total gain depends strongly on the high voltage V supplied to the dynode chain. Assuming a linear dependence $\delta = K\Delta V$ with a constant voltage drop ΔV, a gain of

$$G = \delta^n = K^n \Delta V^n$$

results. It should be noted that $\delta(\Delta V)$ can deviate strongly from the assumed linear behavior for some dynode materials and that many voltage divider designs do not foresee a constant voltage drop between the dynodes. As a rule of thumb, however, it remains true that the tube gain varies with a high power of the applied voltage, setting high requirements on the stability of the power supply.

Within a confined parameter space, the photomultiplier signal depends linearly on the intensity of the incident light. A typical source of nonlinear behaviour is a too low supply voltage. In this case, secondary electrons are accelerated so slowly that their space charge field substantially decreases the effective electric field on the dynode surface and therefore impedes the charge collection process. The influence of this space charge effect increases at higher photocurrents, so that higher voltages are necessary for stronger light intensities.

In connection with high currents, the photocathode or dynodes may lose more charge due to secondary emission than can be replaced through the voltage divider. Since this leads to a shift of electric potentials, it causes variations of the PMT gain. A similar effect takes place at excessive anode currents: the voltage drop across the load resistance becomes so high that the potential difference between the last dynode and the anode is changed. For applications where high currents appear only in short pulses, the charge reservoir of the electrodes is usually enlarged by adding capacitors to the voltage divider which can be recharged in the time between pulses.

Photomultipliers are often used as fast and sensitive beam loss monitors because of their very fast signal rise time of few nanoseconds and the adjustable high gain. A number of possible usage scenarios are discussed in the following.

Photomultiplier with bulk scintillator

The most common use of photomultipliers for beam loss detection is in combination with a scintillating material. Scintillators are substances capable of absorbing high-energy ionizing radiation and emitting a part of the absorbed energy in the form of low-energy photons, often in the optical wavelength range. Four main types of scintillators are used for radiation detection:

Organic crystals: A few aromatic hydrocarbons like anthracene and stilbene exhibit fluorescence and can be crystallized into a solid scintillator. They are usually very expensive, only available in small volumes, and rapidly deteriorate after absorption of few kGy of dose [Bro04]. The scintillation efficiency of anthracene, however, is the highest of any known organic scintillator and therefore often quoted as a reference. On average, anthracene emits one photon in the blue wavelength range per 60 eV of absorbed energy.

Organic liquids: Organic scintillators can be dissolved in an organic solvent. Typical solvents are mineral oils or specific aromatic compounds like toluene or xylene. In a liquid scintillation cocktail, the absorption of ionizing radiation is mainly carried by the solvent which transfers the energy to the scintillating solute. Liquid scintillator is usually cheap and fairly radiation-resistant—a drop of the light output to $1/e$ after absorption of several 100 kGy or even MGy has been reported [Gol95]. However, its use in BLMs is often discouraged by safety concerns because the liquids are highly flammable and sometimes toxic.

Plastics: Organic scintillators can also be dissolved in a solid polymer base like polystyrene or polyvinyltoluene. Plastic scintillators are cheap, can easily be produced in big volumes and can be cut to any desired shape. This flexibility makes plastics the most widely used scintillators for BLM applications. A typical plastic scintillator emits one photon in the wavelength range of 400–500 nm (blue) per 100 eV of absorbed energy. With decay times of few nanoseconds, the light pulses are comparatively fast. The major drawback of plastic scintillators is their susceptibility to radiation damage: after absorption of several kGy to few 10 kGy, the light output drops to $1/e$ of its initial value [Ham99, Li05]. Therefore, plastic scintillators are not well suited for use in high-radiation environments.

Inorganic crystals: Many inorganic crystals can be used as scintillators. Typical examples are the alkali halide salts NaI and CsI, or other well-known substances like ZnS or $PbWO_4$. Often, various dopants are used to influence the properties of anorganic scintillators. While light output and radiation hardness vary greatly between different crystals, their high density gives them an advantage in sensitivity compared to organic scintillators. A major drawback for the use in BLMs is that inorganic crystals are usually expensive and hard to handle because of their chemical and mechanical vulnerability.

Photomultipliers coupled to bulk scintillators have wide application in high-energy physics detectors, and are used as beam loss monitors at many facilities. While predominantly plastic scintillators are used (e.g. at FLASH), there are also examples for liquid scintillators (LANSCE, Los Alamos) and inorganic crystals (LEDA, Los Alamos).

Sensitivity example: Assuming a photomultiplier with a gain of 10^5 is coupled to a small piece of plastic scintillator with a volume of $100\,\mathrm{cm}^3$, a density of $1\,\mathrm{g/cm}^3$, and an average light output of 1 photon per 100 eV, a sensitivity of

$$S = 100\,\mathrm{cm}^3 \cdot 1\,\frac{\mathrm{g}}{\mathrm{cm}^3} \cdot \frac{1}{100\,\mathrm{eV}} \cdot 0.6 \cdot 0.3 \cdot 10^5 \cdot e \approx 18\,\frac{\mathrm{C}}{\mathrm{Gy}}.$$

can be estimated. Two additional factors enter the calculation: 0.6 for the collection efficiency, i.e. the fraction of photons created in the scintillator that actually reach the photocathode, and 0.3 for the quantum efficiency of the photocathode. For long scintillator geometries, also the attenuation length of the material must be taken into account. An example of this is found in appendix B.1.5.

The calculation illustrates that photomultipliers with bulk scintillators are the most sensitive BLMs available. The widely adjustable gain and the freedom of scintillator shapes and volumes also gives them a high degree of versatility.

Photomultiplier with Čerenkov fiber

A longitudinally position sensitive beam loss monitor can be constructed by fitting a long quartz fiber to a photomultiplier. Relativistic charged particles traversing the fiber emit Čerenkov light that propagates with a typical velocity of about $\frac{2}{3}c$ because of the high index of refraction of the material. As in the case of the long ionization chamber, the arrival time of the photon pulses at the PMT can be converted into the longitudinal position of the beam loss. Assuming that the fiber is mounted in parallel to the beam line and extending from the PMT into the direction of the beam, the longitudinal beam loss position follows the relation (cf. (3.2))

$$\Delta z \approx \frac{3}{5} c \Delta t.$$

Thus, at a bunch frequency of 1 MHz, the maximum length of the fiber allowing an unambiguous reconstruction of the beam loss position is about 180 m.

Carefully selected quartz fibers can be radiation hard up to doses of several 10 MGy, and have attenuation coefficients as low as 1–3 dB/km. A test setup for the detection of beam losses with long fibers is installed at FLASH. [Jan00, Goe05]

Sensitivity example: The main problem in evaluating the sensitivity of a Čerenkov fiber beam loss monitor is the estimate of the collection efficiency, i.e. the ratio of photons created in the fiber that actually reach the photocathode before being absorbed or refracted out of the material. For a thin fiber, this efficiency is usually

assumed to be of the order of few percent [Gor95]—0.02 in the example. It is further assumed that the charged particles are minimum ionizing with a mass collision stopping power of 2 MeV cm^2/g (cf. 2.1.1).

The Čerenkov light yield in a typical range of photocathode sensitivities, between 350 nm and 550 nm, can be calculated as

$$\frac{\mathrm{d}N}{\mathrm{d}x} \approx 475 \sin^2\theta \, \frac{\text{photons}}{\text{cm}} \approx 280 \, \frac{\text{photons}}{\text{cm}}$$

using 50° as the approximate opening angle of Čerenkov radiation in quartz [Leo94].

Other parameters are a photocathode efficiency of 0.03, a tube gain of 10^5, a fiber length of 1 m, and a radius of 100 µm. Using these values, a sensitivity of

$$S = 31.4\,\text{mm}^3 \cdot \frac{1}{2\,\text{MeV cm}^2/\text{g}} \cdot \frac{280}{\text{cm}} \cdot 0.02 \cdot 0.3 \cdot 10^5 \cdot e \approx 2.4 \, \frac{\mu\text{C}}{\text{Gy}}$$

is obtained.

Bare photomultiplier

Due to the high gain of a photomultiplier tube, it is sometimes not even necessary to couple it to any kind of external light source. In the electromagnetic shower of a high-energy particle beam, a PMT can generate a signal from Čerenkov light and optical transition radiation originating from the glass tube itself. In addition, secondary emission from the photocathode or from the dynodes may contribute. Because of this mixture of effects, it is hard to estimate the sensitivity of bare photomultipliers to electromagnetic dose.

A number of cheap side window photomultipliers are successfully employed in this way as beam loss monitors at the Jefferson Lab FEL [Kra97a]. Only moderate radiation damage by a slow darkening of the tube glass is reported [Jor08].

Aluminum cathode electron multipliers

It is possible to increase the radiation hardness of a tube-based BLM even further by replacing the photocathode with a metal film, e.g. an aluminum layer. Such an *aluminum cathode electron multiplier*—sometimes abbreviated as ACEM or called *nuclear radiation monitor*—is practically a secondary emission monitor with an integrated tube amplifier. ACEMs usually do not belong to the standard product range of tube manufacturers and are therefore relatively expensive. At FLASH, they are used in few locations where a high radiation background would lead to a rapid darkening of scintillator material.

Sensitivity example: Under the assumptions of an average secondary emission yield of 1 %, minimum ionizing particles, and a tube gain of 10^5, the sensitivity of an ACEM with a cathode surface of 10 cm^2 is approximately

$$S \approx 10\,\text{cm}^2 \cdot \frac{1}{2\,\text{MeV cm}^2/\text{g}} \cdot 0.01 \cdot 10^5 \cdot e \approx 5 \, \frac{\mu\text{C}}{\text{Gy}}.$$

4 Dark current transport at FLASH

The term *dark current* describes the set of all charge carriers being emitted from a device or structure unintentionally, or the electrical current formed by these charges. In the context of electron accelerators, this usually refers to electrons that are released into the beamline vacuum without being part of the beam.

At FLASH, two main sources of dark current must be distinguished: The superconducting niobium cavities and the normal-conducting rf gun cavity. Dark current from the rf gun (short: gun dark current) has become a major problem in the operation of the accelerator due to its high intensity. Its most notable effect is the activation of multiple beamline components to dose equivalent rates of several mSv/h. This has necessitated considerable radiation protection effort and impairs maintenance work in the affected sections. When FLASH is operated with few bunches, the power transported by the dark current can exceed the beam power by an order of magnitude or more. Even close to the end of the accelerator, losses of gun dark current contribute significantly to the total absorbed dose in the undulator magnets. It has therefore become necessary to investigate the origin and the transport mechanism of dark current in the FLASH linac.

The discussions in this chapter are largely based on extensive tracking simulations. Similar studies for the TTF linac, FLASH's predecessor, have been published in [Sch01]. Before the detailed analysis of dark current transport mechanisms, a qualitative discussion of the contribution from superconducting cavities is given.

4.1 Dark current from superconducting cavities

For the niobium cavities, the amount of dark current produced at typical field amplitudes is strongly dependent on the surface quality of each individual resonator. While most cavities emit currents of only few nA, individual ones may produce as much as several µA [Lil05].

The energy of a field-emitted electron exiting the cavity can range from practically zero to a value $\Delta E_{\max}$. $\Delta E_{\max}$ denotes the maximum energy gain of a beam particle passing the cavity when it is accelerated on the crest of the rf wave. The energy spectrum of dark current exiting a string of N cavities ranges from zero to a maximum of $E_{\max} = \sum_{i=1}^{N} \Delta E_{\max,i}$. Because the optical lattice of the linac does not permit the transport of electrons with energies substantially lower than that of the beam, most of the dark current from the cryomodules is lost close to its origin. As there is no preferred direction for the emission of dark current, the possibility of transport in the forward (towards the dump) and in the backward direction (towards the gun) has to be considered. Figure 4.1 illustrates these possibilities

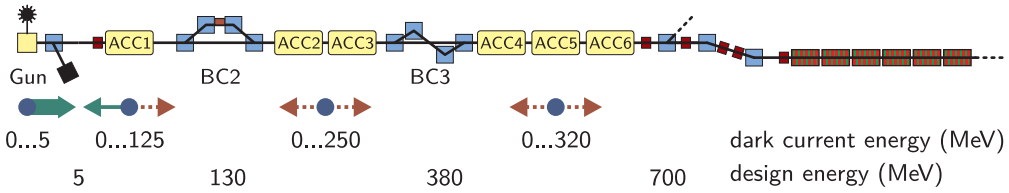

Figure 4.1: Dark current emitters and possible transport directions in the FLASH beamline. The energies given reflect a typical setup for a final beam energy of 700 MeV.

for an exemplary final beam energy of 700 MeV; for the individual acceleration modules, the situation is the following:

ACC1: Backwards transport is possible because no optical elements except a solenoid magnet are present in the very short beamline to the rf gun. In the forward direction, the dark current has an energy of up to 125 MeV. Considering that the beam already has an initial energy of 5 MeV after the gun and is usually accelerated at an off-crest phase of 8° in ACC1, its energy amounts to $(5 + 125 \cdot \cos 8°)$ MeV ≈ 128.8 MeV. Inserting the BC2 scraper (cf. section 1.4.2) in its recommended position prevents transport of electrons more than 2 MeV below the average beam energy and is therefore sufficient to suppress dark current from ACC1 completely. However, in nonstandard modes of operation—if the scraper is opened or if the beam is accelerated at higher off-crest phases—, it cannot be excluded that a tiny fraction of this current gets transported to the entrance of ACC2 and further down the linac.

ACC2–3: Although a part of the dark current emitted in the backwards direction has the correct energy to be transported through the quadrupoles upstream of ACC2, it is lost at the latest upon entering the last (i.e. most downstream) dipole of BC2, where it is deflected into the vacuum chamber. In the forward direction, the energy mismatch with respect to the lattice forbids any transport.

ACC4–6: Transport is prevented by the BC3 dipoles in the backwards direction, and by the huge energy mismatch in the forward direction.

In summary, dark current from any cryogenic acceleration module is only a local problem because it cannot be transported far from the module. The situation for the rf gun is entirely different.

4.2 Dark current from the rf gun

The core part of the FLASH electron gun is a copper cavity with a full and a half cell for the rf frequency of 1.3 GHz as depicted in Fig. 4.2. A circular opening in its

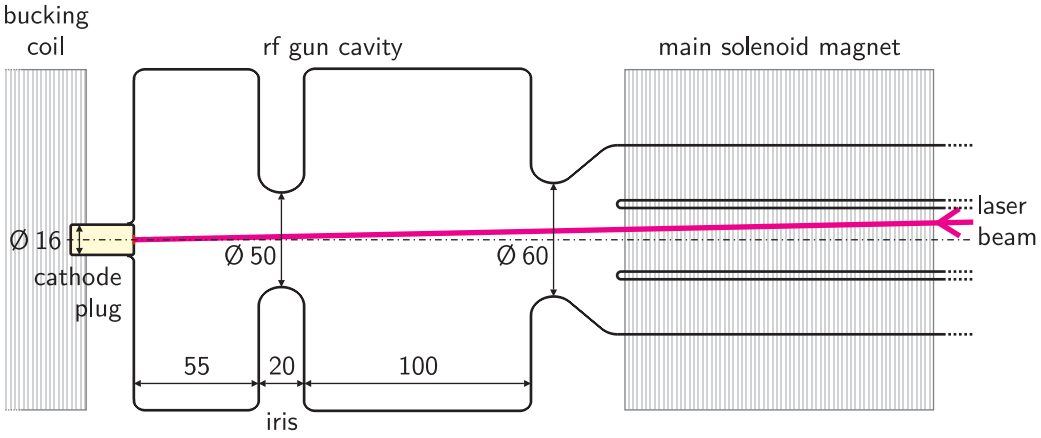

Figure 4.2: Schematic of the core parts of the FLASH rf gun. The photocathode covers a circular area of 5 mm diameter on the cathode plug surface. The gun cavity is rotationally symmetric. All measures are given in mm.

backplane allows the insertion of a cathode plug of 16 mm diameter. The plug itself is usually made of molybdenum and carries a layer of Cs_2Te as a photocathode of high quantum efficiency. This coating covers a circular area of 5 mm diameter on the front surface of the plug.

The rf gun is equipped with a main solenoidal magnet to focus the electron bunches as they are exiting the cavity. To avoid a negative impact on the beam emittance, an additional weak solenoidal *bucking coil* is used to compensate the longitudinal magnetic field on the cathode surface. It is located just behind the cavity backplane and thus surrounds the cathode retention system.

A low beam emittance is crucial to the operation of a free-electron laser. Hence, a main goal in the operation of the rf gun is to limit the emittance blow-up caused by transverse space charge forces while the electrons are still at non- or low-relativistic energies. As a result of this, the gun is operated with electric field amplitudes as high as practically possible—with the current rf source, the limit is at about 44 MV/m. Due to the high gradient, all surfaces in the rf gun cavity become potential field emitters. Depending on roughness, cleanness, and local geometry, any of the copper, molybdenum and cesium telluride surfaces can be strong sources of dark current. The total dark current exiting the rf gun cavity can be measured with two Faraday cups located 1.09 m and 1.23 m downstream of the cathode. Dark currents in the range of few 100 µA during the rf pulse—with large variations between individual photocathodes—are quite typical, as reported e.g. in [Han05].

The gun dark current is thus about two or three orders of magnitude higher than that of a whole cryogenic acceleration module. In addition, its spatial distribution and energy spectrum overlap with that of the beam, which allows a significant fraction to be accelerated and transported through the linac. Losses of gun dark

current can be observed in many places along the machine, most prominently in bunch compressors, collimators, and also in the FEL undulators. In the bunch compressors, activation caused by these losses has occasionally reached critical proportions with dose equivalent rates of more than 10 mSv/h at the surface of the vacuum chambers [Leu08].

A simulational model has been devised to investigate the mechanisms of transport and loss of gun dark current. Because a good modeling of the dynamics inside the rf gun cavity is essential, the tracking code Astra [Flo00] has been chosen as the main simulation tool. However, the original version of this code is not well prepared for the detailed simulation of a complete accelerator with a complex beam line geometry. Therefore, several modifications of the code have been implemented, most importantly:

- parallelization of the existing code to reduce simulation times by the use of a PC cluster
- addition of detailed 3-dimensional aperture modeling capabilities

Further information on these features can be found in appendix A.

4.3 Field emitters in the rf gun cavity

To model the dark current emission from the rf gun, it is necessary to identify the main field emitters in the gun cavity. Some work in this direction has already been reported in [Han05] with the conclusion that the cathode plug is the only emitter of dark current that can be transported into the linac. The investigation discussed here uses a more detailed, 2-dimensional field map [Sek08] of the rotationally symmetric gun cavity.

The amplitude of the electric field is shown in Fig. 4.3. Beside the desired field maxima in the center of the cavity, local maxima exist at the rounded edges of the cathode plug and of the backplane, at both sides of the central iris, and at the iris near the exit of the cavity. In all of these places, the field amplitude is even higher than on the cathode surface. When 42 MV/m are assumed as a typical peak field on the cathode surface, a local amplitude of more than 54 MV/m is predicted near the edge of the plug due to its rounding radius of only 1 mm.

To follow the trajectories of field-emitted electrons from these hot spots, a simulation has been set up using the parallel version of the Astra tracking code. Each potential field emitter listed in Tab. 4.1 is modeled individually.

4.3.1 Field emission model

The current I emitted from a surface by field emission is a function of the applied electric field. The form of this dependency has been established in [Fow28] by solving the Schrödinger equation for the tunneling of an electron through the potential barrier at the surface of a solid. With an additional factor β_e that accounts for

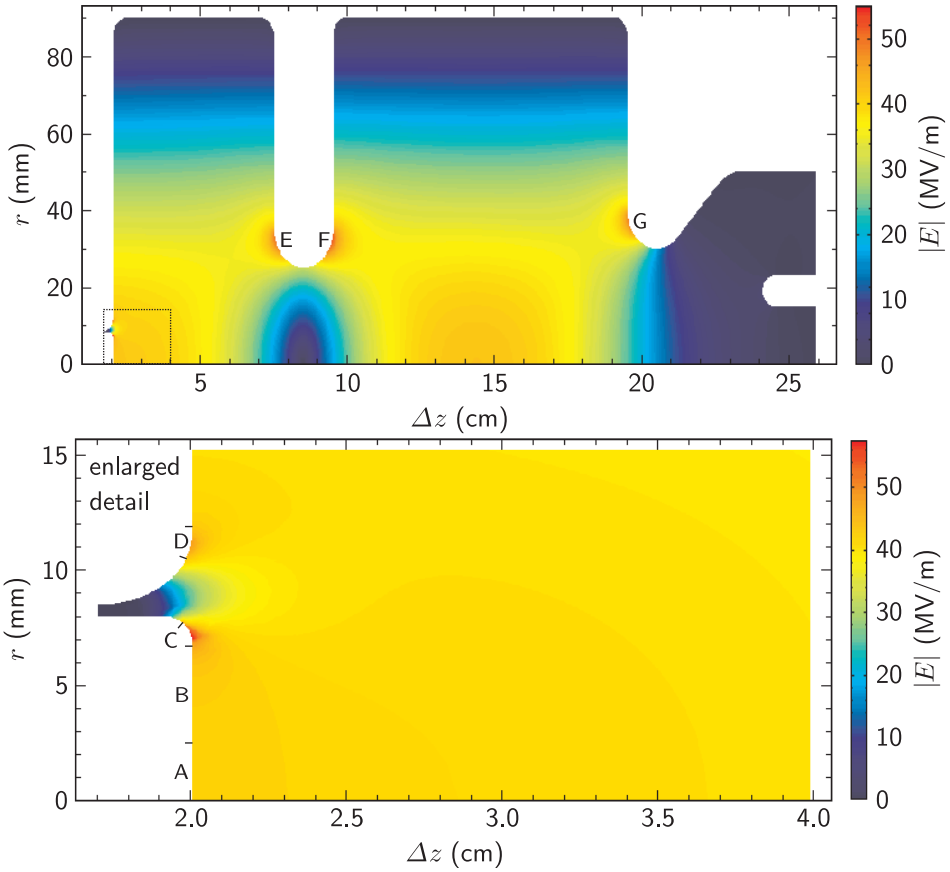

Figure 4.3: Electric field amplitudes in the rf gun [Sek08]. In addition to the field maxima in the center of the cavity, local maxima are found at the rounded edges of the cathode plug (C) and of the backplane (D), at both sides of the central iris (E and F), and at the iris near the exit of the cavity (G).

	Surface	Material	Field amplitude $\hat{E}$ (MV/m)
A	photocathode	Cs_2Te	42
B	plug surface	Mo	42
C	plug border	Mo	54
D	backplane border	Cu	44
E	iris upstream	Cu	48
F	iris downstream	Cu	48
G	iris exit	Cu	48

Table 4.1: Potential field emitters in the gun cavity. Each surface region listed here belongs to a local maximum of the electric field. The photocathode plug is divided into three parts: the border, the plane surface not covered by Cs_2Te, and the photocathode.

Surface material	**Work function** ϕ (eV)	**Field enhancement factor** β_e
Cu (polished)	4.5 [And49]	80 (60 [Wan89], 62 [Pal98], 30, 46, 76 [Sri99])
Mo (polished)	4.6 [CRC03]	170 (164 [Boh03])
Cs_2Te (photocathode)	3.5 [Pow73]	220 (220 [Boh03])

Table 4.2: Material-dependent parameters for the Fowler-Nordheim field emission model. The field enhancement factor used for the initial tracking studies is an upper limit to a number of literature values for comparable surfaces (given in small print).

the enhancement of the electric field due to microscopic surface irregularities, an effective version of this *Fowler-Nordheim equation* reads as follows [Lat81]:

$$I = A_e \cdot 1.54 \cdot 10^{-6} \cdot 10^{4.52\phi^{-0.5}} \cdot \frac{\beta_\mathrm{e}^2 E^2}{\phi} \cdot \exp\left(-\frac{6.53 \cdot 10^9 \phi^{1.5}}{\beta_\mathrm{e} E}\right) \cdot 1\,\mathrm{A} \qquad (4.1)$$

A_e is an effective emitting area in units of m^2, ϕ the work function of the material in units of eV, and E the macroscopic electric field applied to the surface in units of V/m.

For setting up the input distribution of the tracking simulation, an expression for I as a function of time is needed. Because a standing wave is formed in the gun cavity, the momentary electric field at any point within it is given by $E(\boldsymbol{x}, \varphi(t)) = \hat{E}(\boldsymbol{x}) \cos\varphi(t)$ with $\varphi(t) = \omega t$. Therefore, maximum field emission is expected for $\varphi = 0$.

For an evaluation of $I(\varphi)$, appropriate values of ϕ, β_e, and A_e have to be determined. The work function ϕ, i.e. the minimum energy needed to remove an electron from a solid, can be found in the literature. The field enhancement factor β_e depends both on the material and on the preparation of the surface; a number of suitable values have been reported and are summarized together with the work function in Tab. 4.2. The value of β_e chosen as input for the tracking simulation is generally an upper limit to the literature values. In this way, field emission is overrather than underestimated. Finally, there is no a priori knowledge of the effective emitting area A_e. Therefore, $I(\varphi)$ is normalized to an arbitrary amplitude in the process of setting up the input distributions for the tracking code. Choosing an arbitrary current or charge density for the simulation is permissible because space charge forces are negligible for the dark current.

Using the values from Tab. 4.2, the dark current distributions $I(\varphi)$ can be calculated. As shown in Fig. 4.4, the curves are almost Gaussian with rms widths between 12° and 20°. Particles from the Cs_2Te photocathode cover the broadest range of phases, those from the backplane border the narrowest one.

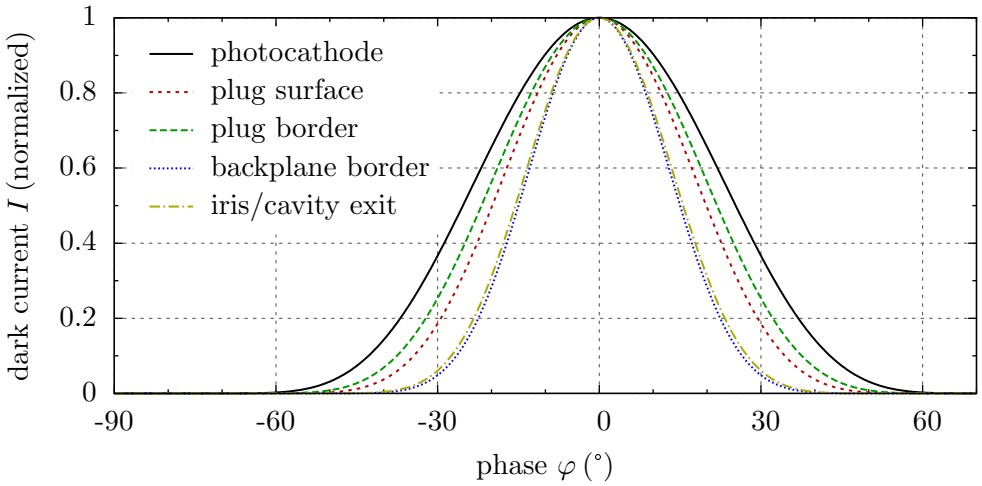

Figure 4.4: Dark current emission vs. phase of the electric field. The curves follow the Fowler-Nordheim formula with parameters according to Tab. 4.1 and 4.2. The amplitude has been normalized to 1.

4.3.2 Tracking studies in the photoinjector

Once the temporal distributions have been calculated, the modeling of the field emitters can be concluded. For each of the high field regions, one million particles are uniformly distributed over the respective surface area. Afterwards, the tracking is started with time steps below 1 ps. It continues until all of the particles have either hit an aperture or have reached the entrance to cryomodule ACC1 about 2.6 m downstream of the cathode. The aperture model includes the gun cavity, beam pipes, and the vacuum chamber of a spectrometer dipole magnet, but excludes the gun collimator described in section 1.4.1; a detailed discussion of this removable collimator is given later.

Because the field of the main solenoid magnet may have a substantial influence on the particle trajectories close to the exit of the cavity, the simulation is repeated for peak fields between 145 and 205 mT. The value in the center of this range, 175 mT, is the usual operating point for the main solenoid. The electric field amplitude on the cathode surface is set to a value of 42 MV/m that is typical for the normal operation of the FLASH rf gun.

Figure 4.5 shows the trajectories of 30 randomly picked particles from each field emitter for a solenoid peak field of 175 mT. It is obvious that electrons emitted in the backwards direction, i.e. from the upstream face of the iris or from the cavity exit, cannot reach the beam pipe. From the downstream face of the iris, only a few particles are transported into the region of the solenoid, but are then lost on the beam pipe walls.

This rough assessment is confirmed by the transmission percentages given in Tab. 4.3. The numbers indicate the fraction of particles from each field emitter that

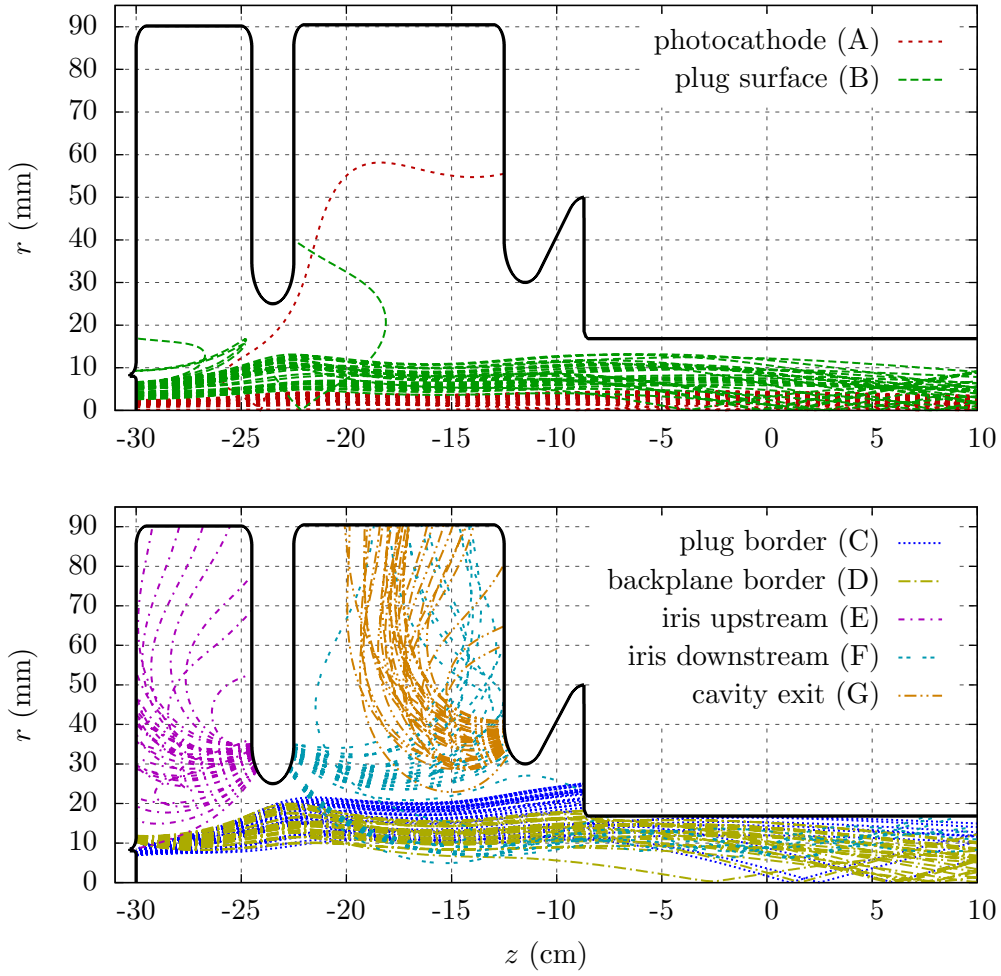

Figure 4.5: Trajectories of electrons from various field emitters in the rf gun cavity for a main solenoid peak field of 175 mT. The aperture model has been simplified by excluding the geometry of the input coupler from about −8 cm onwards.

reach a given position downstream of the gun cavity, namely, the first Faraday cup 2GUN, and the entrance of cryomodule ACC1. As a main result, electrons from the photocathode area have by far the highest probability of being transported into the beamline and, subsequently, into the cryomodule. It is also apparent that electrons from all emitters on the cathode plug or the backplane have at least a non-zero probability of reaching the module.

However, the energy of an electron must be sufficiently close to that of the beam to allow its transport over a serious distance. If just particles with a momentum above 1 MeV/c are counted, it turns out that only electrons from the photocathode or the surrounding surface of the molybdenum plug are thus transportable. Table 4.3 further indicates that practically all of the dark current from the Mo surface arriving at ACC1 is lower than 2 MeV/c in momentum unless the solenoid is tuned to a far weaker field than normal. In an analysis of dark current transport along the complete linac, all field emitters but the photocathode itself can therefore be neglected.

4.3.3 Measurement of the field enhancement parameter

The simulation of dark current dynamics in the gun cavity also allows to estimate the field enhancement parameter β_e from a measurement of the dark current against the electric field amplitude $\hat{E}$. In the FLASH injector, the measuring device closest to the gun cavity is the Faraday cup 2GUN, located about 1.2 m downstream of the photocathode. As set out before, only a fraction of the total dark current emitted at the photocathode can reach this Faraday cup. This transmission τ depends on the electric field amplitude itself and can be estimated by tracking simulations as before. Figure 4.6 summarizes the results and shows that $\tau(\hat{E})$ increases monotonically from about 11 % to 58 % when the field amplitude is increased from 30 to 43 MV/m while all other simulation parameters stay constant.

In preparation of the measurement, both solenoid magnets were set up according to the field strengths in the simulation, and a valve in front of cryomodule ACC1 was closed to prevent dark current from that module from reaching the Faraday cup. As the FLASH rf gun lacks a field probe, the field amplitude has to be determined indirectly by measuring the rf power sent into the cavity *(forward power P_f),* and the *reflected power* P_r. From electromagnetic calculations, the field amplitude can then then be approximated as [Bae03]

$$\hat{E} = 23.336 \, \frac{\mathrm{MV/m}}{\sqrt{\mathrm{MW}}} \sqrt{P_\mathrm{f} - P_\mathrm{r}}.$$

The resonance frequency of the copper cavity is tuned solely by its temperature, which can be regulated by the cooling water system. Because a variation of the forward power changes the heat load during the rf pulse, the cooling water temperature needs to be adjusted for each different setting of P_f. This prevents the cavity from detuning and thus increasing the reflected power. By observing this,

	Main solenoid peak field (mT)						
Emitter	**145**	**155**	**165**	**175**	**185**	**195**	**205**
Faraday cup 2GUN (percentage of all particles)							
photocathode	72.90	67.63	56.57	54.34	43.95	31.63	28.18
plug surface	10.04	4.93	2.88	2.33	2.13	2.23	2.47
plug border	15.96	16.22	11.00	4.04	0.27	0.19	0.22
backplane border	6.49	3.28	1.53	1.23	1.15	1.30	0.74
iris upstream	0	0	0	0	0	0	0
iris downstream	0	0	0	0	0	0	0
cavity exit	0	0	0	0	0	0	0
Faraday cup 2GUN (percentage of particles above 1 MeV/c)							
photocathode	72.35	67.00	56.02	53.71	43.25	30.99	27.48
plug surface	9.54	4.34	2.29	1.68	1.39	1.48	1.65
plug border	0	0	0	0	0	0	0
backplane border	0	0	0	0	0	0	0
iris upstream	0	0	0	0	0	0	0
iris downstream	0	0	0	0	0	0	0
cavity exit	0	0	0	0	0	0	0
Module entrance ACC1 (percentage of all particles)							
photocathode	53.58	43.83	32.55	20.99	14.68	11.02	8.82
plug surface	3.43	1.09	1.12	1.17	1.28	1.35	1.47
plug border	11.44	5.62	0.07	0.07	0.08	0.09	0.10
backplane border	0.16	0.24	0.33	0.41	0.49	0.58	0.65
iris upstream	0	0	0	0	0	0	0
iris downstream	0	0	0	0	0	0	0
cavity exit	0	0	0	0	0	0	0
Module entrance ACC1 (percentage of particles above 1 MeV/c)							
photocathode	53.30	43.49	32.21	20.64	14.28	10.62	8.41
plug surface	3.20	0.83	0.83	0.85	0.91	0.96	1.04
others	0	0	0	0	0	0	0
Module entrance ACC1 (percentage of particles above 2 MeV/c)							
photocathode	52.43	42.57	31.24	19.61	13.19	9.49	7.21
plug surface	2.51	0.07	0.02	$\ll$0.01	0	0	0
others	0	0	0	0	0	0	0

Table 4.3: Transmission of dark current from field emitters in the rf gun. The values specify the percentage of macroparticles that reach one of the given positions in the beam line after having been started from a field emitting surface. Faraday cup 2GUN is located 1.2 m downstream of the cathode, the entrance of ACC1 is at about 2.6 m. The values in each column represent a different setting of the main solenoid field. 175 mT is the standard value for operation of the rf gun.

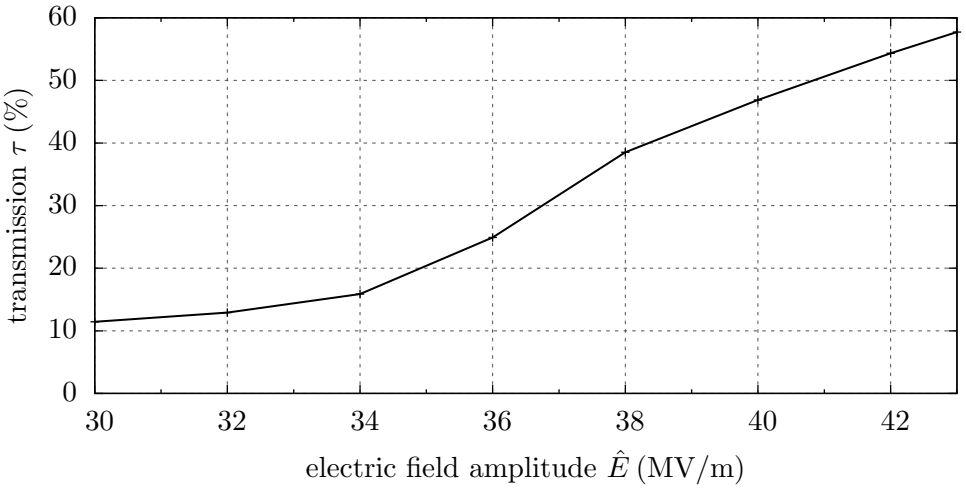

Figure 4.6: Transmission of dark current from the photocathode to Faraday cup 2GUN as a function of the electric field amplitude

the mean reflected power could be kept to about $P_\mathrm{r} \approx 25\,\mathrm{kW}$ during the measurement, which is almost negligible compared to 1.6–3.4 MW of forward power.

As shown in Fig. 4.7, the dark current $\bar{I}_\mathrm{meas}(\hat{E})$ measured at the Faraday cup ranges from 3 to about 220 µA. The total current emitted from the photocathode is higher and can be estimated with the simulated transmission function τ as $\bar{I}_\mathrm{meas}(\hat{E})/\tau(\hat{E})$; it ranges from 23 to 380 µA. To obtain a value for the field enhancement parameter β_e, these data are fitted with the mean dark current expected from the Fowler-Nordheim relation (4.1).

In its given form, (4.1) can only be used to calculate the momentary field emission current $I(E)$ for a value of the time-dependent electric field strength $E(t)$. What is measured, however, is an average current

$$\bar{I}(\hat{E}) = \int_0^{2\pi} I(\hat{E}\sin(\varphi))\mathrm{d}\varphi. \tag{4.2}$$

There is no closed analytic form for the exact averaging of the Fowler-Nordheim equation. Using a number of approximations, a widely used formula is derived by [Wan89]:

$$\bar{I}(\hat{E}) \approx 5.7 \cdot 10^{-12} \cdot 10^{4.52\phi^{-0.5}} \cdot \frac{A_\mathrm{e}(\beta_\mathrm{e}\hat{E})^{2.5}}{\phi^{1.75}} \exp\left(-6.53 \cdot 10^9 \frac{\phi^{1.5}}{\beta_\mathrm{e}\hat{E}}\right) \tag{4.3}$$

Results

If a work function of 3.5 eV for Cs_2Te is assumed as before, the only free parameters are A_e and β_e. A direct numerical integration of (4.2) and the approximation (4.3) are matched to the data points by a least-squares fit. The results are summarized

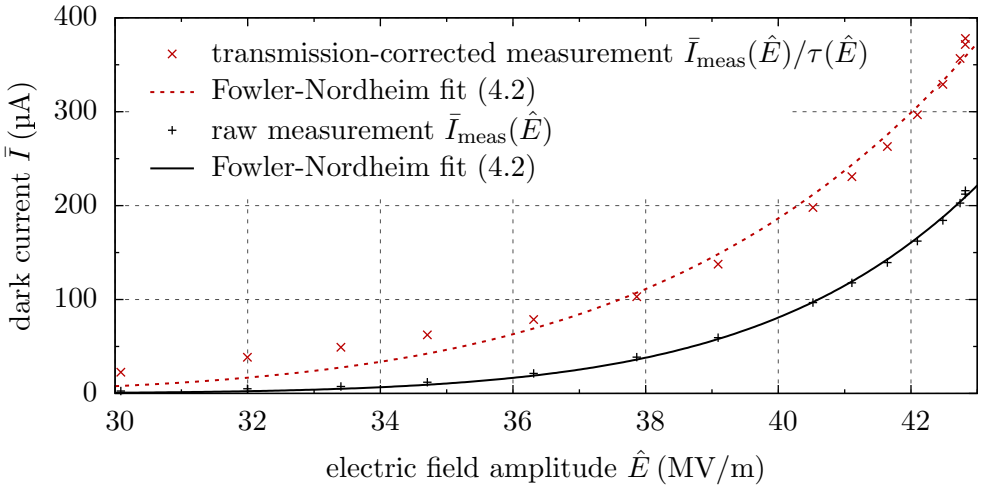

Figure 4.7: Gun dark current measurements for determination of the field enhancement parameter β_e. The *raw measurement* data series indicates the dark current measured at a Faraday cup, the *transmission-corrected* data series estimates the total dark current emitted from the photocathode.

Data set	Fit function	β_e	A_e (m^2)
corrected	exact (4.2)	142	$7.4 \cdot 10^{-16}$
corrected	approx. (4.3)	145	$6.9 \cdot 10^{-16}$
uncorrected	exact (4.2)	89	$8.5 \cdot 10^{-14}$
uncorrected	approx. (4.3)	90	$8.9 \cdot 10^{-14}$

Table 4.4: Fowler-Nordheim fit parameters for the gun dark current measurement.

in Tab. 4.4 and indicate a field enhancement factor of 142. Because it reproduces this value within 2 %, formula (4.3) can be considered a valid approximation for the conditions of the FLASH rf gun.

The surface quality of photocathodes can vary considerably. Hence, this measurement of the field enhancement factor cannot be representative for all photocathodes used at FLASH. However, it gives a useful exemplary value for β_e for the tracking simulations.

It should be noted that the fit does not reproduce the transmission-corrected data set as well as the uncorrected data set, which follows Fowler-Nordheim behavior much more faithfully (cf. Fig. 4.7). The evaluation of the raw measurement data yields a field enhancement factor of about 90. However, because of the long distance between detector and field emitter, the influence of the transmission $\tau(\hat{E})$ cannot be neglected. As long as measurements closer to the field emitter are not possible, tracking simulations give the best estimate of $\tau(\hat{E})$.

4.4 Start-to-end simulation setup

To study the transport of dark current along the whole accelerator, a detailed three-dimensional aperture model has been set up using the new features of the Astra tracking code described in appendix A. Apart from ordinary beam pipes, the most important elements included in the model are

- the gun cavity,
- the vacuum chamber of a spectrometer dipole behind the rf gun,
- collimators,
- cryomodules with cavities,
- bunch compressor vacuum chambers, and
- the undulator vacuum chambers.

The geometry of these components and their influence on the transport of dark current will be discussed in the following sections. First, a number of preparatory steps are outlined.

4.4.1 Setup of the accelerator model

The transport of electrons along an accelerator is influenced by a huge number of components. Apart from a description of the aperture, the main element classes to be incorporated into the accelerator model are the following:

- cavities
- solenoid magnets
- quadrupole magnets
- dipole magnets

The number of individual elements (about 70 quadrupoles and 50 cavities) makes a manual setup impractical. Therefore, a semi-automatic method has been chosen to assemble all the necessary information in an input file for the tracking code *Astra.*

An algorithm gathers information on the position, rotation, and effective magnetic length of all quadrupole magnets from the machine control system, and complements the information with the focusing strengths k from the design optics [Bal08]. A similar procedure is used for the main dipole magnets and their nominal bending radii ρ.

Because the superconducting cavities have individual quench limits, and due to beam dynamics considerations, they are operated with different accelerating gradients. For the simulation, the standard gradient distribution for high energy

operation from [Bal08] is applied to single cavity field profiles. Detailed information on the gradient distribution, the used field profiles, and on the conversion between accelerating gradients and electric field amplitudes can be found in appendix B.2. To reflect the conditions of typical FEL operation, acceleration module ACC1 is set up with an off-crest phase of $-8°$, ACC2 and ACC3 with $-20°$. ACC4–6 are set to accelerate the beam on the crest of the rf wave so that it reaches an end energy of about 980 MeV. Finally, the accelerator model uses a value of 175 mT for the peak field of the main gun solenoid, reflecting the typical operation point at FLASH.

4.4.2 Tracking of beam

Although the ultimate goal is an investigation of dark current transport, the accelerator model is first used to track the main electron beam. During this run, *Astra* determines a number of model- and energy-dependent values that are needed for the tracking of dark current. Before proceeding to the results of the beam tracking, this mechanism deserves some explanation.

Fixing of model- and energy-dependent parameters

In the initial input file, quadrupole field strengths are specified by the energy-independent focusing strength k. Because the tracking code needs to calculate the actual magnetic field B, it first tracks a probe particle and determines its momentum p upon entering the magnet. This information allows to calculate the field gradient $g = kp/e$, and, consequently, the well-known quadrupole field components

$$B_\mathrm{x} = gy \quad \text{and} \quad B_\mathrm{y} = gx.$$

This procedure works in a similar way for dipole magnets. Here, the energy-independent parameter is the bending radius ρ, and it is converted to the dipole field

$$B = \frac{p}{e\rho}.$$

Another important result of the beam tracking step is the definition of the phases of the rf fields. As set out before, the cavities are initially configured with a value φ for the phase of the beam with respect to the rf wave, in a way that $\varphi = 0$ corresponds to maximum energy gain of the beam. Because these phase values are relative to the a priori unknown beam arrival time, the tracking code needs to transfer them to a set of internal, absolute phases Φ before starting the main simulation. This is done by tracking a number of probe particles through the accelerator model until they have passed the first cavity. The tracking is repeated under variation of Φ until the probe particles obtain the maximum exit energy. The internal phase of the cavity is then fixed by adding the relative value φ configured in the input file. This *auto-phasing* procedure is iteratively applied for all cavities.

To summarize, the initial tracking with beam uses a number of automatisms to transform the relative quantities k, ρ, φ to the absolute quantities g, B, and Φ. As the dark current should be tracked under exactly the same conditions as the electron beam, the accelerator model must be set up using these absolute quantities. Formally, this is ensured by the following substitutions in the input file:

- quadrupole focusing strenght $k \longrightarrow$ magnetic field gradient g
- dipole bending radius $\rho \longrightarrow$ magnetic field strength B
- relative rf phase $\varphi \longrightarrow$ absolute phase Φ

Initial distribution

The initial particle distribution is mainly defined by the parameters of the injector laser system. The temporal shape—i.e., the distribution of emission times from the photocathode—is Gaussian with a width of 4 ps [Ayv06]. In accordance with operational procedures at FLASH, the phase of photoemission is chosen as $\varphi \approx -60^\circ$ with respect to the rf field as indicated in Fig. 4.8. This emission phase marks a compromise between minimization of the energy spread and maximization of the momentum of the electron bunch. For comparison, the figure also shows the simulated bunch charge arriving at a position 0.5 m downstream of the gun cavity as a function of φ.

Spatially, the particles are uniformly distributed over the area of a circle of 2 mm diameter, centered on the photocathode. The charge of the photoemitted electron bunch is set to 0.65 nC, a value typical for operation of the FEL. A total of 10^5 particles are simulated.

Main results

While the tracking of the beam itself is not the objective of these studies, it allows to verify the correctness of the accelerator model. Space charge effects—i.e. the expansion of the bunch due to Coulomb repulsion between the electrons—are included in the simulation only up to the entrance of the first bunch compressor. The influence of the emission of coherent synchrotron radiation in the dipole magnets is neglected. Therefore, the simulation cannot reproduce fine details in the bunch structure like other start-to-end simulations that use multiple specialized tracking codes (e.g. [Doh04]). However, only little difference should be expected in the collective properties of the electron bunch.

There is no prior experience in using *Astra* to model a complete accelerator of this complexity. It therefore lacks a number of output options that are common in other tracking codes—most notably, a calculation of the Twiss parameters. To circumvent this, the phase space of the beam is saved to disk in longitudinal intervals of about 15 cm. A custom external application is used to calculate the collective properties as set out in appendix A.

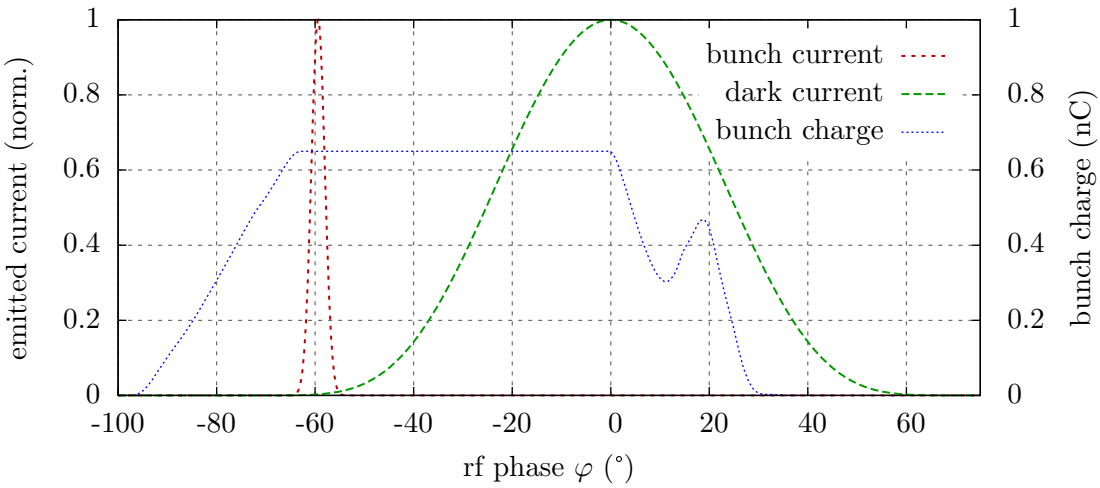

Figure 4.8: Phase dependence of field- and photoemission current. The red curve shows the bunch current for the typical phase of laser-induced photoemission, the green curve indicates dark current as predicted by the Fowler-Nordheim field emission model. For comparison, the plot also shows the result of a simulated phase scan of the laser pulse: the blue curve shows the bunch charge arriving at a position 0.5 m downstream of the gun cavity for a given photoemission phase.

Figure 4.9 shows some of the main results of the tracking: First, the trajectory of the center of mass of the beam lies well within the defined apertures. Second, the beam is accelerated to the expected beam momenta for a high-energy FEL run at FLASH, i.e. 4.9 MeV after the gun, 126 MeV after ACC1, 450 MeV after ACC3, and 980 MeV after ACC6. Third, the beam size is well-behaved with rms values well below 500 µm after BC3, and predominantly below 200 µm in the FEL undulators.

It is worth noting that the focusing strengths of five quadrupoles just downstream of BC2 have been adjusted to match the optics of the model to the design values from [Bal08]. This is one of the matching options commonly used for setting up the machine. The bottom plot of Fig. 4.9 shows that the beta functions yielded by the simulation are predominantly below 30 m behind BC2. They also reproduce two main features of the design optics, namely, the sharp focusing in the energy collimation section, and—albeit imperfectly—the periodicity in the undulator section.

4.5 Start-to-end tracking of dark current

The tracking simulation for the gun dark current is based on the same accelerator model as the simulation for the main beam; particularly, it uses the magnetic field gradients, dipole field strengths, and internal rf phases determined in the previous

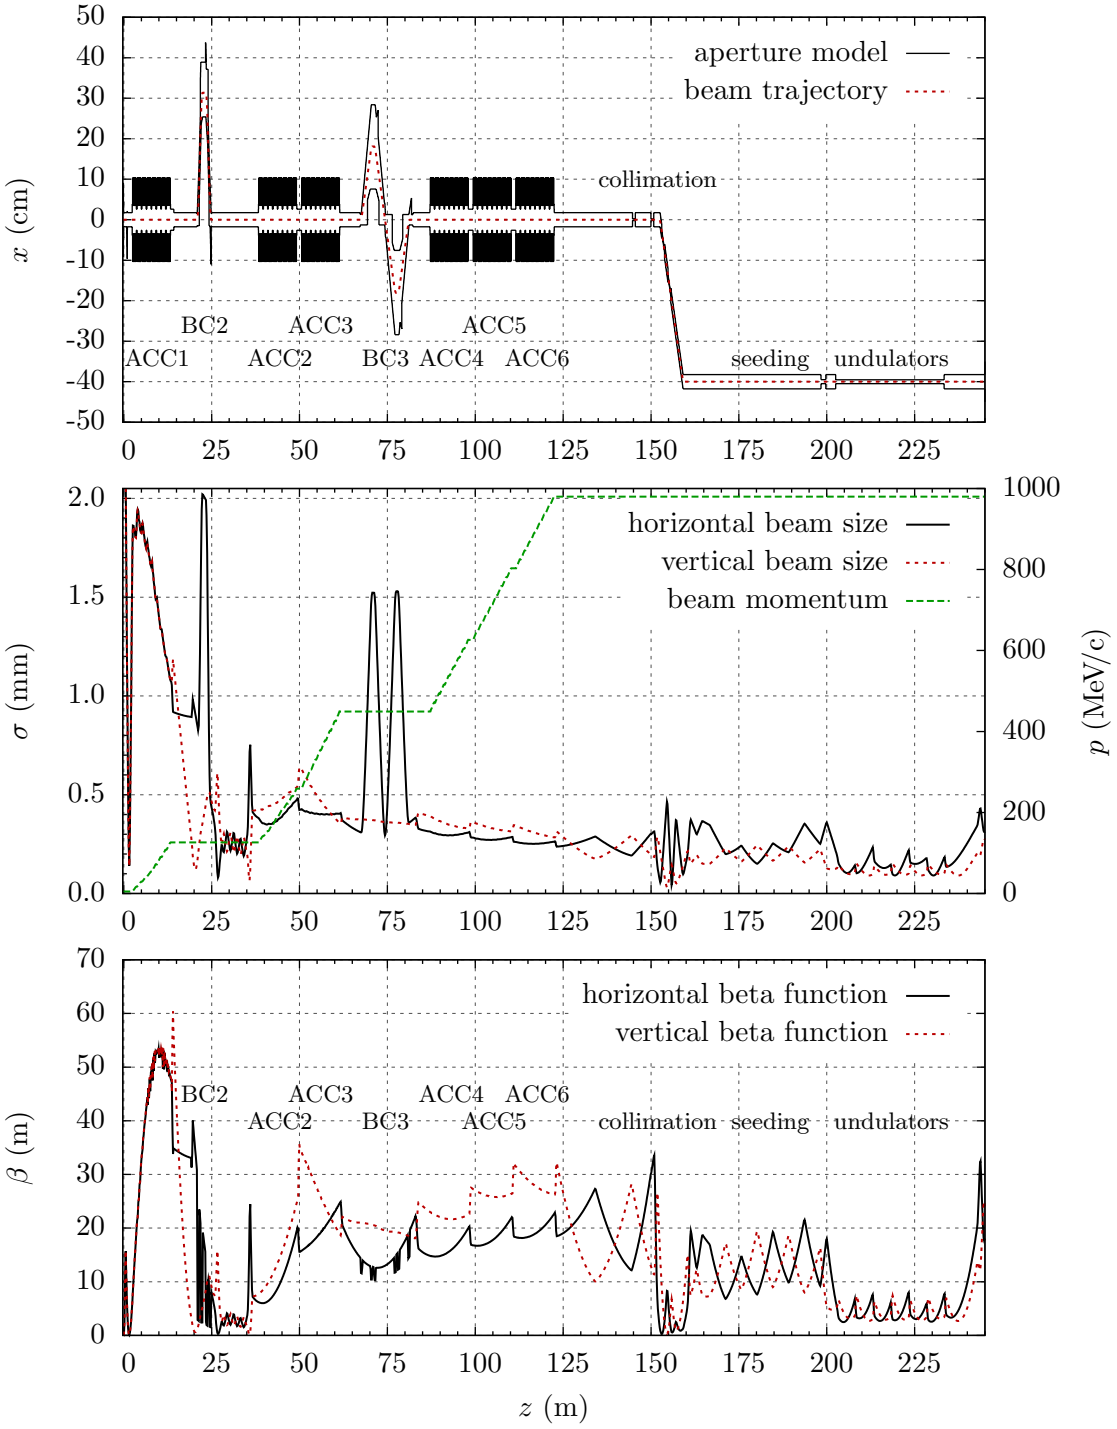

Figure 4.9: Main results of the beam tracking simulation. *Top:* horizontal section of the aperture model overlaid with the beam trajectory. *Center:* rms beam sizes and average beam momentum. *Bottom:* beta functions; distortions are visible in the bunch compressor and collimation sections because the algorithm for the calculation of β produces artefacts in dipole fields.

step. The decisive difference is the definition of the initial particle distribution.

The temporal and spatial distribution of dark current emission is the same as assumed above for the photocathode. As Fig. 4.8 indicates, the overlap with the typical emission phase range of the bunch is quite small.

The aperture model is the same as discussed before; for the initial tracking, the gun collimator and the BC2 collimator are removed while the fixed transverse and energy collimators are in place with the standard apertures described in section 1.4.3. The number of simulated macroparticles is chosen as 10^6 to obtain a sufficient resolution for small dark current losses.

4.5.1 Overview of simulation results

The number of macroparticles reaching a given beamline position can be extracted from the simulation output. In a similar way, it is possible to determine the number of macroparticles that have hit an aperture in a given interval along the longitudinal axis. To obtain more descriptive values, these particle numbers are converted to a current; this is done by normalizing the number of particles reaching the position of Faraday cup 2GUN to the measured dark current of 160 µA (cf. Fig. 4.6).

Before entering the discussion of individual beam line sections, the most important simulation results will be reviewed. Figure 4.10 illustrates that there are four locations with major dark current losses:

- the beamline between gun and ACC1 *(GUN)*
- bunch compressor BC2
- bunch compressor BC3
- the transverse collimators *(TCOL)*

The dark current losses are most intense in the GUN section. However, particle energies are still on the order of few MeV and thus significantly below the threshold for the giant dipole resonance. Hence, almost no radio-activation of components is caused (cf. section 2.2.4). The generated radiation field is nevertheless substantial and poses a hazard to electronic components in the vicinity.

The situation is more problematic in the sections behind the first cryomodule. About 70 % of the 85 µA of gun dark current accelerated in ACC1 are lost in the vacuum chamber of bunch compressor BC2. With electron energies above 100 MeV—and, consequently, bremsstrahlung photons up to the same energy—, both the giant dipole resonance and the quasi-deuteron effect contribute to the activation of materials, making the vacuum chamber and surrounding magnet yokes major sources of residual radioactivity. At higher electron energies, the bremsstrahlung spectrum extends to shorter wavelengths, and more kinds of photonuclear reactions become possible. Therefore, even the relatively small loss of 8 µA of dark current in BC3 can cause substantial activation. The transverse and energy

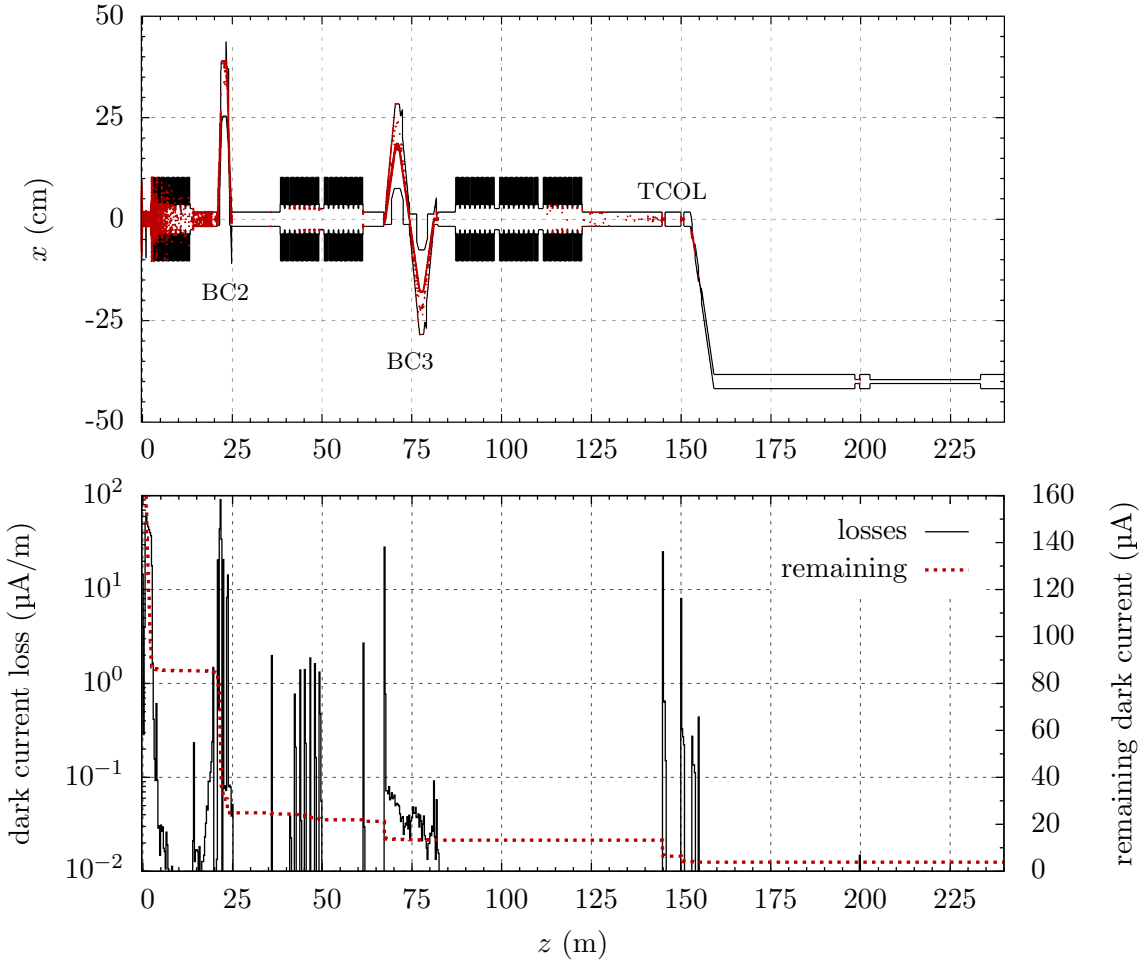

Figure 4.10: Dark current losses along the accelerator. The upper plot shows the aperture model of the FLASH linac and the position of lost dark current particles. The lower plot shows the relative beam loss per path length.

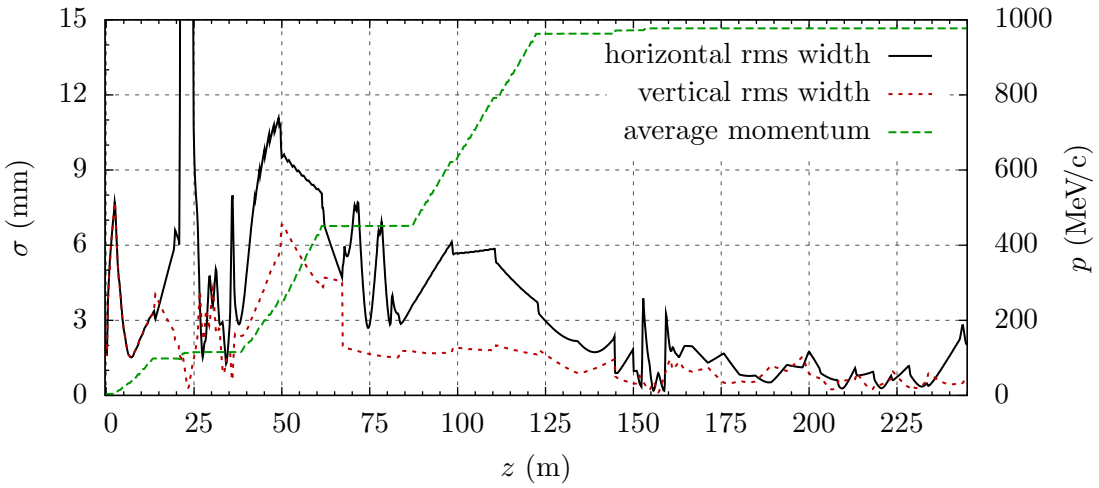

Figure 4.11: Transverse dimensions of gun dark current. The rms widths given here include only particles traveling in forward direction.

collimators reduce the remaining current from 13.2 µA to 3.9 µA. In the simulation, only minor losses are observed in front of the undulator.

To set the result of 3.9 µA of remaining dark current in perspective, it is helpful to calculate the total charge carried within a single rf pulse with a (flat top) length of 800 µs:

$$3.9\,\mu\text{A} \cdot 800\,\mu\text{s} \approx 3\,\text{nC}$$

If FLASH is thus operated at its nominal rf pulse length, the gun dark current reaching the undulator carries the same charge as three typical electron bunches.

Beside the transmission and the loss locations of the dark current, it is instructive to examine its transverse dimensions along the accelerator. In Fig. 4.11, the horizontal and vertical rms widths of the particle ensemble are shown together with its average momentum. The width of the dark current spot is of the order of several millimeters except in very few places. This is to be compared to the small dimensions of the electron beam shown in Fig. 4.9. It is also noticeable that the average momentum of the dark current rises in sections that do not contain any accelerating fields, e.g. around 25 m and 150 m. This effect is readily explained by the substantial loss of low-energy particles at obstacles—in these cases, at the BC2 vacuum chamber and at the collimators.

4.5.2 Comparison with measured radioactivity

At this point, it is useful to compare the simulated loss pattern from Fig. 4.9 with measured activities. Figure 4.12 shows the equivalent dose rate at the surface of various beamline elements along the accelerator. It represents the average of

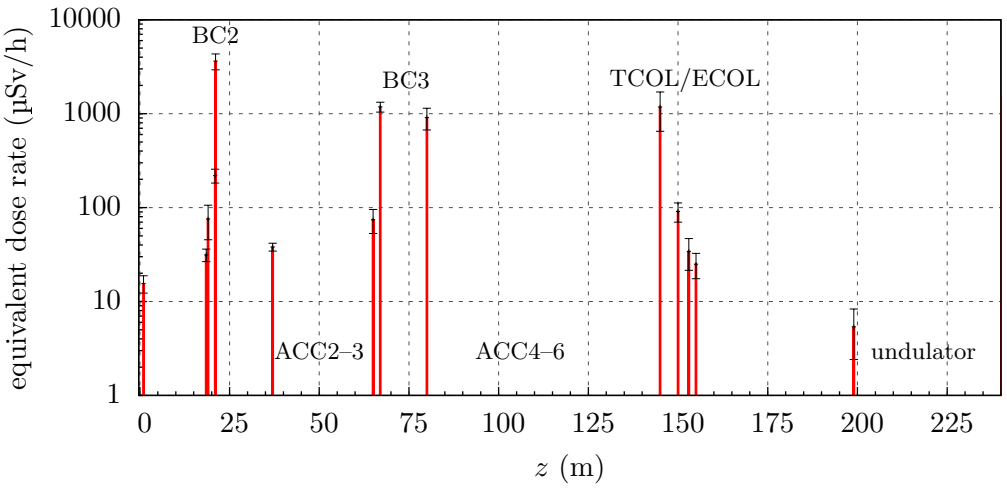

Figure 4.12: Measured activation of the FLASH beamline. The plot shows the residual equivalent dose rate few minutes after switching off the accelerator on the surface of selected beamline components. The data represent an average over measurements performed from July to December 2008 [TTF08].

11 measurements performed from July to December 2008 [TTF08]. The measurements were usually taken few minutes after the shutdown of the machine.

As predicted by the simulation, BC2 exhibits the highest activation along the accelerator with an equivalent dose rate of several thousand µSv/h. The loss peaks at 36 m, at the entrance of BC3, and at the transverse and energy collimators also appear as hot spots in the measurement. The dark current losses between 40 m and 50 m lie within cryomodule ACC2, and the loss at about 61 m is within cryomodule ACC3. These places are therefore not accessible to measurement.

A few places with high measured activities are not predicted by this simulation. First, the beam line between the gun and ACC1 shows an equivalent dose rate of more than 10 µSv/h, which can be explained with losses of dark current from the cryomodule. Second, the dose rate at the exit of bunch compressor BC3 is almost as high as at its entrance. Certainly, dark current emitted in the backwards direction from ACC4–6 contributes to the activation at this point; a detailed discussion of this section is given below. Third and last, minor activation of few µSv/h is measured at the undulator absorber at 199 m. To the most part, this can be attributed to the practice of using it as a temporary beam dump during machine setup.

The following sections give a more detailed analysis of the behavior of the gun dark current in various locations along the accelerator.

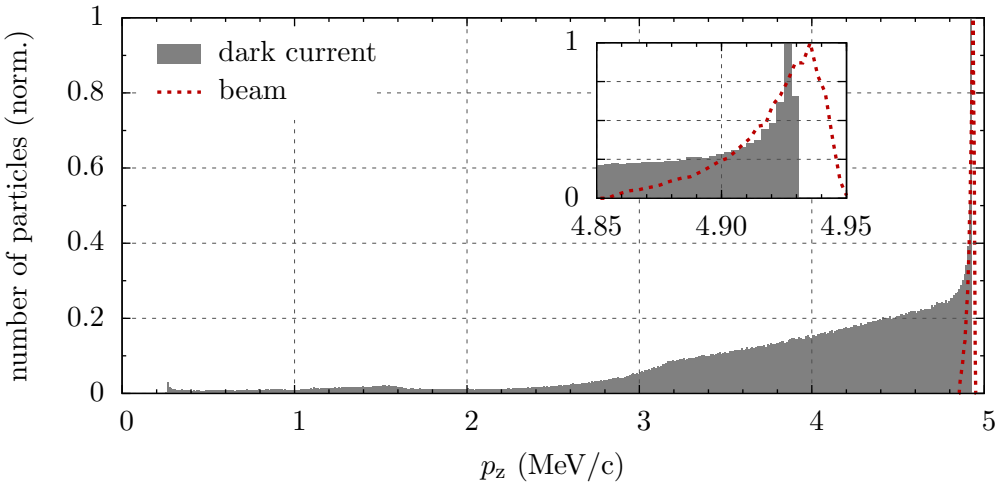

Figure 4.13: Momentum spectrum of dark current and beam 80 cm behind the gun cavity. The curves have been normalized to a maximum of 1.

4.5.3 Injector (rf gun and ACC1)

Just as the injector is of critical importance to the formation of the electron beam, it also defines the properties of the gun dark current that is transported downstream.

The main differences between beam and dark current are the size of the source area and the distribution of electron energies. Figure 4.13 compares the longitudinal momentum spectrum of beam and dark current after all particles have left the gun cavity. The dark current covers an extremely broad range of momenta from 200 keV/c to 4.93 MeV/c while the beam has a mean momentum of 4.92 MeV/c with an extremely small energy spread of about 20 keV. So, although the phase ranges for photo- and field emission hardly overlap at all (cf. Fig. 4.8), the resulting momentum spectra *do* have a considerable overlap.

The huge energy spread of the dark current also implies big velocity differences between individual electrons in the non- or low-relativistic regime. Thus, the particle ensemble must be expected to disperse longitudinally. This is illustrated in Fig. 4.14 which shows the longitudinal phase space of the dark current. Behind the main dark current "bunch" with an rms bunch length of more than 4.5 cm, two faint after-bunches can be distinguished between $z = 0.5$ m and $z = 0.6$ m. They consist of particles that were trapped in the half- or full cell of the gun cavity, respectively, for a full period of the rf wave.

Having established these characteristics, the spatial distribution of dark current losses can now be analyzed. Figure 4.15 gives an overview of the injector section. The top two plots illustrate location and intensity of the losses, the bottom plot compares the transverse dimensions of the beam and of the dark current.

As discussed before, the main solenoid magnet has a dominant influence on the

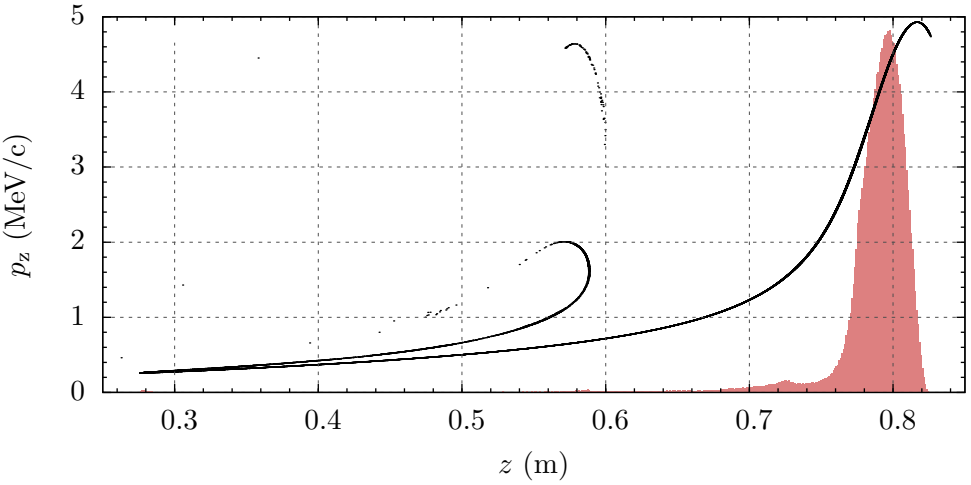

Figure 4.14: Longitudinal dark current phase space in the injector. The colored area indicates the longitudinal charge density on an arbitrary linear vertical scale.

transport of dark current. Located at about $z = 0$, i.e. 30 centimeters downstream of the cathode, the solenoid provides the appropriate focusing to obtain an electron beam of low emittance after the first acceleration module. Compared with the beam, most of the dark current has a substantially lower energy and gets strongly overfocused with a focal point at about $z = 16\,\text{cm}$. Beyond this point, the particle ensemble expands transversely and soon fills the complete beam pipe, causing almost constant losses from $z = 50\,\text{cm}$ to the entrance of ACC1.

A minor distortion is caused by the vacuum chamber of the *IDUMP* dipole magnet at about 1 m. The flat chamber is horizontally slightly wider than the neighboring beam pipes of 35 mm diameter, but has a height of only 27.5 mm. This creates a peak of losses at the entrance of the chamber, and a second one at its exit because of the reduction in horizontal aperture. Vertically, a piece of the downstream beam pipe is in the shadow of the vacuum chamber.

Due to the huge longitudinal extent of the dark current, only a part of it arrives at the right phase to be captured and accelerated by the rf in ACC1. Instead, many electrons are decelerated and deflected, hitting the irises of the first cavity or neighboring vacuum components. A small fraction is repelled completely and travels backwards—few particles even as far as the backplane of the gun cavity.

Figure 4.16 illustrates the bunching of dark current in ACC1. It shows the longitudinal phase space behind the acceleration module. The dark current electrons occupy several rf buckets. 98 % of the particles are concentrated in the first bucket, and about 1.5 % in the second. Consequently, the charge transported in later buckets is negligible.

One final observation on the transport within the cryomodule is that the dark

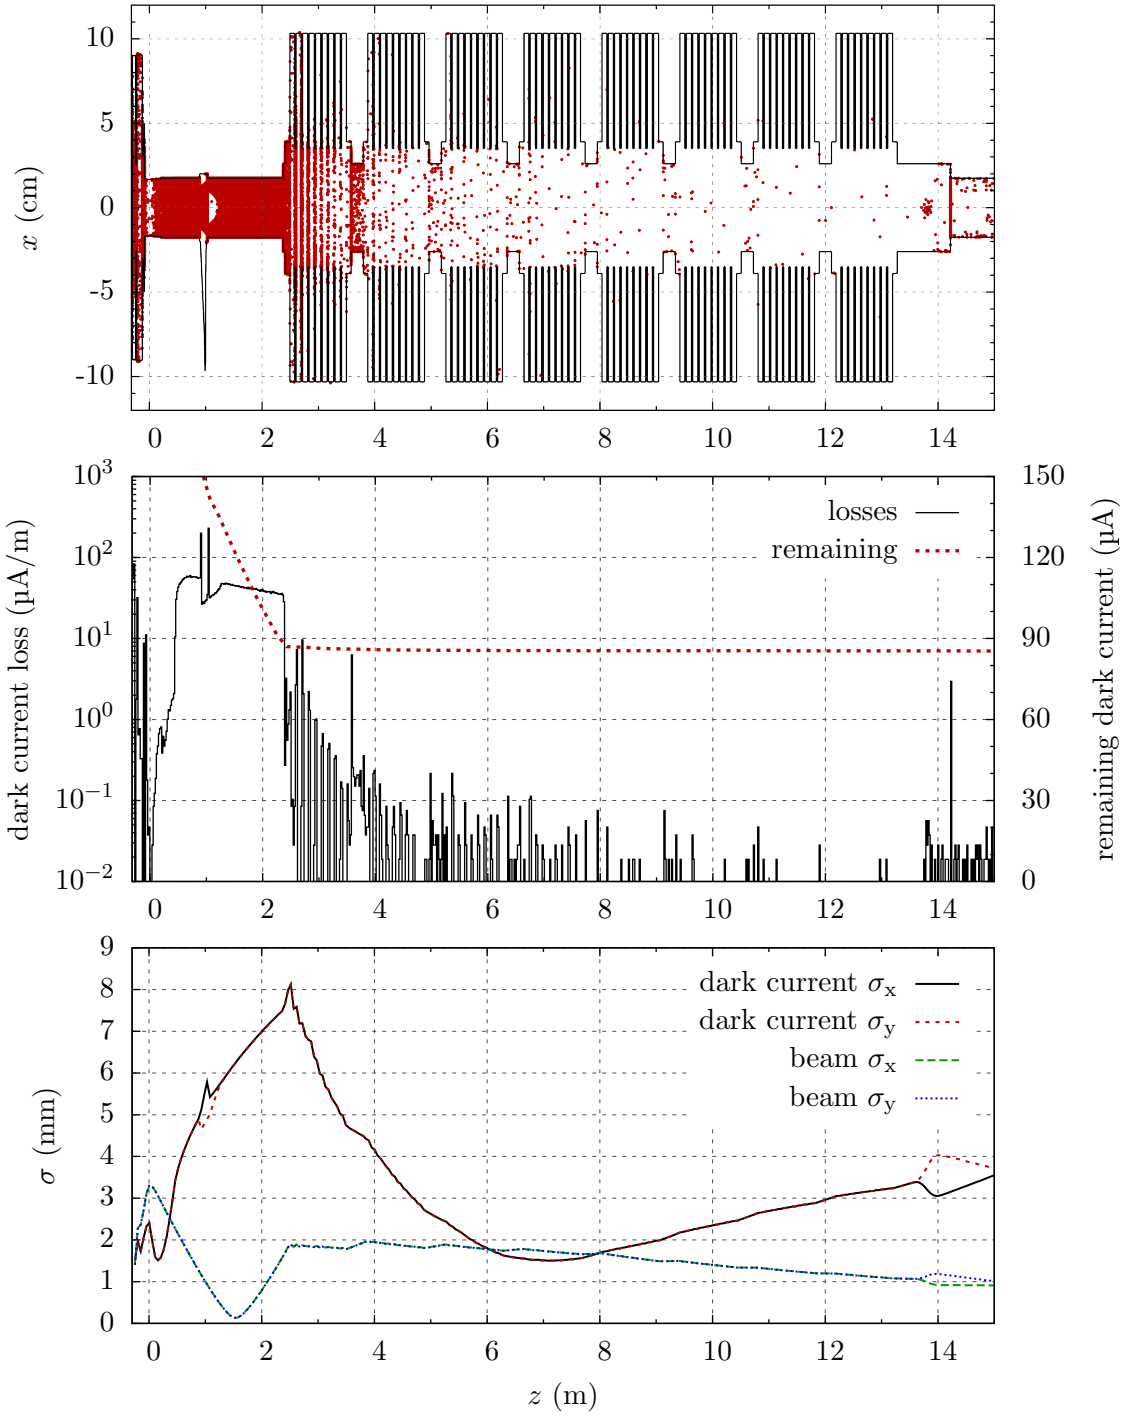

Figure 4.15: Dark current losses in the injector. *Top:* Positions of lost particles within the aperture model. *Center:* Dark current loss per beamline lenght and total amount of transported dark current. *Bottom:* Transverse rms width of beam and dark current.

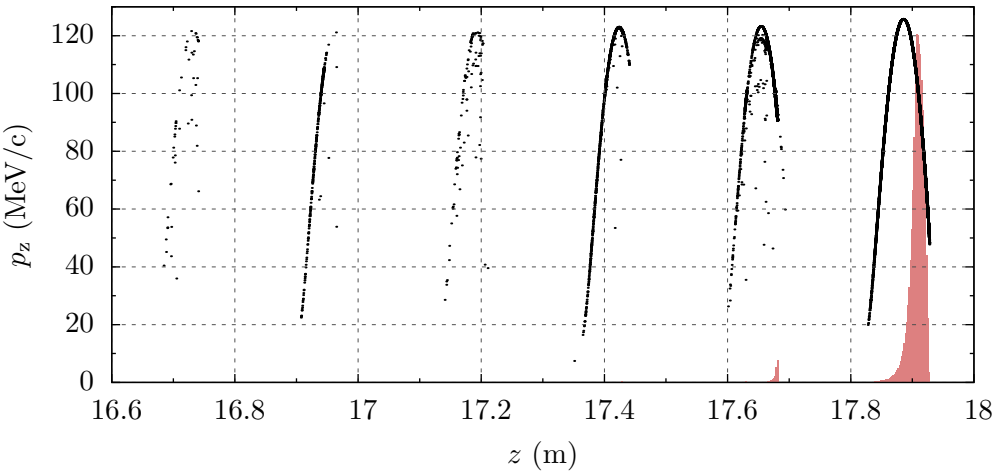

Figure 4.16: Bunching of dark current in ACC1. The plot shows the longitudinal phase space of the dark current behind the acceleration module. The colored area indicates the longitudinal charge density on an arbitrary linear vertical scale.

current is heavily influenced by cavity focusing. The electrons experience strong kicks towards the axis of the first cavities because of their huge transverse offset and low energy. The transverse momenta are big enough to lead to another over-focusing, with the focal point inside the cryomodule at about 7 m. Afterwards, the particle ensemble slowly expands to a width of slightly more than 3 mm.

To conclude the discussion of the injector section, it should be remarked that the quadrupole doublet at the exit of ACC1 ($z \approx 14$ m) is the first optical element that changes the transverse aspect ratio of both dark current and beam. This has a major impact on the transport of dark current in the following bunch compressor section.

4.5.4 Bunch compressor BC2

After exiting ACC1, the dark current is strongly divergent in the horizontal plane. Losses of low intensity occur up to 19.5 m, where a strong quadrupole triplet causes a minor increase by a further horizontal defocusing. Shortly after, the particles reach the vacuum chamber of bunch compressor BC2.

The bunch compressor consists of four dipole magnets arranged around a flat vacuum chamber with an internal height of 8 mm. Its geometry is detailed in Fig. B.3. The aperture model used for tracking includes the full geometry information except for minor simplifications at the synchrotron radiation ports. In the central dispersive section, the horizontal aperture amounts to 135 mm.

As shown in Fig. 4.17, 5.3 µA of dark current are lost at the transition from the circular beam pipe to the flat vacuum chamber. This is, however, only a minor

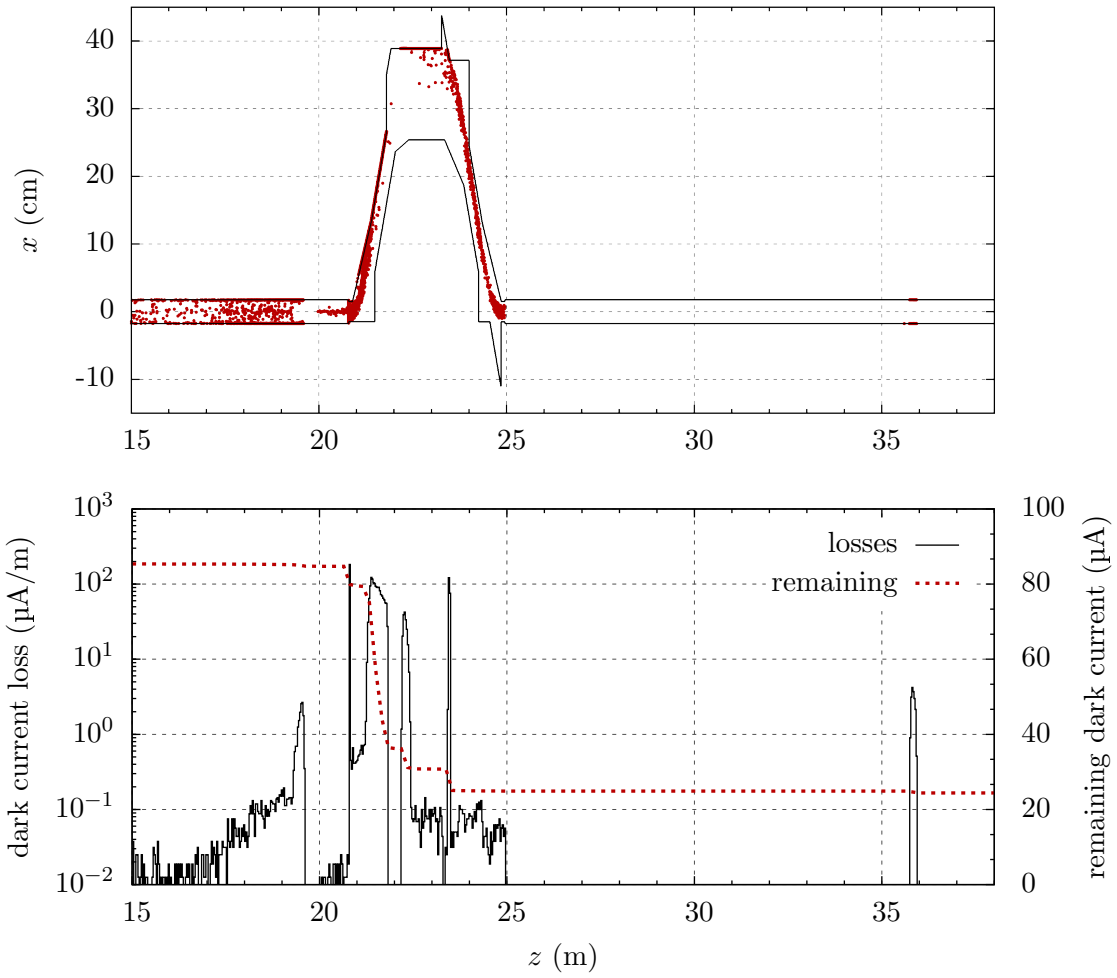

Figure 4.17: Dark current losses in BC2. *Top:* Positions of lost particles within the aperture model. *Bottom:* Dark current loss per beamline lenght and total amount of transported dark current.

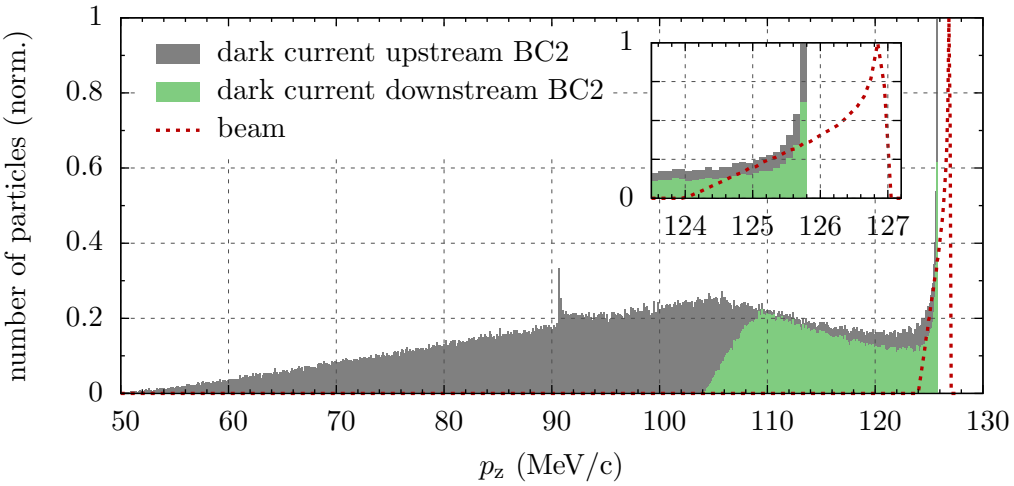

Figure 4.18: Momentum spectrum of dark current and beam before and after passing BC2. The curves have been normalized to a maximum of 1. The inset plot shows the high energy range in more detail.

fraction of the total loss in BC2. The major part of the electrons, about 43 µA, hits the left wall of the chamber after the first bend. This is a consequence of the low energy of the dark current. For the same reason, another 11 µA are lost on the outer walls of the chamber between the two inner dipoles. After the third dipole magnet, only minor losses occur, mainly at the bottom and the top of the chamber.

The bunch compressor acts as a momentum filter for the dark current by letting only particles above a certain threshold energy pass. Figure 4.18 shows the longitudinal momentum spectrum of dark current and beam before and after BC2. For the standard settings of the bunch compressor used in the simulation, the threshold energy is about 104 MeV while the beam covers the range between about 124 and 127 MeV. As mentioned before, the BC2 scraper can be moved into the central part of the chamber from the outside. It has a huge potential to raise the threshold energy and intercept the biggest part of the dark current; a quantitative discussion is given in section 4.5.9.

Behind the bunch compressor, the dark current is transported without further losses up to about 36 m. There, the final three quadrupoles of this beamline section create a huge horizontal bulge of more than 8 mm rms width, pushing a significant part of the current into the left and right wall of the beam pipe. Afterwards, the remaining dark current of about 24 µA enters acceleration module ACC2.

4.5.5 Acceleration modules ACC2–3 and bunch compressor BC3

Unlike at ACC1, the dark current is already bunched when entering ACC2. However, its longitudinal phase space is strongly influenced by the large momentum

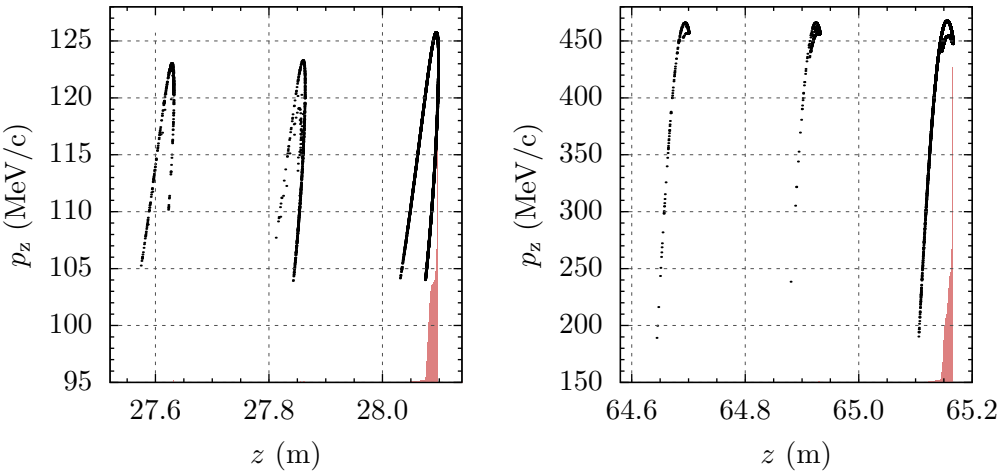

Figure 4.19: Longitudinal dark current phase space before and after acceleration in ACC2–3. The colored areas indicate the longitudinal charge density on an arbitrary linear vertical scale.

compaction of BC2. A rough estimate of the path length difference between high and low energy particles can be obtained by a quick linear beam optics calculation. The energy spectrum of the dark current covers the range from 114 to 126 MeV, and the momentum compaction factor of BC2 amounts to $\alpha_c = 180.6\,\text{mm}$ (cf. section 1.4.2). Therefore, particles at the low end of the spectrum will lag about

$$\Delta s \approx \alpha_c \cdot \frac{126 - 114}{126} \approx 2\,\text{cm}$$

with respect to those of highest energy. As a consequence, the longitudinal phase space gets considerably skewed. This effect is visible in Fig. 4.19, which also shows the phase space after acceleration in ACC2–3. Again, the $\sin(\varphi)$ dependence of the rf wave is imprinted on the energy profile of the particle ensemble. But due to the skewness of the initial distribution, two distinct energy levels now coexist at the same longitudinal position. Later rf buckets are almost completely negligible at this stage, as more than 99 % of the remaining particles are now contained within the first bucket.

After these initial considerations, the focus can be put on the simulated dark current losses shown in Fig. 4.20. In ACC2, the dark current undergoes a quick transverse expansion. Because of the prior vertical collimation by the BC2 vacuum chamber, losses occur only in the horizontal plane in the most narrow parts of the cavity string. The dark current is focused by two quadrupoles between the cryomodules. No more losses occur until the exit of ACC3, where a part of the dark current is stripped off horizontally. In total, about 3.2 µA are lost in ACC2 and ACC3. Assuming a worst-case scenario where all the energy of the lost electrons

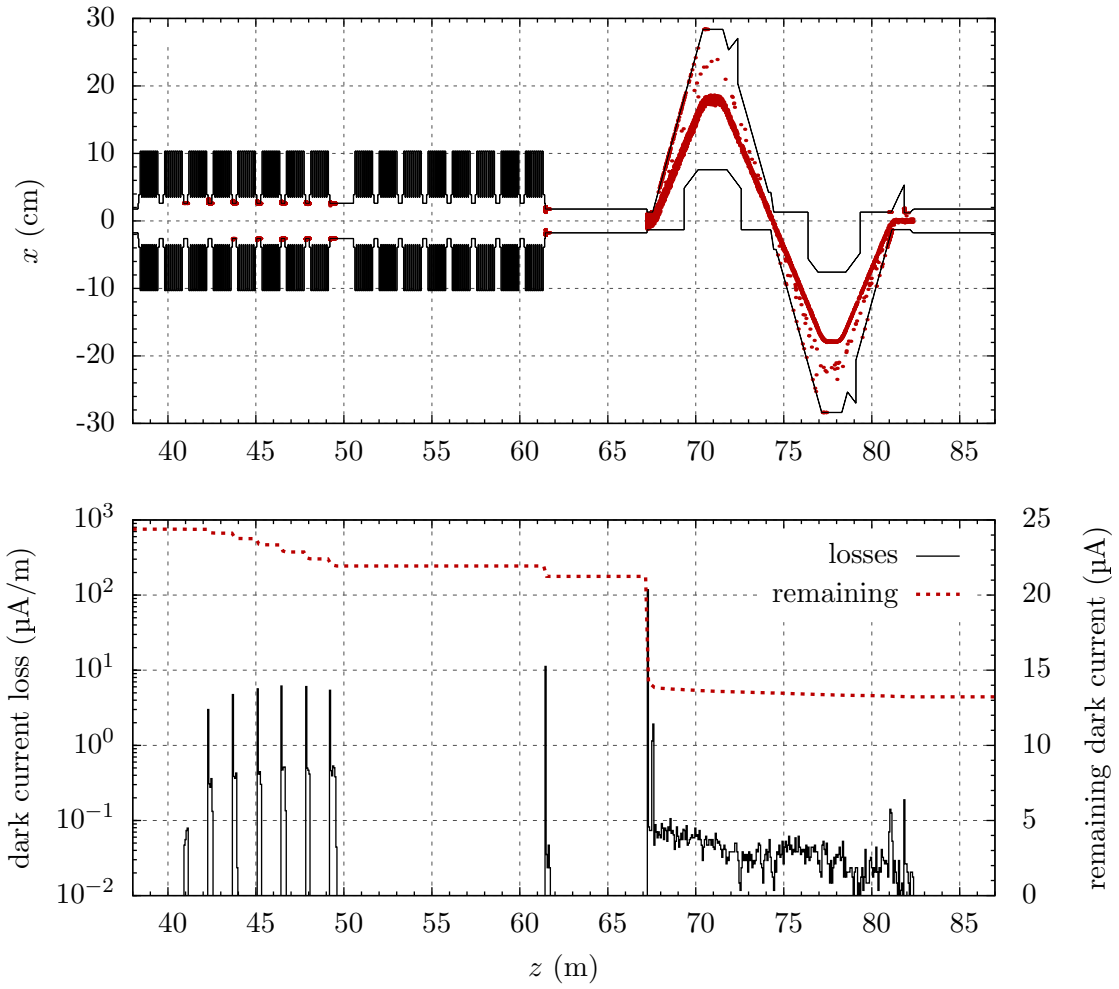

Figure 4.20: Dark current losses in ACC2–3 and BC3. *Top:* Positions of lost particles within the aperture model. *Bottom:* Dark current loss per beamline lenght and total amount of transported dark current.

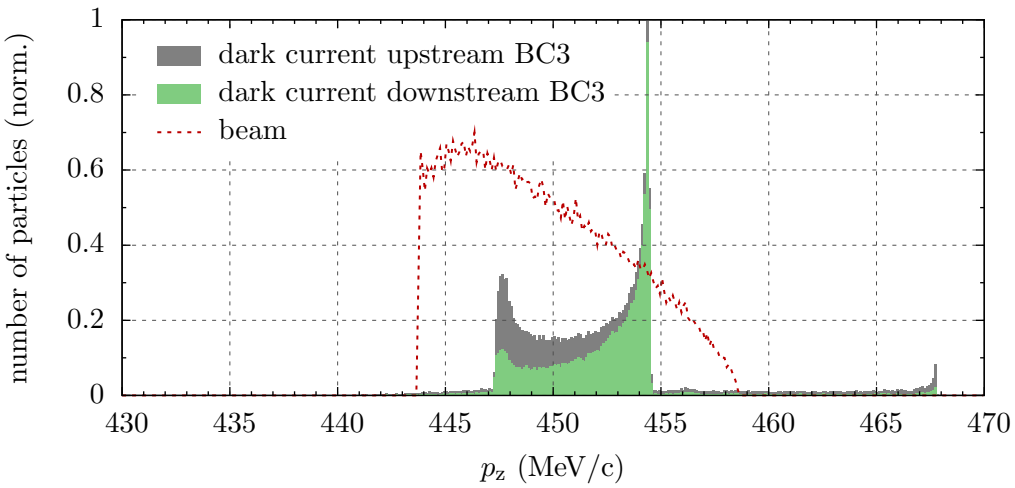

Figure 4.21: Momentum spectrum of dark current and beam before and after passing BC3.

is absorbed within the cryomodules, this could generate a maximum heat load of

$$P = 3.2\,\mu\text{A} \cdot 450\,\text{MeV} \cdot 800\,\mu\text{s} \cdot 10\,\text{Hz} \approx 12\,\text{W}$$

for the design parameters of FLASH. While this is well within the capacity of the cryogenic system, it constitutes a notable contribution to the average load.

The next major obstacle for the transport of dark current is bunch compressor BC3. Like BC2, it has a flat vacuum chamber of 8 mm internal height. The deflection of the beam is provided by six dipole magnets, in an arrangement that is often described as an "S"-shaped chicane. A detailed drawing of the vacuum chamber is found in Fig. B.4. As for BC2, the aperture model accurately reproduces the geometry of the chamber except for minor simplifications at the synchrotron light ports.

BC3 hardly acts as an energy collimator for the dark current. Instead, almost the complete loss of 7.5 µA occurs at the transition from the beam pipe to the low vacuum chamber. Low but almost constant vertical dark current losses continue until the end of the chamber. The trajectory of the main dark current part lies very close to the outer wall of the vacuum chamber in the last dipole. A minor horizontal missteering could increase losses at this point strongly and thus contribute to the observed activation of the beam line.

Before entering the discussion of the next acceleration stage, it is helpful to compare the momentum spectrum of the dark current with that of the beam (Fig. 4.21). While the vast majority of dark current electrons has an energy between 447 and 455 MeV, a weak branch extends up to 468 MeV. The spectrum of the beam covers energies between 443 and 459 MeV, and thus overlaps with more than 90 % of the dark current.

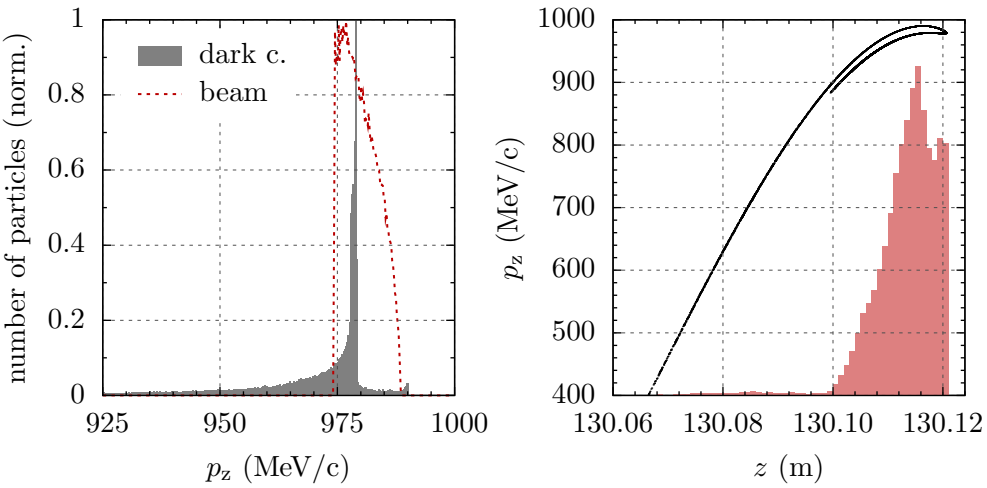

Figure 4.22: Momentum spectrum and longitudinal phase space of dark current after ACC4–6. The colored area in the right plot indicates the longitudinal charge density on an arbitrary linear vertical scale.

A number of effects contribute to this result. First, the acceleration of the beam at an off-crest phase of 20° has two major consequences: it increases the energy spread of the beam itself (cf. Fig. 1.5), and it allows part of the dark current to get a higher energy gain than the beam due to acceleration on the crest of the rf wave. The second effect reminds of the synchrotron oscillation in circular accelerators: the bulk of the dark current is accelerated on a falling slope of the rf field in ACC1 and thus imprinted with a negative correlation of energy versus longitudinal coordinate ($\mathrm{d}E/\mathrm{d}z$). As seen in Fig. 4.19, this correlation is inverted by the momentum compaction of bunch compressor BC2. The acceleration in ACC2–3 again happens on the falling slope of the field so that the existing positive correlation can partly cancel the newly imprinted, negative one. As a result, the net energy spread is reduced.

4.5.6 Acceleration modules ACC4–6 and collimation section

The final acceleration in cryomodules ACC4, ACC5, and ACC6 brings both the beam and the dark current to an energy of about 980 MeV. While the beam and the main part of the dark current are accelerated on-crest, the backmost dark current electrons are transported far off the crest of the rf wave, creating a huge tail in the energy spectrum (Fig. 4.22). Single electrons even get decelerated and are lost in ACC6 or in the adjacent beam line.

The particles enter the main collimation section at 145 m. At this point, the dark current has an rms spot size of about 2.3 mm in the horizontal and 1.4 mm in the vertical plane. Hence, the first transverse collimator with an inner radius of 2 mm is quite effective. As Fig. 4.23 shows, it removes roughly half of the incoming

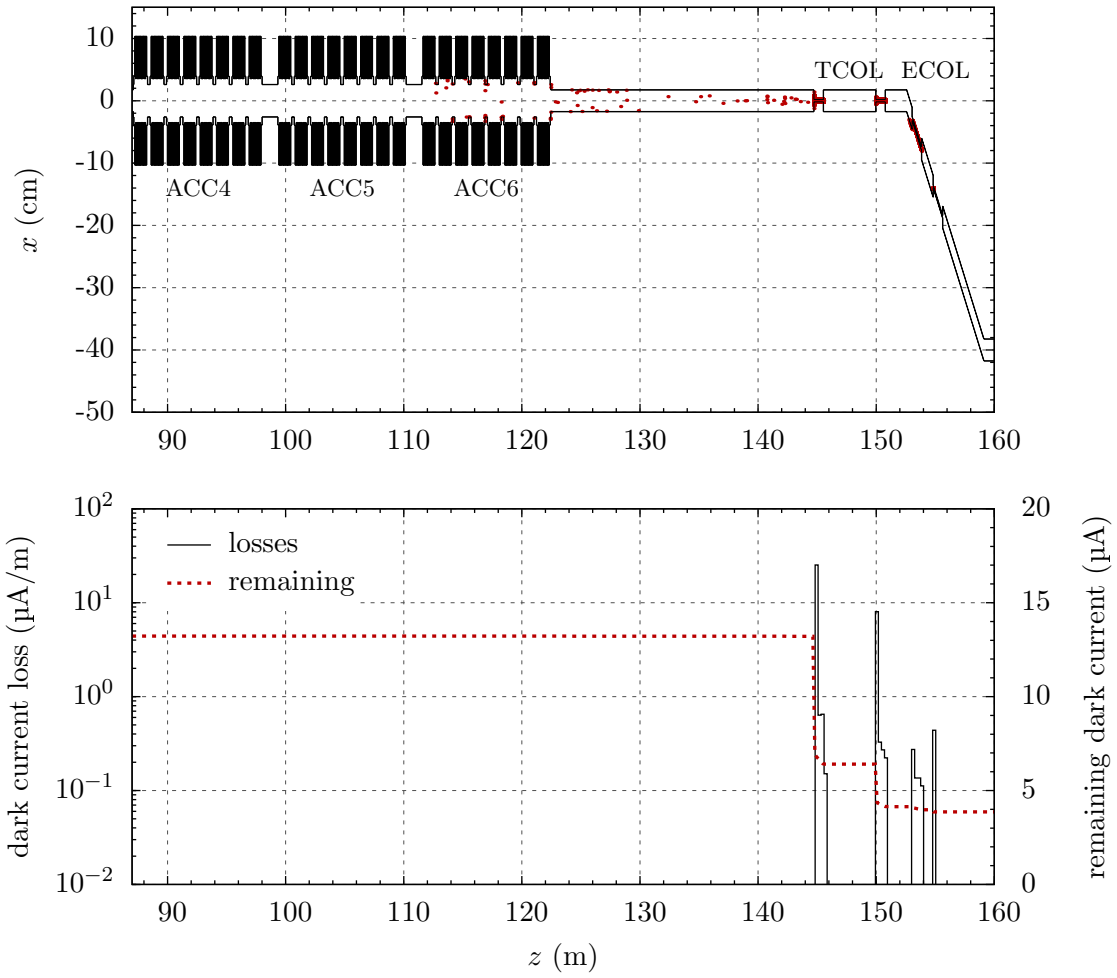

Figure 4.23: Dark current losses in ACC4–6 and in the collimation section. The TCOL section contains two transverse collimators, the ECOL section contains two energy collimators. *Top:* Positions of lost particles within the aperture model. *Bottom:* Dark current loss per beamline lenght and total amount of transported dark current.

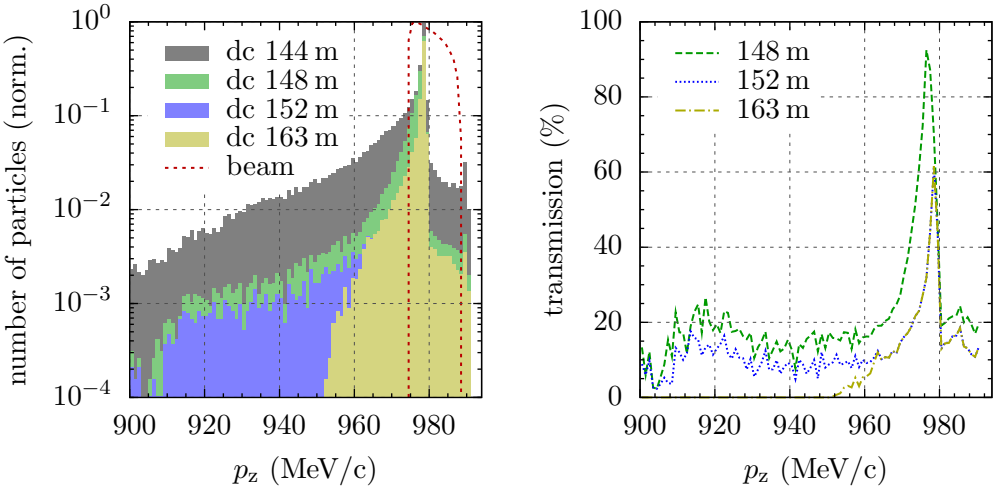

Figure 4.24: Energy dependence of dark current collimation. Longitudinal positions: 144 m—upstream first collimator, 148 m—downstream first collimator, 152 m—downstream second collimator, 163 m—downstream all collimators. *Left:* Longitudinal momentum spectra. *Right:* Transmission through the collimators relative to the amount of dark current before collimation.

dark current (∼7 µA). The second transverse collimator intercepts another 2.3 µA, and the energy collimators show almost no influence.

The reason for this is illustrated in Fig. 4.24 which compares the longitudinal momentum spectra of the dark current before collimation, downstream of the first collimator, downstream of the second collimator, and downstream of all collimators. The transmission through the transverse collimators is energy dependent. This is a consequence of the fact that the dark current is heavily subjected to chromatic effects, especially to energy-dependent focusing in quadrupole fields, which cause a correlation between the transverse offset of a particle and its energy in the collimator. The effect is so strong that the dark current spectrum has almost reached its final shape downstream of the second collimator. The energy collimators with their nominal energy aperture of ±3 % only cut off a minor tail of electrons below 950 MeV. If the *collimation efficiency* of the section is defined as

$$\eta = \frac{\text{incoming dark current} - \text{transmitted dark current}}{\text{incoming dark current}},$$

a value of $\eta = 71\,\%$ is obtained. A collimation efficiency of 100 % is not achievable due to the considerable overlap between beam and dark current in the transverse plane and in the energy spectrum.

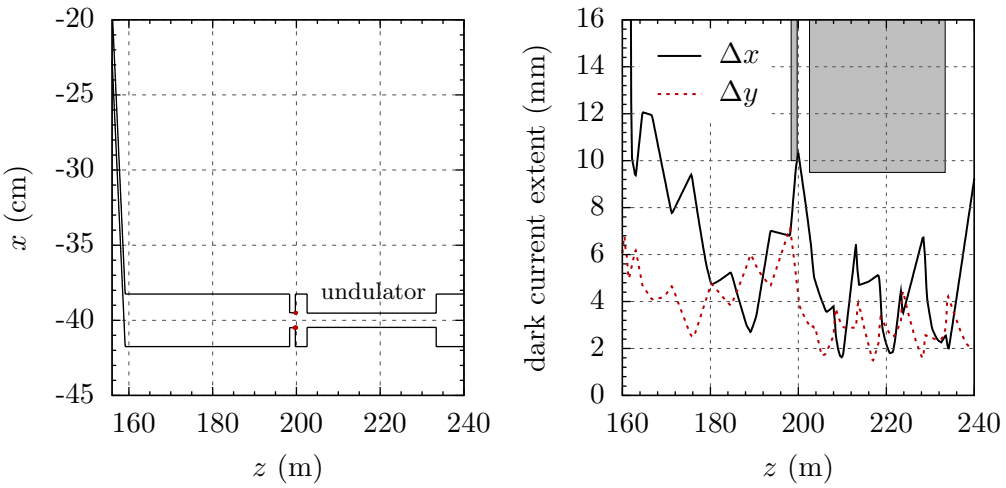

Figure 4.25: Aperture model and transverse extent of dark current in the FEL beamline. The shaded areas in the right plot indicate limited beam line and vacuum chamber apertures.

4.5.7 FEL beamline

Behind the collimators, almost no losses of dark current occur in the simulation. The only exception is a beam pipe of 10 mm diameter at about 200 m, in front of the FEL undulators, where a few particles narrowly exceed the available transverse aperture. Figure 4.25 shows the maximum transverse extent of the dark current particle ensemble, i.e.

$$\Delta x = x_{\max} - x_{\min}$$
$$\Delta y = y_{\max} - y_{\min},$$

as a function of z. The plot indicates that the losses are caused by a horizontally defocusing quadrupole[1] just in front of the narrow beamline. Although the nominal current loss in this scenario amounts to only 5 nA, it will be increased by any missteering or misalignment.

Except for dedicated collimators, the vacuum chambers of the FEL undulators constitute the smallest aperture along the whole accelerator. If the inner diameter of 9.5 mm is compared to a horizontal dark current extent of up to 7 mm, it is clear that the room for steering of the beam is extremely limited. Steering in the vertical plane is slightly more relaxed due to the maximum extent of 5 mm. This situation is problematic because of the radiation sensitivity of the undulator magnets. Especially during the setup of an FEL run, the overlap between the photon- and electron beam has to be established by steering. Because dark current losses occur much earlier than beam losses, the efficiency of this setup procedure

1 Q16SEED at 198.16 m

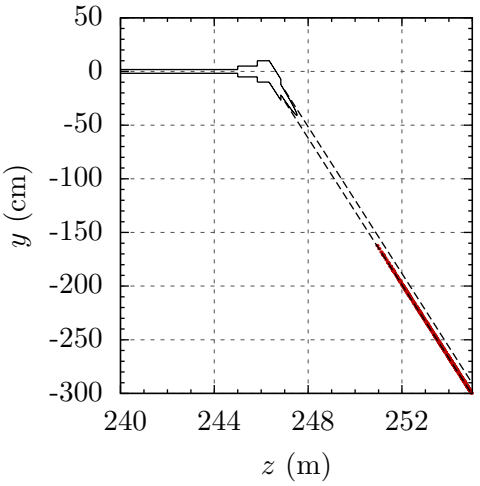

Figure 4.26: Aperture model of the dump line. The figure shows a vertical section through the aperture model at $x = -40\,\mathrm{cm}$; as the dump line is tilted towards positive x, the beam pipe leaves the section plane and is only displayed incompletely. Dark current losses are indicated in red.

and, ultimately, the performance of the FEL are impaired. The following chapter will resume the discussion of dark current losses in the undulator.

4.5.8 Dump

The electrons are deflected downwards into the dump by a strong dipole magnet with a nominal bending angle of 19° at the end of the FEL beamline. As the dipole is rotated by 7° around the z axis, the beam experiences an effective bend of 2.4° in the horizontal plane, and a bend of slightly less than 18.9° in the vertical plane. The main beam pipe with its diameter of 10 cm follows this direction, and the aperture model is set up accordingly. Because no exact geometry data are available on single vacuum components of the dump line, the model of the dipole chamber has been assumed to be bounded by two parallel, vertical planes spaced at 6 cm with a vertical aperture of 20 cm. No further geometric details are included.

Figure 4.26 shows a vertical section through the aperture model at $x = -40\,\mathrm{cm}$, i.e. through the axis of the FEL beamline. Because the dump line is tilted towards positive x, the beam pipe leaves the section plane and is thus only displayed incompletely. In the dump line, the dispersion created by the dipole is not compensated. Only two quadrupoles at 250.2 m and 251.5 m can be used to defocus the beam in order to reduce the power density at the exit window separating the accelerator vacuum from the beam dump. In the default mode of operation at the time of writing—and hence, in the simulation—, the quadrupoles are switched off. In this case, the dark current is left to expand due to the increasing vertical dispersion. The low-energy part of the dark current collides with the lower part of the beam pipe starting from about 251 m.

Up to the dump dipole, the simulation includes no element that could introduce vertical dispersion. Hence, the geometric arguments discussed in appendix C apply and the vertical dispersion can be calculated from the bending angle and the effective length of the magnet using (C.4). The result is plotted in Fig. 4.27; the

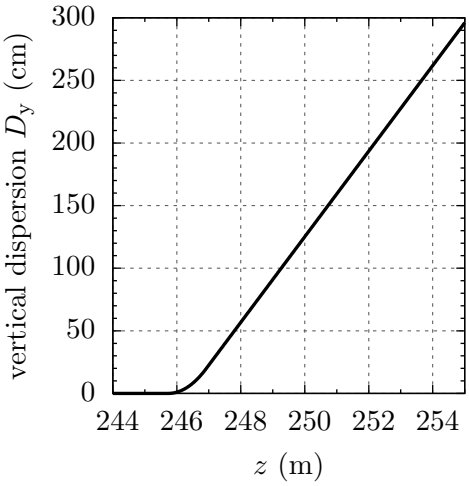

Figure 4.27: Vertical dispersion in the dump line. It is assumed that the quadrupole magnets downstream of the dipole are switched off.

dispersion grows from zero to about 3 m at the dump window. Recalling that the dark current spectrum stretches to the lower end of the main collimators' energy aperture, the maximum energy deviation is $\delta = -3\,\%$ (or, in absolute values, 950 MeV compared to a beam energy of 980 MeV). Considering the radius of the beam pipe, $r_{\max} = 5\,\mathrm{cm}$, the dark current must therefore collide with the beam pipe when

$$r_{\max} \leqslant D_{\mathrm{y}}\delta \quad \Leftrightarrow \quad D_{\mathrm{y}} \geqslant r_{\max}/\delta,$$

i.e. at a vertical dispersion exceeding 166 cm. Consistent with simulated losses, this is the case from $z = 251\,\mathrm{m}$ onwards. Although the first quadrupole could reduce the dispersion starting at about 250 m, it has already reached a value of 130 cm at this point, corresponding to a maximum excursion of 3.9 cm. In terms of dark current losses, the dump line is therefore extremely sensitive to missteering and to energy variations.

4.5.9 Counter measures

Mainly because of the activation of beam line components, dark current losses remain a major problem in the operation of FLASH. Apart from the fixed collimators discussed above, three devices have been implemented with the purpose of restricting its transport along the linac:

Dark current kicker: This is a fast kicker magnet located just behind the main gun solenoid. It is driven by an alternating current with a frequency of 1 MHz. The phase of this current is synchronized to the bunch clock so that the electron bunches pass the kicker on the zero crossing of the magnetic field and thus experience little or no influence. The gun dark current arriving in the time between the bunches, however, is deflected transversely to a varying degree. Because of insufficient phase stability, the device has found relatively little use during normal operation of the FEL. In addition, a huge part of

the kicked dark current is still transported to the first bunch compressor and adds to the activation caused there. Recent developments including a major improvement of the phase stability and the change of the drive signal from a sine to a rectangular wave have shown promising results and could solve most of these problems [Obi09].

Gun collimator: The main goal of the kicker is to remove as much of the dark current as possible before it can be accelerated by ACC1. The same is true for the removable gun collimator, located about 1.5 m downstream of the cathode. Since it is also located downstream of the kicker, it can support its operation by intercepting a big part of the kicked dark current.

BC2 collimator: As described before, the BC2 collimator or scraper is a copper block that can be inserted into the vacuum chamber of bunch compressor 2. Its position can be varied freely.

This section focuses on the influence of the two collimators. Since they are removable, FLASH is often operated without them. An analysis of their efficiency is therefore crucial to improve the operational procedures of the accelerator. For this purpose, the tracking simulation described above is repeated with an altered aperture model for each combination of collimator settings.

Gun collimator

Insertion of the 8 mm diameter gun collimator allows for undisturbed passage of the beam but removes almost two thirds of the incident dark current (i.e., 87 µA out of 136 µA). However, a huge part of the electrons intercepted this way would anyhow be lost on the beam pipe because of the strong transverse expansion of the dark current. This is illustrated in Fig. 4.28 which shows that the final difference in the current transported by ACC1 amounts only to about 37 µA out of 85 µA, or 43 %. The spot size at the exit of the acceleration module decreases moderately from 3.2 to about 2 mm (rms).

BC2 collimator

The BC2 collimator is located approximately in the longitudinal center of the bunch compressor chamber. At this point, the dark current completely fills the horizontal aperture from about $x = 30$ cm to the outer wall of the vacuum chamber due to its huge energy spread. As shown in Fig. 4.29, the beam occupies only a comparatively small region between 31 and 32 cm. Hence, the collimator has a huge potential of intercepting dark current before cutting into the beam.

For the simulation, the BC2 collimator is inserted just so far that none of the simulated beam particles are lost while passing through the magnetic chicane. In this setting, the scraper covers all the space with $x \geqslant 32.11$ cm. The beam trajectory, i.e. the center-of-mass trajectory of all beam electrons, has a distance of 7.5 mm to the edge of the collimator.

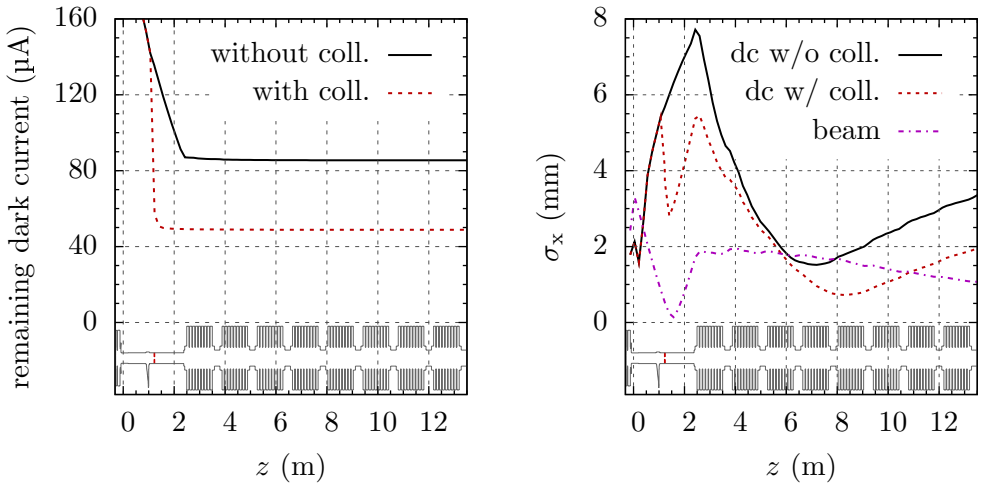

Figure 4.28: Influence of the gun collimator on dark current in the injector. For orientation, both plots show a miniature of the aperture model with the gun collimator marked in red. *Left:* Dark current transported along the injector. *Right:* Transverse rms spot size of dark current and beam.

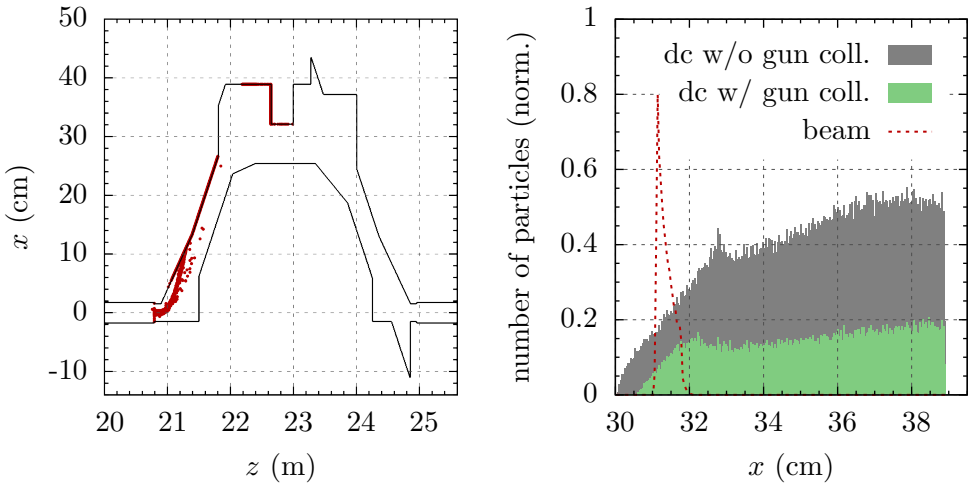

Figure 4.29: Dark current losses at the BC2 collimator. *Left:* Horizontal section of the BC2 aperture model with inserted BC2 collimator. Dark current losses are marked in red. *Right:* Transverse distribution of dark current and beam at the position of the BC2 collimator. The curves for dark current with and without gun collimator have the same vertical scale scale.

One of the benefits of the BC2 collimator directly visible in Fig. 4.29 is the complete elimination of further dark current losses in the vacuum chamber. The copper block intercepts all electrons that would otherwise be lost in various places along the second half of the chamber (cf. Fig. 4.17). With its length of 35 cm, corresponding to roughly 25 radiation lengths of copper, the collimator can absorb most of the energy of the incident particles which would otherwise be released externally as an electromagnetic shower. This contributes greatly to the protection of electronic devices and optical setups located in the vicinity of the bunch compressor.

Collimation efficiency

A suitable figure of merit needs to be chosen to quantify the efficiency of the two removable collimators. This choice should be based on the main hazards presented by the dark current to the accelerator. These are, first, the activation of beamline components, and second, demagnetization of the undulator magnets. For both cases, the total energy deposited in the respective component is of decisive importance. It is therefore reasonable to present the simulated losses along the beamline not in terms of lost current, but of power.

The required normalization includes a number of assumptions; to convert the number of macroparticles N to the momentary dark current (i.e., the dark current averaged over one period of the rf wave or over the length of the rf pulse), the measured value of $I = 160\,\mu\text{A}$ at Faraday cup 2GUN has been used. This current corresponds to 889515 macroparticles in the simulation:

$$I = CN \quad \text{with} \quad C = \frac{1.80 \cdot 10^{-10}\,\text{A}}{\text{macroparticle}}$$

The total time-averaged current depends on the length of the rf pulse τ_{rf} and on the repetition rate ν_{rf}. For further calculations, the design parameters of FLASH are used:

$$\bar{I} = \nu_{\text{rf}}\,\tau_{\text{rf}}\,I \quad \text{with} \quad \nu_{\text{rf}} = 10\,\text{Hz},\ \Delta T = 800\,\mu\text{s}$$

For each simulated particle lost on an obstacle, the tracking code reports the impact momentum p_i. To account for low- and nonrelativistic particles correctly, the kinetic energy is computed using the rest mass of the electron m:

$$T_i = \sqrt{p_i^2 c^2 + m^2 c^4} - mc^2$$

The average power released by N lost macroparticles is then given by

$$\bar{P} = \frac{\bar{I}}{e} \cdot \frac{1}{N} \sum_{i=1}^{N} T_i = \frac{\nu_{\text{rf}}\,\tau_{\text{rf}}\,C}{e} \cdot \sum_{i=1}^{N} T_i. \tag{4.4}$$

To visualize the losses along the accelerator, the lost macroparticles are sorted into bins of 1 m length according to their longitudinal position. For each bin, the

	collimators inserted			
section	**none**	**gun**	**BC2**	**gun+BC2**
gun	92.5 µA/ 2.0 W	130.9 µA/ 3.1 W	92.5 µA/ 2.0 W	130.9 µA/ 3.1 W
BC2	59.8 µA/43.3 W	40.2 µA/27.0 W	81.6 µA/63.1 W	47.6 µA/33.8 W
BC3	8.0 µA/28.8 W	1.5 µA/ 5.4 W	2.1 µA/ 7.6 W	0.6 µA/ 2.3 W
coll.	9.4 µA/71.6 W	4.4 µA/33.2 W	0.9 µA/ 7.4 W	0.5 µA/ 3.5 W
dump	3.9 µA/30.2 W	2.6 µA/19.9 W	0.1 µA/ 0.9 W	0.1 µA/ 0.9 W

Table 4.5: Dark current losses in most affected sections for various combinations of gun and BC2 collimator settings. *Gun:* gun cavity and beam line up to ACC1 entrance. Because only the photocathode is simulated as a field emitter, losses are underestimated. *BC2/BC3:* vacuum chambers of the bunch compressors. *Coll.:* transverse and energy collimators, and the beam line between them. *Dump:* dump line and dump. This is equivalent to the amount of dark current traversing the undulators.

released power is determined using (4.4). The resulting *power loss per meter* is an often-used quantity, but its absolute value has to be interpreted with care because of the dependence on the chosen bin size, especially in the case of longitudinally point-like losses.

Figure 4.30 shows the power loss per meter for various combinations of the gun and BC2 collimator settings. In all cases, dark current losses at BC2, BC3, and at the transverse collimators are dominant and of an order of magnitude around 10 W/m. The gun section is special because of two reasons: On the one hand, the obtained values around 1 W/m are underestimated as only dark current from the photocathode is simulated. Other field emitters may create local losses of higher intensity. On the other hand, these losses are qualitatively different due to the low electron energies; the penetration depth in material is short, photonuclear reactions are practically excluded, and the radiation field consists mainly of comparatively soft bremsstrahlung photons.

As the gun collimator is limiting the transverse phase space of the dark current, it has the biggest impact in locations where losses occur due to the transverse extent of the particle ensemble. This is the case in ACC2–3, BC3, and at the transverse collimators. In contrast, it is only moderately effective or ineffective in limiting losses caused by the huge energy spread of the dark current: in BC2, at the energy collimators, and in the dump line. Hence, the energy-collimating scraper in BC2 is a reasonable complement; when positioned close to the beam as described above, it can practically eliminate all downstream losses except the ones at BC3 and at the transverse collimators.

The absolute power loss in the most affected sections of the accelerator is summarized in Tab. 4.5. The values underline that the gun collimator is helpful in limiting the power dumped in bunch compressor BC2; however, about 30 W of losses remain in any case and inevitably contribute to the activation of surround-

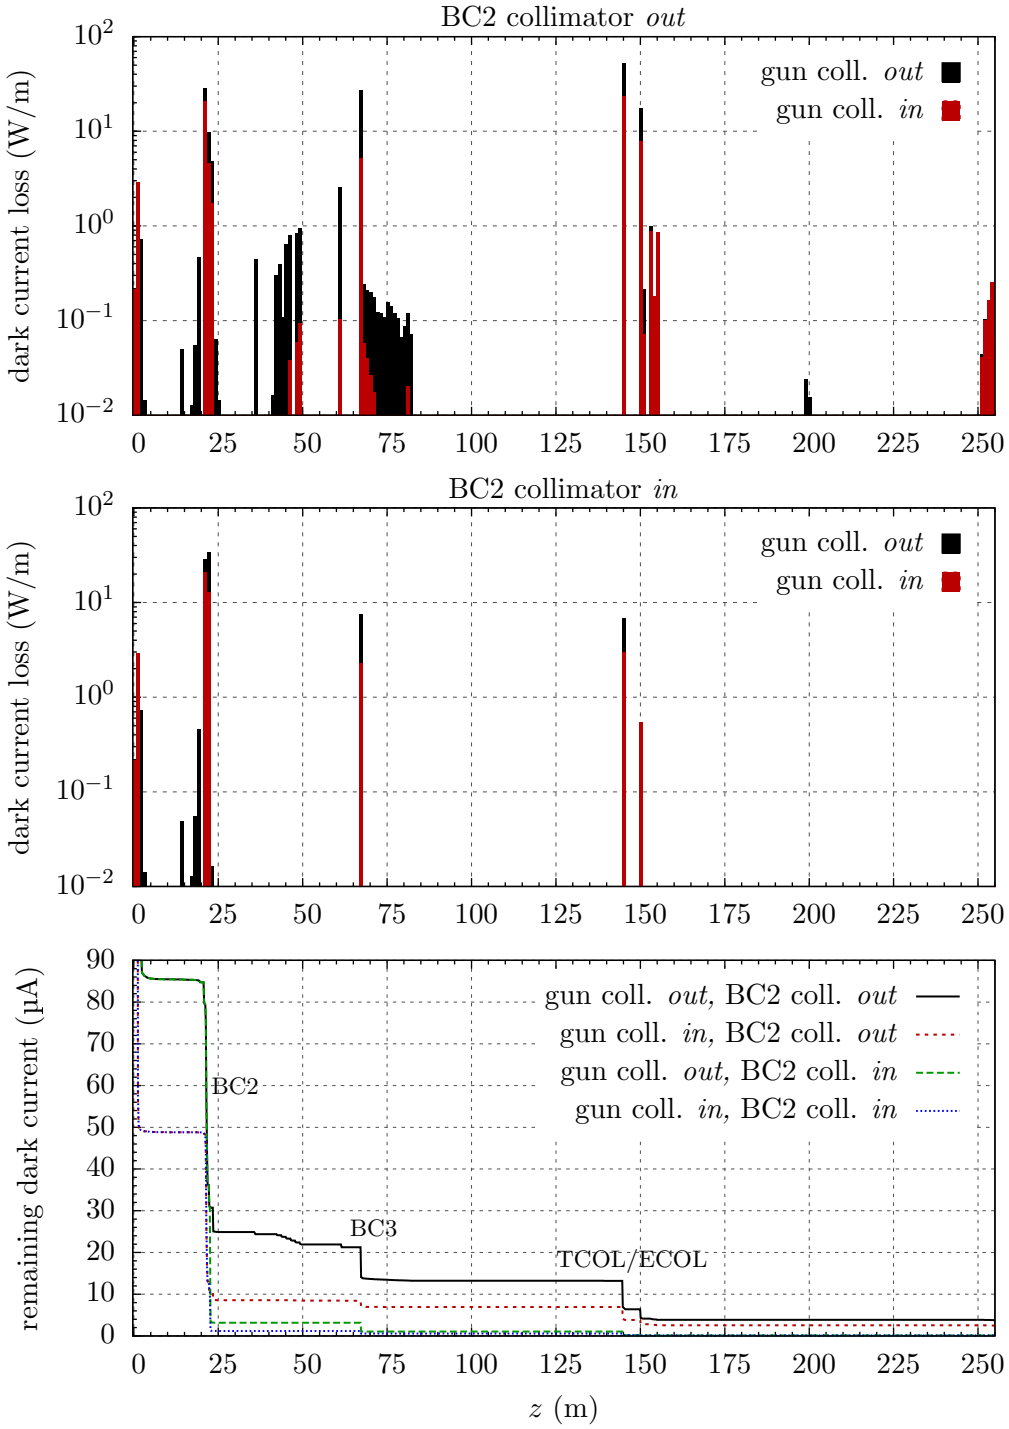

Figure 4.30: Dark current losses along the accelerator for various combinations of gun and BC2 collimator settings.

ing components. With the current setup of the beamline, only an active element like the dark current kicker or a reduction of the rf pulse length can mitigate the problem.

Without collimators, BC3 is also subjected to a bombardment of about 30 W. However, the gun collimator can reduce this power to 5 W, and with the BC2 scraper it is possible to achieve a further mitigation by 50 %.

If the dark current is left unchecked in the upstream sections, the main collimators absorb roughly 70 W of it. Here, the BC2 collimator can reduce the exposure by about an order of magnitude while the gun collimator brings an improvement of only 50 %.

As the final indicator for the efficiency of the collimation system, the table lists the total power of the dark current traversing the undulator and reaching the dump line. This power can be released partially or completely into the undulator by a missteering or misfocusing. Even the tiny remaining current carries a power of 30 W because of the energy of almost 1 GeV per particle. While the gun collimator can only reduce this value by one third, use of the BC2 collimator almost eliminates the dark current, leaving less than 1 W.

4.6 Final remarks

The transport of gun dark current is a major problem for the operation of the FLASH linac. It is not only responsible for the strong residual radioactivity in the bunch compressors and in the transverse collimators, but it also limits the effective aperture available for steering and focusing of the beam.

Because FLASH is not equipped with monitors for the measurement of dark current except for few Faraday cups behind the gun, tracking simulations are the only way to obtain a coherent picture of the transport mechanism. In this chapter, it has been shown that under normal operating conditions no field emitters except the photocathode have a significant contribution to the gun dark current entering the linac. It has also been shown that geometrical details of beamline elements—such as the bunch compressor chambers—have a decisive influence on the transport of dark current. To allow a detailed representation of these complex apertures, it has been necessary to implement a geometric modeling system superior to those available in conventional tracking codes.

This study shows that a consequent use of the available collimation systems can significantly reduce the radiation exposure especially in the last part of the accelerator. With the current setup, however, a complete elimination of dark current losses is not possible. Especially the huge activation of bunch compressor 2 is inevitable as long as additional active systems like the dark current kicker have not been commissioned.

5 Damage scenarios for the FLASH undulators

The free-electron laser process depends strongly on a good field quality of the wiggler or undulator. Field errors cause the electron beam trajectory to deviate from that of the photon beam, thereby reducing the overlap between both. They also change the longitudinal velocity of the beam and move it out of resonance. Detailed discussions of these effects can be found, e.g., in [Yu92] and [Wal93]. A recent calculation for FLASH [Sku08] has shown that a deviation of the magnetic field of the order of 0.5 % close to the entrance and exit of the individual undulator segments would decrease the radiation power by 10 %. It is therefore obvious that the undulator magnets must be protected from damage caused by the beam.

After a brief review of available data on radiation damage in permanent magnets, this chapter will focus on the analysis of typical beam loss scenarios for the undulator beamline at FLASH. The efficiency of a shielding installation is evaluated with the help of Monte Carlo simulations.

5.1 Radiation damage in permanent magnets

Like most modern insertion devices, the FLASH undulators are based on permanent magnets made of an alloy of neodymium, iron, and boron ($Nd_2Fe_{14}B$ or, short, NdFeB). The radiation sensitivity of this material has been an object of extensive experimental studies. All particle types prevalent in the radiation environment of an electron accelerator—electrons and positrons, gamma photons, and neutrons—have been shown to cause damage to specific magnet samples. However, the results of these experiments are often difficult to interpret and the observed demagnetization depends strongly on the parameters of the individual sample.

For example, [And07] reports relative magnetization losses of $3.5 \cdot 10^{-4}\,\mathrm{Gy}^{-1}$, $1.6 \cdot 10^{-4}\,\mathrm{Gy}^{-1}$, $9.9 \cdot 10^{-6}\,\mathrm{Gy}^{-1}$, and zero for four magnet samples of identical dimension after irradiation with fast neutrons under identical conditions. The experiment covered doses up to 220 Gy and neutron fluences up to $3 \cdot 10^{14}\,\mathrm{cm}^{-2}$. The authors attribute the difference in radiation resistance to the choice of dopants in the NdFeB alloy.

For the low energy regime, [Oku94] reports a magnetic field loss of $3.5{\cdot}10^{-8}\,\mathrm{Gy}^{-1}$ for irradiation with electrons of 17 MeV, and less than $1.8{\cdot}10^{-9}\,\mathrm{Gy}^{-1}$ for irradiation with a ^{60}Co source, i.e. with photons below 1.2 MeV. The maximum absorbed doses in this reference are specified as 2.6 MGy and 2.8 MGy, respectively. The authors

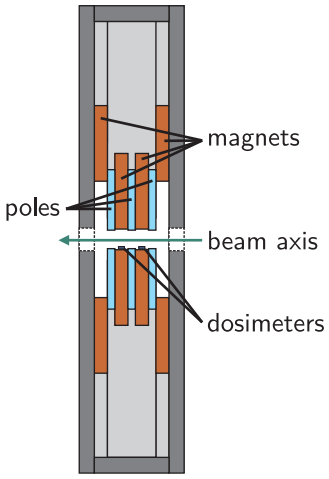

Figure 5.1: Schematic of the diagnostic undulator installed at FLASH. The magnetic field integral along the beam axis is zero.

arrive at the conclusion that the radiation damage caused by gammas is equivalent to the damage caused by the released secondary electrons.

These observations are confirmed for higher beam energies by [Lun89]. For direct irradiation with an 85 MeV electron beam, a remanence loss of $4.2 \cdot 10^{-5}$ Gy^{-1} at an absorbed dose of 360 Gy is measured. In the bremsstrahlung-dominated radiation field created by dumping the beam on a 2.5 mm thick plate of tungsten alloy, three different magnet samples show relative demagnetizations of $3.1 \cdot 10^{-8}$ Gy^{-1}, $3.0 \cdot 10^{-8}$ Gy^{-1}, and $1.5 \cdot 10^{-9}$ Gy, at absorbed doses up to 25 MGy.

While the available data support a rough hierarchy by indicating that neutrons are more damaging than electrons which are, in turn, more damaging than photons, the underlying microscopic damage mechanism is still under discussion. A comprehensive overview of related work can be found in [Spe03].

5.1.1 Measurements at FLASH and TTF

In view of the huge spread of published data, it is desirable to measure the radiation damage under realistic conditions. For this purpose, a removable *diagnostic undulator* has been installed at FLASH. As shown in Fig. 5.1, it consists of a hybrid magnetic structure with four NdFeB magnets and three soft-iron pole pieces per side. The end magnets are removed from the beam axis so that the resulting magnetic field integral across the device is zero. The structure uses the same components as the main undulators and is located just in front of them, where a comparatively high radiation background is expected. Two thermoluminescence dosimeters (TLDs) placed on the lower magnets allow a periodic measurement of the accumulated dose close to the beam pipe.

After an initial measurement of the magnetic field profile, the diagnostic undulator has been removed from the beamline twice for remeasurement. Table 5.1 shows the relative peak field change of the central pole and the accumulated dose at the time of measurement as published in [Sku08]. A linear fit to the data gives

Date	Dose (kGy)	$\Delta B/B$ (%)
2004-08-13	0	0.0
2006-03-21	37	−2.3
2007-09-29	61	−3.1

Table 5.1: Accumulated dose and measured peak field change of the central pole of the dosimetric undulator as published in [Sku08]

a demagnetization of $5.4 \cdot 10^{-7}\,\mathrm{Gy}^{-1}$.

In the same reference, older measurements [Pfl03] of the first undulator segment from phase 1 of the TESLA Test Facility are re-evaluated. This undulator used similar magnet pieces, and accumulated a local dose of more than 12 kGy during three years of operation. For this device, a demagnetization of $2 \cdot 10^{-7}\,\mathrm{Gy}^{-1}$ is estimated.

These values can be used to derive thresholds for the operation of FLASH. If the maximum tolerable deviation of the magnetic field is assumed to be of the order of 0.5 %, the measurements of the diagnostic undulator indicate a local dose limit of about 10 kGy. Assuming further an integrated FEL operation time of 200 days per year and a projected undulator life time of 10 years, the dose rate should be limited to 5 Gy/d.

In order to enforce this tight limit, the machine protection system implements special safety measures for the undulators. These measures will be discussed in chapter 6.

5.2 FEL beamline description

The FEL beamline of FLASH is a heterogeneous assembly of beam pipes with different diameters and various surrounding components. The particularities of this section are explained in the following.

A number of removable optical transition radiation (OTR) screens for transverse beamsize measurements are installed in an early part of the FEL beamline. As these screens are essentially 350 µm thick silicon wafers, they absorb very little of the beam energy. They do, however, cause a significant blowup of the transverse emittance by Coulomb scattering. The resulting beamsizes are too big for the aperture of the undulator vacuum chambers, and would therefore cause substantial beam losses and, consequently, radiation damage to the permanent magnets.

To prevent this, two strong electromagnetic dipoles have been installed downstream of the OTR stations, at about 196.5 m. The machine protection system allows the use of the OTR screens only if these *suppressor dipoles* are powered at their maximum current of 3.5 A, corresponding to a field of 312 mT. With the magnetic length of 10 cm per dipole, a deflection angle of 1.1° (or 19 mrad) in the horizontal plane can be obtained for a beam energy of 980 MeV according to (C.1). Considering the beam pipe radius of 17 mm, this makes the beam hit the pipe at about 197.4 m.

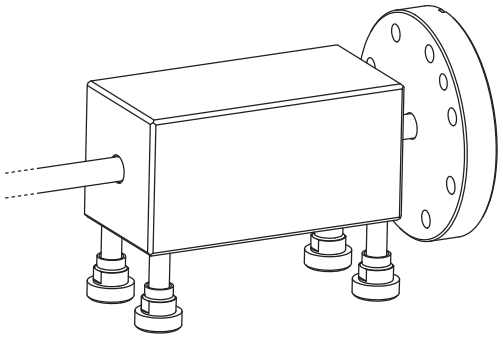

Figure 5.2: Drawing of the undulator absorber [Sie08]. The absorber consists of a massive copper block soldered to the beam pipe just behind a flange. It is prepared for water cooling.

This does not pose a threat to the beam pipe itself because the machine protection system automatically limits the beam to a maximum of 2 bunches per macropulse upon activation of the dipoles. However, the created electromagnetic shower endangers downstream components and the permanent magnets of the first undulator segment. Although this kind of operation is limited to short time spans, an analysis of the expected radiation field is essential and has therefore been chosen as the subject of the first simulation presented here.

Another peculiarity of the section is the insertion of a smaller beam pipe with 10 mm instead of 34 mm diameter between 198.38 m and 199.74 m. The reason for this was to facilitate the installation of a number of experimental beam position monitors (BPMs). As discussed earlier, the small aperture of this beam pipe limits the freedom for beam steering severely. Especially at the initial transition from 34 to 10 mm, frequent beam and dark current losses have to be expected. The shower produced under these conditions is much more detrimental to the undulator magnets because its direction is practically identical with the beam axis. This is investigated in the second simulation scenario.

A simple shielding setup has been implemented based on these analyses. It consists of two parts. The first part is a copper absorber installed just behind the flange joining the thick and the thin beam pipes (Fig. 5.2). Due to severe space constraints, it is of relatively small dimensions and covers a volume of $7 \times 7 \times 14\,\mathrm{cm}^3$. Its length therefore corresponds to about 10 radiation lengths of copper. Just behind the absorber, at 198.6 m, a 5 cm thick wall of lead bricks covers an area of about $40 \times 40\,\mathrm{cm}^2$ around the beam pipe. Evaluation of the shielding efficiency of this setup is one of the main objectives in the discussion of the first scenarios.

Behind the experimental BPM insert, the regular 34 mm beam pipe continues up to 202.55 m, where another transition to 10 mm takes place. The diagnostic undulator is installed close to this transition. At 203.44 m, the small beam pipe leads into the vacuum chamber of the first undulator segment which, again, has a marginally smaller aperture of 9.5 mm diameter. In total, there are six undulator segments, each of about 4.5 m length. These segments are separated by 60.7 cm long so-called *intersections* based on a 10 mm beam pipe. Each intersection provides strong focusing by two electromagnetic quadrupoles. In addition, it is equipped with a wirescanner for the determination of the transverse beam

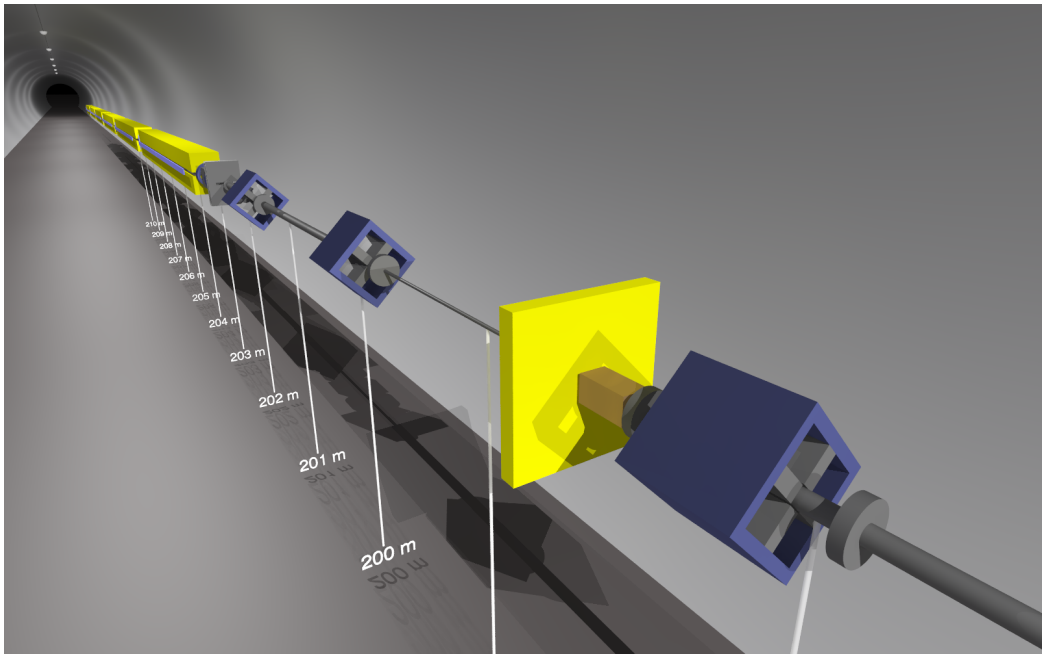

Figure 5.3: Three-dimensional view of the FEL beamline model

profile, beam position monitors, and other diagnostics. The setup of all intersections is identical; with minor modifications, the same arrangement of components is found in front of the first and behind the last undulator segment.

5.3 Modeling of beamline and components

The various beam loss scenarios have been simulated with the multi-particle transport code *Fluka* [Fas03, Fer05]. Fluka is a general purpose Monte Carlo code that can simulate the propagation and interaction of many particle types in matter, including photons, electrons, positrons, and neutrons. The problem description is formulated in a combinatorial geometry language: First, a list of geometrical primitives *(bodies)* is defined in a cartesian coordinate system. Typical examples are cylinders, cuboids, or half-spaces. Afterwards, these bodies are combined to *regions* using a combination of the three logical operations AND, AND NOT, OR. Each region can be associated with a specific material and, if necessary, with a user-written routine to simulate magnetic fields.

For the shower simulation, the FLASH FEL beamline has been modeled starting from the longitudinal position of 197 m, which is about 6.5 m upstream of the first undulator segment. Beam losses further upstream are not expected to have a significant contribution to the accumulated dose in the undulator magnets. Because the magnets are located very close to the beam axis and well shielded to the outside, only components close to the beamline are included in the geometry.

A three-dimensional rendering of the implemented model is shown in Fig. 5.3. The main visible components include

- beam pipes,
- flanges,
- quadrupole magnets,
- a copper absorber,
- a lead wall,
- the diagnostic undulator, and
- the 6 main undulator segments.

In the following, the most important details of the model will be described.

Environment: All defined regions are embedded in a number of external vacuum regions. Test runs have shown that the simulation of scattering in air has no significant contribution to the dose deposition in beamline elements.

Beam pipes: As indicated before, two types of steel beam pipes with different diameters are used in the FEL beamline. The standard pipe has an inner diameter of 34 mm and a wall thickness of 2 mm, the thin pipe has an inner diameter of 10 mm and a thickness of 1 mm. The material of all beam pipes is modeled as pure iron. Also the flanges at the connections of individual beam pipe segments are modeled as solid iron cylinders with an inner cut-out according to the respective pipe diameters.

Quadrupole magnets: The massive iron yokes of the quadrupole magnets have a significant shielding effect and are therefore included in the model. Each quadrupole is associated with a longitudinal region corresponding to its effective magnetic length. Within this region, the magnetic field is simulated in the beam pipe material and in the enclosed vacuum region as

$$\boldsymbol{B} = (-gy, -gx, 0).$$

In effect, this is the hard-edge model of a perfect magnetic quadrupole. The gradient g and the effective magnetic length are identical to the ones used in the previous tracking simulations.

Shielding elements: In simulations including the shielding setup, the absorber is modeled as a solid copper block of $7 \times 7 \times 14\,\mathrm{cm}^3$ volume with a cylindrical cutout of 10 mm diameter. The lead wall is implemented as a solid lead block of $40 \times 40 \times 5\,\mathrm{cm}^3$, placed symmetrically around the beam pipe. For historic reasons, the wall has an oversized cutout of 38 mm diameter which is included in the model.

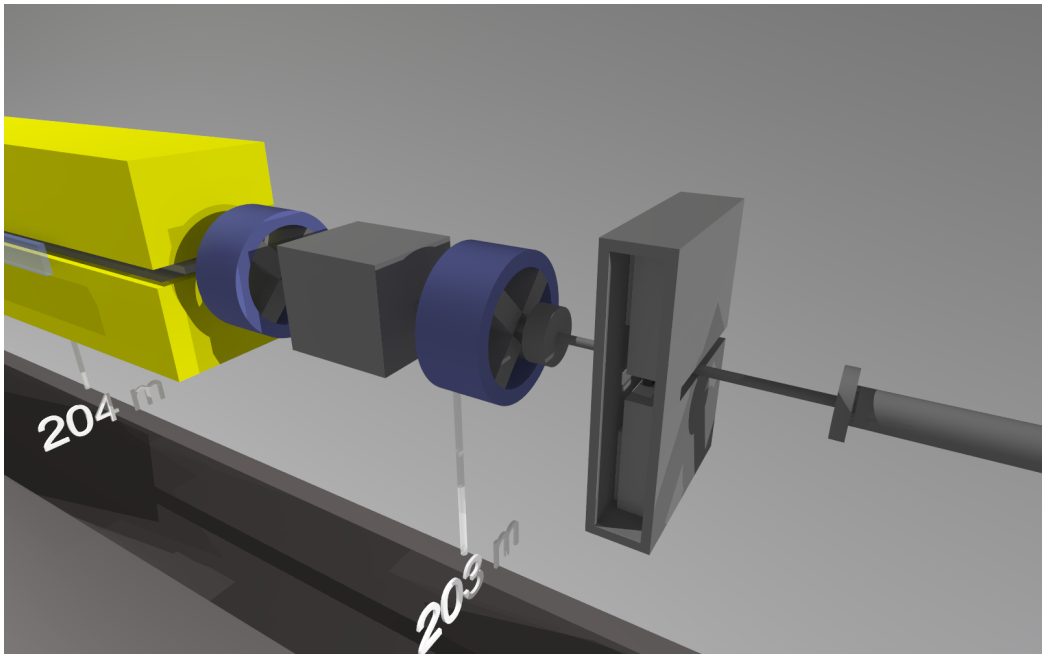

Figure 5.4: Detail view of diagnostic undulator and entrance to the first main undulator segment.

Diagnostic undulator: A coarse description of the diagnostic undulator has already been given in section 5.1.1. Except for the slight displacement of the end magnets, the arrangement of magnets, poles, and holders is identical to the main undulators and will be described in detail in the following section. The outer support structure, however, is unique; it consists of 1 cm thick iron plates and has a substantial shielding effect. Transversely, the structure covers an area of $22 \times 30.9\,\mathrm{cm}^2$; a slit of 1.5 cm height and 17 cm width allows to slide the diagnostic undulator sideways over the beam pipe (Fig. 5.4). The effective magnetic field region of the diagnostic undulator is very small compared to that of other magnetic components in the beamline and is therefore neglected.

Undulator intersections: Figure 5.4 also shows the components present in an undulator intersection—in this case, in front of the first undulator. Apart from two quadrupole magnets and surrounding flanges, the central wirescanner chamber is included in the model. It is realized as a hollow iron cube of 5 mm wall thickness with outer dimensions of $14.5 \times 14.5 \times 16.3\,\mathrm{cm}^3$. In the real installation, the wirescanner assembly extends further outside in a very complex geometry. However, it is assumed that these external parts have only little influence on the deposition of dose in the undulator.

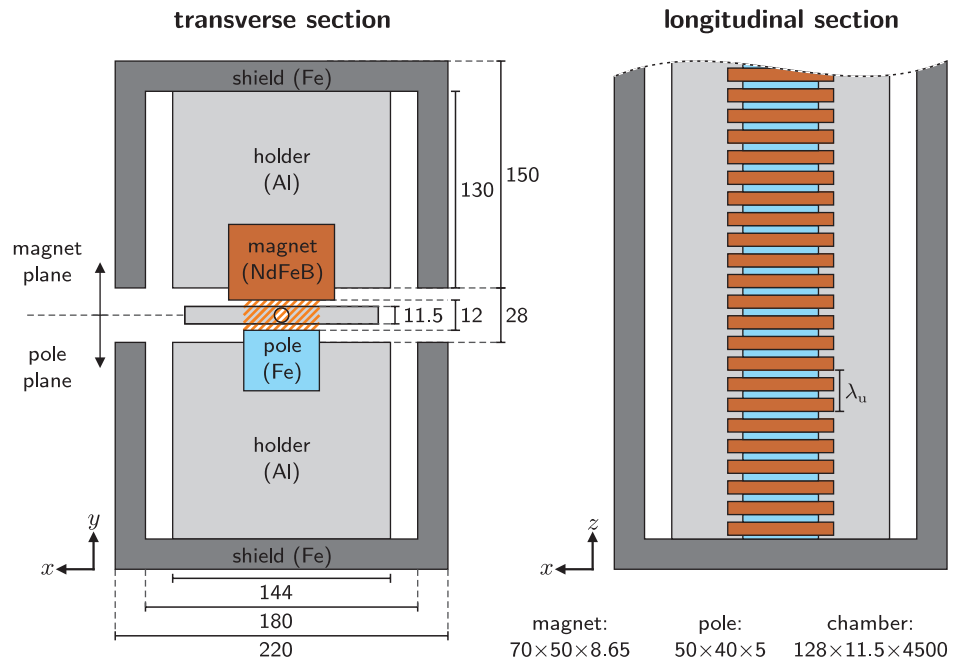

Figure 5.5: Geometry of the undulator model. The hatched area indicates the region in which the magnetic field of the undulator is simulated. The upper and lower half of the transverse section correspond to different z coordinates. The drawings are true to scale; the gap between magnets/poles and vacuum chamber has been enlarged for clarity. All dimensions are in units of mm.

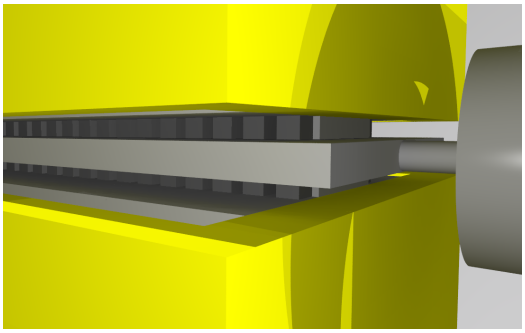

Figure 5.6: Detail view of the undulator model. The external iron shield is shown in yellow, NdFeB magnets are dark gray.

Undulator model: The undulator has the highest complexity of all simulated components. The beam passes through a flat aluminum vacuum chamber of 11.5 mm height with an inner diameter of 9.5 mm. As shown in Fig. 5.5 and 5.6, the chamber is inserted into the 12 mm high gap of a hybrid magnet structure consisting of alternating pieces of magnets and soft iron poles. A full undulator period consists of two magnets and two poles with a full length of $\lambda_\mathrm{u} = 2.73\,\mathrm{cm}$. The structure is retained by a system of aluminum clamps that is modeled as a massive block of $144 \times 130\,\mathrm{mm}^2$ cross sectional area with appropriate cut-outs for magnets and poles. On the outside, the undulator is supported by a massive iron shield of 2 cm thickness. Further design details can be found, e.g., in [Hah00] and [Pfl03a].

Modeling of magnetic fields

The maximum energy deposition by electrons and positrons in matter takes place at energies below 1 MeV, where the trajectories of these particles are strongly influenced by magnetic fields. Therefore, an adequate modeling of the main magnetic fields in the undulator region is necessary to obtain meaningful estimates of deposited dose.

The configuration of magnetic field regions is illustrated in Fig. 5.7. The magnetization of the permanent magnets is directed along the z axis, with alternating polarity from one piece to the next. The soft iron poles in between direct the magnetic flux into vertical direction. A simple calculation given e.g. in [Doh08] yields the following solution for the field in the enclosed space between these hybrid structures:

$$\boldsymbol{B} = -B_0\big(0,\ \cosh(k_\mathrm{u} y)\sin(k_\mathrm{u} z),\ \sinh(k_\mathrm{u} y)\cos(k_\mathrm{u} z)\big) \tag{5.1}$$

Here, B_0 is the peak field on the axis of the undulator, and $k_\mathrm{u} = 2\pi/\lambda_u$ is the undulator wave number. While also horizontal components appear in the lateral fringe fields, (5.1) is a good approximation for the central part of the undulator which is marked as hatched area in Fig. 5.5.

Similar to the sacrificial undulator, the main undulators have a short transition zone for tapering of the magnetic field at the undulator entrance and exit. For simplicity, this transition is not modeled in detail. Instead, the undulator model starts and ends with a pole of half thickness, covering a total length of 163 full undulator periods. Correspondingly, a hard-edge model of the undulator field has been implemented. This model must ensure that a particle entering the undulator on axis will also leave it on axis with no transverse deflection.

If the origin of the coordinate system is located on the undulator axis such that $z = 0$ at the entrance of the field region, this criterion is fulfilled only if (5.1) is shifted by λ_u:

$$\boldsymbol{B} = -B_0\big(0,\ \cosh(k_\mathrm{u} y)\cos(k_\mathrm{u} z),\ \sinh(k_\mathrm{u} y)\sin(k_\mathrm{u} z)\big) \tag{5.2}$$

To validate this, only the field on the undulator axis ($y = 0$) is considered:

$$B_\mathrm{y}(z) = -B_0\cos(k_\mathrm{u} z)$$

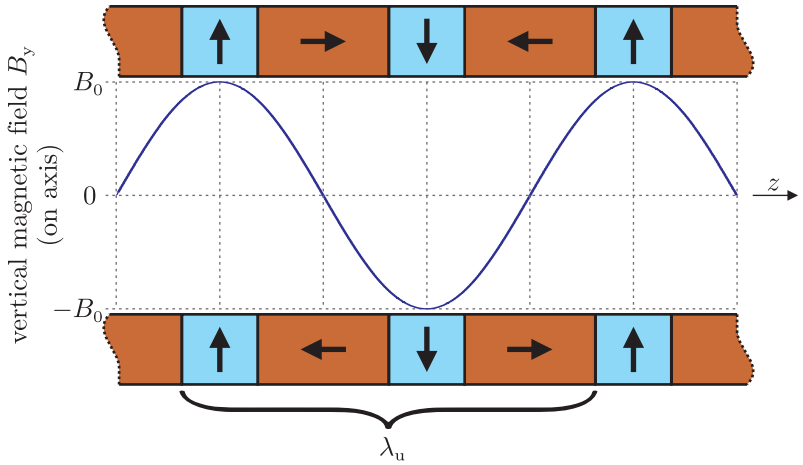

Figure 5.7: Model for undulator fields. Magnets are shown in red, poles in light blue. The polarity of the magnetic field in the material is indicated by arrows. The enclosed plot shows the vertical magnetic field on the beam axis.

The horizontal offset at position z can now be calculated with a simple integration. Assuming that the longitudinal velocity $v_z = \mathrm{d}z/\mathrm{d}t$ is approximately constant, the Lorentz force is

$$m\ddot{x} = ev_z B_y,$$

and therefore:

$$\begin{aligned}
\dot{x}(z) &= \frac{e}{m}\int_0^z B_0 \cos(k_u z')\,\mathrm{d}z' \\
&= \frac{eB_0}{m\,k_u}\sin k_u z \\
x(z) &= \frac{eB_0}{m\,k_u\,v_z}\int_0^z \sin(k_u z')\,\mathrm{d}z' \\
&= \frac{eB_0}{m\,k_u^2\,v_z}(1-\cos k_u z)
\end{aligned}$$

At the end of the field region, the desired result of $\dot{x} = x = 0$ is obtained because $z = 163\lambda_u = 163 \cdot 2\pi/k_u$.

If, for comparison, $B_y(z=0) = 0$ had been demanded, (5.1) would have had to be used to model the field. The integration for this case yields

$$\begin{aligned}
B_y(z) &= -B_0 \sin(k_u z) \\
\dot{x}(z) &= \frac{eB_0}{m\,k_u}(1-\cos k_u z) \\
x(z) &= \frac{eB_0}{m\,k_u\,v_z}\left(z - \frac{\sin k_u z}{k_u}\right),
\end{aligned}$$

Table 5.2: Step size settings for Fluka simulations. In addition to the shown step sizes, the accuracy for determining a region boundary intersection is set to 50 µm, and the maximum direction change per tracking step is limited to 10°.

Regions	Max. step size
Quadrupole fields	50 µm
Dipole fields	50 µm
Undulator fields	20 µm

and the result would have been a strong horizontal trajectory drift in the undulator.

The implemented magnetic field in the various regions is specified as follows:

Undulator center: This corresponds to the hatched area in Fig. 5.5, comprising those parts of internal and external vacuum regions and of the vacuum chamber falling within the lateral extent of the pole pieces. The field is calculated using (5.2), with the nominal undulator peak field of $B_0 = 486\,\mathrm{mT}$.

Magnet pieces: Within the NdFeB material, the nominal remanent field of $B_\mathrm{z} = 1.15\,\mathrm{T}$ is oriented longitudinally. The polarity follows Fig. 5.7.

Pole pieces: Within the poles, the magnetic flux is approximated as purely vertical. The flux density is calculated by evaluating (5.2) at the tip of the pole; with $y = 6\,\mathrm{mm}$, the value of $B_\mathrm{z} = 1.03\,\mathrm{T}$ is obtained. Again, the polarity is implemented according to Fig. 5.7.

Other regions within the undulator are considered field-free.

5.4 Simulation parameters

Before the analysis of the various beam loss scenarios, the most important simulation settings that differ from the default setup of Fluka need to be mentioned. In this case, most modifications are concerned with the accuracy of tracking within magnetic fields. Because Fluka, in its original state, has only rudimentary support for magnetic fields, they have been implemented as a user routine according to the descriptions given above. The default algorithm for automatic adaptation of the tracking step size is not aware of the resulting deflection for charged particles. Hence, a simulation with default parameters produces inaccurate or plainly wrong results. As a remedy, the step size in magnetic field regions has been explicitly limited.

Forcing the tracking algorithm to smaller step sizes than necessary results in excessive computation times. Therefore, the step sizes are determined using an empirical approach. An ensemble of $4 \cdot 10^5$ particles at $z = 197\,\mathrm{m}$ from the beam tracking with Astra is used as the initial particle distribution in Fluka. This distribution is repeatedly tracked with decreasing step sizes until a straight trajectory through all quadrupoles and undulator segments is achieved. Table 5.2 summarizes the final step size settings which are used in all of the following simulations.

Material	Range R (mm)
Aluminum	0.215
Iron	0.083
Copper	0.076
Lead	0.008

Table 5.3: Mean range in matter for electrons of 200 keV kinetic energy. Calculated in CSDA using [Ber05].

Physics model

Additionally, a number of options controlling Fluka's physics model are changed from their default setting. First, the production and transport thresholds for electrons and positrons are lowered to a value of 200 keV. This means that particles with a lower kinetic energy are neither produced nor tracked by the code; instead, their energy is deposited at the point of interaction. To evaluate the impact this has on the results, it is useful to calculate the *mean range* R of the particles, i.e. the average penetration depth in matter. In rough approximation, R can be calculated by integrating over the inverse of the stopping power,

$$R \approx \int_0^T \frac{\mathrm{d}x}{\mathrm{d}E} \mathrm{d}E,$$

with T denoting the kinetic energy of the particle. This *continuous slowing down approximation* (CSDA) neglects effects such as the microscopic path length increase by multiple scattering. Nevertheless, it is sufficient to validate the chosen parameters. For a kinetic energy of 200 keV, the CSDA ranges given by [Ber05] are listed in Tab. 5.3. They are well below 250 µm for all materials used in the model. Since the finest structures in the model, the magnet poles, have a thickness of 5 mm, the chosen threshold allows an accurate simulation of energy deposition.

Also the thresholds for photon production and transport are lowered from their default setting. As photons do not cause a quasi-continuous energy deposition along their path, the concept of range cannot be used to determine a suitable value. Instead, a conservative low energy of 10 keV is chosen. Considering the background of high-energy bremsstrahlung (cf. Fig. 2.4), photons below this energy are not expected to have a significant contribution to the deposited dose.

Furthermore, photonuclear reactions are enabled. This explicitly includes simulation of the giant dipole resonance and the quasi-deuteron effect, but also of the Δ resonance[1] and of high-energy effects such as photospallation. The chosen settings therefore allow an estimate of expected neutron fluences.

1 The Δ baryons consist solely of u and d quarks and decay under production of protons and neutrons, among other particles. The lowest resonance energy for Δ production is 1232 MeV, with a full resonance width of 118 MeV [PDG08]. Therefore, a channel for production of Δs will open up with the planned energy upgrade of FLASH to 1.2 GeV.

Radiation quantities

To provide an overview of intensity and distribution of the shower, a meaningful physical quantity for a plot must be chosen. A traditional choice for this is *particle fluence* F, i.e. particles per unit area or, more precisely, the number of particles entering a sphere of unit cross-sectional area:

$$[F] = 1\,\mathrm{m}^{-2}$$

One typical way of evaluating F in a Monte Carlo code is in the form of *track-length fluence,* i.e. the sum of all particle track lengths falling within a volume element $\mathrm{d}V$, normalized to that volume. However, if not differentiated by particle types and energies, fluence is a rather abstract quantity and it is hard to judge the damage potential of a radiation field from it.

Instead, *ambient dose equivalent* $H^*(10)$ is chosen as the main quantity for overview plots. This is an operational quantity for use in radiological protection introduced by the International Commission on Radiation Units and Measurements (ICRU). It is defined as the quality-weighted dose that would be deposited by the radiation field at a depth of 10 mm inside the *ICRU sphere*, a virtual body of a mass composition equivalent to human tissue:

$$[H^*(10)] = 1\,\mathrm{Sv} = 1\,\frac{\mathrm{J}}{\mathrm{kg}}$$

The quality-weighting is based on a dimensionless quality factor Q that is high for particularly damaging kinds of radiation, and low for less damaging ones. The complete definition of $H^*(10)$ can be found in [ICR93]. In Fluka, this quantity is calculated using tables of fluence-to-ambient-dose-equivalent conversion factors as published in [ICR97] and [Pel00].

Being a quantity for radiological—i.e., personnel—protection, it can be argued that ambient dose equivalent is not a good representation of the danger a radiation field presents to components such as permanent magnets. For electromagnetic showers, however, it is a good approximation of the actually absorbed dose in material. And, compared to fluence, it is a more intuitive unit.

Statistics

A final remark concerns the statistics of the Monte Carlo simulations. Each scenario is associated with an ensemble of $4 \cdot 10^5$ particles describing the initial beam, and the simulation consists of 10 independent Fluka runs with different seeds for the code's pseudorandom number generator. This means that, although each run starts from the same set of $4 \cdot 10^5$ primary particles, it provides results that are (practically) statistically independent from the other runs. This approach makes it possible to calculate statistical errors. All simulation results quoted in the following sections—including plots—, are based on the average of these 10 runs. Given errors specify the standard error of the mean, i.e. the standard deviation of the single results divided by $\sqrt{10}$.

5.5 Beam loss scenarios

5.5.1 Operation with activated suppressor dipoles

As discussed before, FLASH is occasionally operated in a mode in which two suppressor dipoles deflect the beam onto the beam pipe in order to protect the undulators from local beam losses. To simulate this situation, the dipoles—located at about 196.5 m—are included in the input file for start-to-end tracking of the beam with the Astra code (cf. section 4.4). The dipoles are initialized with their maximum field of 312 mT. An ensemble of $4 \cdot 10^5$ beam particles is tracked to 197.0 m, where the beam has left the magnetic fringe fields of the dipoles but has not yet collided with the beam pipe. Coordinates and momenta of the beam particles are saved to a file which is then used as the input distribution for the shower simulation with Fluka.

Figure 5.8 shows an overview of the shower distribution for the cases with and without shielding—i.e., with just the narrow beam pipe in the upper plot, and with both copper absorber and lead wall installed at 198.5 m in the lower plot. The quantity shown is dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$ for single bunch operation with the design parameters of FLASH (bunch charge 1 nC, repetition rate 10 Hz). The simulated beam energy is 980 MeV. The shown dose rate is averaged over the vertical range from −25 cm to +25 cm; the plot is overlaid with a cut through the model geometry at $y = 0$.

A huge amount of radiation is released at about 197.4 m, where the deflected beam hits the pipe. At 25 cm distance from the beam pipe, dose rates are of the order of 100 Sv/h. To set this in perspective, the median lethal dose to humans is about 4.5 Sv (as set out e.g. in [Art08]). The most intense part of the shower escapes to the right (towards negative x), but a strong radiation field develops also to the opposite side and inside the beam pipe. Six meters further downstream, close to the entrance of the undulator, the dose equivalent rate is still of the order of 0.1–1 Sv/h.

In the shadow of the shielding setup, dose equivalent rates are reduced by about an order of magnitude; at the undulator entrance, values are in the range of 0.01–0.1 Sv/h. Still, radiation sensitive components clearly need additional local shielding, especially if they are to be placed close to the shield wall.

Figure 5.9 shows the dose rate absorbed by beamline components. Again, the color-code indicates values that are averaged over the vertical range from −25 cm to +25 cm. This makes the plot slightly misleading because, in contrast to the ambient dose equivalent which is defined at any point in space, absorbed dose is only defined in matter. Therefore, the averaging favors parts of the geometry with more material along the vertical axis. This is plainly visible in the gaps between the outer undulator shield and the inner magnet holder structure at $z > 203.5$ m and $7.2\,\mathrm{cm} < |x| < 9\,\mathrm{cm}$. Nonetheless, the plot allows relative comparisons. It indicates that, apart from the initial beam pipe, the components most affected by the shower are the first quadrupole magnet (Q16 at 198.2 m) and its surrounding

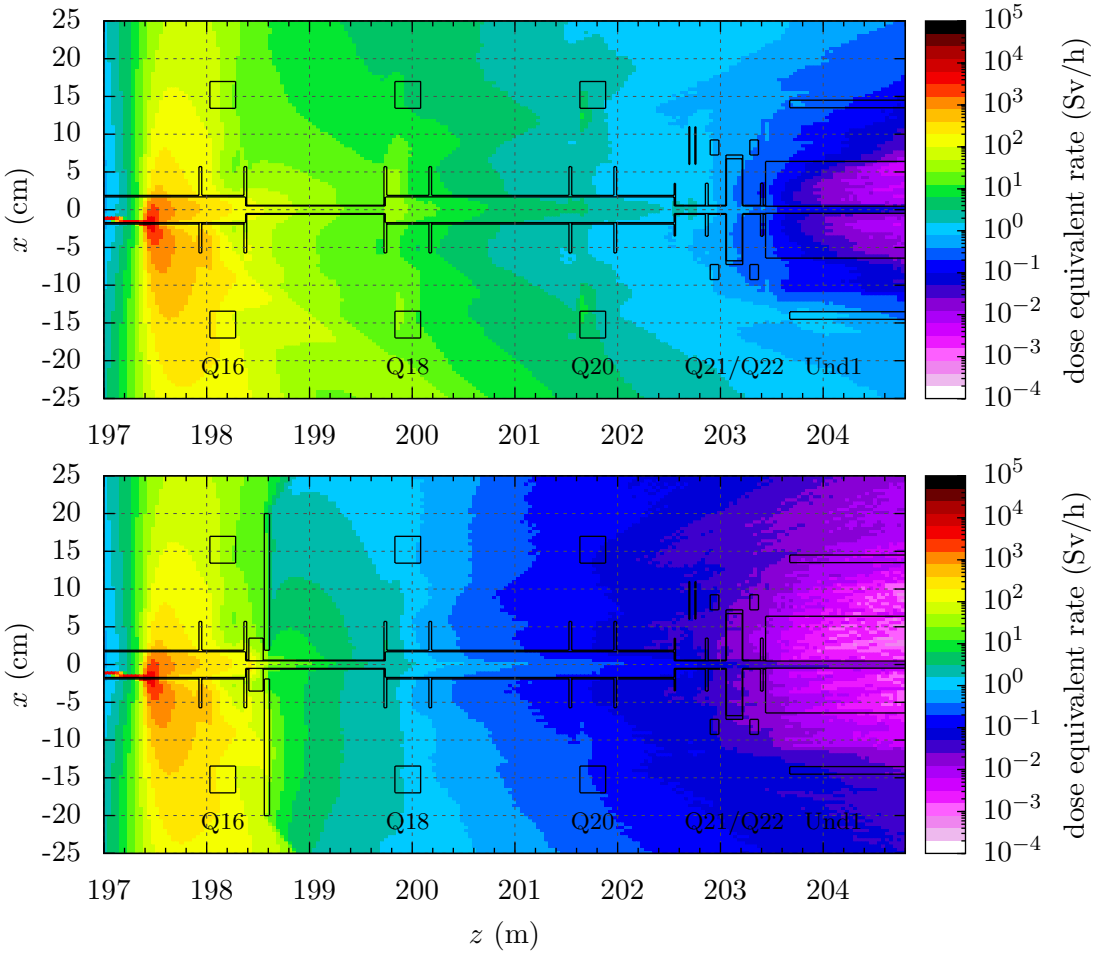

Figure 5.8: Ambient dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$ during operation with suppressor dipoles. The plot shows the section upstream of the main undulators. Quadrupole magnets are indicated, and about one third of the first undulator chamber is visible on the right side. Normalized to 1 bunch at 1 nC and 10 Hz. *Top:* no shielding. *Bottom:* shielding by copper absorber and lead wall at 198.5 m.

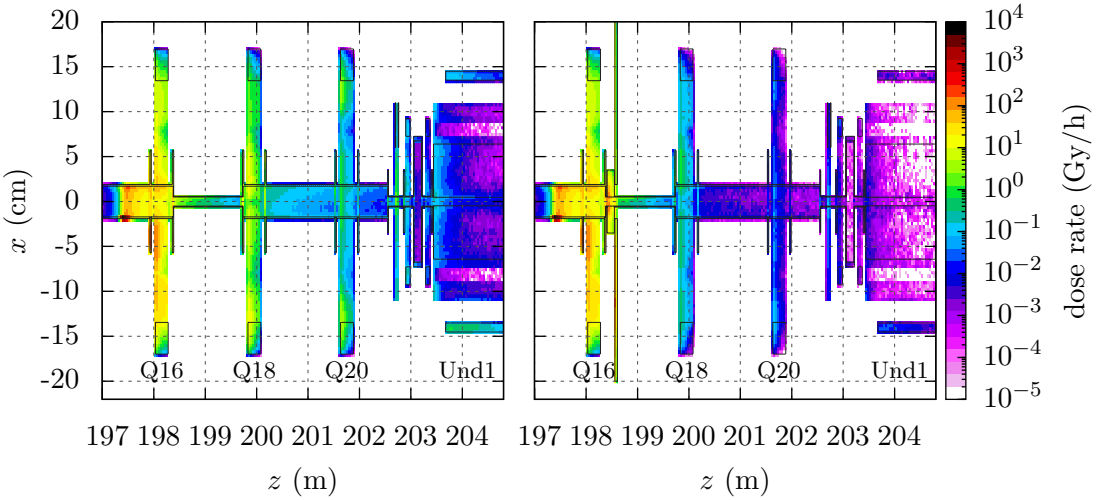

Figure 5.9: Absorbed dose rate in beamline components during operation with suppressor dipoles, averaged over the vertical range of $[-25; 25]$ cm. Normalized to 1 bunch at 1 nC and 10 Hz. *Left:* no shielding. *Right:* shielding by copper absorber and lead wall at 198.5 m.

Table 5.4: Dose deposition in the first undulator magnets during operation with suppressor dipoles. Normalized to 1 bunch at 1 nC and 10 Hz.

	Dose rate (Gy/h)	
	Unshielded	**Shielded**
Top	0.185 ± 0.011	0.010 ± 0.002
Bottom	0.212 ± 0.013	0.010 ± 0.002

flanges. The next two quadrupoles absorb one or two orders of magnitude less energy, and profit from the installed shielding.

Figure 5.9 also indicates that the highest dose rates in the undulator appear in the first few centimeters of the device. Hence, the most affected magnets are the ones close to the undulator entrance. Table 5.4 summarizes the deposited dose rates for the first (or most upstream) magnets of the undulator structure. In the case of no shielding, these magnets accumulate about 5 Gy/d, which coincides with the lifetime dose rate limit discussed before. Although operation with activated suppressor dipoles is limited to short time spans, the reduction to 0.24 Gy/d achieved by the shielding setup provides a welcome safety margin.

Power absorption

With active suppressor dipoles, the beam power is not properly contained in a beam dump anymore. It is therefore important to clarify how much power is dissipated in individual components. To answer this, the energy balance of each modeled component is determined with Fluka. From these data, the net power absorption is calculated with the same beam parameters as above. The beam

Component	Power deposition (W) Unshielded	Shielded
Lead wall	—	1.1008 ± 0.0005
Copper absorber	—	0.3940 ± 0.0004
Beam pipe u/s Q16	3.0311 ± 0.0004	3.0315 ± 0.0004
Flange u/s Q16	0.4820 ± 0.0002	0.4816 ± 0.0002
Quadrupole Q16	1.5945 ± 0.0005	1.5961 ± 0.0005
Flange d/s Q16	0.1241 ± 0.0001	0.1254 ± 0.0001
Quadrupole Q18	0.2211 ± 0.0003	0.0131 ± 0.0001
Escaped	4.0748 ± 0.0008	2.9125 ± 0.0005

Table 5.5: Power deposition in various beamline components during operation with suppressor dipoles. The assumed total beam power is 9.8 W. *Escaped* subsumes all particles leaving the simulation volume. *u/s:* upstream (towards lower z), *d/s:* downstream (towards higher z)

carries a total power of

$$P = 10\,\mathrm{Hz} \cdot 1\,\mathrm{nC} \cdot 980\,\frac{\mathrm{MeV}}{e} = 9.8\,\mathrm{W}.$$

Table 5.5 lists the most affected beamline components. The biggest single part of about 3 W is deposited in the beam pipe at and around the point of impact. This causes mainly a moderate and therefore tolerable heating of the pipe.

As observed before, quadrupole Q16 is the next big absorber of about 1.6 W of radiation power. Since the magnet consists only of radiation-hard materials, this is unproblematic. If the operation with deflected beam is extended over long periods of time, activation of both the beam pipe and of the iron yoke might become significant.

Because they constitute comparatively massive bodies, flanges also absorb a certain amount of power. This is especially dangerous because the mechanical stress introduced by thermal expansion and contraction can ultimately create a vacuum leak in these components. In the simulation, the flange upstream of Q16 overlaps with the main shower emanating from the beam pipe. It therefore absorbs a maximum power of about half a watt, which is considered noncritical.

When the shielding components are included in the simulation, it is found that they absorb a total of 1.5 W. While this reduces the load on downstream components, it obviously cannot improve the situation for quadrupole Q16 and its surroundings. Because of severe space limitations, only a major redesign of the section could remedy this problem.

Finally, the high *escaping* power of 4 W in the unshielded and 3 W in the shielded case, i.e. the total power of all particles leaving the simulation volume, is an indicator for the amount of radiation released to the immediate environment of the accelerator. It is therefore clear that the installed shielding setup does not replace

a beam dump. Nevertheless, it effectively limits the dose absorbed by downstream components during operation with activated suppressor dipoles. Moreover, the shielding is sufficient to remove any threat this mode of operation poses to the undulator magnets.

5.5.2 Missteered beam at the experimental BPM beamline

In the second scenario, an accidental beam loss at the flanged transition between the 34 mm and the 10 mm beam pipe just behind Q16 is simulated. Because the beam sizes in this part of the beamline are still much bigger than in the undulator itself (cf. Fig. 4.9), the transition is comparatively often hit by a missteered beam during linac setup. Even more frequently, and for longer periods of time, parts of the dark current are lost on the flange. The copper absorber is the main device for protection of the undulators from the generated shower.

The simulation is set up in the same way as described above. Again, the suppressor dipoles are used to deflect the beam horizontally; this time, however, only a kick of 4.3 mrad is induced. This is just sufficient to guide all beam particles onto the flange downstream of Q16.

The resulting dose equivalent rates are shown in Fig. 5.10. In the unshielded case, the biggest part of the beam is transmitted through the flange, travels alongside the small beam pipe, and then hits the next flange in front of Q18, where it is dispersed to a wide solid angle. At both of these main interaction points, huge quantities of electrons, positrons, photons, and neutrons are released. Because these hot spots are much closer to the undulator than in the last scenario, also the dose equivalent rate at the undulator entrance is higher—of the order of 10 Sv/h.

Introduction of the absorber changes the picture completely: The beam is still transmitted through the flange, but most of its energy is now deposited in the copper. The lead shield serves as an additional barrier for the created radiation field. Although the radiation field behind the shielding is still significant, the dose equivalent rate at the undulator entrance is reduced by two orders of magnitude to about 0.1 Sv/h.

The distribution of absorbed doses shown in Fig. 5.11 indicates that the yoke of quadrupole Q18 is one of the main absorbers of radiation energy in the unshielded case. Also the undulator receives a significant dose rate; as shown in Tab. 5.6, the first magnets absorb about 8.2 Gy/h for the loss of a single bunch. The loss of macropulses with 30 bunches is the worst case allowed by the MPS. In this case, the dose rate could reach 250 Gy/h for some time before the MPS intervenes by disabling the beam (cf. section 6.2). The shielding setup removes this hazard completely by reducing these dose rates to 30 mGy/h for a single bunch or 0.9 Gy/h for 30 bunches.

The evaluation of the power deposition in various beamline elements (Tab. 5.7) shows that the flange on which the beam is dumped absorbs only 0.17 W from the loss of a single-bunch beam. This is the direct consequence of the low stopping power for electrons of high energy; while bremsstrahlung is emitted by the electrons

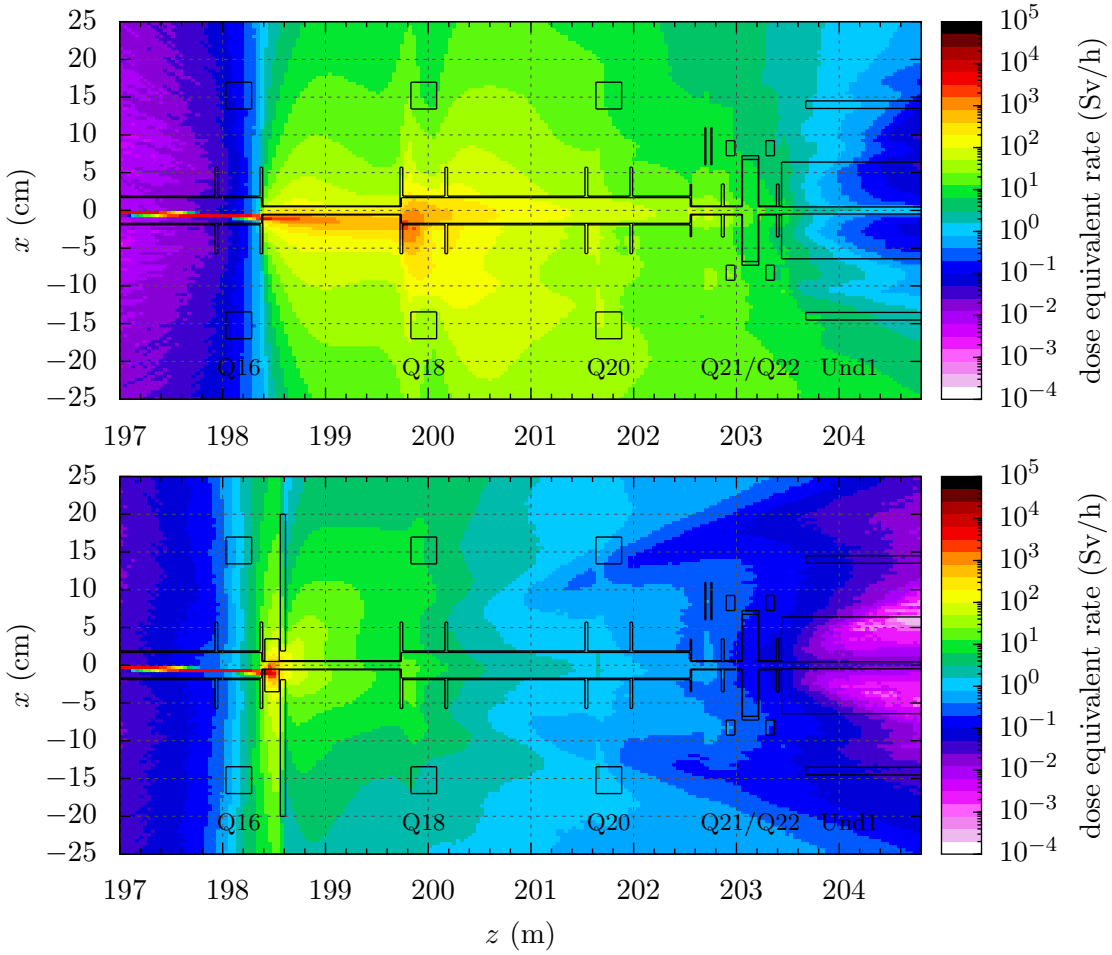

Figure 5.10: Ambient dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$ during an accidental beam loss at the beginning of the small beam pipe behind Q16. Normalized to 1 bunch at 1 nC and 10 Hz. *Top:* no shielding. *Bottom:* shielding by copper absorber and lead wall at 198.5 m.

Table 5.6: Dose deposition in the first undulator magnets for accidental beam loss behind Q16. Parameters: 1 bunch, 980 MeV, 10 Hz.

	Dose rate (Gy/h)	
	Unshielded	**Shielded**
Top	8.17 ± 0.06	0.031 ± 0.004
Bottom	8.21 ± 0.10	0.030 ± 0.004

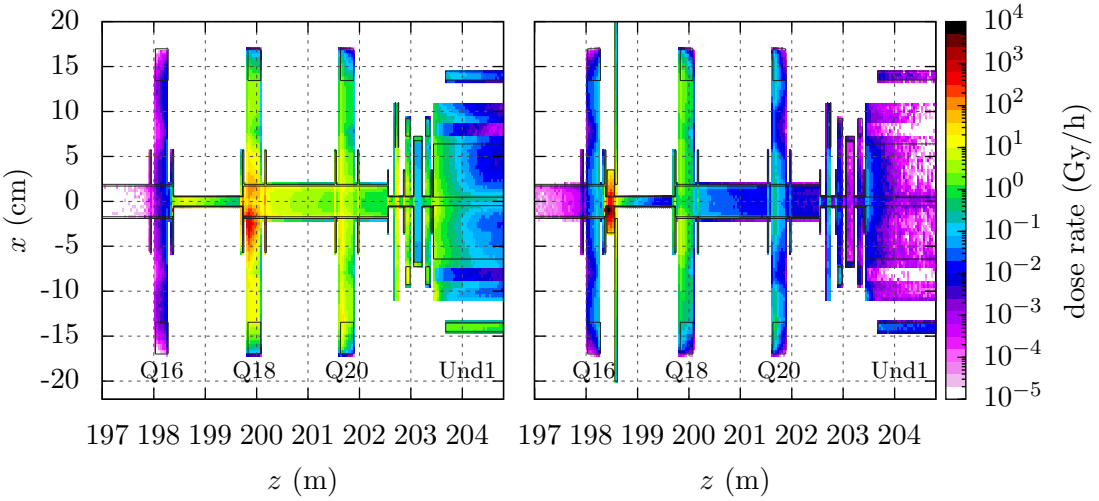

Figure 5.11: Absorbed dose rate in beamline components for accidental beam loss behind Q16, averaged over the vertical range of $[-25; 25]$ cm. Normalized to 1 bunch at 1 nC and 10 Hz. *Left:* no shielding. *Right:* shielding by copper absorber and lead wall at 198.5 m.

	Power deposition (W)	
Component	**Unshielded**	**Shielded**
Lead wall	—	0.58888 ± 0.00014
Copper absorber	—	8.30645 ± 0.00032
Flange d/s Q16	0.1669 ± 0.0001	0.17563 ± 0.00004
Flange u/s Q18	0.9251 ± 0.0002	0.02080 ± 0.00003
Quadrupole Q18	3.3865 ± 0.0014	0.09475 ± 0.00011
Flange d/s Q18	0.3038 ± 0.0002	0.00273 ± 0.00001
Beam pipe d/s Q18	0.6711 ± 0.0004	0.00337 ± 0.00002
Quadrupole Q20	0.3960 ± 0.0004	0.00752 ± 0.00002
Escaped	2.9625 ± 0.0008	0.52463 ± 0.00022

Table 5.7: Power deposition in various beamline components for accidental beam loss behind Q16. The assumed total beam power is 9.8 W. *Escaped* subsumes all particles leaving the simulation volume. *u/s:* upstream (towards lower z), *d/s:* downstream (towards higher z)

passing the flange material, its thickness is neither sufficient for the re-absorption of a significant part of the photons nor for the production of a significant number of electron-positron pairs. Instead, the photon-electron beam develops into a full electromagnetic shower in the copper absorber which can thus intercept about 85 % of the initial beam power; another 6 % are absorbed by the lead wall. Hence the shielding setup is an efficient protection against accidental beam losses.

On a final note, the full loss of 30 bunches at 10 Hz may cause a problematic power deposition of 5.3 W in the flange. This problem is in part remedied by the fact that, up to now, the FLASH rf systems allow only a maximum repetition rate of 5 Hz. Nevertheless, operational care is advised until a redesign of the section can address this issue.

5.5.3 Dark current losses in front of the diagnostic undulator

Because the first 10 mm beam pipe with its shielding installation acts like a collimator, the second transition from 34 mm to 10 mm diameter in front of the diagnostic undulator is practically free of accidental beam losses. This is supported by the optics setup which uses quadrupoles Q18–Q22 to reduce the beta functions to an average value of less than 5 m starting from about 202.5 m.

Nevertheless, a special situation can create strong dark current losses in this location. During the setup of the accelerator, the quadrupole gradients are adjusted to match the beam energy. When the setup is finished, the electromagnets are *cycled* in order to bring their remanent magnetization into a well-defined state; this procedure typically works by ramping the magnet current to its minimum, then to its maximum, and finally to the desired setpoint—or in a similar manner, depending on the magnet type. As this would obviously cause beam losses, the operation procedure for FLASH requires a manual disabling of the injector laser before cycling. However, it leaves any measures against the transport of dark current to the discretion of the operators. If no such measures like the disabling of an acceleration module are taken, the defocused gun dark current can be dumped close to the undulator, usually triggering an emergency shutoff by the BIS.

The magnets can be cycled in groups or individually; since their cycling procedures vary as well, there is no single representative model for dark current losses caused this way. Instead, a well-defined exemplary scenario is chosen by assuming normal dark current transport up to $z = 199\,\mathrm{m}$. The following two quadrupoles Q18 and Q20 are set to their maximum gradient of $+15.2\,\mathrm{T/m}$. The transport up to the end of the magnetic field region of Q20 is done with Astra as described in the previous chapter. To obtain better statistics, the number of started macroparticles is increased by a factor of ten, yielding about $2 \cdot 10^5$ remaining dark current particles after the quadrupole. This ensemble is then used as the input distribution for the Fluka simulation in which all remaining undulator quadrupoles are set up according to the design optics. The simulation results are normalized using the following parameter set, describing the worst-case scenario for FLASH:

- dark current behind Q20: 3.9 µA (cf. Tab. 4.5)

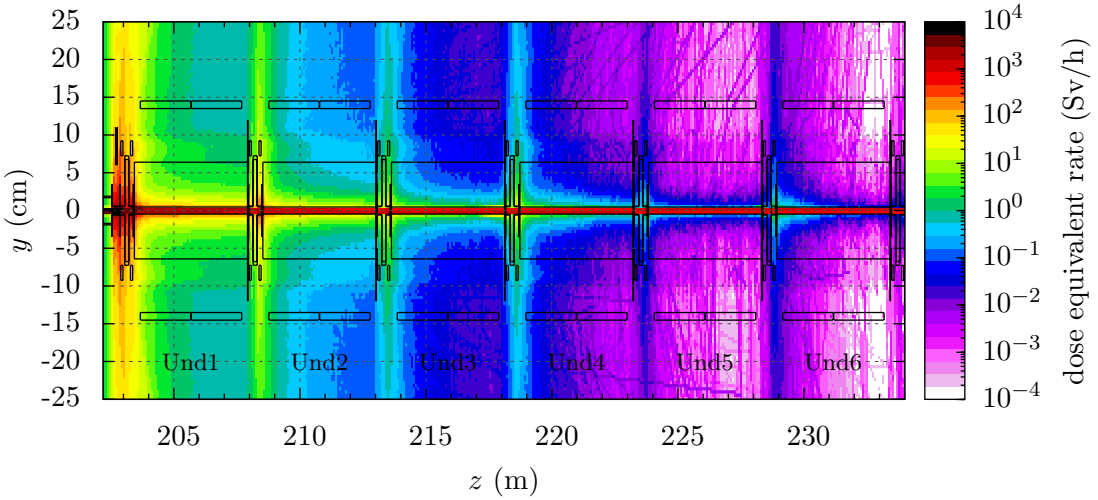

Figure 5.12: Ambient dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$ for dark current losses caused by quadrupole cycling. Averaged over $y \in [-25; 25]$ cm. Normalized to a current of 3.9 µA, 800 µs rf pulse length and 10 Hz repetition rate.

- rf pulse length: 800 µs
- repetition frequency: 10 Hz

In passing Q18 and Q20, the dark current suffers a strong vertical defocusing and hits the transition to the 10 mm pipe. As shown in Fig. 5.12, grazing losses continue throughout the whole undulator section with pronounced maxima at the positions of the quadrupoles, i.e. at the maxima of the beta function. Because of the lack of shielding, these locations are also the only ones in which a significant part of the shower can escape transversely. Close to the entrance of undulator 1, dose equivalent rates are of the order of 1000 Sv/h, and in front of undulator 2 still of 100 Sv/h.

With the parameters given above, the dark current carries a total power of 30.4 W. Only 16.3 % of this power are transported all the way to the exit of undulator 6 without interaction with beam pipes or vacuum chambers. Another 8.2 % are released as radiation into the environment. The remaining 23.0 W are deposited in various beamline elements; especially, a power of 3.9 W is absorbed by permanent magnets, 2.9 W of this by the magnets in undulator 1.

Figure 5.13 shows the corresponding distribution of absorbed dose rates. Unlike in previous plots, the averaging is restricted to the height and depth of the internal magnet structure. In general, the dose rate decreases with increasing distance from the beam axis. Hence, the deposition of energy in the magnet pieces is highly inhomogeneous; for undulator 1, the mean dose rate varies from more than 1 kGy/h in the vicinity of the beam axis to less than 5 Gy/h at the outer edges of the magnets.

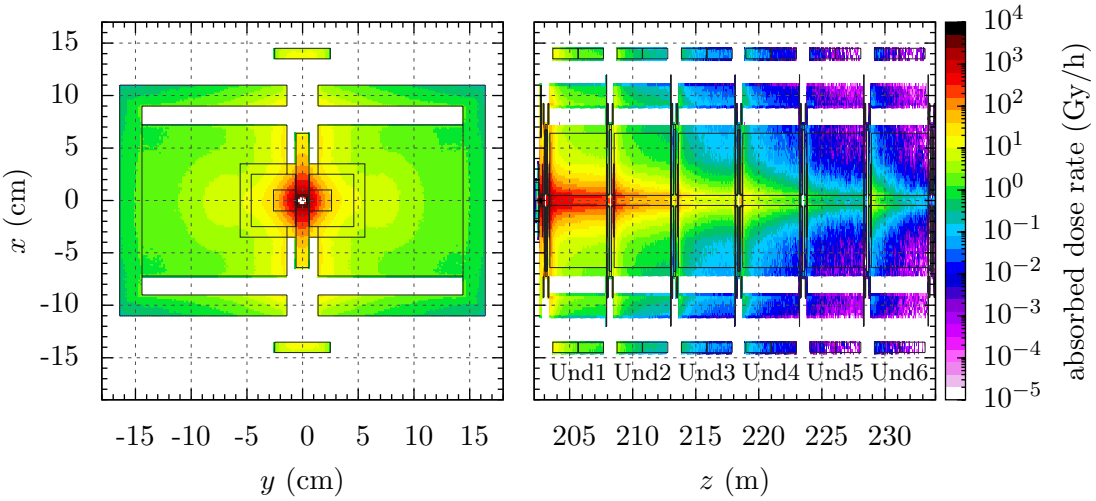

Figure 5.13: Absorbed dose rate for dark current losses caused by quadrupole cycling. *Left:* transverse distribution, averaged over the full interior of undulator 1, $z \in [203.48; 207.87]$ m. *Right:* longitudinal distribution, averaged over the vertical range covering both lower and upper magnets, $y \in [-5.6; 5.6]$ cm.

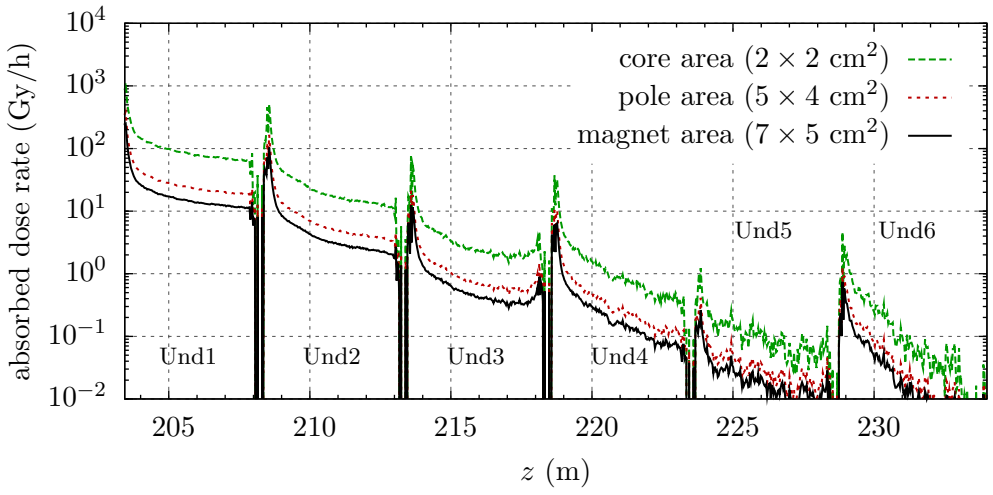

Figure 5.14: Absorbed dose rate in permanent magnets vs. z for dark current losses caused by quadrupole cycling.

	Dose rate (Gy/h)						
	Diag.	**Und1**	**Und2**	**Und3**	**Und4**	**Und5**	**Und6**
Top	4508 ± 7	1057 ± 2	369 ± 1	45 ± 1	26.6 ± 0.4	0.9 ± 0.1	2.3 ± 0.1
Bottom	4478 ± 6	1065 ± 2	365 ± 1	44 ± 1	27.1 ± 0.4	1.2 ± 0.1	2.5 ± 0.1

Table 5.8: Average dose deposition in the first undulator magnets for dark current loss during quadrupole cycling.

The dose also decreases with increasing distance from the entry of each undulator segment. As observed before, the strongest losses occur in the intersections, close to or in the quadrupole magnets; since the generated showers are directed towards positive z, the following undulator is the main absorber. This can be seen more clearly in Fig. 5.14 which shows the dose rates deposited in the permanent magnets as a function of the longitudinal coordinate. To reflect the inhomogeneity of the energy deposition, the dose rate is calculated for three different transverse dimensions: first, the complete cross-sectional area of the permanent magnet pieces; second, the area of the pole pieces; and third, a core area of $2 \times 2\,\mathrm{cm}^2$ centered horizontally and aligned with the edge of the magnet pieces facing the vacuum chamber. Because the geometry is completely symmetric, only the results for the upper magnet structure are shown.

From the beginning of the undulator section to its end, the dose rates decrease by five orders of magnitude. Within each segment, the first magnets absorb at least 30 times more energy than those close to the end. To estimate the worst-case impact of the cycling procedure, it is therefore sufficient to consider the absorbed dose rates for the first magnets of each segment as summarized in Tab. 5.8. In addition, the table lists the average dose rate for the two inner magnet pieces of the diagnostic undulator which are more than four times higher than those expected for undulator 1.

The cycling of quadrupoles at FLASH is a relatively slow procedure. For a very rough estimate, it can be assumed that losses comparable to this simulation occur for an integrated time of about 30 seconds per cycle. Without intervention by the machine protection system, this would accumulate an average dose of 37 Gy in the magnets of the diagnostic undulator and 9 Gy in the first magnets of undulator 1. This means that the daily allowance of 5 Gy for the undulator magnets could be exceeded with a single cycling of the quadrupoles at full rf pulse length.

These high energy depositions should also be expected to generate significant neutron flux. Figure 5.15 shows the simulated neutron flux in the most affected region around undulator 1. In this case, the main sources of neutrons are ferrous components hit by electromagnetic showers, namely the diagnostic undulator at 202.7 m, the flange upstream of Q21, and the entry regions of undulators 1 and 2. In the vicinity of the magnets, neutron fluxes are of the order of $10^7\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. According to [And07], fluences (i.e. time-integrated fluxes) of $10^{13}\,\mathrm{cm}^{-2}$ can cause first measurable demagnetization; in terms of figures, a fluence of this magnitude

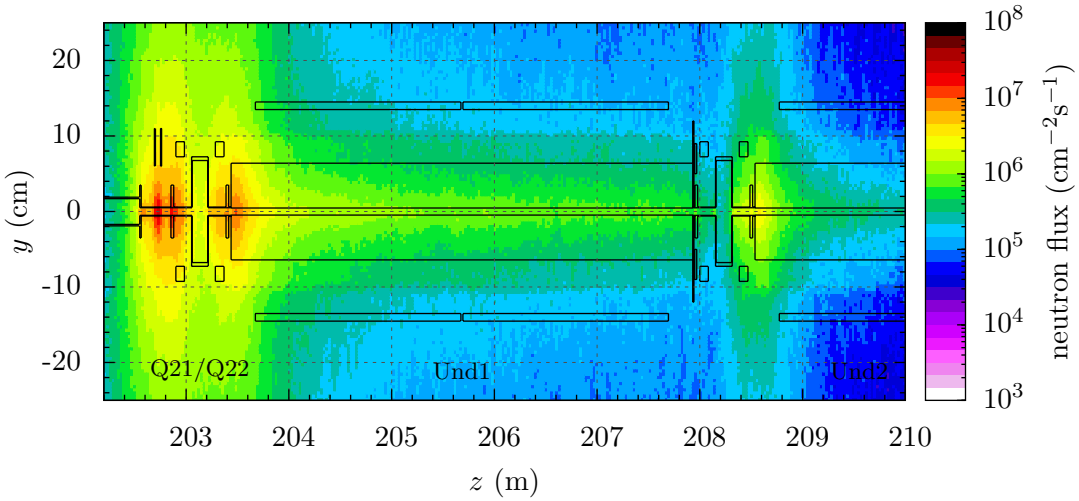

Figure 5.15: Neutron flux from dark current losses caused by quadrupole cycling. Averaged over $y \in [-25; 25]$ cm. Normalized to a current of 3.9 µA, 800 µs rf pulse length and 10 Hz repetition rate.

could be accumulated after about 10 days of continuous irradiation according to the given parameters.

Although this scenario simulates a rather extreme condition, it yields a lot of useful information: First, it quantifies the damage potential held by even the small amount of gun dark current reaching the undulator section. Second, it shows several characteristics of the dose distribution that can be considered universal—e.g. the big inhomogeneity of dose within a single magnet piece, or the much stronger exposure of the front part of the undulator segments as compared to the back part. And last, although these incidents are practically always handled by the MPS before dangerous doses are accumulated, dark current losses due to quadrupole cycling do happen occasionally. Therefore, these calculations provide an input for further discussion on the improvement of operational procedures at FLASH.

5.5.4 Horizontal beam loss in the undulator

Similar to the previous scenario, the last two simulations deal with extreme situations. In this section, the consequences of a complete horizontal beam loss in undulator 1 are analyzed; a corresponding vertical beam loss is examined in the following section.

To produce a sufficient deflection, the last horizontal corrector coil in front of the undulators (at 200.56 m) is set to its maximum field of 32 mT in the Astra simulation. After the beam has left the magnetic field region of the coil, the particle distribution is saved at 201.0 m and imported to Fluka for further tracking. The

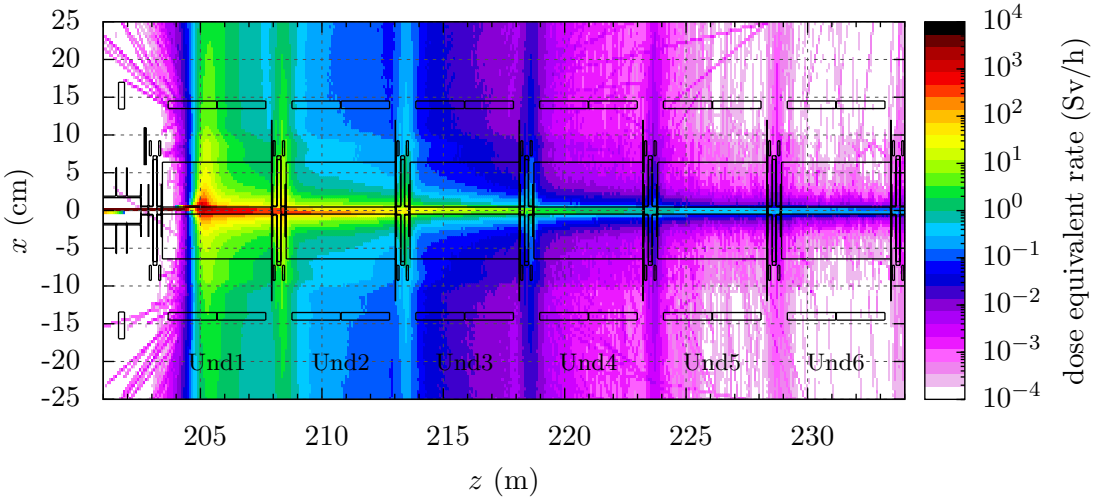

Figure 5.16: Ambient dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$ for a horizontal beam dump in undulator 1. Parameters: 1 bunch at 1 nC and 10 Hz.

beam enters the quadrupole doublet Q21/Q22 at an offset of 2.3 mm; because Q22 is horizontally defocusing, it increases the deflection even more so that the well-focused beam hits the vacuum chamber at 204.8 m. This scenario therefore represents a point-like loss in the middle of an undulator segment.

The dose equivalent rate overview in Fig. 5.16 exhibits characteristics similar to the previous scenario. The main qualitative difference is that practically none of the primary particles are transported all the way through the beam line to the exit of the last undulator. Furthermore, it is noteworthy that the created showers are highly contained within the undulator. At a horizontal distance of 25 cm, equivalent dose rates are below 5 Sv/h which is to be compared with about 100 Sv/h or more for previous scenarios.

The absorbed dose rates shown in Fig. 5.17 confirm that a major part of the beam energy is deposited in the vacuum chamber and in the magnet structure. The left plot shows the transverse distribution of dose rates, averaged over a narrow longitudinal range of 30 cm length shortly behind the beam impact. In this worst affected region, the deposition within the magnet ranges from 5 Gy at the outer edges to more than 10 kGy/h close to the undulator axis. While the distribution is still notably asymmetric at this position, the shower tends to symmetrize with increasing z.

The longitudinal distribution of dose rates as shown in Fig. 5.18 is very similar to the one caused by the dark current loss. In absolute values, the average dose rates are slightly lower because of the small beam power of 9.8 W as compared to the 30.4 W carried by the dark current. The maximum dose rate in a single magnet reaches 400 Gy/h if averaged over the full magnet volume, or 2 kGy/h in the core region. It should be kept in mind, however, that beam losses are automatically

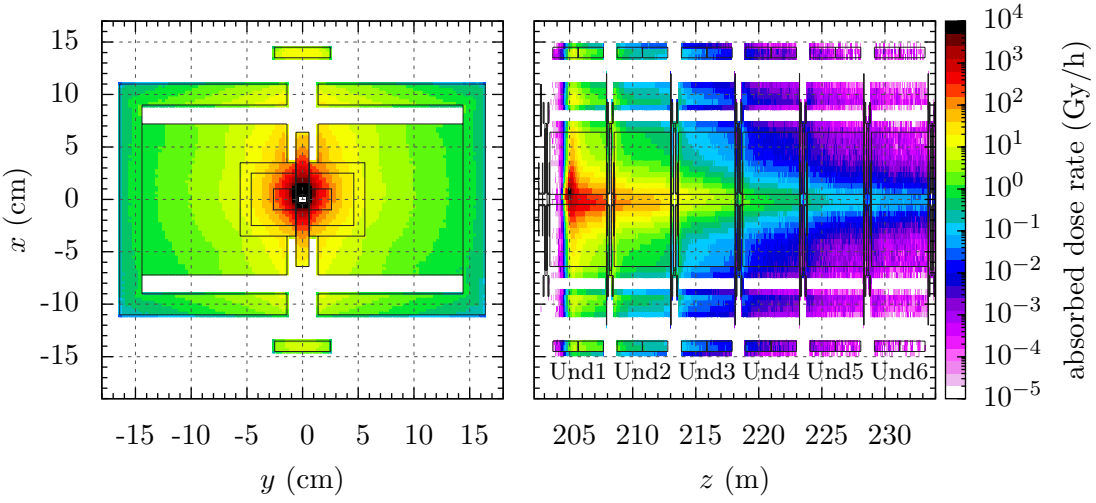

Figure 5.17: Absorbed dose rate for a horizontal beam dump in undulator 1. The color scale is capped at 10^4 Gy/h. Parameters: 1 bunch at 1 nC and 10 Hz. *Left:* transverse distribution, averaged over the vicinity of the beam impact, $z \in [204.9; 205.2]$ m. *Right:* longitudinal distribution, averaged over the vertical range covering both lower and upper magnets, $y \in [-5.6; 5.6]$ cm.

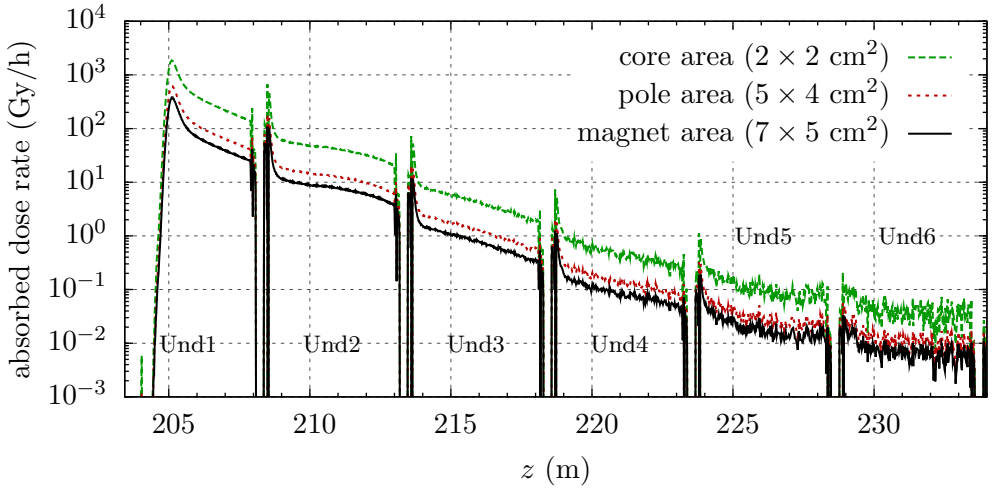

Figure 5.18: Absorbed dose rate in permanent magnets vs. z for a horizontal beam dump in undulator 1. The color scale is capped at 10^4 Gy/h. Parameters: 1 bunch at 1 nC and 10 Hz.

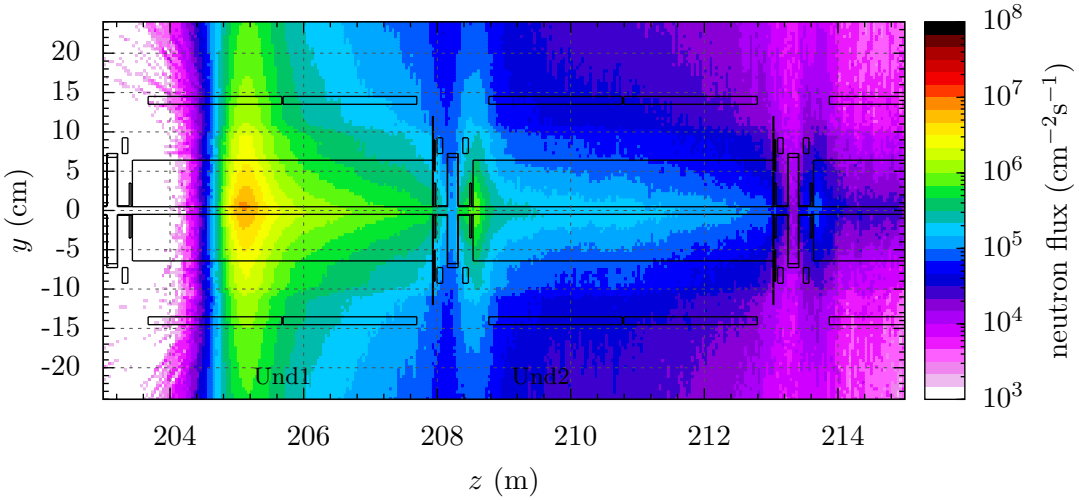

Figure 5.19: Neutron flux caused by horizontal beam dump in undulator 1, averaged over $y \in [-24; 24]$ cm. Parameters: 1 bunch at 1 nC and 10 Hz.

accompanied by dark current losses. For the parameters chosen here, this would increase dose rates by roughly a factor of four.

The same remark is valid for the neutron flux depicted in Fig. 5.19. Although the average flux is below $10^7\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ throughout the whole geometry, this may be modified by a factor of up to four by inclusion of dark current losses. The loss of more than a single bunch per macropulse increases the number of neutrons accordingly.

To obtain information on the typical neutron energies within the undulator, the flux through the upper and lower surfaces of the vacuum chambers can be analyzed. Figure 5.20 shows the resulting neutron spectrum. The number of photonuclear reactions in the simulation is not sufficient to resolve the low energy part of the spectrum; what can be observed, however, is that the spectrum extends almost to the beam energy of 980 MeV, and that these high-energy neutrons are much less prevalent than fast neutrons with energies up to few MeV.

In the same way, the spectra of electrons/positrons and photons can be determined. Here, the lower energy boundary is determined by the chosen 200 keV production threshold. Both spectra extend to the beam energy of 980 MeV and show a superior number of low energy particles as expected for an electromagnetic shower. Finally, two distinct features are discernible. First, there is a substantial increase of the $\mathrm{e}^{\pm}$ flux close to the maximum energy. This is caused by the primary electrons of the beam which can cross the boundary surfaces multiple times due to scattering events. Second, a line at 511 keV in the photon spectrum is caused by the annihilation of $\mathrm{e}^{+}/\mathrm{e}^{-}$ pairs.

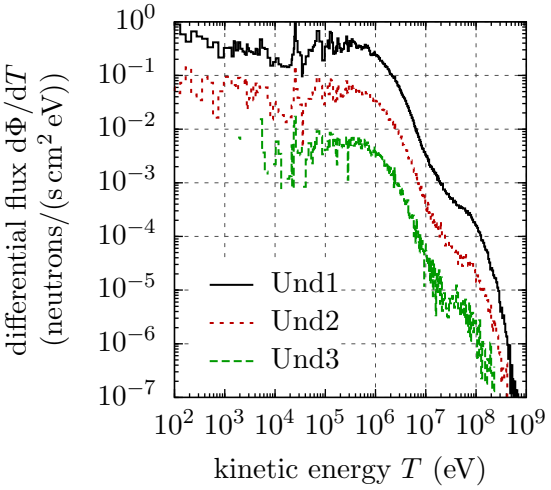

Figure 5.20: Neutron spectrum caused by horizontal beam dump in undulator 1.

5.5.5 Vertical beam dump in the undulator

To complete the discussion, the last scenario examines a vertical beam loss in undulator 1. Like in the previous case, a corrector coil at 200.76 m is set to its maximum field of 32 mT, deflecting the beam downwards. To overcome the vertical focusing of quadrupole Q22 just in front of the undulator, the incident beam is additionally displaced by −1.2 mm. This causes a small fraction of the particles to be lost at the transition to the vacuum chamber of undulator 1. Afterwards, the outer parts of the beam are grazing along the lower wall of the chamber; close to the end of the undulator segment, between 207 and 208 m, the core of the beam comes into contact with the aluminum chamber and creates the primary shower.

The results of the shower simulation are summarized in Fig. 5.22. As usual, all results are normalized to an incident beam of a single bunch per macropulse with a charge of 1 nC at a repetition rate of 10 Hz. It can be seen that, although the point of impact is located within undulator 1, the highest energy deposition is caused by the electromagnetic shower in undulator 2. Again, it can be observed that the initial asymmetry of the shower decreases towards higher z. While the absorbed dose rates for corresponding upper and lower magnets differ by factors of 2–4 in undulator 1 (or slightly more for the first few periods), the ratio varies between only 1.5 and 3 in undulator 2; from the third segment onwards, the difference is barely noticeable.

The highest dose rate in a single magnet slightly exceeds 1 kGy/h if averaged over the full magnet volume, or 8 kGy/h in the core region. A vertical beam loss therefore has a more detrimental effect on the undulator than a horizontal one. Even without considering the additional contribution by gun dark current, a beam loss of 10 seconds will deposit an average dose of 3 Gy in some of the most affected magnets—or more than 20 Gy in their core region.

Since this scenario is very similar to the last one, neutron fluences and energy spectra differ only slightly in absolute magnitude and warrant no further discus-

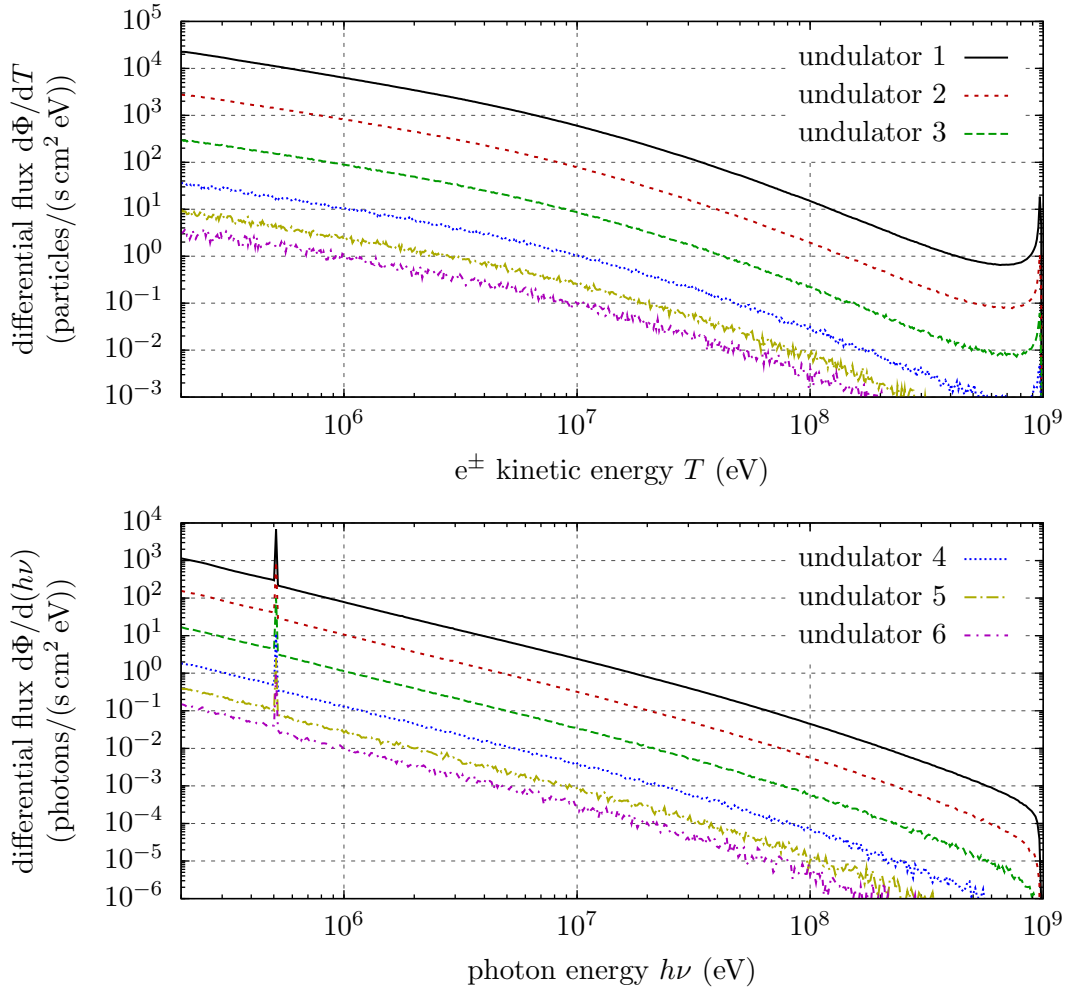

Figure 5.21: Electron/positron and photon spectra caused by horizontal beam dump in undulator 1. The photon spectrum shows a line due to e^+/e^- annihilation at 511 keV.

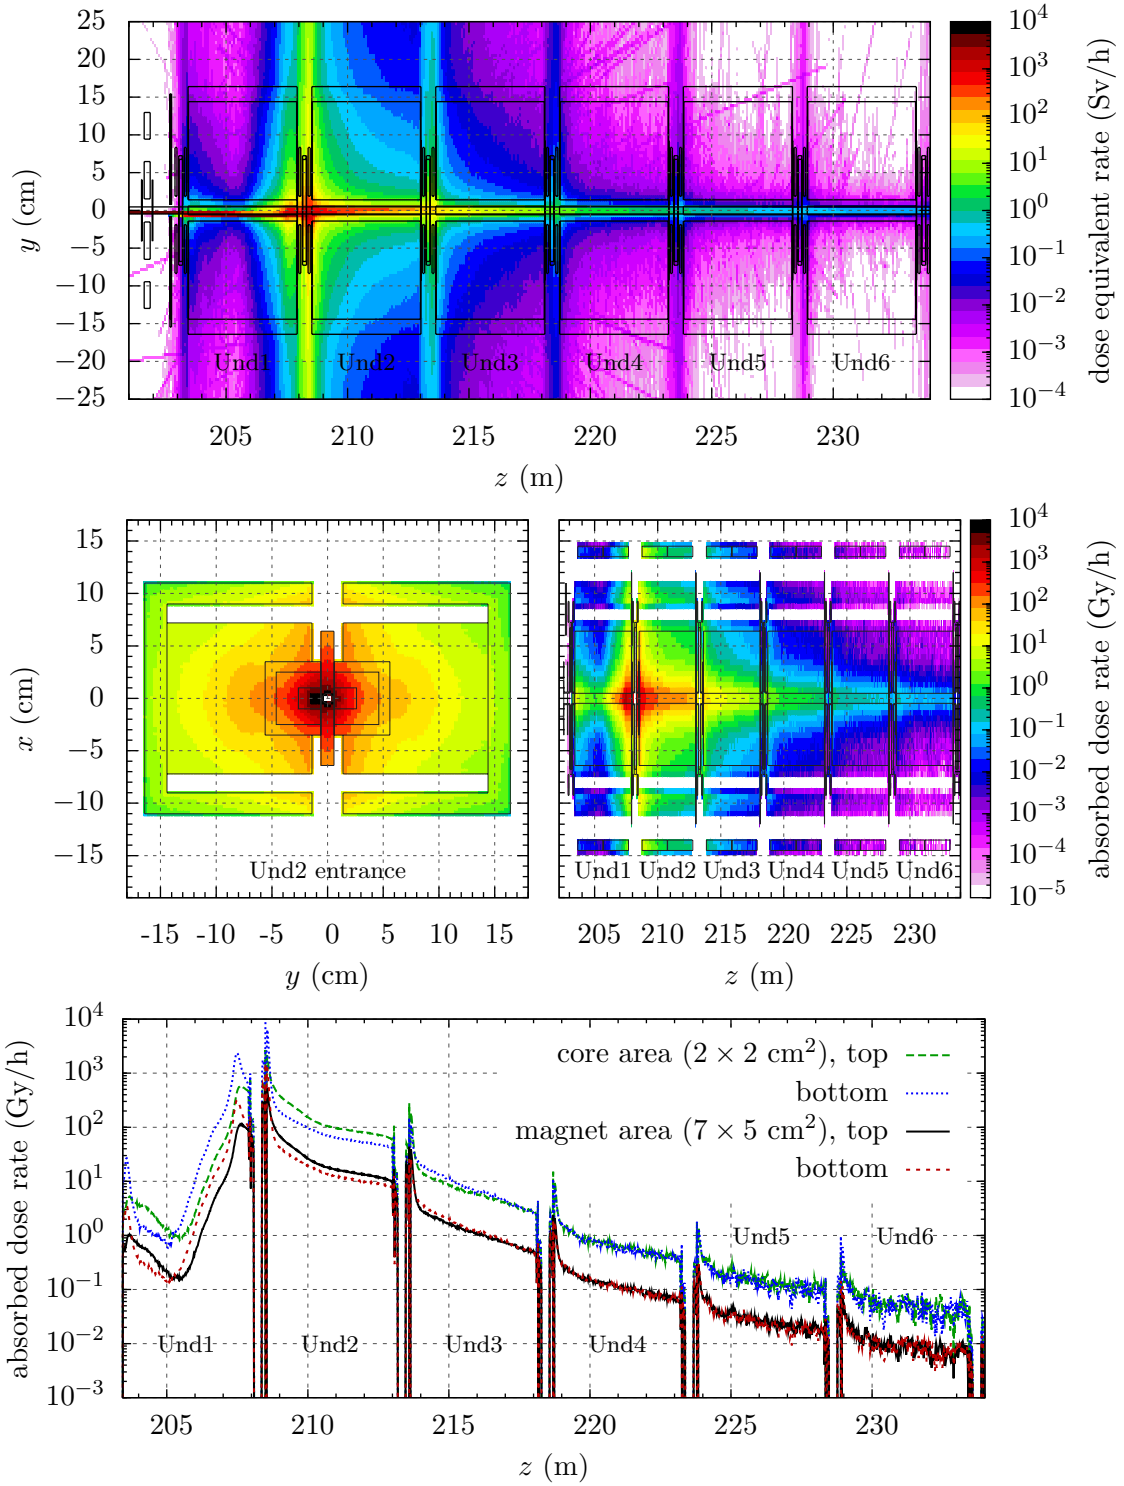

Figure 5.22: Dose rate plots for vertical beam loss in undulator 1. *Top:* Ambient dose equivalent rate $\mathrm{d}H^*(10)/\mathrm{d}t$. *Mid-left:* Absorbed dose rate, projection over $z \in [208.57; 208.67]$ m. *Mid-right:* Absorbed dose rate, projection over $y \in [-5.6; 5.6]$ cm. *Bottom:* Absorbed dose rate vs. z for various transverse areas.

sion.

5.6 Final remarks

Present-day insertion devices, especially those for ultraviolet and X-ray free-electron lasers, have to fulfill demanding requirements in terms of magnetic field quality. In most cases, this precludes the use of electromagnets; therefore, in the foreseeable future the vast majority of wigglers and undulators will continue to be constructed in permanent magnet technology.

It is by now a well-established fact that permanent magnets are prone to demagnetization under irradiation. This does not apply only to the often-used NdFeB alloy found in the FLASH undulators, but also to industrially available substitute materials like samarium cobalt. It is therefore necessary to limit the dose absorbed over the lifetime of the magnets. There are a number of ways to achieve this:

First, sufficient collimation and shielding are essential; as shown in this chapter, even minimalistic setups can achieve a considerable reduction of dose rates.

Second, the accumulation of radiation dose in case of accidental beam losses can be restricted by an active machine protection system that disables the beam in case of excessive losses. All of the accident scenarios discussed in this chapter can be—and have been, in fact—triggered by a single operating error. It has been shown that without the intervention of an MPS, unacceptable doses can be accumulated within seconds. A final calculation can further illustrate the problem:

In the vertical beam loss scenario, a single-bunch beam with a repetition rate of 10 Hz has been shown to deposit a dose rate of roughly 1 kGy/h in one of the permanent magnets. The underlying figure is a dose contribution of about 30 mGy per electron bunch of 1 nC. Using the FLASH design parameters for maximum beam power, a single bunch train contains 7200 of these bunches; this means that, within a single pulse of 800 µs length, a total dose of 200 Gy would be deposited—or, taking into account the repetition rate, a dose rate of 2 kGy/s.

Conversely, in order to stay within the dose rate limit of 5 Gy/d, the relative average beam loss in the 7200-bunch scenario would have to be limited to $3 \cdot 10^{-8}$, i.e. 30 aC or 180 electrons per bunch. For the current maximum operation parameters of FLASH, 800 bunches per macropulse at 5 Hz, the limit for relative loss is $5 \cdot 10^{-7}$, corresponding to 520 aC or 3250 electrons per bunch. This makes it clear that also very small beam losses must be controlled tightly; in this respect, the protection of insertion devices is one of the most challenging tasks for the MPS of a superconducting accelerator.

6 Machine protection system for FLASH

The discussion of beam loss scenarios for the undulator section has already illustrated the damage potential of the FLASH electron beam. The accelerator is capable of transporting high average beam powers that can easily cause mechanical damage, e.g. by evaporating material. With its demonstrated operation parameters of 800 bunches per macropulse, 5 Hz repetition rate, and 1 GeV beam energy, FLASH can sustain an average beam power of 4 kW at the bunch charge of 1 nC, or up to three times more at higher charge. This power will increase to 94 kW or more if the design parameters of the accelerator (7200 bunches, 10 Hz repetition rate) and an already scheduled energy upgrade to 1.3 GeV are realized.

An accelerator with this damage potential needs an active machine protection system that can identify hazardous conditions and react before harm is caused. As pointed out before, the loss of a full bunch train can already cause substantial damage. Hence, the MPS must be able to interrupt the beam *within* a macropulse. The high possible bunch frequency of 1 MHz (or future 9 MHz) necessitates a reaction time in the range of microseconds.

This chapter provides a detailed description of the machine protection system of FLASH and of its subsystems. The MPS is a substantially enhanced version of the system used in the first phase of the TESLA Test Facility. Its development can be traced in [Sch02] and [Noe03]. First results from the commissioning for FLASH are found in [Fro06a].

6.1 Machine protection system overview

The MPS consists of multiple subsystems and has connections to several external devices as shown in Fig. 6.1. Each of these entities logically belongs to one of three categories represented as vertical layers in the drawing.

The top layer comprises *signal providers* such as beam loss monitors or magnet power supplies. Devices in this category send alarms or status information to the MPS logic subsystems. Only simple digital signals (e.g. on–off, OK–alarm) are transmitted.

The central part of the machine protection system, the *MPS logic* layer, is functionally divided into a slow and a fast subsystem. On the one hand, the Beam Interlock System (BIS) is fully programmable and handles complex logic tasks at relatively low speed. Its main task is to ensure safe operating conditions by monitoring the status of critical devices and by imposing limits on the beam power. On the other hand, the Beam Interlock Concentrators (BICs) have a simple hardwired logic capable of handling alarm conditions with sub-microsecond reaction

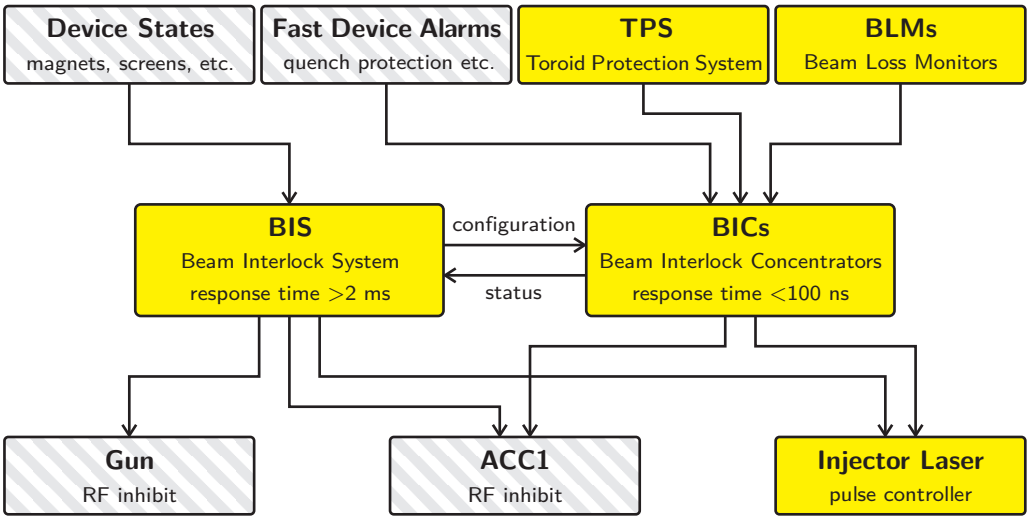

Figure 6.1: Overview of the FLASH machine protection system. Yellow (solid) boxes indicate MPS subsystems, gray (striped) boxes indicate external devices. From top to bottom, the shown systems are classified as: signal providers, MPS logic, actuators.

times. Their main purpose is the fast interruption of beam production within a macropulse.

The bottom layer shows the main *actuators*. This category comprises all points where the MPS logic may act on the operation of the machine, e.g. by preventing beam to be produced or transported. The main actuator for both the fast and the slow logic systems is the injector laser.

6.2 BIS—Beam Interlock System

The core part of the *Beam Interlock System* is a Siemens S7-300 programmable logic controller (PLC) with about 400 digital input channels. The cycle time time for the MPS program—and thus the typical response time—is about 1.2 ms. [Sta08]

The main function of the BIS is to ensure that all accelerator components are in a well-defined and appropriate state to allow safe beam transport. For this purpose, devices that are vital or possibly hazardous have a connection to the PLC. These connections carry on/off signals from magnet power supplies, open/closed signals from valves, in/out signals from screens, as well as a plethora of digital signals from over-temperature switches, water flow controllers, position switches for collimators and wirescanners, and many others. From this comprehensive overview of the state of the machine, the BIS derives a so-called operation mode.

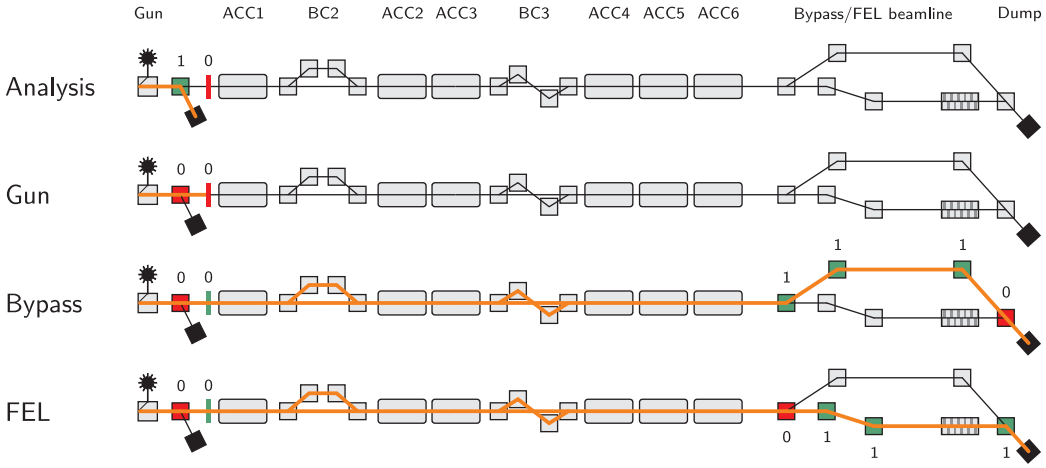

Figure 6.2: Operation modes for FLASH. The colored accelerator components have a required state for the respective operation mode: Red (0) – dipole off/valve closed. Green (1) – dipole on/valve open.

6.2.1 Operation mode

An *operation mode* is an abstraction of a valid beam path through the machine. The BIS prohibits beam production by disabling the gun rf and the injector laser unless the requirements for one of the predefined modes are fulfilled. There are four main operation modes, each one defined by a unique list of critical components with their valid states (Fig. 6.2):

Analysis mode The beam is deflected directly behind the gun cavity into a spectrometer beam line. Requirements: spectrometer dipole on, ACC1 entry valve closed.

Gun mode The beam is dumped on a valve before entering the first acceleration module. Requirements: spectrometer dipole off, ACC1 entry valve closed.

Bypass mode The beam is guided through all acceleration modules, deflected upwards into the bypass line, and finally downwards into the main dump. Requirements: all valves in main linac open, all valves in bypass and dump line open, spectrometer dipole off, linac-to-bypass dipoles on, bypass-to-dump dipole on, FEL-to-dump dipole off.

FEL mode The beam is guided through all acceleration modules, deflected through the collimation section, passes the FEL undulators, and is finally deflected into the main dump. Requirements: all valves in main linac open, all valves in FEL and dump line open, spectrometer dipole off, linac-to-bypass dipoles off, collimation section dipoles on, FEL-to-dump dipole on.

The FEL and bypass modes each require that a specified set of quadrupoles is switched on. Additional requirements exist, but are omitted here for the sake of simplicity.

Since only basic on/off signals are collected, a valid operation mode does not ensure correct beam transport. It merely guarantees that all devices are in working order. If, for instance, a critical quadrupole fails during FEL operation, the requirements for FEL mode are no longer fulfilled (and also not those for any other mode), and the BIS will stop beam production in the injector.

6.2.2 Beam mode

A superconducting linac like FLASH is routinely used with large variations on the beam current. Because the repetition rate and the bunch charge are more or less fixed parameters, the current is mainly defined by the number of bunches per macropulse. The *beam mode* is a concept designed to adjust the strictness of the active machine protection to the desired number of bunches. Three modes are defined:

Single pulse mode allows up to 2 bunches/macropulse. There is no fast reaction on beam losses within a bunch train.

Short pulse mode allows up to 30 bunches/macropulse. There is no fast reaction on beam losses within a bunch train.

Long pulse mode does not impose a limit on the number of bunches, but enables fast intra-train protection by the BIC network as explained in the following section.

The selection of beam modes is a two-way process: The operator can choose between short and long pulse mode based on the desired beam power, and the BIS can limit to single or short pulse mode based on the status of devices or relevant events. As an example, to protect OTR (optical transition radiation) screens from damage when they are inserted in the beamline, the BIS will limit the beam mode to *short* or *single* depending on the screen type. No beam is allowed while a screen is moving or in an otherwise undefined position.

If the machine is operated in single or short pulse mode, the fast intra-train protection provided by the BICs is disabled. This means that the machine protection system generally does not react on excessive beam losses—although exceptions to this rule exist, the most important being the undulator protection explained below. On purpose, the single and short pulse modes allow to dump the complete electron beam at an arbitrary location, as is often required for dedicated experiments or for calibration of diagnostics. In addition, they greatly simplify the initial setup of the machine.

The limit of 30 bunches is mainly a historical one, based on the following consideration: The total response time of the fast BIC network is of the order of 3 µs. When operating with long bunch trains at the design bunch frequency of 9 MHz,

this means that even after the detection of beam losses approximately 30 more electron bunches will be lost in the machine before the MPS can react. Therefore, all beamline components should anyhow be designed to withstand the temporary impact of this number of bunches.

On the other hand, at typical beam parameters of 1 nC charge, 5 Hz repetition rate, 1 GeV energy, 30 bunches carry a beam power of 150 W. At maximum, this figure could increase to 390 W (3 nC, 10 Hz, 1.3 GeV). This beam power usually does not pose a threat to thin beam pipes because only a small fraction is deposited in the material. However, the released amount of radiation is considerable and has often been observed to cause single event effects in computers located near the beamline. It can also interfere with personnel radiation safety directly, by triggering emergency shutdowns, or indirectly, because of huge amounts of induced radioactivity. Therefore, preparations have been made to impose an additional, stricter limit on the charge per macropulse in short pulse mode.

6.2.3 Undulator protection

As the $Nd_2Fe_{14}B$ permanent magnets of the FEL undulator are prone to demagnetization under comparatively low radiation doses (cf. chapter 5), they need active protection even in the single and short pulse modes. Therefore, the undulator is equipped with 38 beam loss monitors. Because of the dense placement of the BLMs, any electromagnetic shower is always observed on more than one monitor. A special BIS subroutine takes advantage of this by counting the number of simultaneous BLM alarms in the undulator section. If 2 alarms are pending over 120 s, 4 over 30 s, or 8 over 5 s, the injector laser is stopped and has to be restarted manually. In addition, module ACC1 is cut from rf power to suppress the transport of gun dark current.

To facilitate tuning, the tolerance of this routine can be increased for a time span of 30 minutes, limiting the beam to 2 bunches per macropulse. A similar routine protects the dump region from missteering of the beam and subsequent activation.

6.3 BICs—Beam Interlock Concentrators

As mentioned before, a network of BICs serves as fast intra-macropulse beam interlock logic. The function of a BIC module can be described as a fast logical OR concentrating 16 alarm inputs to two outputs within a processing time of less than 50 ns. A configurable mask allows to suppress any of these input channels. To allow the use of the BICs as a distributed system, all external connections use V.11 differential signaling ([ITU96], better known as RS-422) for maximum interference resistance. The logic functionality is implemented in an FPGA and a redundant hard-wired TTL network. A Profibus interface to the BIS allows configuration of the BIC masks and a remote readout of the inputs. [Wer08]

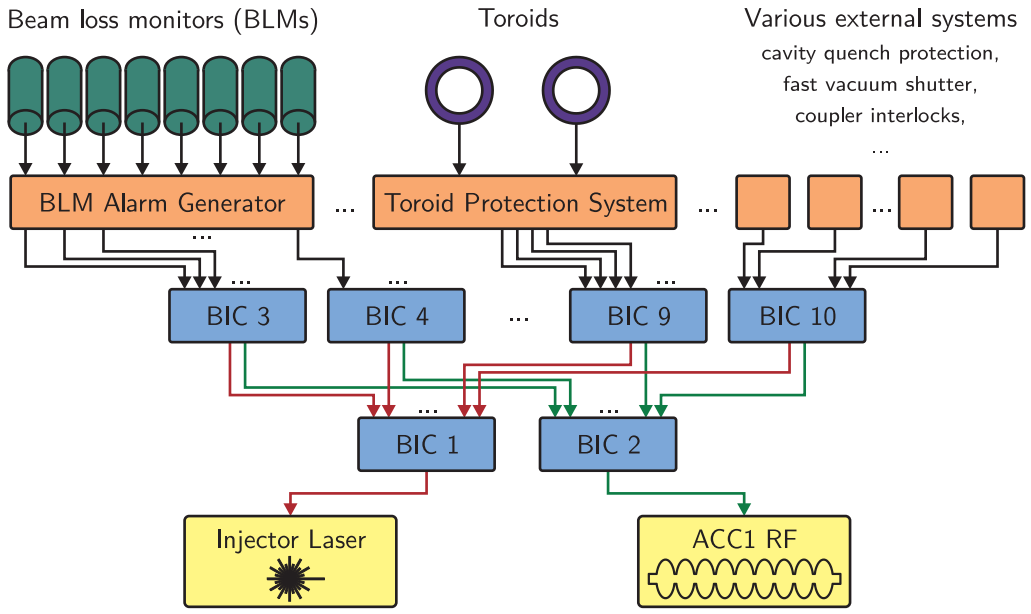

Figure 6.3: Organization of beam interlock concentrators (BICs). BIC 3–10 collect alarms from various sources, perform a logical OR, and relay the result to BIC 1–2 which switch off the injector laser and the rf power of acceleration module ACC1.

Ten BIC modules are in use at FLASH, connected in a tree-like fashion as illustrated in Fig. 6.3. The biggest part of the input channels is used by 83 BLM alarm signals, but also the toroid protection system, the cavity quench detection, a fast vacuum shutter, and rf coupler interlock signals are connected.

When any of the non-masked inputs signal an alarm status, the top-level BICs switch off the injector laser to suppress the production of new bunches. Additionally, acceleration module ACC1 is cut from rf power as a precaution against further transport of dark current from the rf gun. The BIC outputs stay in this alarm state for at least 2 ms. As a consequence of this, a beam loss alarm is guaranteed to suppress the whole remaining macropulse—whose maximum length is 800 µs—, but it does not interfere with the next macropulse that comes 200 ms later.

Measurements have confirmed that the total reaction time of the BIC network to a beam loss alarm on any input channel is below 4 µs. This time is dominated by cable delays. Therefore, up to three more bunches enter the machine after detection of the loss at the standard bunch frequency of 1 MHz.

6.4 BLMs—Beam Loss Monitors

FLASH is instrumented with more than 80 beam loss monitors for the detection of electromagnetic showers (Fig. 6.4). 17 of them are aluminum cathode electron

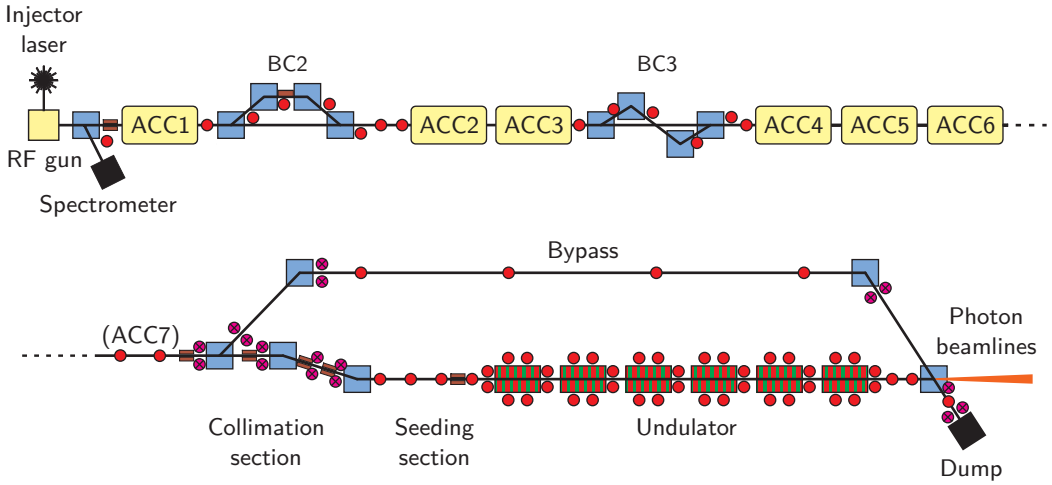

Figure 6.4: Overview of BLMs installed in FLASH. Red circles indicate photomultipliers with a mounted scintillator, purple (crossed) circles indicate aluminum cathode electron multipliers.

multipliers, located near collimators and main dipoles where high radiation backgrounds are expected. The remaining BLMs are photomultipliers with a plexiglass light guide and a piece of plastic scintillator as active medium. Detailed information on the geometry of the scintillation panels, the types of scintillator material, and on the used tubes can be found in appendix B.1.

All tubes are powered by a central high voltage power supply that allows to control the voltage for each BLM separately. The anode signals are transferred out of the accelerator tunnel to a central point via long cables. Due to the intrinsic high gain of the photomultiplier/electron multiplier tubes, preamplifiers in the tunnel are not necessary. The readout chain is based on a number of custom-built boards using the VME standard. Figure 6.5 illustrates the connections between the five main types of boards:

Main timing board: This is a standardized timing board for the FLASH/TTF facility. Its main function in the BLM system is to supply the clock distributor board with a 9 MHz and a 1 MHz beam-synchronous clock, with trigger signals for the beginning of the rf pulse and the start of the ADC sampling interval, and to ensure synchronization to the macropulse repetition rate.

BLM clock distributor: This board generates timing information for the alarm generators and transmits it via a flat cable. It also forwards the 1 MHz clock to a delay board.

Delay board: This is a standardized delay board that receives 1 MHz clock pulses on the input side and forwards them to 8 output channels with adjustable delay times. The resolution of the delay stages is 4 ns, the maximum delay is about 1 µs.

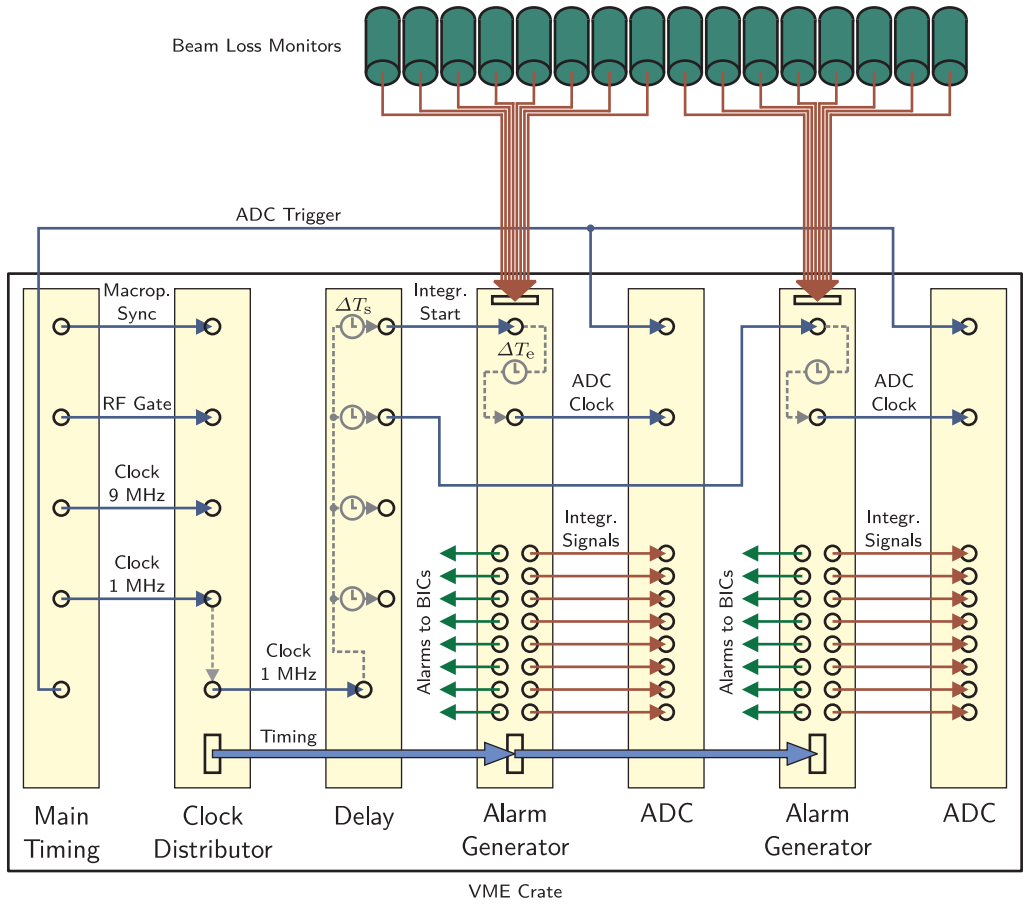

Figure 6.5: Simplified schematic of the FLASH BLM electronics. Connections: blue—timing, red—analog signal, green—digital alarm signal.

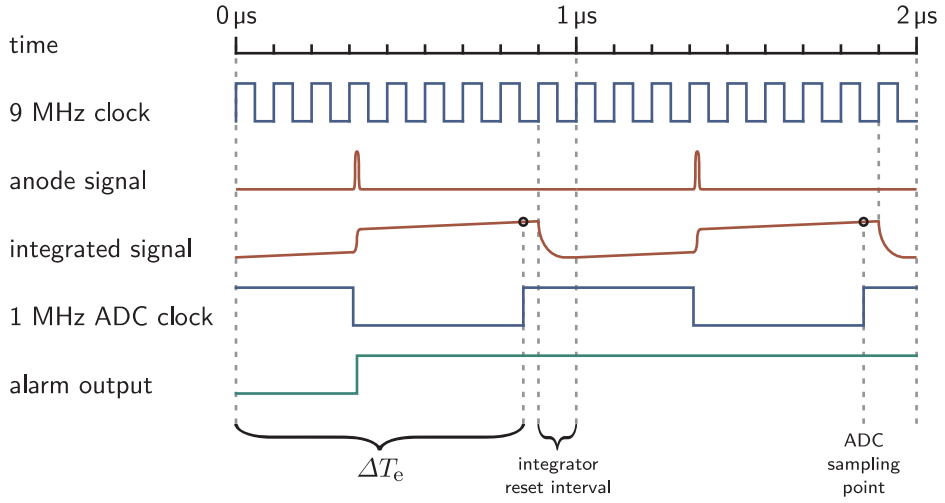

Figure 6.6: Timing diagram for FLASH BLM signals. The minimum pulse width for the alarm output is about 800 µs.

BLM alarm generator: The alarm generator is the central device for beam loss detection via photomultiplier/electron multiplier tubes. It performs an analog comparison of the PMT signals to various programmable thresholds. If one of the thresholds is exceeded or if an error condition is detected, it generates alarm signals on a number of output channels which are connected to the BICs. For monitoring purposes, the BLM signals are integrated over intervals of 1 µs—the standard bunch spacing of the accelerator—and forwarded to an external ADC. Each BLM alarm generator can handle eight input channels. Its functionality is described in more detail in the following sections.

ADC: This standardized 14 bit ADC board [And01] samples the BLM signals from the alarm generator with a clock frequency of 1 MHz in sampling intervals of 2048 µs length. These sampling intervals are started by a trigger from the main timing board with the repetition rate of the machine, typically 5 Hz. The trigger point is set to 700 µs before the first bunch of the bunch train.

Technical documentation on the BLM clock distributor and alarm generator boards can be found in [Goe04] and [Ste04b].

6.4.1 Signal processing

The BLM alarm generator board has identical signal processing stages for its eight input channels. The input stage limits the signal voltage to ±5 V and allows to add an offset voltage $V_{\text{offs\_in}}$ that is adjustable in the range of ±10 mV via a DAC. All other stages of the board use this pre-conditioned signal V_{in} or derivatives of it.

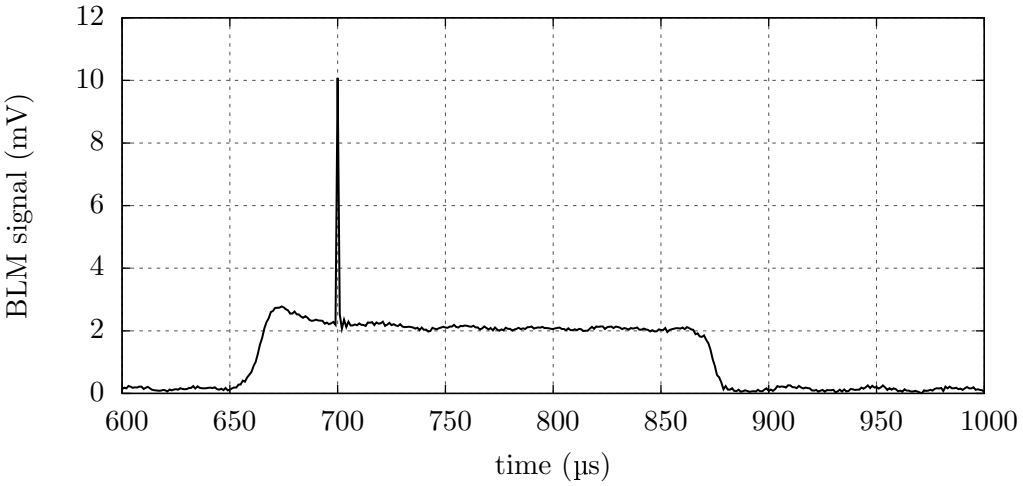

Figure 6.7: Typical signal of a BLM in bunch compressor BC2, measured with an ADC at the respective integrator output of the alarm generator. During the rf pulse of the gun, at about 650–880 µs, dark current losses are visible. In addition, there are losses by a single bunch at 700 µs.

The alarm generator is supplied with a 9 MHz clock by the FLASH timing system. To visualize the BLM signals, an integration stage supplies an output to a 1 MHz ADC. As shown in Fig. 6.6, V_{in} is integrated over eight clock cycles (889 ns) of the alarm generator, and the ADC clock is adjusted in a way that the sampling point is close to the end of this integration period. In the ninth clock cycle, the integrator is reset to zero.

The value shown by the ADC for a given BLM signal depends heavily on the timing. As shown in Fig. 6.5, the start time of the integration can be adjusted by the external delay ΔT_{s}. The additional delay ΔT_{e} generated on the alarm generator board defines the sampling point of the ADC. Procedures to set these delays correctly are introduced in section 7.2.

Figure 6.7 shows the typical signal from a BLM located in bunch compressor BC2 as measured by the respective ADC, averaged over 100 pulses. The signal socket visible in the range of about 650–880 µs is caused by losses of dark current from the rf gun. In addition, the partial loss of a single bunch at the start of the macropulse is registered; by convention, the ADC trigger point precedes this first bunch by 700 µs.

6.4.2 Alarms

Beam loss alarms are generated by analog comparison of DAC output voltages to the BLM signal. For each channel, there are three adjustable beam loss thresholds:

Single bunch: This is a simple threshold designed to react on excessive losses by a

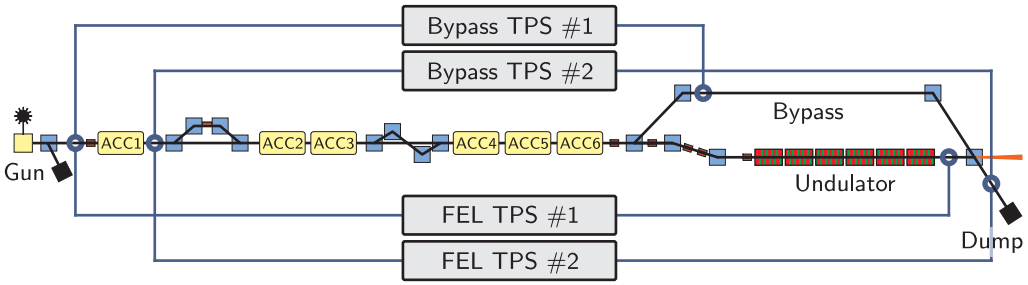

Figure 6.8: Overview of toroid pairs used for the FLASH TPS.

single bunch. It triggers when $V_{\text{in}} > V_{\text{thr\_single}}$.

Multi bunch: The multi bunch threshold allows to react on repeated losses within a macropulse. For every microsecond in which $V_{\text{in}} > V_{\text{thr\_multi}}$, a counter is incremented. When the counter exceeds a configurable threshold, the alarm is triggered. The counter is reset to zero at the beginning of each macropulse.

Integration: This is a threshold for the total losses within a macropulse. An integrator with time constant τ provides a voltage that is proportional to the accumulated input signal over the 2 ms long rf gate; the alarm triggers when $\frac{1}{\tau}\int V_{\text{in}}\mathrm{d}t > V_{\text{thr\_int}}$.

There are a number of additional alarm conditions that signalize errors in the beam loss monitoring system. The most important two are:

High voltage failure: At the BLM voltage divider, a small fraction of the high voltage is added to the signal voltage as a DC offset. The alarm generator constantly compares this offset to an adjustable threshold, and an alarm is triggered when the high voltage is too low.

Negative result: This alarm is raised if the output voltage of the ADC integrator drops below zero. This is usually an indication of faulty or swapped connections, or of electromagnetic interference.

For each channel, all alarms and error conditions are combined in a logical OR gate that controls the alarm output.

6.5 TPS—Toroid Protection System

The *Toroid Protection System* is a microprocessor- and FPGA-controlled system for differential current monitoring. Its purpose is to monitor the global transmission through the linac and to raise alarms in case of excessive beam loss. As illustrated in Fig. 6.8, four independent TPS units are in use at FLASH, two for bypass and two for FEL operation. Each unit is connected to a pair of toroids

and constantly monitors the charge loss between them. As additional inputs, the TPS units receive the 9 MHz machine timing and the *bunch gate,* a signal from the injector laser marking the desired bunch pattern in the current macropulse.

The charge comparison performed by the TPS is a relative one: When the linac is operating with reasonably low beam losses, a calibration can be invoked. In this calibration step, the system scans for the toroid signals and determines their respective timing and average amplitude. The average transmission found during this procedure is defined as 100 %. To be more exact, the ratio of the mean up- and downstream signal amplitudes, $k = \bar{V}_\mathrm{u}/\bar{V}_\mathrm{d}$, is stored as a correction factor for the downstream toroid. Subsequently, the absolute bunch charge is derived using the fixed calibration constant of 500 mV/nC:

$$Q_\mathrm{u} = V_\mathrm{u} \cdot \frac{1}{500\,\mathrm{mV/nC}}$$
$$Q_\mathrm{d} = V_\mathrm{d} \cdot \frac{k}{500\,\mathrm{mV/nC}}$$

Because of this relative calibration, a moderate difference in signal amplification or cable attenuation between the up- and downstream toroids can be ignored. However, the decision to regard a machine state as loss-free and hence suitable for calibration must be based on a different piece of instrumentation—in this case, on the beam loss monitors.

Each TPS unit is able to raise four types of alarms within a processing time of few 10 ns. For the description of the alarms, the charge on the upstream and downstream toroid will be denoted by $Q_{\mathrm{u},i}$ and $Q_{\mathrm{d},i}$, respectively. The index i indicates the number of the bunch within the macropulse, and n is the number of the last bunch whose signal has just reached the TPS:

Charge validation: The charge validation alarm is triggered when the charge measured at the upstream toroid is below 0.1 nC although the bunch gate indicates that a bunch is desired in this position of the macropulse. This is a safety measure against broken cables, timing errors, or a defective input stage of the TPS.

$$Q_{\mathrm{u},n} < 0.1\,\mathrm{nC}$$

Single bunch: The single bunch alarm is triggered when the transmission for one bunch is below 75 %:

$$Q_{\mathrm{d},n} < 0.75 Q_{\mathrm{u},n}$$

Slice: During the bunch train, the TPS constantly calculates the average transmission for the last 30 bunches. The alarm is triggered if the mean loss exceeds 10 %:

$$\sum_{i=n-29}^{n} Q_{\mathrm{d},i} < 0.9 \sum_{i=n-29}^{n} Q_{\mathrm{u},i}$$

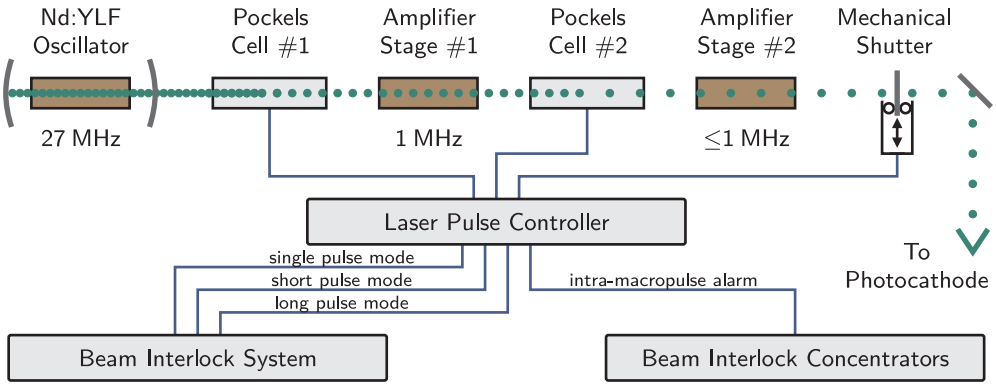

Figure 6.9: Schematic of the FLASH injector laser system.

Integration: The integration alarm triggers when 30 nC of charge loss are accumulated during the bunch train.

$$\sum_{i=1}^{n} (Q_{\mathrm{u},i} - Q_{\mathrm{d},i}) > 30\,\mathrm{nC}$$

All TPS alarms are connected to one of the BICs. Depending on the operation mode of the machine, either the FEL or the bypass pair are activated (unmasked) by the BIS. The use of two toroid pairs per mode guarantees a limited redundancy against the failure of one TPS unit or of one toroid. The fact that bypass TPS #2 and FEL TPS #2 are actually connected to the same two toroids is explained by a time-of-flight difference of few nanoseconds between both beam paths—if only a single TPS unit was used here, it would need to be recalibrated on each change between FEL and bypass mode to adapt to the changed signal timing.

Extensive information on the toroid protection system can be found in [Ham04], [Ham04a] and [Ham07].

6.6 Laser pulse controller

The *pulse controller* is the device that controls number and spacing of bunches in a macropulse by picking single laser pulses out of a pulse train. It is also the main actuator for beam inhibits issued by the BICs or by the BIS. The controller is designed as a VME board carrying a fully programmable FPGA. It has inputs for the beam mode defined by the BIS, for a BIC alarm signal, for the 9 MHz machine timing, and for a macropulse trigger. There are two control outputs for Pockels cell drivers, and one for a mechanical shutter. [Wer08]

Figure 6.9 shows a simplified schematic of the FLASH injector laser. A pulsed laser oscillator produces linearly polarized light pulses with a frequency of 27 MHz. The outcoupled light then enters the first Pockels cell which can turn the plane of polarization if a voltage is applied. Together with a polarizer, it acts as a fast

pulse picker forming a pulse train with a reduced frequency of 1 MHz. This train then traverses an amplification stage and enters the second cell which allows only pulses matching a user-configurable pattern. After final amplification, the laser light is guided onto the rf gun photocathode. This beam path can be blocked by inserting a mechanical shutter.

In the context of machine protection, the laser pulse controller fulfils the following basic functions:

- Enforce the limit on the number of bunches as given by the currently selected beam mode—i.e. let a maximum of 2 or 30 pulses pass in single or short mode.
- Close the laser shutter on request of the BIS. This may happen when there is no valid operation mode or when the slow undulator or dump protection routines are triggered.
- Block the Pockels cell based pulse pickers as long as the BIC input is in an alarm state.

7 Setup and operation of the FLASH MPS

The FLASH machine protection system is a fairly complex aggregate of multiple subsystems and a huge number of monitors. Great care must be taken to ensure that all components work as expected and that failures do not go unnoticed.

A number of procedures for calibration and setup of various system parameters have been developed for the beam loss monitoring system. They are described in the following sections; a few remarks on the operational experience with the system are given at the end of the chapter.

7.1 Cabling verification

All of the more than 80 BLMs are controlled and read out from a central location outside of the accelerator tunnel, close to the last acceleration module. Cables fan out from the high voltage power supply to distribution boxes in the accelerator tunnel, to which the BLMs are connected with patch cables. In a similar fashion, the BLM signal cables are pre-collected at patch panels in the tunnel. Long cables connect these to central patch panels at the back of the electronics racks which are in turn wired to the inputs of the alarm generators. From each alarm generator, eight short cables transfer the integrated signals to a neighboring ADC. Another eight cables send alarms to the BICs.

It is obvious that this kind of setup provokes swapped cables and other kinds of erroneous connections. This is even more problematic in the environment of an accelerator that is subject to frequent upgrades and redesigns that ask for an adjustment of the BLM coverage. As a further complication, the high voltage power supply, the alarm generator, and the ADC are each configured individually and have custom mappings of connector numbers to symbolic BLM names.

An automated procedure for testing and verification is indispensable to avoid false mappings between channels and to ensure basic functionality of the safety-critical BLM system. It consists of a passive and an active part.

Passive plausibility check

A central database contains manually updated information about the cabling. Amongst others, it includes the following fields:

> *BLM name* (“1GUN”), *signal cable name* (“25A2”), *HV cable name* (“25HV2”), *rack side patch panel connection* (“TOP C1”)

While the cable names are essential for practical work on the system, a simple algorithm can establish a one-to-one mapping from the patch panel connection to a triple of numbers identifying the channel in the control system:

alarm generator number, alarm generator channel, ADC number

A comparison of these numbers, deduced from the cabling database, to those stored in the control system for the respective BLM name identifies inconsistencies.

Active cable test

An inherent problem of most BLM designs is that a failure of the monitor can silently produce a signal level of zero—and thus give the false impression that no beam losses are present. The most likely failure mode for photomultipliers and electron multipliers is a high voltage drop out. This can be caused by a variety of mechanisms, e.g. failure of the power supply, cable breaks, or defects in the voltage divider circuit. Hence, the system design provides several ways of detecting this kind of failure, and the switching of the high voltage for individual BLMs can be used for an active cable test. The system design stipulates several independent indications of a missing tube voltage:

HV readback: The high voltage power supply provides a readback of the voltage on each of its channels.

Baseline drift: A tiny fraction of the tube high voltage is added to the signal voltage as a DC offset. Therefore, the signal baseline drops significantly when the high voltage is turned off. The ratio is roughly 5 mV/kV [Goe03].

HV alarm: The BLM alarm generator monitors the DC offset on the signal lines. A low readback triggers a “no HV” alarm.

BIC alarm: The “no HV” alarm from the alarm generator must show up on the BIC channel assigned to the respective BLM.

Since each of these indicators comes from a different device (power supply, alarm generator, BIC, ADC), a test checking all of them is a solid validation of the whole beam loss monitoring chain. It reliably points out swapped and faulty cables as well as many kinds of electronic malfunctions.

Because the algorithm for this test is quite obvious, only the criterion for detection of the baseline drift will be mentioned. With the high voltage switched on, an ADC trace a_i consisting of 2048 samples is read out (cf. Fig. 6.7). The first $N = 500$ samples are well before the start of the rf pulse and therefore free of beam loss signals. They are used to calculate the average signal baseline and its

Value	Symbol
$x_{\text{offs\_in}}$	Input offset DAC setting
$V_{\text{offs\_in}}$	Input offset voltage
$x_{\text{offs\_adc}}$	ADC offset DAC setting
$V_{\text{offs\_adc}}$	ADC offset voltage
τ_{single}	Single threshold DAC setting
τ_{multi}	Multi threshold DAC setting
τ_{int}	Integration threshold DAC setting
y	ADC sample (counts)
U	ADC sample (voltage)

Table 7.1: List of quantities and symbols used in the discussion of setup and calibration of the BLM system.

standard error:

$$\bar{a} = \frac{1}{N}\sum_{i=1}^{N} a_i$$

$$\sigma_{\bar{a}} = \sqrt{\frac{1}{N(N-1)}\sum_{i=1}^{N}(a_i - \bar{a})^2}$$

After switching off the high voltage, the baseline is measured a second time, resulting in $\bar{b} \pm \sigma_{\bar{b}}$. A drop of the baseline by 10σ is considered significant. In detail, the check *fails* if

$$\bar{b} > \bar{a} - 10\sigma_{\bar{a}} \qquad \vee \qquad \bar{a} < \bar{b} + 10\sigma_{\bar{b}}.$$

7.2 Electrical setup of the BLM system

The operation of the BLM system is influenced by a number of user-configurable parameters. Where possible, algorithms have been developed to facilitate a reproducible and correct setup of the system. After the discussion of these algorithms, a procedure for the calibration of the alarm threshold values will be explained. The symbols used are summarized in Tab. 7.1.

7.2.1 Input offset

As stated before, the input offset $V_{\text{offs\_in}}$ is a voltage that is added to the BLM signal in the input stage of the alarm generator, resulting in the voltage V_{in} that all other stages refer to. A change of the input offset therefore affects both the level of the alarm thresholds relative to the signal baseline and the ADC output. $V_{\text{offs\_in}}$ is generated by an 8-bit DAC like the threshold voltages.

Because of its comparatively long integration time span of about 2 ms, the integration alarm stage is most sensitive to changes of the input offset. The integration alarm is therefore used in the setup procedure for the input offset. In condensed form, an algorithm ensures that $x_{\text{offs\_in}}$ *is adjusted to the lowest possible value that*

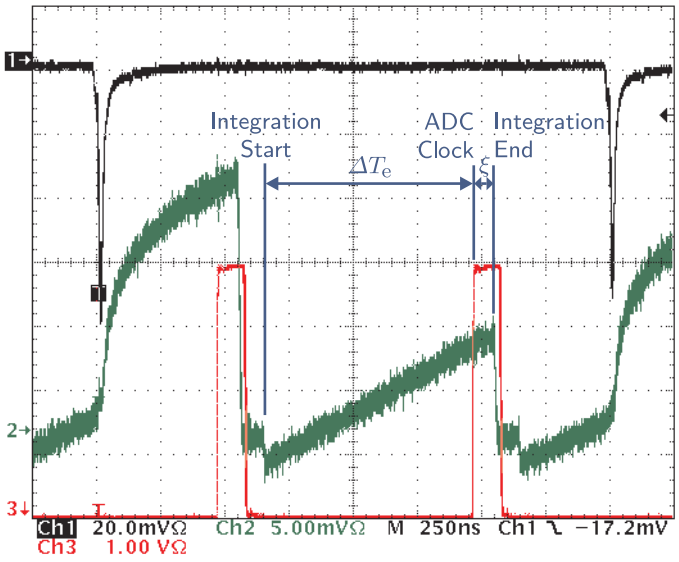

Figure 7.1: Oscilloscope trace of in- and outputs of the ADC integration stage of the BLM alarm generator. Black (Ch1)—photomultiplier signal V_{in}; green (Ch2)—Integrator output; red (Ch4)—ADC clock signal.

triggers an integration alarm while the threshold is set to zero, $\tau_{\mathrm{int}} = 0$. Consequently, the alarm vanishes when the threshold is raised to one or higher.

As additional constraints, the BLM must be connected and powered (because of the dependence of signal baseline on high voltage), and must be free of beam loss signals that would enter into the integration result. In practice, this means that the accelerator should be off during the setup of the input offset.

7.2.2 ADC sampling point

The values sampled by the ADC depend on four parameters of the alarm generator: the input offset voltage $V_{\mathrm{offs\_in}}$, an additional offset voltage for the ADC output $V_{\mathrm{offs\_adc}}$, and the timings of the integration cycle start ΔT_{s} and of the sampling point ΔT_{e}, which are illustrated in Figs. 6.5 and 6.6.

An example of actual beam loss signals is shown in Fig. 7.1. The picture represents an oscilloscope trace of three signals from the alarm generator board: In black (Ch1), the output of a photomultiplier tube (V_{in}), in green (Ch2), the output of the ADC integration stage, and in red (Ch3), the ADC clock. The horizontal axis spans a time of 2.5 µs, which corresponds to 2.5 cycles of the ADC integrator.

Because the machine was operated at a bunch frequency of 500 kHz at the time of measurement, the distance of the two visible peaks in the PMT signal is 2 µs. The peak voltage is about 70 mV which corresponds to a current of 0.7 mA into the input resistance of 100 Ω. The output of the integration stage always has a positive slope between the short reset intervals. While the integrator reaches

23 mV with beam loss, it end up at 10 mV even in a completely loss-free cycle. This is a cumulative effect of the input offset and the ADC offset which also enter the integration.

The rising edge of the ADC clock approximately marks the sampling point. To obtain the best signal-to-noise ratio, it should be located near the end of the integration period. However, moving it too close to this end increases the risk of sampling in the reset interval due to timing jitter.

Delay calibration

The sampling point is mainly defined by the delay ΔT_e relative to the start of the integration cycle. It is set for all channels of an alarm generator board by an 8-bit digital delay. To obtain a calibration of the delay component, the interval ξ between the rising edge of the ADC clock and the reset of the integrator has been measured with an oscilloscope for multiple delay settings. Table 7.2 and Fig. 7.2 show the measured points and the result of a linear regression of the data; it is found that ξ is related to the delay setting x by

$$\xi = (27 + 0.552x)\,\mathrm{ns}.$$

Because the range of the delay is limited to $0 \leqslant x \leqslant 255$, it is guaranteed that the rising slope of the ADC clock occurs at least 27 ns before the integrator reset. However, an additional hidden delay of about 45 ns [And01] is added in the ADC between the clock signal and the actual sampling point. This can still cause the sampling to fall into the reset period.

Timing scan

A simple timing scan ensures the correct sampling of the integrator signal. The alarm generator's delay setting x is ramped from 0 to 255, and at each step the baseline of a BLM signal—the average of the first 500 samples—is read out from the connected ADC. Because this moves the sampling point over the repetitive integrator signal, the result is a fine scan of the signal shape (Fig. 7.3). The sampling point falls into the reset period for $x \leqslant 40$. The digital delay becomes unstable at settings above 190 because the difference to the zero position approaches 111 ns, one period of the 9 MHz clock.

With some margin, safe settings for the alarm generator delay are therefore found in the range of about 50–175. A setting of 50 yields the highest signal-to-noise ratio.

7.2.3 Integrator timing

When the sampling point has been set up correctly using the ADC clock delay, it must still be ensured that the beam loss signals themselves do not fall into the reset period of the integrator. For this, the integration start time ΔT_s can be varied for

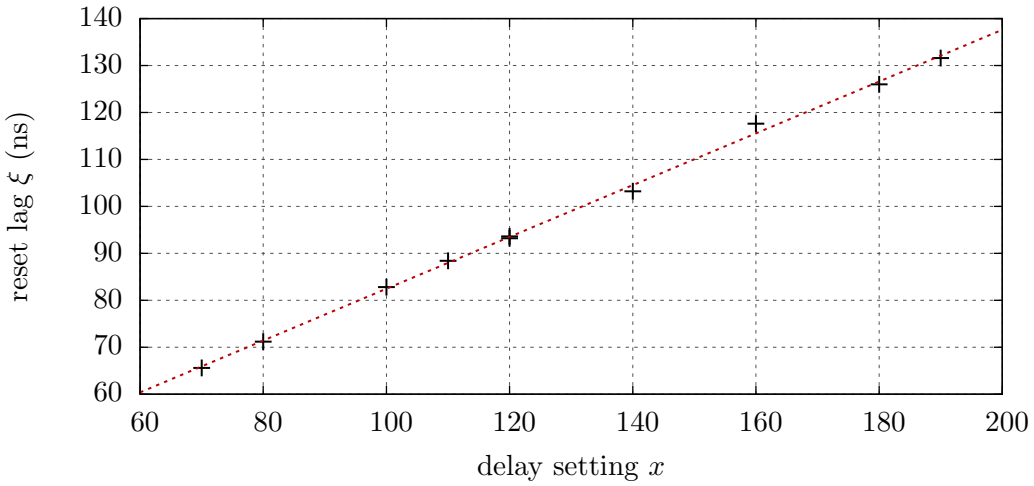

Figure 7.2: Calibration measurement for the internal delay of the BLM alarm generator. For various settings of the delay, the time ξ between the ADC sampling point and the reset of the integrator was measured with an oscilloscope.

delay steps x	**reset lag** ξ (ns)
70	65.6
80	71.2
100	82.8
110	88.4
120	93.2
120	93.6
140	103.2
160	117.6
180	126.0
190	131.6
Regression: $\xi = a + bx$	
$a = (27.3 \pm 1.0)$ ns	
$b = (0.5516 \pm 0.0078)$ ns	

Table 7.2: Calibration measurement for the internal delay of the BLM alarm generator. For various settings of the delay, the time ξ between the ADC sampling point and the reset of the integrator was measured with an oscilloscope. The parameters at the bottom of the table are the result of a linear regression of the data.

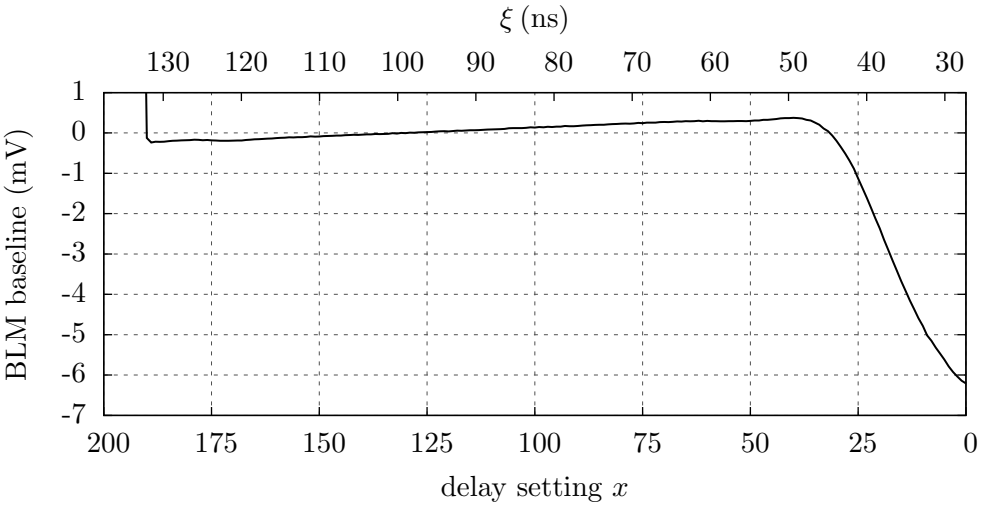

Figure 7.3: ADC clock timing scan. The plot shows the BLM signal baseline read by an ADC in dependence of the alarm generator's digital delay setting—the result is a fine scan of the integrator output signal. The integrator reset is visible at $x \leqslant 45$, for $x \geqslant 190$ the digital delay does not work correctly (note that the horizontal scales are reversed for clarity).

each BLM alarm generator with an external delay board. The desirable situation corresponds to the one shown in Fig. 7.1 where the photomultiplier signals arrive just in the middle of the integration cycle.

Because the integrator timing is the same for all eight BLM channels on a board, care must be taken to use similar cable lengths for all of the connected monitors. To have some safety margin against single-period jumps of the 9 MHz clock signal, it is reasonable to demand that all eight signals arrive within a time span of about 600 ns. At a signal speed of $0.9c$, this corresponds roughly to a maximum difference in cable length of 160 m. For the existing cabling, this is not negligible, and an appropriate grouping of channels has to be observed.

Timing scan

The correct setup of the integration start time is verified by a scan of the external delay ΔT_s over the full range of 1 µs. While constant losses by a single bunch are induced, the ADC is read out for a configurable number N of macropulses per delay step. These N ADC traces are averaged, resulting in a signal shape y_i comparable to Fig. 6.7. During the timing scan, it is not certain that the signal from the bunch occurs exactly at the designated position of $i = 700$, i.e. 700 µs after the ADC trigger. Because of this, the bunch loss signal b is calculated as the

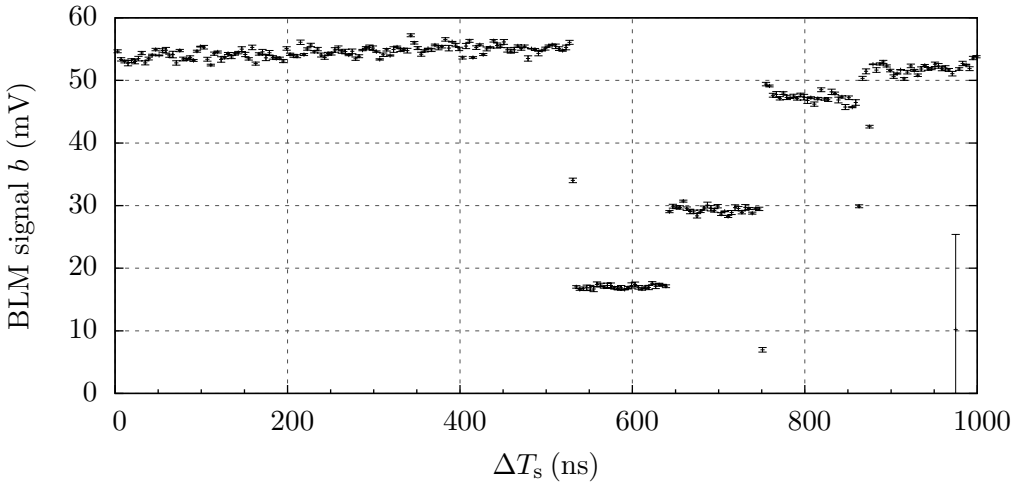

Figure 7.4: ADC integrator timing scan. The plot shows the signal sampled by the ADC for a roughly constant beam loss as a function of the integrator delay ΔT_s. Between 530 and 640 ns, the photomultiplier pulse falls into the reset period of the integration stage and is not properly processed.

maximum of three samples, and corrected by the baseline of the signal:

$$b = \max\left(y_{699}, y_{700}, y_{701}\right) - \frac{1}{500}\sum_{i=1}^{500} y_i$$

The exemplary scan result $b(\Delta T_\mathrm{s})$ shown in Fig. 7.4 exhibits a number of steps of 111 ns width. This illustrates that the alarm generator fixes the start of the integration to the next edge of its 9 MHz clock—in other words, although the external delay ΔT_s is varied in fine steps of few nanoseconds, the actual start time only changes in full 9 MHz periods.

The plot shows clearly that parts of the photomultiplier pulse fall into the integrator reset period for delay settings of 530–640, or 640–750 ns. If the reset interval is moved to 750–850 ns, it still interferes with the long tail of the PMT pulse. It is worth noting that some of the signal amplitude is caused by dark current losses that—unlike bunch losses—occur with a frequency of 1.3 GHz and therefore always contribute to the integration. For completeness, it should also be remarked that delay settings close to the border between two clock periods can cause unexpected or unstable behavior of the integrator (e.g. at $\Delta T_\mathrm{s} = 751$ or 975).

Because of unequal cable lengths, it is recommended to perform this scan separately for each of the BLMs connected to an alarm generator board. Afterwards, the delay ΔT_s must be chosen in a way that guarantees the maximum distance to the reset interval for all of these BLM channels. The need for induced beam losses

prevents a complete automatization of the procedure; in the future, the setup could be simplified by a fully automated BLM test pulse system.

7.2.4 ADC offset

The last part of the electrical setup for the BLM system is concerned with the offset voltage $V_{\text{offs\_adc}}$, a per-channel setting of the BLM alarm generator. Its main purpose is to adapt the integrator output level to the dynamic range of the connected ADC. More precisely, it is used to move the baseline of all BLMs to the same number of ADC counts.

The standard ADCs used in the BLM system have a resolution of 14 bits; their digital output y ranges from 0 to $16383 = 2^{14} - 1$, where 0 counts correspond to $+1\,\text{V}$ and 16383 counts to $-1\,\text{V}$. The sampled voltage U is therefore converted into ADC counts y as

$$y = 8192\,\text{V}^{-1} \cdot (1\,\text{V} - U) \quad \Leftrightarrow \quad U = 1\,\text{V} - \frac{y}{8192\,\text{V}^{-1}}. \tag{7.1}$$

Since the integration stage always produces a positive signal level if the input offset is set up correctly, it is found that 7800 counts—roughly $+50\,\text{mV}$—is a reasonable target count that each BLM should reach when it detects no beam loss.

In a simple algorithm, the setting of the 8-bit DAC generating the offset, $x_{\text{offs\_adc}}$, is increased stepwise until the mean counts of the corresponding ADC channel fall just below 7800. Only the first 500 samples of the ADC trace are considered to avoid influence by beam loss signals. During the procedure, the BLM must also be powered to account for the baseline shift caused by high voltage variations.

7.2.5 Cross-calibration of offsets and thresholds

The adjustment of the ADC offsets concludes the electrical setup of the BLM system. However, the two offsets $V_{\text{offs\_in}}$ and $V_{\text{offs\_adc}}$ and—more importantly—the alarm thresholds are all set by DAC values that have not yet been related to real voltages. A particular question of high practical relevance is "to which signal U on the ADC does a threshold of τ correspond?" A cross-calibration algorithm has been developed is an attempt to answer this.

The whole calibration requires the absence of beam loss signals—ideally, the accelerator should be switched off. To obtain meaningful results, the ADC timings must have been set up according to the instructions given before.

The first part of the calibration is based on a scan of the ADC offset $x_{\text{offs\_adc}}$ across the complete range of 0 to 255 while the input offset is set to zero, $x_{\text{offs\_in}} = 0$. For each scan step, an ADC trace is read out and the mean number of ADC counts y is calculated. The result is a linear relation

$$y(x_{\text{offs\_in}} = 0, x_{\text{offs\_adc}}) = y_0 + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_adc}}} x_{\text{offs\_adc}}. \tag{7.2}$$

The same scan is repeated for the input offset, this time with $x_{\text{offs\_adc}} = 0$. Before the start of this scan, however, the single bunch, multi bunch, and integration thresholds are set to zero, and the counter limit for the multi bunch alarm is reduced to one, making it basically another variant of the single bunch alarm. Again, the result is a linear dependence

$$y(x_{\text{offs\_in}}, x_{\text{offs\_adc}} = 0) = y_0 + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_in}}} x_{\text{offs\_in}}. \tag{7.3}$$

The topmost plot in Fig. 7.5 shows the primary results of these two scans for an exemplary BLM channel.

During the scan of $x_{\text{offs\_in}}$, it is also registered when one of the alarms is triggered. After such an event, the respective threshold is increased until the alarm vanishes, and the scan continues. This is illustrated in the bottom plot of Fig. 7.5, which shows that the integration alarm is triggered first at an input offset of 161 for a threshold of zero—remembering the setup procedure for the input offset from section 7.2.1, this is exactly the operating point for the BLM channel. For a threshold setting of $\tau_{\text{int}} = 119$ or higher, the integration alarm cannot be triggered by the input offset anymore.

In the same way as for the integration alarm, the single and multi bunch alarms are registered. Because they are far less sensitive to the input offset, only few data points can be collected this way. Usually, the alarms can only be triggered for thresholds of 0, 1, or 2. In order to obtain better statistics, the scan is therefore repeated 20 times in an interval around the offset where the alarm was first detected. As the central plot in Fig. 7.5 shows, the variation in $x_{\text{offs\_in}}$ is relatively big because any noise spike on the BLM channel can trigger the single or multi alarm. Hence, a considerable error has to be expected when extrapolation from these points to high threshold values is attempted.

In summary, measurements for the following relations have now been obtained:

$$\begin{aligned}
\tau_{\text{single}}(x_{\text{offs_in}}) &= \tau^0_{\text{single}} + \frac{\mathrm{d}\tau_{\text{single}}}{\mathrm{d}x_{\text{offs_in}}} x_{\text{offs_in}} \\
\tau_{\text{multi}}(x_{\text{offs_in}}) &= \tau^0_{\text{multi}} + \frac{\mathrm{d}\tau_{\text{multi}}}{\mathrm{d}x_{\text{offs_in}}} x_{\text{offs_in}} \\
\tau_{\text{int}}(x_{\text{offs_in}}) &= \tau^0_{\text{int}} + \frac{\mathrm{d}\tau_{\text{int}}}{\mathrm{d}x_{\text{offs_in}}} x_{\text{offs_in}} \\
y(x_{\text{offs_adc}}, x_{\text{offs_in}}) &= y_0 + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs_adc}}} x_{\text{offs_adc}} + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs_in}}} x_{\text{offs_in}}
\end{aligned}$$

The last equation is a contraction of (7.2) and (7.3). The unknown quantities can be determined by a linear regression of the data points; a comprehensive compilation of calibration results and further details is found in section B.1.6.

It is now possible to calculate the ADC count that corresponds to a given threshold setting τ_*, where $*$ stands for the threshold under consideration:

$$y_* = y_0 + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_adc}}} x_{\text{offs\_adc}} + \frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_in}}} \cdot \frac{\mathrm{d}x_{\text{offs\_in}}}{\mathrm{d}\tau_*} \cdot \left(\tau_* - \tau^0_*\right) \tag{7.4}$$

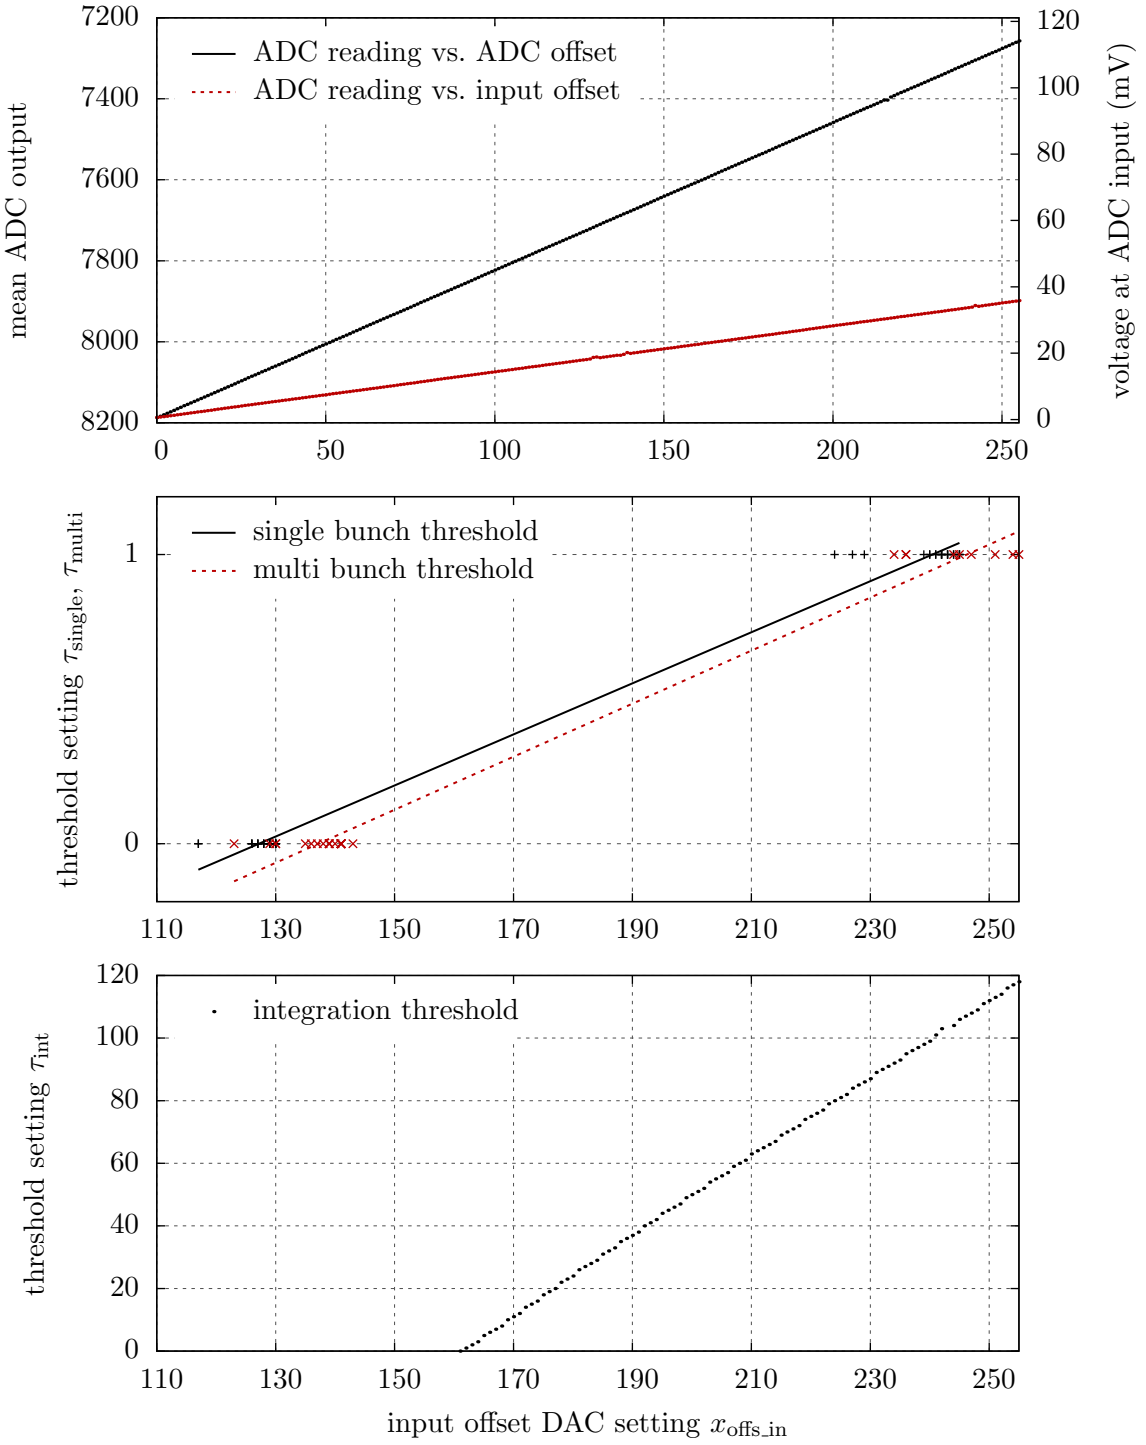

Figure 7.5: Electrical calibration data of a BLM channel. In the calibration procedure, the input offset is scanned through its complete range, and the response of the ADC and of the beam loss alarms is registered.

These ADC counts can be readily converted into voltages using (7.1). Applying error propagation and introducing the notation $\sigma[X]$ for the standard error of quantity X yields an uncertainty of

$$\sigma\left[y_*\right] = \Bigg\{ \sigma^2\left[y_0\right] + x_{\text{offs\_adc}}^2 \sigma^2\left[\frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_adc}}}\right] + \left(\tau_* - \tau_*^0\right)^2 \left(\frac{\mathrm{d}x_{\text{offs\_in}}}{\mathrm{d}\tau_*}\right)^2 \cdot \left(\sigma^2\left[\frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_in}}}\right] + \left(\frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_in}}}\frac{\mathrm{d}x_{\text{offs\_in}}}{\mathrm{d}\tau_*}\right)^2 \sigma^2\left[\frac{\mathrm{d}\tau_*}{\mathrm{d}x_{\text{offs\_in}}}\right]\right) + \left(\frac{\mathrm{d}y}{\mathrm{d}x_{\text{offs\_in}}}\frac{\mathrm{d}x_{\text{offs\_in}}}{\mathrm{d}\tau_*}\right)^2 \left(\sigma^2\left[\tau_*\right] + \sigma^2\left[\tau_*^0\right]\right) \Bigg\}^{\frac{1}{2}}.$$

Discussion

With equation (7.4), a means of relating an ADC reading to a certain threshold setting has been found. However, this result is in need of interpretation.

As set out before, the *integration threshold* pertains to the accumulated losses during the rf gate, which falls entirely into the 2048 µs long sampling interval of the ADC. From the timing setup, the start and end of the rf gate relative to the ADC trigger can be determined. If the gate thus spans from sample m to n, the condition for an alarm to be triggered can be formulated in terms of ADC counts y_i:

$$\frac{1}{n-m+1}\sum_{i=m}^{n} y_i < y_{\text{int}}(\tau_{\text{int}}, x_{\text{offs\_adc}}, x_{\text{offs\_in}})$$

If this inequality is expressed in voltages instead of ADC counts, the *less-than* changes into a *more-than* condition, which might be more intuitive. Two remarks have to be added: First, the electrical characteristics of the rf gate and ADC integrators are different—for instance, both constitute low passes with different cut-off frequencies. Second, the ADC integrator only covers about 8/9 of the total time because of the periodic reset. Since the calibration is done with a DC voltage, pulsed signals will appear larger by a factor of 9/8 on the ADC display relative to their real influence on the alarm integrator.

The *single* and *multi bunch* alarms do not refer to an integrated signal at all. Instead, these alarm thresholds pertain directly to the input signal—they are triggered by the amplitude of a pulse, not by its integrated charge. As the calibration procedure uses a DC voltage, (7.4) can be applied only for DC-like signals. This is the case for dark current losses because they appear with the rf frequency of 1.3 GHz which is far above the bandwidth of the photomultipliers and of the alarm generator electronics.

Losses of electron bunches, on the other hand, cause short current pulses with comparatively high amplitudes. The pulse width depends on a number of factors—cable length, scintillator geometry, saturation effects in photomultiplier tubes and

electronics, longitudinal distribution of beam losses. Hence it is not possible to establish a fixed ratio between the amplitude of a pulse and its integrated charge. In short, for pulsed losses there is no universal way to tell from the ADC reading whether the single or multi bunch alarms would be triggered.

In spite of these limitations, the cross-calibration procedure yields valuable information about the status of the individual BLM channels. It has allowed to identify several setup and hardware faults.

7.3 BLM response to beam losses

It has already been stated that the signal from a photomultiplier depends linearly on the intensity of the incident light—within a confined parameter space. The most important parameters in this context are the high voltage applied to the tube and the intensity of light reaching the photocathode. In addition, the final output signal is influenced by the signal processing electronics. It is therefore necessary to characterize the response of the whole readout chain—BLM, alarm generator, ADC—to beam losses.

7.3.1 High voltage dependence

At first, the dependence between signal and high voltage will be examined. An exemplary photomultiplier showing a special beam loss pattern is selected for the following discussion. The ADC traces in the top plot of Fig. 7.6 show strong losses by a single bunch at sample 700, and a background of dark current losses of varying magnitude between samples 670 and 1030. Because this BLM mainly registers the electromagnetic shower emanating from the energy collimators, the signal variation can be attributed to an inhomogeneous energy distribution along the pulse, which is in turn caused by non-uniformities of the rf regulation. Because the fluctuation of the loss pattern from pulse to pulse is small, it can be considered a static distribution of light intensities.

The high voltage of the photomultiplier is increased from 0 to 1500 V in steps of 25 V. At each step, the average of 20 successive ADC traces (20 pulses) is calculated. The upper plot in Fig. 7.6 shows exemplary traces for a number of high voltage settings. Based on this data, a number of representative sample positions are selected. Then, the BLM signal U at these positions is plotted against the high voltage setting S. Because the loss pattern is static, several curves for $U(S)$ representing individual light intensities have thus been obtained. To extend the range of intensities, a second measurement with the same BLM, but strong beam loss signals from two individual bunches has been taken. Following the considerations of section 3.2.5 on the photomultiplier gain, an approximate power law dependence is expected for $U(S)$. Hence, a double logarithmic scale is chosen for the plot (Fig. 7.6, bottom).

In fact, for small light intensities the curves do represent straight lines; it is

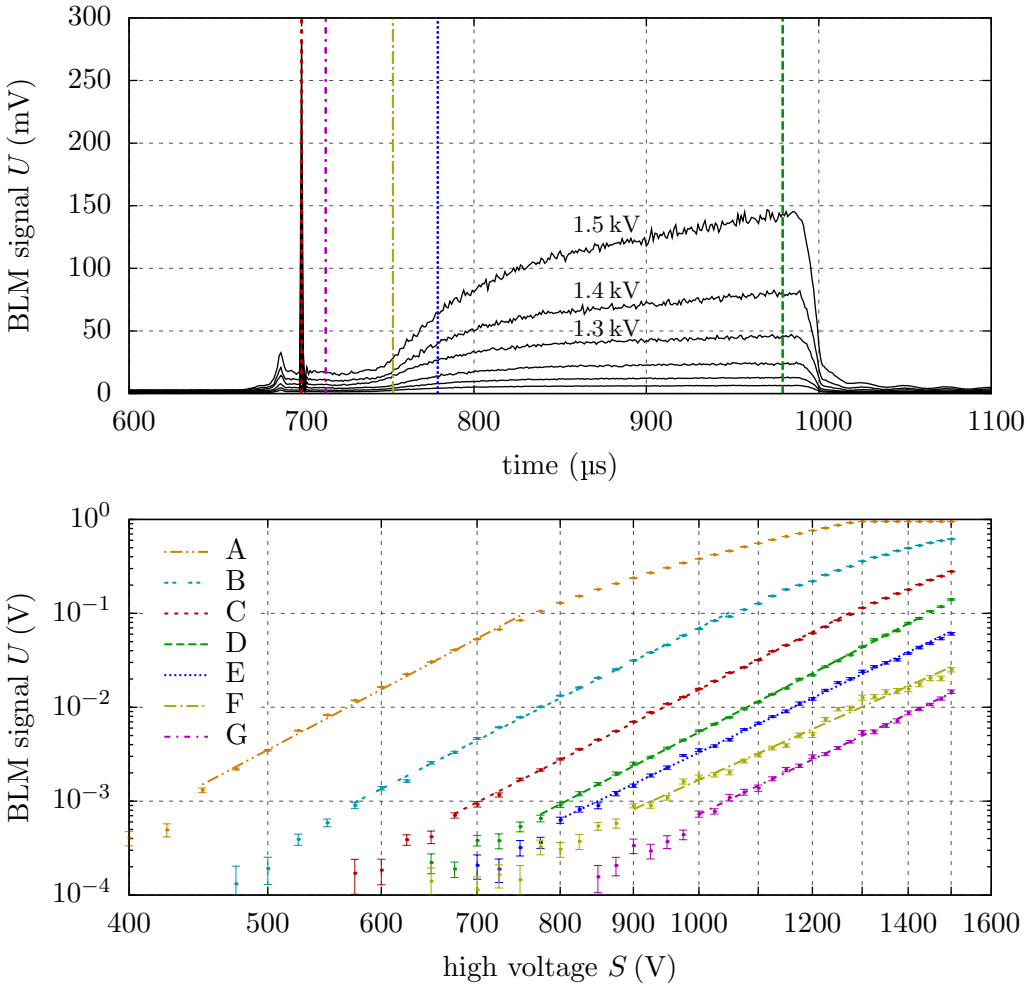

Figure 7.6: Photomultiplier response for various light intensities and high voltage settings. *Top:* ADC traces of stationary losses for high voltage settings of 1.0, 1.1, 1.2, 1.3, 1.4, and 1.5 kV. The vertical lines indicate samples with different signal strenghts that are used in the second plot. *Bottom:* BLM signal as shown by the ADC for various light intensities and high voltage settings. Curves C–G correspond to the samples marked in the upper plot, the lowest four of them represent dark current losses. Data series A and B have been obtained in a separate measurement. Data points represent the mean value of 20 successive pulses, error bars indicate the statistical error of the mean. The plot also shows power law fits to the data in selected ranges.

	coefficient $a/10^{-26}$	**exponent** b	**goodness-of-fit** χ^2/dof	
			fit range	full range
A	45 ± 38	8.11 ± 0.13	180	$6.5{\cdot}10^6$
B	46 ± 19	7.73 ± 0.06	6.1	$2.2{\cdot}10^4$
C	9 ± 2.3	7.75 ± 0.04	3.6	110
D	7.1 ± 4.3	7.34 ± 0.09	0.4	2.4
E	550 ± 570	6.83 ± 0.15	2.1	3.4
F	37 ± 13	7.32 ± 0.05	1.7	3.5
G	1.01 ± 0.25	7.91 ± 0.03	0.9	2.6

Table 7.3: Fit parameters for the photomultiplier high voltage scan according to (7.5).

therefore justified to describe them by the power law

$$U(S) = a \left(\frac{S}{1\,\mathrm{V}} \right)^b \cdot 1\,\mathrm{V}. \tag{7.5}$$

This is especially valid for all points corresponding to dark current losses (data series D–E). At higher intensities, deviations become apparent, especially at signal amplitudes of $U > 0.1\,\mathrm{V}$. To obtain estimates for a and b, (7.5) has been fitted to a selected range of values from the the data series. The results are summarized in Tab. 7.3 and displayed as line segments in Fig. 7.6. While the value of coefficient a is not of particular interest, knowledge of the exponent b allows to scale BLM signals from one high voltage setting to another. The weighted average of the exponents is

$$\bar{b} = 7.64 \pm 0.32. \tag{7.6}$$

To quantify the deviation from linearity, the reduced chi-square statistic χ^2/dof of the fit is listed in the table. If N designates the number of points in the data series and $\Phi = 2$ is the number of fit parameters, the problem has $\mathrm{dof} = N - \Phi - 1$ degrees of freedom. Therefore, the reduced chi-square figure becomes

$$\frac{\chi^2}{\mathrm{dof}} = \frac{\chi^2}{N - \Phi - 1} = \frac{1}{N - \Phi - 1} \sum_{i=1}^{N} \frac{(U_i - U(S_i))^2}{\sigma^2[U_i]}. \tag{7.7}$$

U_i and $\sigma[U_i]$ denote the measured signal and its uncertainty, S_i is the corresponding high voltage setting. Obviously, smaller values of χ^2 indicate better agreement of the data with the power law (7.5). The table lists χ^2/dof for the fit range indicated in Fig. 7.6 and for all recorded data points except those that evidently exceed the range of the ADC.

Discussion

It has been shown that major deviations from the expected power law dependence of $U(S)$ occur for high beam loss intensities, and are especially pronounced for

values of $U > 0.1\,\mathrm{V}$. It is useful to add that this behavior is observed for photomultiplier and aluminum cathode electron multiplier tubes alike. With a few calculations, it can be pointed out that the alarm generator electronics is the main reason for this kind of signal saturation:

As set out before, the input stage of the alarm generator limits the voltage V_{in} of any incoming signal to about $\pm 5\,\mathrm{V}$. The documentation of the ADC integrator [Goe03] specifies an output of

$$U = 0.5\,\frac{\mathrm{mV}}{\mathrm{Vns}} \int V_{\mathrm{in}} \mathrm{d}t. \tag{7.8}$$

Thus, an ADC reading of $U = 0.1\,\mathrm{V}$ corresponds to an integral of $\int V_{\mathrm{in}} \mathrm{d}t = 200\,\mathrm{Vns}$. Assuming a typical pulse width of 30 ns, an amplitude of about 6–7 V can be estimated. The top of the pulse would therefore be subject to *clipping* (i.e. truncation at the 5 V limit); this effect obviously gains even more influence for higher pulses.

To illustrate this effect, the temporal profile of a photomultiplier pulse was measured with an oscilloscope under normal operating conditions of the accelerator. An amplitude of 270 mV and an integral of 8300 mV ns were determined. In Fig. 7.7, this profile is shown after scaling to various amplitudes $\hat{V}_{\mathrm{pulse}}$ between 0 and 50 V, and after (simulated) clipping at 5 V in the alarm generator input stage. The lower plot shows the resulting output of the ADC integrator according to (7.8). The ADC response is only linear up to a pulse amplitude of 5 V; above, the slope of the response curve decreases rapidly. It does, however, stay strictly monotonic because of the effective broadening of the clipped pulse.

On top of this purely artificial limitation, it must be taken into account that the linearity of the photomultiplier response at high pulsed currents is not guaranteed due to various effects discussed earlier. In fact, for the voltage divider setup used in the FLASH BLMs, the tube manufacturer specifies a maximum pulse current of 100 mA for a maximum deviation of 5 % from linear behavior. However, with the input impedance of 100 Ω, this corresponds to an amplitude of 10 V. It can be concluded that the multiplier tube is only a secondary source of nonlinearities in the BLM readout chain.

7.3.2 Calibration

The calibration of a beam loss monitor falls into two parts: The adjustment of the display units, and the adjustment of the alarm thresholds.

As set out before, the natural units for a BLM display would be absorbed dose or dose rate. Taking, for instance, an ionization chamber, a good estimate of the absorbed dose corresponding to a measured charge pulse can be derived just from the knowledge of the active gas volume. For photo- and electron multipliers, matters are more complicated. The main reason is the huge gain spread between the tubes; it is quite typical to find that the gains of PMTs from the same production batch form a distributions of width $\sigma_G/\bar{G} \approx 0.5$ at the same high voltage

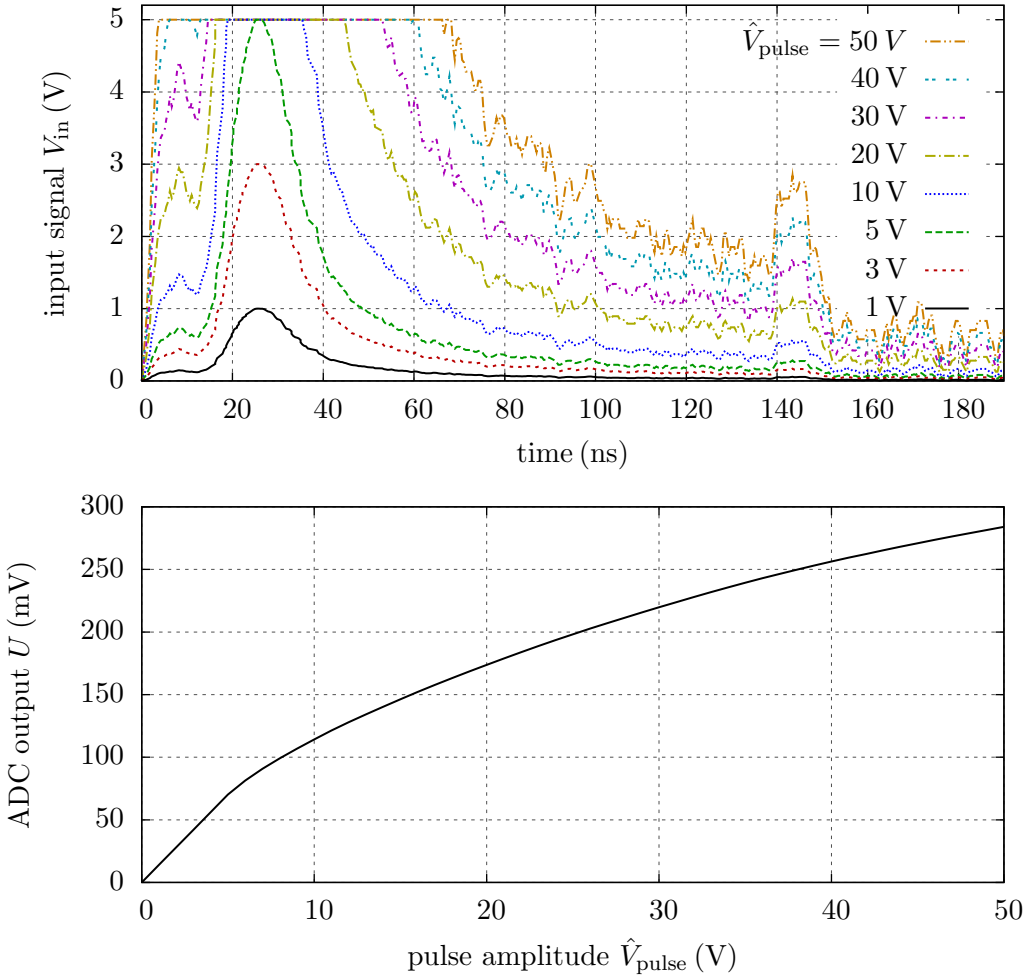

Figure 7.7: Influence of pulse clipping in the BLM readout chain. The top plot shows a measured photomultiplier pulse that has been artificially scaled to various amplitudes. The pulses are clipped at 5 V to simulate the influence of the BLM alarm generator input stage. The lower plot shows the resulting voltages at the ADC integrator output.

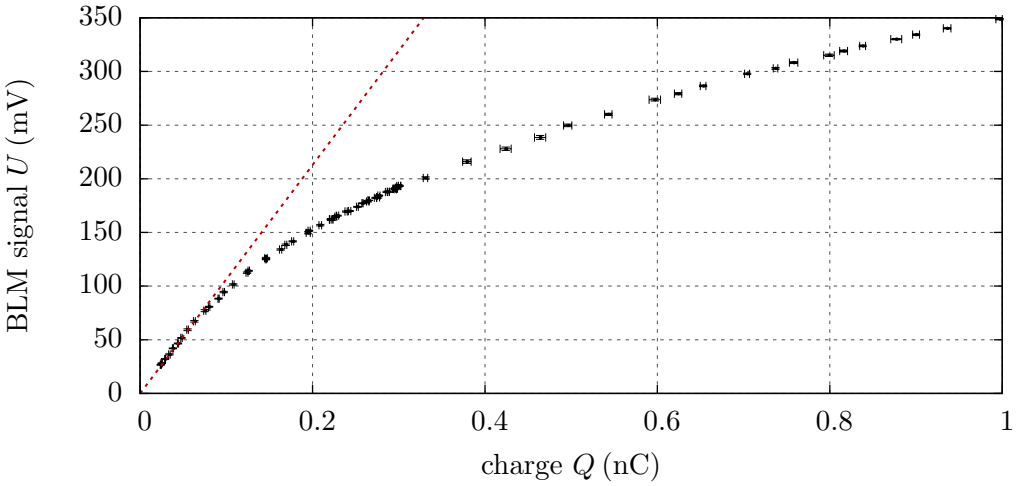

Figure 7.8: Beam loss monitor signal as a function of locally dumped charge

setting, as reported in [Lan05]. It is therefore not uncommon for the output signal of two tubes to differ by a factor of four under exactly the same external conditions. In addition, the light collection efficiency of the various types of scintillator panels used is an unknown quantity. The only way to obtain a good calibration of photomultiplier signal versus dose would therefore be to expose each BLM to a well-known radiation field; if there is no strict requirement to use the monitors for dosimetry, this effort is prohibitive. For practical purposes, a much simpler calibration scheme is used.

Calibration by induced beam loss

At FLASH, most of the beam loss monitor displays are adjusted to roughly reflect *local charge loss*. For this, the beam is deliberately driven into the vacuum chamber a few meters upstream of the monitor. The intensity of the electromagnetic shower is then varied by changing the charge of the electron bunches. The result of such a scan, the BLM signal as a function of the dumped charge, is shown in Fig. 7.8.

The curve clearly shows the impact of signal clipping discussed before, but is linear for small charges. Because any contribution from dark current is removed by subtracting the signal at zero charge, a linear regression through the origin can be used to determine the slope ζ of this part. The high voltage S is then adjusted to achieve a slope of 1, i.e. a direct correspondence between BLM signal and lost charge. Using the exponent b determined in (7.6), the new high voltage setting may be calculated as $S' = S\zeta^{-1/b}$.

Sometimes, the adjustment of the high voltage is not possible for practical reasons. Among other causes, this is the case when the desired threshold settings would exceed the range of the DACs at the new multiplier gain. In these cases the high voltage is kept unchanged, and instead the display units are multiplied by a

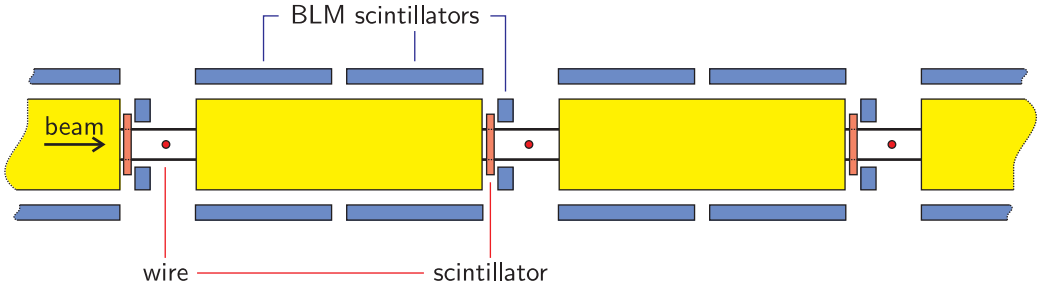

Figure 7.9: Simplified schematic of the FLASH undulator section. The picture shows a top view of two full undulator segments. Each segment is surrounded by six beam loss monitors and a wirescanner installation.

coefficient.

This calibration can only provide a rough estimate of the lost charge. Changes of beam loss position, angular distribution, or of beam energy can cause substantial variations of the local dose rate at the BLM. If the monitors are placed sufficiently close to each other, however, the procedure gives a satisfactory calibration for operational purposes.

Calibration by wirescans

Because the permanent magnets of the FEL undulators are susceptible to radiation damage, it is not advisable to calibrate the beam loss monitors in this section by dumping the complete beam even at low charge. Instead, advantage can be taken of the wirescanners that are present in the space between the undulator segments. Figure 7.9 shows a simplified schematic of the section. Each undulator segment is preceded by a wirescanner and surrounded by six beam loss monitors. Four of the scintillation panels have the shape of long strips extending parallel to the beam line, the other two are small blocks of vertical orientation placed near the exit of the undulator.

Each wirescanner installation consists of a set of three vertical and three horizontal wires of various materials and diameters (10 µm carbon, 10 and 50 µm tungsten) that can be moved through the beam individually, and of a photomultiplier with a small scintillation panel at the far end of the corresponding undulator segment. The integrated output pulse of this photomultiplier is commonly referred to as the *wirescanner signal*. While this monitor is practically equivalent to a BLM, its scintillator is smaller in volume and its light output is strongly attenuated to guarantee a linear response of output signal to scattered charge.

The six BLMs of an undulator segment are calibrated simultaneously by a slow wirescan with the horizontal or vertical 10 µm tungsten wire. In this procedure, the wire takes several seconds to cross the beam. During the scan, the following data are recorded for each pulse:

- bunch charge Q_i
- wire position x_i
- wirescanner signal W_i
- BLM signal U_i for each BLM

Figure 7.10 compares the beam profiles obtained with the wirescanner and with the loss monitor for an exemplary measurement. The profile shown by the BLM appears much broader because of the well-known saturation at high signal levels caused by pulse clipping.

Since the wirescanner signal is proportional to the amount of charge scattered at the wire, it can easily be corrected for charge jitter during the measurement:

$$W_i' = W_i \cdot \frac{\bar{Q}}{Q_i} \quad \text{with} \quad \bar{Q} = \frac{\sum_{i=1}^{N} Q_i}{N}$$

N is the number of data points lying within a reasonable region of interest around the bunch center, typically eight times the rms width of the beam profile. The goal is now to convert the arbitrary units of this wirescanner signal to the transverse linear charge density across the bunch ρ, i.e. to units of charge per length. For this purpose, the normalization condition for the discrete charge profile can be formulated:

$$\begin{aligned}\bar{Q} &= \sum_{i=1}^{N-1} \frac{\rho_i + \rho_{i+1}}{2} (x_{i+1} - x_i) \\ &= \frac{1}{2} \left(\rho_1(x_2 - x_1) + \rho_N(X_N - X_{N-1}) + \sum_{i=2}^{N-1} \rho_i(x_{i+1} - x_{i-1}) \right)\end{aligned}$$

Inserting the basic assumption of proportionality between the wirescanner signal and the linear charge density, $\rho_i = CW_i'$, it is possible to derive the coefficient C and to convert the profile to units of nC/mm.

The BLM signal U can now be calibrated against ρ. The lower plot in Fig. 7.10 shows $U(\rho)$ for a typical measurement. The curve is linear at low charge densities and falls off towards higher signal levels due to clipping of the photomultiplier pulse. A simple linear regression through the origin gives the slope s of the linear part.

The calibration step now consist of forcing this slope to $s = 1$ so that for sufficiently low signals, the BLM display corresponds to the linear charge density scattered at the wire under the given conditions. This can either be done by changing the high voltage of the photomultiplier tube or by adjusting a display coefficient.

Apart from the reduced production of radiation, this calibration procedure has the advantage of a more well-defined source of the electromagnetic shower than

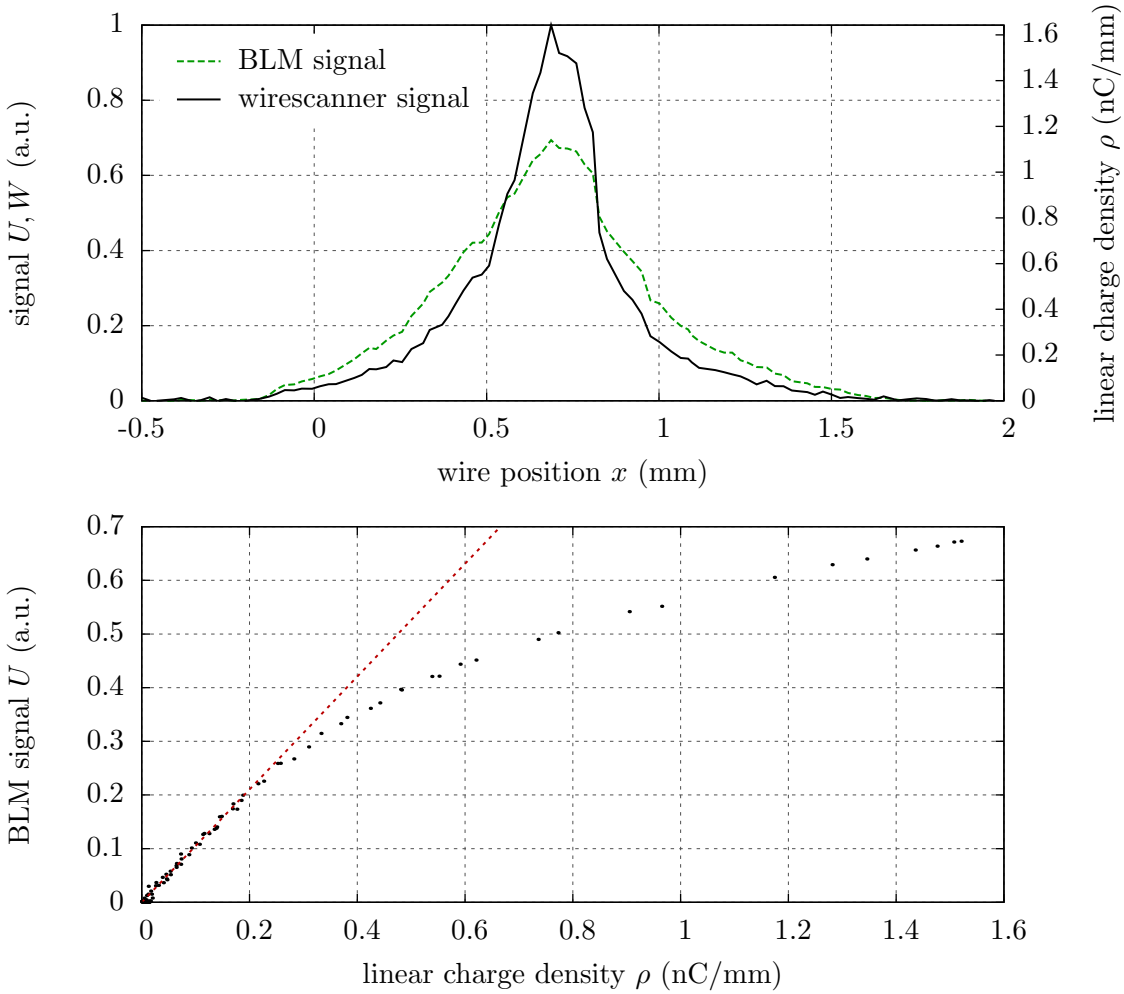

Figure 7.10: BLM calibration by wirescan. The top plot shows the beam profile as measured with the wirescanner photomultiplier and with one of the corresponding BLMs. The wirescanner signal is normalized to a maximum of 1, and can be converted to a linear charge density as indicated by the scale to the right. In the bottom figure, the BLM signal is plotted as a function of this linear charge density. The calibration involves adjusting the BLM gain in a way that the linear part of this curve has a slope of 1.

in the case of a mere dumping of the beam. Because of the periodicity of the undulator section, it can also ensure that the BLM calibrations are equal along its whole length. Nevertheless, it still has the deficit of being strictly valid only for one beam energy.

Adjustment of alarm thresholds

It has been shown that there is no universal relation between the threshold settings for the BLM alarm generators and the signal read by the ADCs. As a practical consequence of this, the precise impact of a given threshold setting can only be evaluated when beam losses are present. The initial setup is therefore done together with the *calibration by induced beam loss* described before, i.e. when the complete beam is dumped in front of the BLM. Thresholds are adjusted according to the following rules:

- The *single bunch alarm* triggers at a bunch charge of 0.1 nC .
- The *multi bunch alarm* triggers for 10 bunches of 0.05 nC.
- The *integration alarm* threshold is set just high enough to ensure that it cannot be triggered by 10 bunches of maximum charge.

These thresholds are only initial values and have to be adapted to specific requirements during the operation of the accelerator.

The alarm thresholds for BLMs in the undulator sections are more restrictive. The average dose rate deposited in the undulators should be limited to 5 Gy/d as discussed in section 5.1.1. However, it has been agreed to increase the operational limit to 10 Gy/d in order to provide more freedom for the tuning of the FEL [Sch06b].

The FLASH undulators are equipped with a fiber dosimetry system [Hen03]. It measures the attenuation of light in glass fibers that are embedded in the vacuum chambers, and thus provides a reading of absorbed dose along the undulator during the operation of the FEL. This dose measurement is used to manually adjust the thresholds of the undulator BLMs. Initially, all thresholds are set to zero. Subsequently, they are increased to allow an alarm-free passage of beam as long as measured dose rates are below 10 Gy/d. In this way, the thresholds are as restrictive as possible without constraining the normal operation of the accelerator.

7.4 Operational issues

Since its commissioning, the machine protection system of FLASH allows to operate the accelerator with full beam power. On numerous occasions, long bunch trains with up to 800 bunches—the maximum for the current setup of laser and rf systems—have been safely transported through both the FEL and the bypass beamline (see e.g. [Fro06a]). However, routine operation is still limited to few

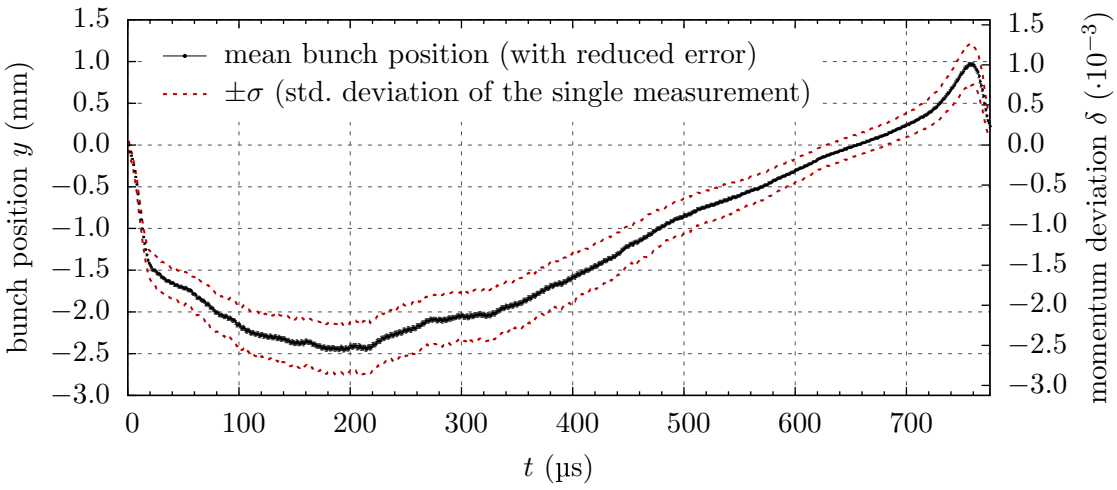

Figure 7.11: Vertical bunch positions along a macropulse of 775 bunches in the dispersive dump line. The average of 139 successive macropulses is shown. The second vertical scale indicates the relative momentum deviation according to the design dispersion.

bunches. This can in part be explained by the relatively small demand for high average power from the photon users. More importantly, the linac setup for a high-power FEL run is much more demanding and thus time-consuming than for a low-power run.

The short pulse mode permits the transport of up to 30 bunches per macropulse with a minimum of interference from the MPS. This allows quick changes to the accelerator setup and the delivery of a low-power beam, often with losses in the range of several percent. In long pulse mode, on the other hand, beam losses are strictly controlled and a much more careful setup is required. Long bunch trains also put a much higher demand on the rf regulation.

The regulation problem is illustrated by Fig. 7.11 which shows the measured vertical bunch positions along a macropulse of 775 bunches in the dump line at $z = 249.12\,\mathrm{m}$. Considering just the bending angle of 19° of the dump dipole, the vertical dispersion in this position amounts to $D_\mathrm{y} = 96\,\mathrm{cm}$, which means that (cf. appendix C, Fig. 4.27)

$$y(\delta) = D_\mathrm{y}\delta$$

with δ denoting the relative energy deviation of the individual bunch. Because of the big value of D_y, other contributions to the vertical offset can be neglected. Furthermore, since the center position of the beam position monitor is uncalibrated, the origin has been chosen at the position of the first bunch.

In this example, the bunch train covers a vertical range of about 3.5 mm, corresponding to a peak-to-peak energy spread of $3.6 \cdot 10^{-3}$ (relative) or 2.5 MeV (absolute) considering the average beam energy of 690 MeV. The energy profile

across the macropulse is highly inhomogeneous due to inadequate regulation of phase and amplitude during the rf pulse: Initially, the field energy stored in the superconducting cavities decreases because it is converted into kinetic energy of the electron bunches—a process called *beam loading*. Consequently, the energy of the first bunches decreases almost linearly.

About 20 µs after the beginning of the macropulse, the rf feedback starts to compensate this power drain. Because the feedback alone is both too weak and too slow, a number of beam loading compensation schemes for the proper supply of additional rf power have been devised; in the shown measurement, the beam loading compensation for acceleration module ACC1 has been set up manually while ACC2–6 have been operated with an adaptive feedforward algorithm described in [Bra07]. Considering that the extent of the static energy pattern is several times bigger than the pulse-to-pulse jitter (indicated by dashed lines), it is clear that there is ample room for improvement of the rf regulation.

This issue obviously makes the loss-free transport of long bunch trains more complicated because it enlarges the effective beam size that has to fit through the apertures of the accelerator. It also has other side effects: For example, the non-uniformity of acceleration phases causes different degrees of bunch compression in the magnetic chicanes, leading to inhomogeneous lasing across the bunch train. It can be concluded that the development of a reliable and fully automated rf control for superconducting cavities is still one of the major challenges in view of future machines with tens or hundreds of rf stations such as the European XFEL or the International Linear Collider.

8 Machine protection system for the European XFEL

In many respects, the design of the planned XFEL facility is based on experience from FLASH. This is also true for the proposed machine protection system; however, the designated average beam power of 600 kW and beam spot sizes down to few micrometers give the XFEL a substantially higher damage potential. The XFEL also features a novel beam distribution system that requires a somewhat more complex MPS logic. Therefore, a new MPS architecture and the redesign of individual components are required. The development of the new system is driven by a group of scientists, engineers and technicians[1] in a collaborative design effort. Because more detailed publications are pending, only a rough outline is given in the present thesis.

8.1 Requirements

The machine protection system for the XFEL must of course fulfill the basic requirements set out in the introduction—protect the accelerator from damage, minimize downtime, limit activation. The system should also provide a high degree of flexibility. Experience with FLASH has shown that the accelerator operation profits from an MPS whose behaviour can be changed or extended in a simple way. For this purpose (and because of the necessity for short reaction times), it has been decided to implement the basic logical functions of the system in a field-programmable gate array (FPGA).

Since most of the MPS electronics will be located in the accelerator tunnel, an elevated radiation background must be expected. An analysis on the behavior of FPGAs has been carried out at DESY [Ryb04] and shows that neutron-induced single event upsets (SEUs) are the major source of malfunctions. Therefore, the design of the electronics should ensure as far as possible that SEUs do not lead to unsafe or uncontrolled behavior of the system. Fault tolerant design schemes like triple modular redundancy (TMR) constitute a possible remedy. In addition, periodic integrity checks of the FPGA configuration should allow the detection of radiation-induced errors.

Power outages or simple cable breaks must put the system into a safe state—typically, by inhibiting beam operation. On the other hand, any alarm condition

1 I. Cheviakov, L. Fröhlich, S. Karstensen, T. Külper, T. Lensch, D. Nölle, K. Rehlich, F. Schmidt-Föhre, M. Staack, M. Werner (DESY Hamburg)

Beam loss location	Distance from injector (m)	Distance from dump kicker (m)	Min. number of lost bunches
Photoinjector	0	−1970	0
Bunch compressor 1	160	−1810	7
Bunch compressor 2	360	−1610	15
Main linac center	1040	−930	44
Main linac end	1650	−320	69
Beam distr. kicker	2010	40	2
Last undulator	3010	1040	44

Table 8.1: Minimum reaction times to MPS alarms for various locations along the XFEL beamline. A signal velocity of $2/3\,c$ is assumed.

should be maskable (i.e., overridable by manual intervention) as long as the main parts of the MPS are functional.

8.1.1 Reaction times

In the XFEL, the distance from the injector lasers to the last undulators is approximately 3 km. Hence, a signal needs about 10 µs to travel from one end of the accelerator to the other at the vacuum speed of light. At the maximum bunch frequency of 5 MHz, up to 50 bunches are present in the beamline. Assuming that a beam loss occurs at the farthest position from the injector and is detected by a BLM without delay, the signal still needs considerable time to reach the injector laser in order to switch it off; this time would be of the order of 15 µs for a fiber-optic transmission line with a signal propagation speed of $2/3\,c$. As a consequence, at least 125 bunches would be lost before the MPS could take any counter measures. These bunches would carry a total energy of about 2.2 kJ, enough to melt about 4 g of copper from room temperature.

To improve reaction times, the MPS needs an additional way of stopping the beam. The only component fast enough for this is the dump kicker; it can send an arbitrary number of bunches safely into the first main beam dump. By using this kicker, the MPS reaction time to a beam loss at the end of the accelerator can be reduced to the equivalent of 44 lost bunches. Table 8.1 lists the minimum number of lost bunches for various possible beam loss locations. Losses in in the collimation section require the longest reaction times (the equivalent of 69 lost bunches). Hence, the design criteria for the collimators demand resistance to the impact of 80–100 full electron bunches [Agh07].

Nonetheless, the average power deposition by a continuous loss would be excessive in most parts of the accelerator. Therefore, a dynamic reduction of the macropulse length is proposed: After detection of a beam loss, the number of bunches for the following pulses would be limited by the MPS. If these limited macropulses are transmitted without further losses, the macropulse length could be increased again, either automatically or manually.

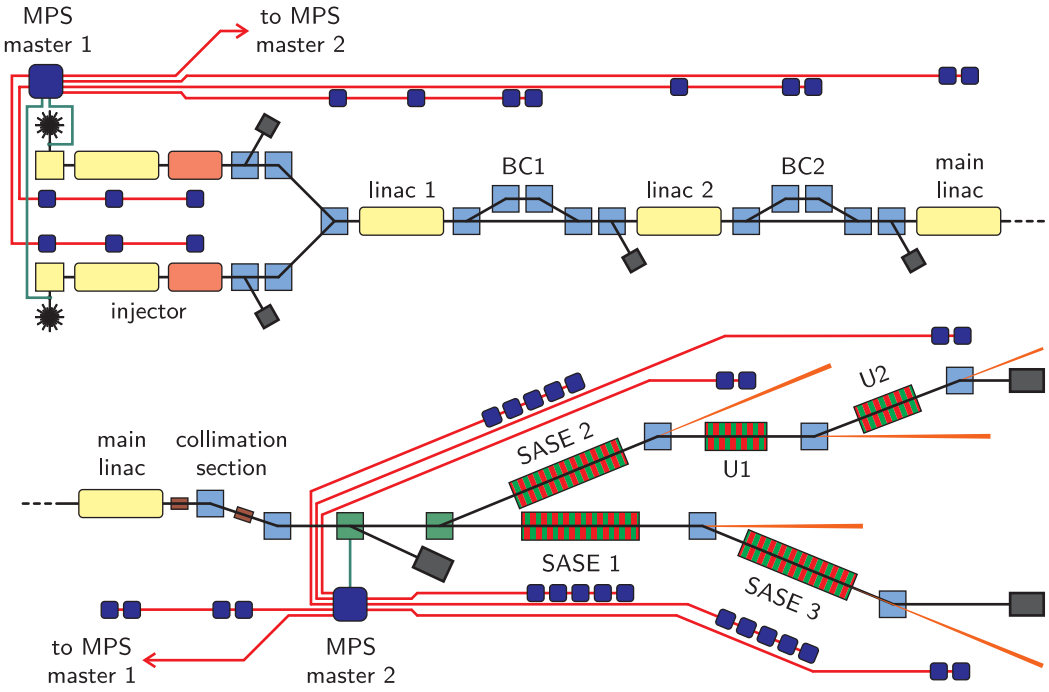

Figure 8.1: Schematic of the proposed MPS architecture. Data from BLMs and components is gathered by distributed MPS slave modules. These slaves are connected to two master modules via multiple optical communication links.

8.2 Architecture

The backbone of the XFEL machine protection system will consist of about 100 MPS slave modules *(slaves)* that are distributed along the machine. Similar to a BIC, each of these modules has at least 12 inputs for digital 1-bit signals (such as on/off or OK/alarm) from critical subsystems, accelerator components, and beam loss detection hardware. In addition, it has four or more output ports to interface with nearby equipment, e.g. to inhibit the operation of a klystron. The main function of each slave is to forward the status of its input channels to another slave or to one of two MPS master modules. For this, each slave is equipped with fast serial input and output ports; the connections are realized with fiber optic cables to avoid electromagnetic interference.

The two MPS master modules *(masters)* are located near the injector and near the dump kicker as indicated in Fig. 8.1. Each one has a multitude of serial input and output ports for communication with the slaves. While the slaves are mainly data gatherers, the programming of the masters defines the behavior of the MPS. They have direct connections to the injector laser(s) and to the dump kicker, allowing them to stop the production of new bunches and to dump bunches that are already in the machine.

8.2.1 Communication between MPS modules

The serial connections between MPS masters and slaves form loops carrying a steady data stream of short *telegrams.* Although the details of the protocol are not yet fixed, a possible communication scheme for a loop with two slaves is illustrated in Fig. 8.2: The communication is initiated by the master which sends a 48-bit telegram to slave 1. This telegram consists of three blocks of 16 bit length. Each of these blocks carries 12 bits of information and a checksum of 4 bits that allows to detect transmission errors. While it is proposed to use a cyclic redundancy check (CRC, [Pet61]) algorithm to calculate this checksum, the implications on latency still have to be investigated. The first block is reserved for general purpose data, e.g. information on the current operation and beam modes. The other blocks of the telegram contain the desired status of the slaves' output ports.

When the telegram is received by slave 1, it sets up its digital outputs according to the information in the respective data block and replaces the block with the status of its input ports. Afterwards, the modified telegram is transmitted to the next slave, which acts accordingly. At the end of the loop, the telegram reaches the master again. It now contains the complete information on all of the slaves' input ports, which can be used as an input to the main MPS algorithm to take necessary protective measures.

The number of slave modules in each loop is mainly limited by the desired latency. Because inevitable cable delays are already of the order of several microseconds, the additional latency from electronics should be considerably lower than 2 µs (corresponding to 10 bunches at 5 MHz). First simulations and test setups have shown that the latency per module—the time for receiving a telegram, modifying it, and sending it—is of the order of 100 ns. Small loops with up to five slaves therefore promise a latency of less than 1 µs plus cable delays.

8.3 Functionality

Each of the two master modules is supplied with the status of all input channels via its serial connections. Additionally, the masters receive a notification on the bunch pattern of the next macropulse from the timing and beam distribution system. Based on this information, the FPGA program can take the necessary steps to ensure the safety of the machine.

8.3.1 Operation modes

From the status of valves and magnets, the MPS determines an operation mode for each of the photoinjectors and for the dump kicker. These operation modes describe valid electron paths through the machine in a similar way as the operation mode of FLASH. The proposed modes for the XFEL MPS are shown in Fig. 8.3. To give an example, if all necessary magnets in front of the photoinjector 1 dump are

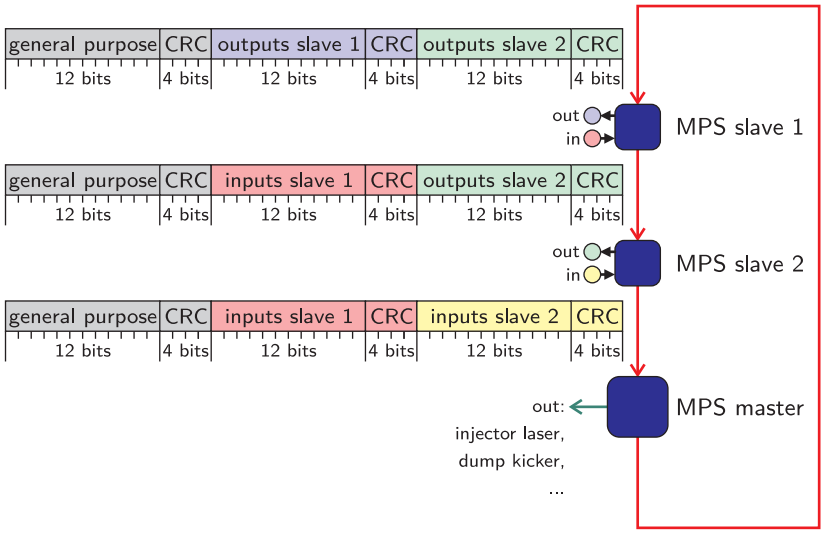

Figure 8.2: Possible communication scheme between MPS modules. The master issues a steady stream of telegrams that are modified by the slaves and finally arrive again at the master.

switched on, the valves in this section are open, the dump dipole itself is switched on, and the valve to the linac is closed, photoinjector 1 is in injector mode.

The production of electron bunches is only allowed when all requirements of an operation mode for the respective photoinjector have been fulfilled. This has the effect that the beam is automatically stopped if a valve along the beam path closes or if a magnet fails. The independent modes for the two injectors allow to operate the linac from one of them while the other is used for studies or switched off for maintenance.

The operation mode for the dump kicker has a slightly different effect: If there is no mode, all bunches are kicked into the main dump, and the injectors are limited to 50 % of the maximum beam power because of the 300 kW dump power limitation. In SASE1 or SASE2 mode, one of the undulator beam lines is not in working order, and all bunches scheduled for this beam line are dumped. In SASE1+2 mode, both beam lines are able to accept bunches and there is no restriction.

8.3.2 Power limit

The MPS can limit the average beam power by restricting the number of bunches per macropulse that are produced in the rf guns. The operators may choose from several predefined power limits; choosing those with small beam power facilitates startup and tuning of the accelerator by disabling unnecessary alarm channels. These limits may also be selected automatically by certain conditions like the insertion of screens into the beam pipe.

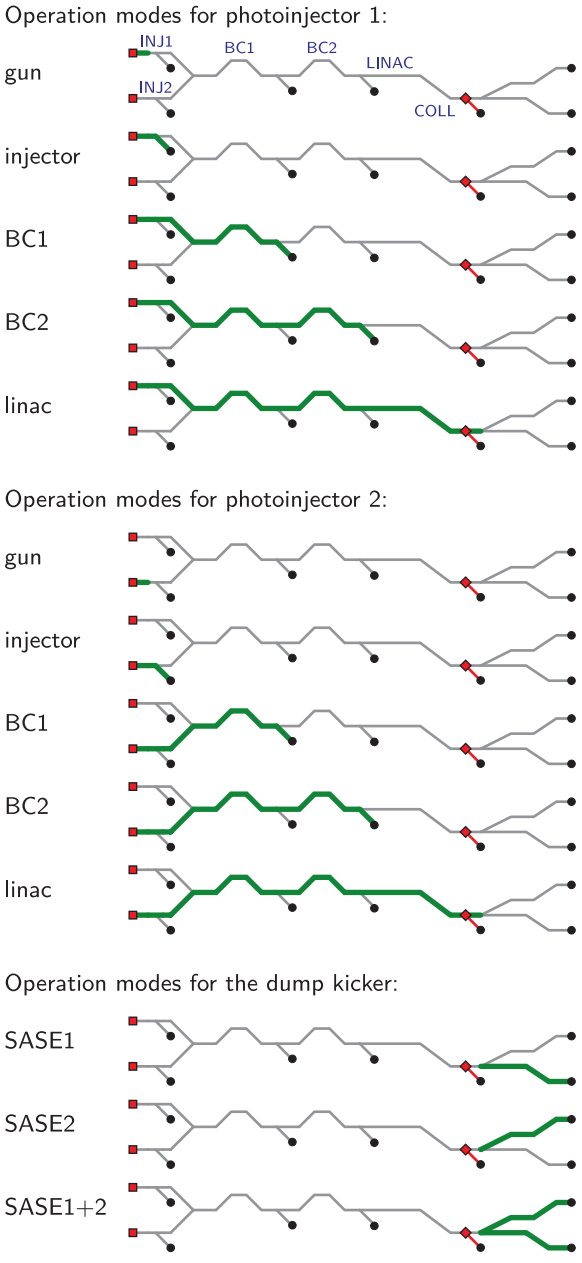

Figure 8.3: Operation modes for the XFEL MPS. There are separate modes for photoinjector 1, photoinjector 2, and the dump kicker. Green lines indicate the designated electron beam path. Black dots represent beam dumps.

As suggested before, the MPS might also dynamically reduce the number of bunches after beam loss alarms, or increase it again when the operation is loss-free. The algorithm for this behavior will have to be based on operational experience. In addition, each of the undulator lines will have a separate power limit to avoid mutual interference: if a screen is inserted in the SASE 1 line, the SASE 2 line should still be able to get full beam power.

8.3.3 Alarm cutoff

An arbitrary number of input channels may be marked as *alarm channels* for the current operation mode of the machine. If any of these channels shows an alarm signal, the beam is immediately inhibited for the rest of the current macropulse by blocking the active photoinjector laser and by activating the dump kicker. Typical alarm sources would include beam loss monitors, exceptions from the rf system, or beam position monitors.

8.4 Reliability and availability

The European XFEL will be a user facility with high demands on the availability of the accelerator. In this respect, a two-fold responsibility falls on the MPS: First, it must avoid accelerator downtime due to damage caused by the beam. Second, the MPS itself should not cause more downtime than necessary. The second point includes soft aspects like the usability of the system and the ease to diagnose and correct faults. It also includes the hard aspect of hardware lifetime in a potentially radiative environment.

The following discussion focuses on the latter; it pertains only to faults that require a repair or replacement of a faulty piece of hardware. It is assumed that soft errors like single event upsets—if not prevented by fault tolerant circuit design—can be resolved quickly by a remote reset of the faulty unit.

The hardware availability of the MPS should be as high as reasonably achievable. To get a rough estimate of the requirements, an availability goal of

$$A_{\mathrm{MPS}} = \frac{\mathrm{uptime}}{\mathrm{uptime} + \mathrm{downtime}} = 0.999$$

is set. For a year of operation (365 full days), this corresponds to a downtime of about 9 hours caused by MPS hardware faults. Instead of uptime and downtime, the well-defined quantities *mean time to failure* (MTTF) and *mean time to repair* (MTTR) are commonly used to define availability:

$$A = \frac{\mathrm{MTTF}}{\mathrm{MTTF} + \mathrm{MTTR}}$$

MTTF is the average time a unit operates until it breaks down, MTTR is the average time needed to repair or replace the unit. Detailed discussions of these quantities and their basic relations are found in a variety of text books, e.g. [Kor07].

If an MTTR of four hours for a fault is assumed—this implies very good fault diagnostics and the capability to swap hardware easily—, the MTTF requirement for the MPS as a whole becomes

$$\mathrm{MTTF}_{\mathrm{MPS}} = \frac{A_{\mathrm{MPS}}}{1 - A_{\mathrm{MPS}}} \mathrm{MTTR}_{\mathrm{MPS}} \approx 4000\,\mathrm{h} \approx 167\,\mathrm{d}.$$

On average, slightly more than two hardware failures per year can be tolerated.

8.4.1 Statistical quantities

To deduce the requirements for individual MPS modules from these considerations, it is necessary to establish a reliability model of the MPS. Before this can be done, some statistical quantities and their relations have to be introduced. This section follows the notation used in [Kor07].

Reliability $R(t)$ is defined as the probability that a component operates correctly for a time span t. For all practical cases, this implies $R(0) = 1$ (the component is not broken at the beginning) and $\lim_{t\to\infty} R(t) = 0$ (the component will break down after a sufficiently long time).

A closely connected quantity is the *failure rate* $\lambda(t)$, describing the statistic frequency of hardware failures for components that have operated for a time t. It is related to the reliability by

$$\lambda(t) = -\frac{1}{R(t)} \frac{\mathrm{d}R(t)}{\mathrm{d}t}. \tag{8.1}$$

Because there is no a priori knowledge about the development of $\lambda(t)$ over time, a constant failure rate is assumed here. With $\lambda = \text{const.}$, (8.1) implies an exponential reliability distribution:

$$R(t) = \exp(-\lambda t)$$

The MTTF is the expectation value of the component lifetime. It is therefore also the inverse of a constant failure rate:

$$\mathrm{MTTF} = \int_0^\infty R(t)\,\mathrm{d}t = \lambda^{-1}$$

8.4.2 Reliability model

A model is required to express the reliability $R_{\mathrm{MPS}}(t)$ of the MPS as a whole in terms of the reliability of its single components. For simplicity, the MPS modules are considered as the only building blocks of the model; possible failures of communication lines or other infrastructure are therefore included in the reliability of the modules. Since they are very similar in design, both masters and slaves are assigned the same reliability function $R_1(t)$.

The model is illustrated by Fig. 8.2. It is assumed that the two masters are connected to L communication loops. Each loop contains N slaves. If a single slave

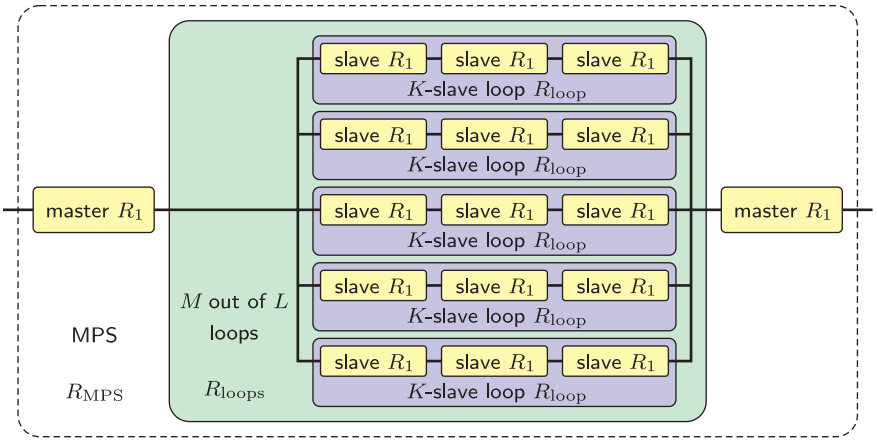

Figure 8.4: Reliability tree for the XFEL MPS. For the MPS to be considered functional, both masters and M out of L loops must be working.

fails, the communication along its loop is interrupted. Therefore, the reliability of a K-slave loop is

$$R_{\text{loop}}(t) = R_1(t)^K.$$

The failure of a loop is not necessarily fatal. It merely means that a certain number of input channels cannot be monitored anymore. The MPS should certainly stop beam operation as soon as such a failure is detected, but the operator may decide that the missing inputs are not vital and can be ignored. On the other hand, if too many input channels fail, safe operation of the accelerator is not guaranteed anymore. The model therefore requires that M out of the total L loops are functional. The reliability of the system of all loops can then be calculated by simple combinatorial arguments:

$$R_{\text{loops}}(t) = \sum_{i=M}^{L} \binom{L}{i} \big(R_{\text{loop}}(t)\big)^i \big(1 - R_{\text{loop}}(t)\big)^{L-i}.$$

The sum extends over all allowed values for the number i of intact loops. The binomial coefficient

$$\binom{L}{i} = \frac{L!}{i!\,(L-i)!}$$

gives the number of possibilities to choose i out of L loops, $(R_{\text{loop}})^i$ is the probability that i loops are still intact and $(1 - R_{\text{loop}})^{L-i}$ is the probability that $L - i$ loops have failed.

The reliability of the MPS as a whole also takes into account the possibility of failures of the master modules. Each of the two masters with individual reliability

$R_1(t)$ is considered a critical component:

$$\begin{aligned} R_{\mathrm{MPS}}(t) &= R_1(t)^2 R_{\mathrm{loops}}(t) \\ &= R_1(t)^2 \sum_{i=M}^{L} \binom{L}{i} \left(R_1(t)^K\right)^i \left(1 - R_1(t)^K\right)^{L-i} \end{aligned}$$

As set out above, a constant failure rate for individual MPS modules will be assumed to evaluate this reliability:

$$R_1(t) = \exp\left(-\frac{t}{\mathrm{MTTF}_1}\right)$$

8.4.3 Case study

The parameters used for a first calculation of the reliabilities are $K = 3$ slaves per loop, $L = 33$ total loops, and $M = 31$ required loops. This implies that 2 arbitrary loops comprising a total of 6 MPS slaves are allowed to fail. This corresponds to 48 dead input channels if an average usage of 8 channels per slave is assumed.

The reliabilities resulting from several values of MTTF_1 are shown in Fig. 8.5. The plot of the reliability $R_1(t)$ of a single module directly demonstrates the influence of the MTTF; for $\mathrm{MTTF}_1 = 1$ a, each MPS module has a probability of only 0.7 % to survive 5 years of operation. For an MTTF of 10 years, this probability increases to 61 %; for 100 years, to 95 %; and for 1000 years, to 99.5 %.

The reliability of a loop with three slaves is correspondingly lower than that of a single slave; the probability that such a loop is still working after 5 years of operation with $\mathrm{MTTF}_1 = 10$ a is only 22 %. The requirement of 31 intact loops out of the total 33 reduces the reliability even further: The MPS slave system cannot be expected to operate for five years without any hardware exchange unless the MTTF of each slave is substantially higher than 100 a. For $\mathrm{MTTF}_1 = 320$ a, the probability for this is 80 %.

The total MPS reliability is completely dominated by that of the slave network. Both masters are single points of failure. However, this can be ignored because the probability of a breakdown of two slave loops is considerably higher than the breakdown of any of the masters. On the premise that the MTTF of each module does not exceeded 100–200 years, a failure of the MPS within the first few years of operation is quite certain if no repairs take place.

If repairs are necessary, the availability of the system is of particular interest. As discussed above, it can be calculated from the MTTF of the system and the MTTR:

$$A_{\mathrm{MPS}} = \frac{\mathrm{MTTF}_{\mathrm{MPS}}}{\mathrm{MTTF}_{\mathrm{MPS}} + \mathrm{MTTR}} \quad \text{with} \quad \mathrm{MTTF}_{\mathrm{MPS}} = \int_0^\infty R_{\mathrm{MPS}}(t)\,\mathrm{d}t$$

Both the system availability A_{MPS} and the corresponding MTTF are shown in Fig. 8.6 as functions of the single-module MTTF. For the plot, $\mathrm{MTTF}_{\mathrm{MPS}}$ has been

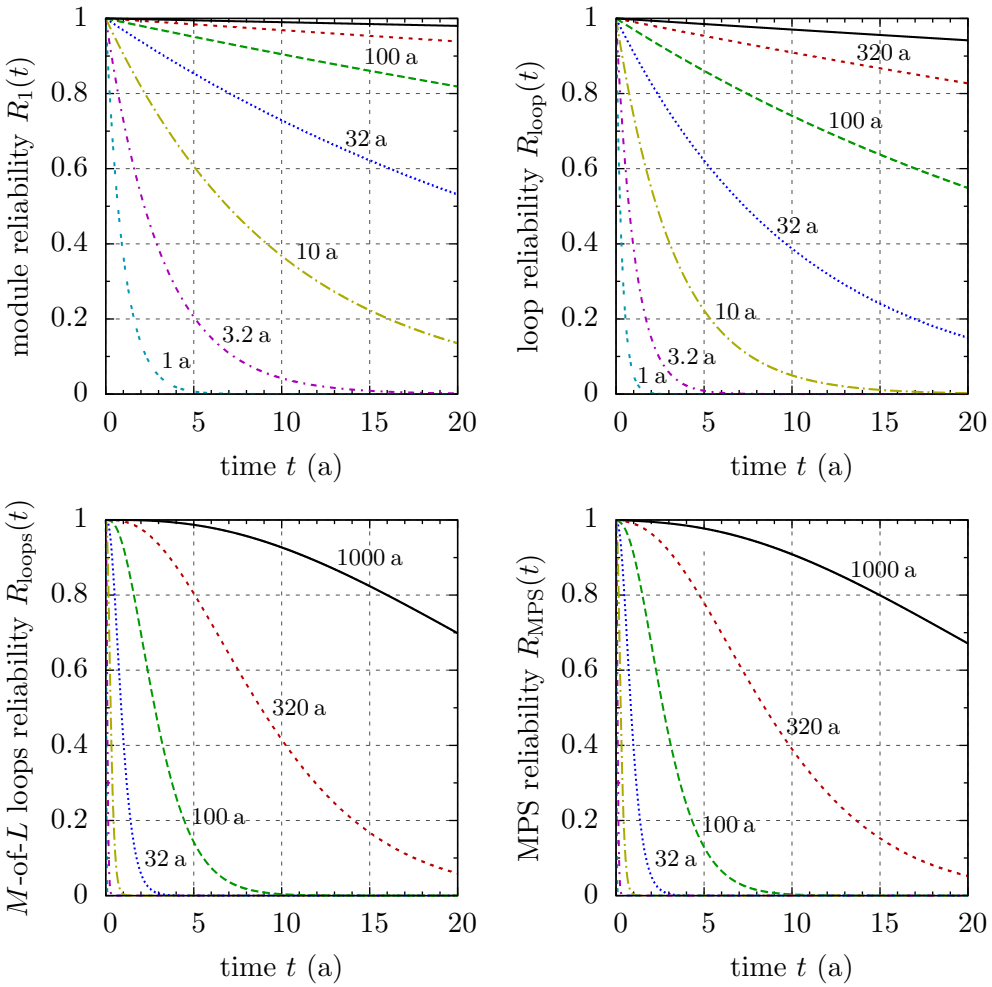

Figure 8.5: Time-dependent reliability of MPS constituents for various values of the MTTF of a single MPS module. Parameters: $K = 3$ slaves per loop, $L = 33$ loops total, $M = 31$ loops required.

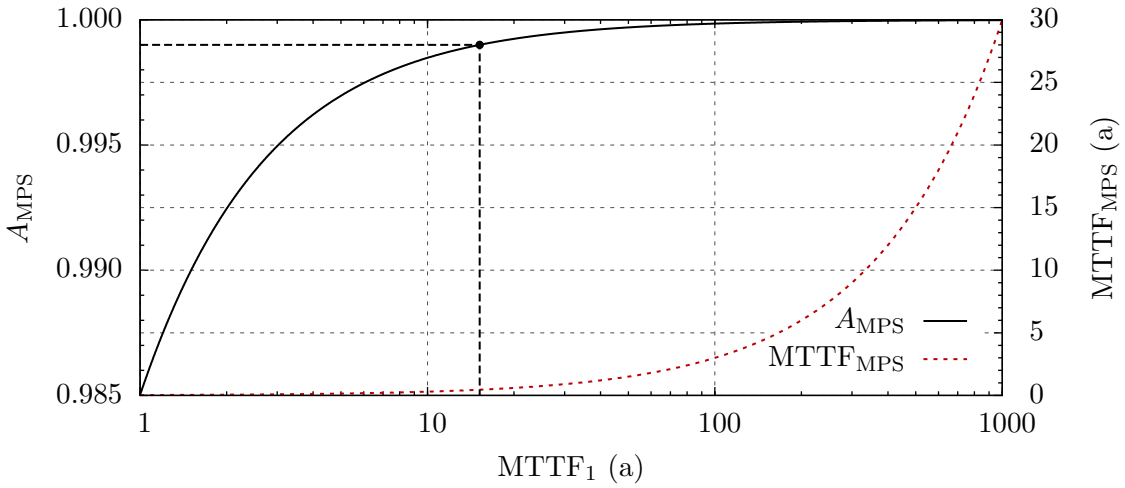

Figure 8.6: MPS availability and MTTF as functions of single-module MTTF. The availability goal of 0.999 is indicated. Parameters: $K = 3$ slaves per loop, $L = 33$ loops total, $M = 31$ loops required, MTTR = 4 h.

determined by numerical integration of $R_{\mathrm{MPS}}(t)$. The figure shows that the single-module MTTF must be higher than about 15 a in order to reach the availability goal of 0.999. To reach an availability of 0.9999, the modules have to exceed an MTTF of 150 a.

8.4.4 Parameter study

Up to now, only a specific configuration with 3 slaves per loop and the allowed failure of 2 loops has been examined. An additional parameter study is needed to quantify the influence of these parameters on the system availability.

Table 8.2 lists a number of possible parameter sets with 1–5 slaves per loop (more are undesirable due to latency). The total number of modules is kept approximately constant at $N = KL \approx 100$. For each parameter set, the single-module MTTF resulting in a system availability of $A_{\mathrm{MPS}} = 0.999$ has been determined.

All cases with $L = M$ are equivalent because they demand that every MPS module is operational; this requires an MTTF of 46.5 a or higher. The dominant parameter for all other cases is the number of required loops, or more precisely the number of loops allowed to fail, $L - M$. For equal values of $L - M$, the MTTF_1 requirement is almost independent of the number of slaves per loop. However, the failure of a 5-slave loop is potentially more hazardous than that of a 2-slave loop.

In the end, deciding upon the number of slaves per loop involves a tradeoff between cost and reliability. If each slave module has an individual connection to the corresponding master, it is likely that its failure can be ignored. If many slaves are connected in a loop, the failure of a single module automatically disables a high number of input channels. On the other hand, increasing the number of loops

slaves/loop K	**loops total** L	**loops req.** M	**slaves required**	**unit MTTF** MTTF_1 (a)
1	100	100	100/100 (100 %)	46.50
		99	99/100 (99 %)	23.36
		98	98/100 (98 %)	15.65
		97	97/100 (97 %)	11.80
		96	96/100 (96 %)	9.49
		95	95/100 (95 %)	7.95
2	50	50	100/100 (100 %)	46.50
		49	98/100 (98 %)	23.25
		48	96/100 (96 %)	15.50
		47	94/100 (94 %)	11.62
3	33	33	99/99 (100 %)	46.04
		32	96/99 (97 %)	22.90
		31	93/99 (94 %)	15.19
		30	90/99 (91 %)	11.33
4	25	25	100/100 (100 %)	46.50
		24	96/100 (96 %)	23.01
		23	92/100 (92 %)	15.18
		22	88/100 (88 %)	11.26
5	20	20	100/100 (100 %)	46.50
		19	95/100 (95 %)	22.89
		18	90/100 (90 %)	15.02
		17	85/100 (85 %)	11.08

Table 8.2: Single-unit MTTF required for a system availability of 0.999. *Slaves required* indicates the number of slaves contained in the M required loops.

requires additional cabling and—equally important—additional input ports at the two master modules. Loops with more slaves therefore reduce the total cost of the system and the complexity of the master hardware.

With the proposed architecture, the MPS also allows more flexible topologies. For example, huge numbers of low-priority channels such as magnet power supplies can be connected to loops with many modules. On the other hand, vitally important channels, e.g. for monitoring of the fast kickers, could be made redundant by connecting the signals to two slaves in independent loops. With a well-planned topology, the MPS should be able to tolerate the failure of individual slave modules until they can be exchanged in a regular maintenance period. Based on the calculations presented here, the MPS modules should have an MTTF of at least 25 a to meet the availability goals and to provide some freedom in choosing the topology.

8.5 Final remarks

The MPS for the European XFEL will have a considerably higher complexity than that of FLASH. The proposed system architecture addresses the possibility of a second photoinjector and the introduction of a fast beam distribution system. In addition, the MPS will be fully programmable and thus allow the adaption to operational needs.

Due to the length and the high bunch frequency of the XFEL, up to 70 bunches out of a single macropulse can be lost before an MPS reaction is possible. This issue can only partly be mitigated by a dynamic reduction of the beam power. In addition, preventive measures like the monitoring of beam positions will be of high importance.

For beam loss detection, the MPS will rely on fast charge comparisons between toroid pairs and on a system of about 200 BLMs. As set out before, high BLM sensitivities are needed to provide sufficient protection of the undulator magnets. The additional requirement of single-bunch resolution can best be met by a photomultiplier-based BLM design.

Like FLASH, the XFEL will not be able to reach its design beam power without a fully functional machine protection system. The MPS will therefore be one of the key elements in making the XFEL a successful user facility.

9 Conclusion

Over the last decades, particle physics and the demand for photon and neutron sources have driven the development of accelerators with increasing stored energies and beam powers. In this process, the field of machine protection has gained in importance.

Especially light sources based on superconducting linacs such as FLASH, the Jefferson Lab FEL [Ben07], or the European XFEL pose special machine protection requirements. On the one hand, the superconducting cavity technology enables these machines to accelerate electron bunches in quick succession, forming beams of high average power. Active machine protection systems with reaction times in the microsecond range are needed to prevent damage to components. On the other hand, these light sources rely on precise insertion devices that are predominantly based on permanent magnets. As these magnets are highly susceptible to demagnetization under irradiation, the tolerance to beam losses is extremely low.

FLASH is in many ways a prototype for future accelerators. It is the first FEL that has reached the extreme ultraviolet spectral range. It also serves as a testbed for the superconducting cavity technology that will be used on a larger scale in the European XFEL and, possibly, in the future International Linear Collider. Therefore, experience gained at this machine is valuable for several other projects. This thesis has discussed some of the key issues for machine protection at FLASH.

Dark current from the rf gun is one of the major problems of the accelerator. It activates components to critical levels in several places along the machine, and it contributes to the dose rate deposited in the FEL undulators. The mechanisms of dark current transport throughout the linac have been discussed with detailed tracking simulations, and practically all places of high observed radioactivity could be explained by the loss of dark current from the rf gun. The investigation has also shown that the consistent use of existing collimators can reduce the activation problem in the biggest part of the accelerator. The simulations were made possible by the parallelization of the tracking code *Astra* and by the implementation of extended three-dimensional aperture modeling capabilities.

The importance of measures for undulator protection has been investigated in a number of shower simulations. To protect the undulators from beam losses in an upstream section of the beamline, a simple shielding setup has been designed and shown to be effective. The dose rates caused by dark current and beam losses in the undulators have been calculated and used to derive practical limits for the operation of the accelerator. For the design parameters of FLASH, the average relative beam loss in the undulators must be controlled down to a level of about 10^{-8}. This emphasizes that reliable undulator protection in a linac with high

average beam power requires either very good collimation of the beam or a highly sensitive and fast beam loss monitor system (or ideally, both).

The FLASH machine protection system satisfies this requirement with a maximum total response time of less than 4 µs and with extremely sensitive photomultiplier-based beam loss monitors. Several algorithms and procedures for the setup of the MPS and for the verification of the functionality of its components have been developed. The complex chain of the BLM readout and alarm generation stages has been characterized in detail and several related problems have been identified and addressed. Finally, the safe operation of FLASH at its current limit of 800 bunches per macropulse has been demonstrated.

Based on experience from FLASH, an architecture for the machine protection system of the future European XFEL has been developed. In contrast to its predecessor, the system will be fully programmable and yet retain a response time that is dominated by cable delays. Several options for the topology of the proposed MPS network have been reviewed, with the result that an availability goal of at least 0.999 for the MPS is achievable with moderate hardware requirements.

For all linear accelerators with high average power, the control of beam losses is a major challenge. While some existing machines are already limited by this today, future light sources based on energy recovery linacs with several hundred megawatts of beam power have been proposed. Beside a good machine protection strategy, these accelerators will require a thorough understanding and control of beam dynamical issues like halo formation. Hence, it is safe to predict that the development and operation of new accelerators will continue to pose a challenge in physics and engineering for a long time.

A Improvements of the Astra tracking code

Tracking codes are indispensable tools for the design and the optimization of modern accelerators. A multitude of these codes exists today, and most of them are specialized on a specific type of accelerator or on the simulation of a specific set of physical effects. In recent years, the exact modeling of space charge effects has drawn particular attention. This development is mainly driven by two communities: on the one hand, the operators of heavy-ion accelerators which transport beams that never reach ultrarelativistic velocities and are thus constantly influenced by their own space charge field; and, on the other hand, the operators of high-brightness electron guns like the photoinjectors of TTF, FLASH, or LCLS.

The development of the TTF photoinjector in particular has spurred the development of Klaus Flöttmann's tracking code *Astra* (an acronym for *A space charge tracking algorithm*) [Flo00]. Astra's main algorithm allows to simulate the propagation of a particle bunch under the influence of its own cylindrically symmetric space charge field and of simple external electric and magnetic fields. Over time, more advanced algorithms and field types have been added, making the code more versatile. Within the wider scope of this thesis, two major changes to the code have been implemented: First, the parallelization of existing algorithms, and second, a novel aperture modeling subsystem.

Appendix A gives an overview of the implemented algorithms and evaluates their efficiency; most of the data are also presented in [Fro07a].

A.1 Parallelization

An increase of computing speed is a bonus for practically all computer simulations. It either cuts down the time one has to wait for results, or, often more importantly, it allows to increase accuracy or to analyze more complex problems. One method of increasing the speed of a simulation is to *parallelize* it, i.e. to split the problem into multiple parts that are suited for concurrent processing on two or more processing units.

While parallel computing systems have been explored since the 1960s, they have only come into significant use for accelerator related particle tracking in the 1980s, usually in the form of vector computers or special multiprocessor systems (see e.g. [Jej88]). In the 1990s, clusters of comparatively inexpensive microcomputers became a suitable platform for parallel computing. This has led to a wide spread of these systems in universities and laboratories, and several tracking codes have

been adapted to take advantage of this. In recent years, the advent of cheap multiprocessor boards and multicore CPUs has accelerated this trend. Owing to this, the target platform for the parallel Astra version is a cluster of PCs or a single multiprocessor PC.

A.1.1 Algorithmic approach

Like most tracking codes, Astra relies on the concept of *macroparticles*. A macroparticle is a point-like entity that is tracked like a single beam particle but represents a higher number of particles to the environment; if a bunch consists of N macroparticles, each one is assigned a charge of Q_i so that the sum equals the desired bunch charge Q:

$$Q = \sum_{i=1}^{N} Q_i$$

While uniform distributions with $Q_i = Q/N$ are common, this is not required.

Astra provides several particle-in-cell (PIC) algorithms to calculate the space charge field generated by the macroparticles. The most frequently used one employs a cylindrically symmetric grid in which the single cells are addressed by a radial and a longitudinal bin number, r and l. The code iterates over all N macroparticles to determine the charge present in each cell, Q_{rl}. Afterwards, the main tracking loop iterates over all particles again, for each one summing up the contributions of all partial charges Q_{rl} to the local space charge field as well as external electric and magnetic fields. The forces exerted on each particle are then integrated with a fourth-order Runge-Kutta algorithm (see e.g. [Atk89]).

Many parallel PIC codes divide the computation work by distributing the grid cells among the available processors, which is especially advantageous in the case of grids with high resolution and with computationally demanding field calculations. However, Astra's cylindrically symmetric algorithm is frequently used with coarse grids (under 1000 cells), and the field calculation is comparatively simple due to the absence of complicated boundary conditions. Investigations with a code profiler also confirm that only a minor part of the execution time is spent in the space charge routines.

On the basis of these considerations, it has been decided to achieve parallelization by distributing the macroparticles among the processors. In a run with N particles and P processors, the initial assignment would follow the pattern:

process 0:	particles	$0 \ldots N/P - 1$
process 1:	particles	$N/P \ldots 2N/P - 1$
...		
process N:	particles	$(P-1)N/P \ldots N - 1$

To exploit this form of data parallelism, a classical *single program multiple data* (SPMD) approach has been chosen; the existing serial source code has been extended by calls to the MPICH2 message passing interface library [MPI09]. It is

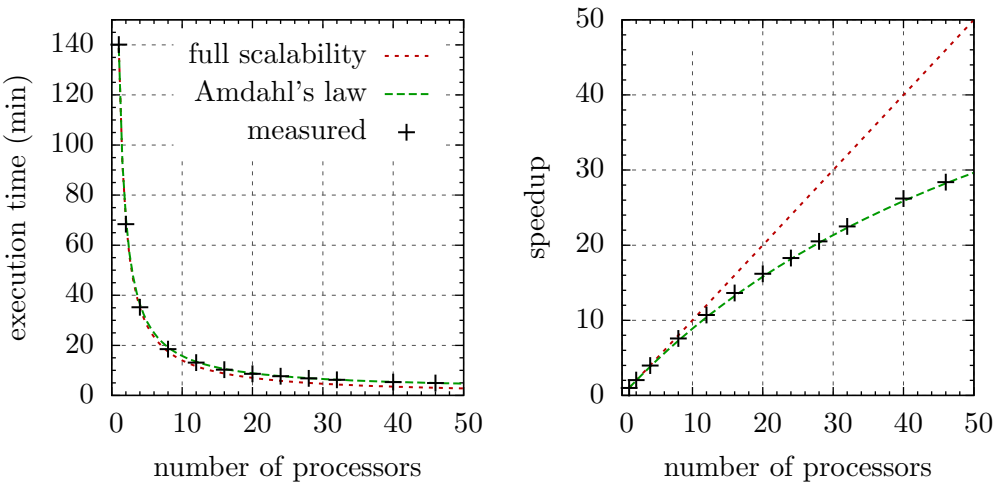

Figure A.1: Astra parallel performance benchmark. The red lines indicate the best theoretically possible execution time based on the measurement for a single processor, i.e. with assumed full scalability. The green lines correspond to Amdahl's law with a sequential code fraction of $\alpha = 1.4\,\%$.

therefore possible to run the executable on a single processor as usual, or to spawn multiple processes within the framework of an existing MPICH2 installation.

A.1.2 Parallel Performance

Figure A.1 shows the measured execution times on a cluster with 23 nodes of two CPUs each. Shared memory is used for communication between the twin CPUs, gigabit ethernet for communication between the nodes. The input file used for benchmarking contains the complete model of the FLASH linac as described in chapter 4, but simulates only the first 20 m of the machine with 10^5 macroparticles including space charge effects. There are no emittance calculations or file outputs except for the saving of one phase space file at the end of tracking.

The speedup ξ achieved by using P processors is defined as the ratio of parallel to sequential execution time, $\xi(P) = \Delta t(P)/\Delta t(1)$. Figure A.1 shows that the measured speedup increasingly deviates from ideal behavior with increasing number of processors—the maximum achieved on the cluster is $\xi(46) = 28.4$. This behavior is described well by *Amdahl's law* [Amd67] which states that if a program can be separated into a parallelizable part and a sequential part that takes a fraction α of the total execution time, the speedup by running it on P processors is given by

$$\xi(P) = \left(\alpha + \frac{1-\alpha}{P}\right)^{-1}.$$

By fitting to the data, a sequential fraction of $\alpha = 1.4\,\%$ is obtained for this

benchmark case. The maximum possible speedup is then given by

$$\lim_{P\to\infty} \xi(P) \approx 71.$$

It should be noted that these figures depend strongly on the specific application—generally, using a higher number of macroparticles will increase the speedup, and a higher number of sequential operations like file saving or emittance calculations will decrease it.

Furthermore, a big part of the sequential execution time is caused by communication between the individual processes, which need to exchange information about basic parameters of the particle distribution and about the space charge field on a regular basis. Therefore, a faster interconnection between the processors (like, e.g., in a pure shared-memory system) can result in a significant performance boost.

A.2 Aperture Modelling

For the dark current simulations described in chapter 4, a good knowledge of the apertures along the machine is essential. It is also necessary to have tools that can reproduce these geometries in an adequate way. Unfortunately, the built-in capabilities for modeling apertures in Astra and in most tracking codes are limited to basic shapes like circular openings or parallel plates.

To facilitate a coherent and exact description of more complex geometries found in the accelerator, a portable aperture library has been developed under the title *ApertureLib.* It is written in C++ under extensive use of the Standard Template Library and allows integration to C and Fortran projects by a set of wrapper functions. The library allows to read aperture models from XML files like the following:

```
<aperture-list>
  <circle z="0" name="drift">
    <r>0.017</r>
  </circle>
  <include z="2.4" name="acc. module">
    <filename>acc_module.xml</filename>
  </include>
</aperture-list>
```

In this example, two aperture elements are specified along with their longitudinal positions in the machine. The first one describes a circular aperture with a radius of 1.7 cm, the second one includes another XML file that defines the geometry of cavities and vacuum chambers inside a cryomodule.

A wide variety of element types is available, allowing to model even complicated geometries:

- Primitive profiles (circle, ellipse, rectangle, ...), rotatable to any direction in space
- Elements delimited by an arbitrary number of planes, specified by position and normal vectors
- Combination of other apertures by a logical AND or OR, e.g. to model alternative beamlines or new shapes like semicircles
- Inclusion of other XML files as building blocks for repetitive structures
- Inclusion of existing radial aperture descriptions in Astra format
- Import of 2-dimensional CAD drawings in DXF format for complex planar geometries

The full description of available input options is found in [Fro08].

ApertureLib is included in the parallel version of the Astra code and can be used instead of Astra's traditional aperture modeling capabilities. In addition to the library, a set of tools has been developed for the validation and inspection of aperture models. The toolkit also allows the generation of plots and checking of phase space files against an aperture model.

A.3 Offline phase space analysis

For the evaluation of collective properties of the particle ensemble in the dark current tracking simulations, Astra's built-in functions are inadequate. The main reason for this is the strict adherence to the initial cartesian coordinate system even after deflection of the beam by dipole magnets; e.g., the horizontal and vertical rms beam width are always calculated along the global x and y axes instead of using a rotated system, resulting in non-intuitively high values.

Instead of changing Astra's documented behavior, an external tool has been developed to calculate the collective properties in a more familiar way. During the Astra run, the complete particle phase space is saved to disk in regular intervals. These phase space files are then read by the tool and evaluated with a set of custom algorithms. For reference, only the calculation of the properties discussed in chapter 4 will be explained here.

Coordinate transformation

The particle ensemble read from a phase space file consists of N particles with the absolute coordinates

$$\boldsymbol{r}_i^* = (x_i^*, y_i^*, z_i^*) \qquad \boldsymbol{p}_i^* = (p_{\mathrm{x}i}^*, p_{\mathrm{y}i}^*, p_{\mathrm{z}i}^*).$$

Additionally, each macroparticle is associated with an individual charge q_i. In an initial step, the average properties of the distribution are calculated:

$$\bar{\boldsymbol{r}} = \frac{1}{NQ}\sum_{i=1}^{N} q_i \boldsymbol{r}_i^* \qquad \bar{\boldsymbol{p}} = \frac{1}{NQ}\sum_{i=1}^{N} q_i \boldsymbol{p}_i^* \qquad \text{with } Q = \sum_{i=1}^{N} q_i$$

Introducing the notation

$$\bar{r} = |\bar{\boldsymbol{r}}| = \sqrt{\bar{x}^2 + \bar{y}^2 + \bar{z}^2} \qquad \bar{p} = |\bar{\boldsymbol{p}}| = \sqrt{\bar{p}_\mathrm{x}^2 + \bar{p}_\mathrm{y}^2 + \bar{p}_\mathrm{z}^2},$$

a unit vector pointing in the direction of motion is derived:

$$\boldsymbol{d} = \bar{\boldsymbol{p}}/\bar{p}$$

The next step is to construct a matrix R that rotates $\boldsymbol{d}$ onto the z axis:

$$\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = R\,\boldsymbol{d}$$

The matrix should therefore correspond to a rotation by an angle α about axis $\boldsymbol{a}$:

$$\alpha = \arccos\left[\boldsymbol{d}\cdot\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}\right] \qquad \boldsymbol{a} = \begin{pmatrix} a_\mathrm{x} \\ a_\mathrm{y} \\ a_\mathrm{z} \end{pmatrix} = \frac{1}{\sin\alpha}\,\boldsymbol{d}\times\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Following this definition, $\boldsymbol{a}$ is a unit vector. Hence, the full rotation matrix reads:

$$R = \begin{pmatrix} 1-C+a_\mathrm{x}^2 C & a_\mathrm{x}a_\mathrm{y}C - a_\mathrm{z}S & a_\mathrm{x}a_\mathrm{z}C + a_\mathrm{y}S \\ a_\mathrm{x}a_\mathrm{y}C + a_\mathrm{z}S & 1-C+a_\mathrm{y}^2 C & a_\mathrm{y}a_\mathrm{z}C - a_\mathrm{x}S \\ a_\mathrm{x}a_\mathrm{z}C - a_\mathrm{y}S & a_\mathrm{y}a_\mathrm{z}C + a_\mathrm{x}S & 1-C+a_\mathrm{z}^2 C \end{pmatrix}$$

with

$$C = 1-\cos\alpha \qquad S = \sin\alpha.$$

The rotation matrix R can now be used to transform the particle coordinates from the absolute cartesian coordinate system to a local one whose z axis is aligned with the average momentum vector of the distribution. The origin of the new coordinate system is chosen as the center of mass of the bunch. So, for each particle, the transformation is done as follows:

$$\boldsymbol{r}_i = R\,(\boldsymbol{r}_i^* - \bar{\boldsymbol{r}})$$
$$\boldsymbol{p}_i = R\,(\boldsymbol{p}_i^* - \bar{\boldsymbol{p}})$$

In addition, the individual propagation angle and momentum deviation are calculated:

$$\boldsymbol{r}_i' = \frac{\boldsymbol{p}_i}{\bar{p}} \qquad \delta_i = \frac{|\boldsymbol{p}_i^*| - \bar{p}}{\bar{p}}$$

Calculation of collective properties

Using the particle coordinates in the new, ensemble-relative system, a number of familiar quantities can be calculated.

Beam sizes, bunch length are given by the weighted standard deviation of the respective particle coordinate:

$$\sigma_\mathrm{x} = \sqrt{\frac{\sum_{i=1}^{N}(q_i x_i^2)}{\frac{Q}{N}(N-1)}} \qquad \sigma_\mathrm{y} = \sqrt{\frac{\sum_{i=1}^{N}(q_i y_i^2)}{\frac{Q}{N}(N-1)}} \qquad \sigma_\mathrm{z} = \sqrt{\frac{\sum_{i=1}^{N}(q_i z_i^2)}{\frac{Q}{N}(N-1)}}$$

Dispersion describes the correlation between the transverse particle coordinates and their momentum deviation, $D_\mathrm{x} := \frac{\mathrm{d}x}{\mathrm{d}\delta}, D_\mathrm{y} := \frac{\mathrm{d}y}{\mathrm{d}\delta}$. It is, of course, a function defined by the lattice of the accelerator. In absence of any knowledge of the lattice structure, however, the linear dispersion can be estimated as the slope of a particle distribution with a linear regression of $x(\delta)$ or $y(\delta)$, respectively. Assuming $x(\delta = 0) = y(\delta = 0) = 0$ and weighting by the charge of the macroparticles,

$$D_\mathrm{x} = \frac{\sum_{i=1}^{N}(\delta_i x_i q_i^2)}{\sum_{i=1}^{N}(\delta_i^2 q_i^2)} \qquad D_\mathrm{y} = \frac{\sum_{i=1}^{N}(\delta_i y_i q_i^2)}{\sum_{i=1}^{N}(\delta_i^2 q_i^2)}$$

is obtained.

Angular dispersion can be estimated in the same way as dispersion by performing a weighted linear regression of $x'(\delta)$ and $y'(\delta)$:

$$D'_\mathrm{x} = \frac{\sum_{i=1}^{N}(\delta_i x'_i q_i^2)}{\sum_{i=1}^{N}(\delta_i^2 q_i^2)} \qquad D'_\mathrm{y} = \frac{\sum_{i=1}^{N}(\delta_i y'_i q_i^2)}{\sum_{i=1}^{N}(\delta_i^2 q_i^2)}.$$

Having thus determined dispersion and angular dispersion, their (linear) contribution to the particle coordinates can be removed:

$$\tilde{\boldsymbol{r}}_i = \begin{pmatrix} \tilde{x}_i \\ \tilde{y}_i \\ \tilde{z}_i \end{pmatrix} = \begin{pmatrix} x_i - D_\mathrm{x}\delta_i \\ y_i - D_\mathrm{y}\delta_i \\ z_i \end{pmatrix} \qquad \tilde{\boldsymbol{r}}'_i = \begin{pmatrix} \tilde{x}'_i \\ \tilde{y}'_i \\ \tilde{z}'_i \end{pmatrix} = \begin{pmatrix} x'_i - D'_\mathrm{x}\delta_i \\ y'_i - D'_\mathrm{y}\delta_i \\ z'_i \end{pmatrix}$$

Emittance is then calculated using the well-known statistical definition found in many text books, e.g. [Wie03], with appropriate weighting by the individual macroparticle charges:

$$\epsilon_\mathrm{x} = \sqrt{\frac{\sum_{i=1}^{N}(q_i {\tilde{x}_i}^2)}{Q}\frac{\sum_{i=1}^{N}(q_i {\tilde{x}'_i}^2)}{Q} - \left(\frac{\sum_{i=1}^{N}(q_i \tilde{x}_i \tilde{x}'_i)}{Q}\right)^2}$$

$$\epsilon_\mathrm{y} = \sqrt{\frac{\sum_{i=1}^{N}(q_i {\tilde{y}_i}^2)}{Q}\frac{\sum_{i=1}^{N}(q_i {\tilde{y}'_i}^2)}{Q} - \left(\frac{\sum_{i=1}^{N}(q_i \tilde{y}_i \tilde{y}'_i)}{Q}\right)^2}$$

Twiss parameters are calculated from the emittances as follows (cf. [Wie03]):

$$\alpha_\mathrm{x} = \frac{-\sum_{i=1}^{N}\left(q_i \tilde{x}_i \tilde{x}'_i\right)}{Q\ \epsilon_\mathrm{x}} \qquad \alpha_\mathrm{y} = \frac{-\sum_{i=1}^{N}\left(q_i \tilde{y}_i \tilde{y}'_i\right)}{Q\ \epsilon_\mathrm{y}}$$

$$\beta_\mathrm{x} = \frac{\sum_{i=1}^{N}\left(q_i \tilde{x}_i{}^2\right)}{Q\ \epsilon_\mathrm{x}} \qquad \beta_\mathrm{y} = \frac{\sum_{i=1}^{N}\left(q_i \tilde{y}_i{}^2\right)}{Q\ \epsilon_\mathrm{y}}$$

$$\gamma_\mathrm{x} = \frac{\sum_{i=1}^{N}\left(q_i \tilde{x}'_i{}^2\right)}{Q\ \epsilon_\mathrm{x}} \qquad \gamma_\mathrm{y} = \frac{\sum_{i=1}^{N}\left(q_i \tilde{y}'_i{}^2\right)}{Q\ \epsilon_\mathrm{y}}$$

A.4 Final remarks

With the integration of the new *ApertureLib* functions, Astra has become a suitable tracking code for the detailed study of beam losses. However, for the simulation of dark current transport with loss rates exceeding 90 %, comparatively big numbers of macroparticles—up to 10 million for the simulations discussed in this thesis—need to be tracked. For an accelerator like FLASH, this results in simulation times of the order of days on the fastest available single-processor computers. Hence, the parallelized version of the code is especially well suited for this task.

It has been demonstrated that simulation times can be reduced drastically by distributing the work load over several processors. The scalability of the code makes it well suited for small PC clusters. Considering the recent trend of increasing the performance of new CPUs by adding more processor cores, the parallel version of Astra seems to be well positioned to take advantage of future CPU developments.

B Technical information

B.1 FLASH beam loss monitors

B.1.1 Photomultipliers

The photomultiplier used in the FLASH BLMs is the 9972B tube manufactured by Electron Tubes Ltd., Ruislip, UK. It is an end-window photomultiplier with a borosilicate window, an S20 multialkali photocathode of 38 mm active diameter, and a 10-stage dynode arrangement. The nominal gain can be varied between 10^3 and 10^6 by adjusting the externally supplied high voltage in the range of 600 V to 1600 V; the corresponding anode sensitivity is quoted as 0.1 A/lm to 200 A/lm. [ET01]

B.1.2 Aluminum cathode electron multipliers

In addition to the photomultipliers, 17 so-called *nuclear radiation detectors* of the type 9841 from the same manufacturer are in use as BLMs at FLASH. The tube is identical to the 9972B photomultiplier except that the photocathode is replaced by a layer of aluminum. The tube is therefore sensitive to ionizing radiation that releases electrons from the aluminum cathode. Its electrical characteristics are identical to the 9972B tube.

B.1.3 High voltage power supply

High voltage for the BLM photomultipliers and aluminum cathode electron multipliers is supplied centrally by a CAEN SY527 multichannel power supply system with several A734N 16-channel high voltage boards. The voltage for each of the more than 80 BLMs is controlled individually to a precision of $\pm 0.3\,\% \pm 2$ V. [CAE00, CAE98]

B.1.4 Scintillator material

The two main types of plastic scintillator used for the FLASH BLMs are polyvinyltoluene-based NE-110/BC-412 and polystyrene-based SCSN-81T, both with peak emission intensity in the blue wavelength range. Table B.1 lists the main properties of these materials.

Property	NE-110/BC-412	SCSN-81T
base polymer	polyvinyltoluene	polystyrene
light output (% anthracene)	60	50
light output (eV/photon)	100	120
wavelength of max. emission (nm)	434	418
attenuation length (cm)	210	110–250
rise time (ns)	1.0	0.9
decay time (ns)	3.3	3.5

Table B.1: Main properties of plastic scintillator materials used for the FLASH beam loss monitors. Data from [Sai08, DET06].

B.1.5 Scintillator shapes

Three typical scintillator shapes are used for beam loss monitors at FLASH: a Y-shaped flat panel, a small block, and a long strip. The active volumes can simply be calculated from the geometries depicted in Fig. B.1:

$$\begin{aligned} V_{\mathrm{Y}} &= 794\,\mathrm{cm}^3 && \text{Y panel} \\ V_{\mathrm{sb}} &= 120\,\mathrm{cm}^3 && \text{small block} \\ V_{\mathrm{sb}} &= 975\,\mathrm{cm}^3 && \text{long strip} \end{aligned}$$

Because the light attenuation length Λ of the NE-110/BC-412 scintillator is about 210 cm [Sai08], it needs to be taken into account for the light output of the "long strip" type which has a length of $L = 195\,\mathrm{cm}$. For ease of comparison, an effective volume will be calculated:

$$\begin{aligned} V_{\mathrm{sb}}^{\mathrm{eff}} &= V_{\mathrm{sb}} \cdot \frac{1}{L} \int_0^L \exp\Big(-\frac{x}{\Lambda}\Big)\,\mathrm{d}x \\ &= V_{\mathrm{sb}} \cdot \frac{\Lambda}{L}\left[1 - \exp\left(-\frac{L}{\Lambda}\right)\right] \\ &= 0.65 V_{\mathrm{sb}} \\ &= 635\,\mathrm{cm}^3 \end{aligned}$$

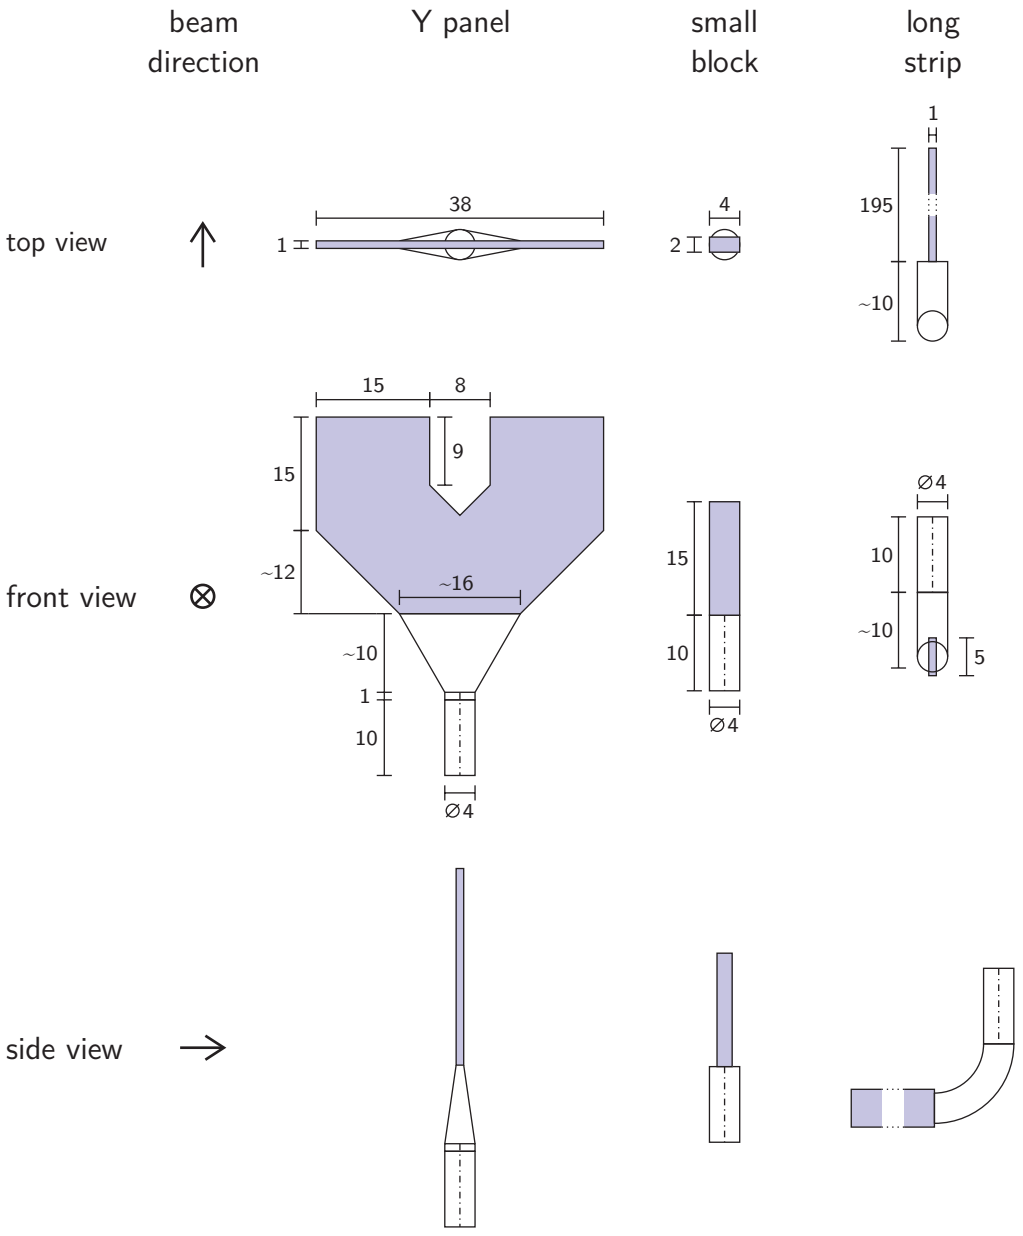

Figure B.1: Typical BLM scintillator shapes used at FLASH. Material: scintillator NE-110/SCSN-81T (blue), light guides plexiglass. Scale 1:10, all dimensions in cm.

B.1.6 Electrical calibration data

For all BLMs used at FLASH, the electrical calibration procedure described in section 7.2 is carried out in irregular intervals. The data given in Tab. B.2 and B.3 reflect the status of the BLM system as of December 1, 2008. With respect to this date, no calibration is more than one month old.

Empty fields for the single and multi bunch alarm parameters indicate that the respective alarm could only be triggered for a single threshold setting and could therefore not be calibrated. All errors with the exception of those given in the last line are the statistical errors from linear regression of the data points. The average values in the last line give the arithmetic mean and the unreduced standard deviation of the respective calibration parameter; these numbers are used as a fallback when no recent calibration data are available for a specific channel.

BLM name	y_0	$\frac{\mathrm{d}y}{\mathrm{d}x_{\mathrm{offs\_adc}}}$ $\cdot 10^{-3}$	$\frac{\mathrm{d}y}{\mathrm{d}x_{\mathrm{offs\_in}}}$ $\cdot 10^{-3}$
1GUN	8145.74 ± 0.03	-3649.80 ± 0.20	-1131.32 ± 0.20
17ACC1	8170.66 ± 0.04	-3608.30 ± 0.27	-1113.60 ± 0.31
1BC2	8165.84 ± 0.03	-3586.29 ± 0.22	-1111.23 ± 0.23
2BC2	8209.88 ± 0.02	-3534.31 ± 0.16	-1101.29 ± 0.21
3BC2	8181.22 ± 0.03	-3587.80 ± 0.20	-1117.43 ± 0.35
4BC2	8253.14 ± 0.07	-1781.11 ± 0.47	-544.21 ± 0.23
4DBC2	8126.15 ± 0.04	-3613.74 ± 0.30	-1121.17 ± 0.39
8DBC2	8147.05 ± 0.03	-3504.76 ± 0.20	-1085.96 ± 0.21
12DBC2	8256.62 ± 0.08	-1859.92 ± 0.52	-576.13 ± 0.30
2UBC3	8268.16 ± 0.03	-1869.66 ± 0.18	-585.81 ± 0.20
1BC3	8233.58 ± 0.01	-1826.54 ± 0.09	-576.62 ± 0.17
5BC3	8269.05 ± 0.03	-1852.61 ± 0.22	-572.11 ± 0.22
11BC3	8262.63 ± 0.04	-1835.65 ± 0.24	-576.74 ± 0.17
14BC3	8236.62 ± 0.06	-1863.21 ± 0.39	-585.34 ± 0.19
2DBC3	8194.17 ± 2.79	-1578.13 ± 18.93	-238.74 ± 13.19
2ACC7	8104.71 ± 0.03	-3644.72 ± 0.23	-1137.12 ± 0.15
10ACC7	8184.28 ± 0.03	-3562.13 ± 0.21	-1098.85 ± 0.20
2.1BYP	8122.30 ± 0.04	-3540.85 ± 0.28	-1095.59 ± 0.17
2.2BYP	8111.42 ± 0.03	-3492.71 ± 0.22	-1102.97 ± 0.24
15.1BYP	8106.49 ± 0.04	-3523.74 ± 0.24	-1100.20 ± 0.26
15.2BYP	8157.47 ± 0.04	-3584.94 ± 0.26	-1110.64 ± 0.20
36BYP	8186.63 ± 0.03	-3648.29 ± 0.22	-1131.80 ± 0.18
59BYP	8185.91 ± 0.03	-3606.74 ± 0.22	-1124.81 ± 0.20
74BYP	8155.07 ± 0.03	-3587.99 ± 0.23	-1113.51 ± 0.29
92BYP	8161.35 ± 0.03	-3611.42 ± 0.19	-1123.56 ± 0.11
2.1TCOL	8141.11 ± 0.03	-3578.12 ± 0.20	-1126.15 ± 0.22
2.2TCOL	8160.01 ± 0.03	-3608.72 ± 0.21	-1115.61 ± 0.21
8.1TCOL	8167.25 ± 0.04	-3556.29 ± 0.28	-1100.89 ± 0.30
8.2TCOL	8115.72 ± 0.04	-3499.91 ± 0.25	-1098.67 ± 0.25
2.1ECOL	8107.76 ± 0.07	-3560.55 ± 0.47	-1103.16 ± 0.24
2.2ECOL	8225.99 ± 0.05	-3561.54 ± 0.31	-1108.93 ± 0.23
3.1ECOL	8112.35 ± 0.04	-3563.43 ± 0.24	-1114.08 ± 0.21
3.2ECOL	8119.81 ± 0.03	-3526.67 ± 0.23	-1097.27 ± 0.17
7MATCH	8193.52 ± 0.04	-3605.09 ± 0.25	-1112.97 ± 0.29

Table B.2: Electrical calibration data of the FLASH BLMs: ADC calibration

BLM name	y_0	$\frac{\mathrm{d}y}{\mathrm{d}x_{\mathrm{offs\_adc}}}$ $\cdot 10^{-3}$	$\frac{\mathrm{d}y}{\mathrm{d}x_{\mathrm{offs\_in}}}$ $\cdot 10^{-3}$
5SUND2	8155.10 ± 0.03	-3548.29 ± 0.22	-1114.17 ± 0.20
3SEED	8239.08 ± 0.03	-3660.60 ± 0.24	-1132.97 ± 0.19
17SEED	8156.59 ± 2.21	-3532.56 ± 14.98	-1112.88 ± 0.34
2EXP	8226.87 ± 0.05	-3597.94 ± 0.31	-1111.94 ± 0.28
10EXP	8160.07 ± 0.03	-3630.38 ± 0.22	-1131.38 ± 0.21
1.1DUMP	8124.70 ± 0.02	-3583.66 ± 0.14	-1114.43 ± 0.07
1.2DUMP	8093.46 ± 0.04	-3583.56 ± 0.30	-1118.08 ± 0.24
6DUMP	8115.28 ± 0.02	-3632.50 ± 0.16	-1125.42 ± 0.10
9DUMP	8090.51 ± 0.03	-3527.76 ± 0.22	-1103.13 ± 0.18
13.1DUMP	8141.31 ± 0.03	-3597.00 ± 0.17	-1130.71 ± 0.11
13.2DUMP	8055.70 ± 7.93	-3079.69 ± 53.78	186.02 ± 66.55
22L.SEED	8148.25 ± 0.04	-3585.84 ± 0.25	-1114.69 ± 0.35
22R.SEED	8128.40 ± 0.04	-3534.81 ± 0.29	-1094.26 ± 0.31
1L.UND1	8127.85 ± 0.04	-3612.48 ± 0.29	-1114.73 ± 0.32
1R.UND1	8192.87 ± 0.04	-3649.52 ± 0.30	-1140.01 ± 0.31
3L.UND1	8087.92 ± 0.05	-3633.57 ± 0.31	-1137.97 ± 0.39
3R.UND1	8225.17 ± 0.04	-3522.91 ± 0.26	-1096.24 ± 0.28
5L.UND1	8143.41 ± 0.03	-3494.67 ± 0.21	-1093.06 ± 0.37
5R.UND1	8155.68 ± 0.03	-3618.28 ± 0.19	-1134.41 ± 0.17
1L.UND2	8154.93 ± 0.03	-3640.70 ± 0.22	-1126.54 ± 0.42
1R.UND2	8150.47 ± 0.04	-3685.28 ± 0.24	-1135.33 ± 0.44
3L.UND2	8107.22 ± 0.03	-3587.15 ± 0.22	-1121.91 ± 0.41
3R.UND2	8173.43 ± 0.07	-3567.39 ± 0.45	-1096.47 ± 0.35
5L.UND2	8115.07 ± 0.04	-3554.03 ± 0.28	-1106.62 ± 0.31
5R.UND2	8208.41 ± 0.04	-3580.37 ± 0.29	-1117.90 ± 0.26
1L.UND3	8111.55 ± 0.04	-3639.12 ± 0.25	-1135.81 ± 0.22
1R.UND3	8153.88 ± 0.03	-3499.69 ± 0.22	-1090.17 ± 0.28
3L.UND3	8184.93 ± 0.04	-3591.62 ± 0.29	-1112.81 ± 0.31
3R.UND3	8136.75 ± 0.01	-3663.36 ± 0.10	-1152.75 ± 0.33
5L.UND3	8209.35 ± 0.05	-3634.23 ± 0.31	-1122.31 ± 0.32
5R.UND3	8129.64 ± 0.03	-3553.69 ± 0.22	-1118.54 ± 0.18
1L.UND4	8129.54 ± 0.05	-3640.18 ± 0.36	-1143.78 ± 0.45
1R.UND4	8157.44 ± 0.05	-3524.23 ± 0.33	-1101.61 ± 0.29
3L.UND4	8127.52 ± 0.04	-3544.34 ± 0.27	-1104.52 ± 0.36
3R.UND4	8113.28 ± 0.03	-3674.88 ± 0.18	-1152.50 ± 0.43
5L.UND4	8181.77 ± 0.03	-3579.44 ± 0.19	-1108.36 ± 0.27
5R.UND4	8118.40 ± 0.03	-3633.07 ± 0.20	-1124.28 ± 0.21
1L.UND5	8194.61 ± 0.04	-3588.43 ± 0.26	-1114.92 ± 0.30
1R.UND5	8152.62 ± 0.05	-3560.28 ± 0.33	-1106.37 ± 0.43
3L.UND5	8214.17 ± 0.04	-3639.75 ± 0.30	-1126.25 ± 0.17
3R.UND5	8102.62 ± 0.03	-3669.72 ± 0.24	-1157.41 ± 0.29
5L.UND5	8103.07 ± 0.04	-3540.26 ± 0.25	-1095.21 ± 0.25
5R.UND5	8148.44 ± 0.04	-3553.26 ± 0.24	-1114.64 ± 0.31
1L.UND6	8121.00 ± 0.04	-3538.70 ± 0.28	-1121.02 ± 0.26
1R.UND6	8159.12 ± 0.03	-3553.02 ± 0.19	-1101.95 ± 0.44
3L.UND6	8177.73 ± 0.02	-3631.02 ± 0.15	-1129.58 ± 0.13
3R.UND6	8162.54 ± 0.03	-3747.99 ± 0.20	-1168.03 ± 0.10
5L.UND6	8129.11 ± 0.04	-3619.85 ± 0.30	-1129.20 ± 0.33
5R.UND6	8153.75 ± 0.06	-3566.18 ± 0.43	-1106.03 ± 0.14
average	8159 ± 46	-3410 ± 532	-1045 ± 224

Table B.2: Electrical calibration data of the FLASH BLMs: ADC calibration (continued)

BLM name	τ^0_{single}	$\frac{\mathrm{d}\tau_{\text{single}}}{\mathrm{d}x_{\text{offs\_in}}}$ $\cdot 10^{-3}$	τ^0_{multi}	$\frac{\mathrm{d}\tau_{\text{multi}}}{\mathrm{d}x_{\text{offs\_in}}}$ $\cdot 10^{-3}$	τ^0_{int}	$\frac{\mathrm{d}\tau_{\text{int}}}{\mathrm{d}x_{\text{offs\_in}}}$ $\cdot 10^{-3}$
1GUN			-0.96 ± 0.02	8.31 ± 0.13	-181.1 ± 0.3	1242 ± 2
17ACC1	-1.001 ± 0.016	8.35 ± 0.09	-0.96 ± 0.02	8.44 ± 0.12	-169.0 ± 0.3	1159 ± 2
1BC2					-186.0 ± 0.3	1064 ± 2
2BC2					-254.7 ± 0.6	1296 ± 3
3BC2	-1.004 ± 0.007	8.51 ± 0.04	-1.22 ± 0.01	9.46 ± 0.07	-194.6 ± 0.4	1324 ± 2
4BC2	-0.793 ± 0.030	8.21 ± 0.18	-0.88 ± 0.04	8.78 ± 0.22	-174.7 ± 0.2	1233 ± 1
4DBC2			-0.97 ± 0.02	8.97 ± 0.09	-190.5 ± 0.4	1316 ± 2
8DBC2	-0.824 ± 0.028	9.30 ± 0.19	-1.03 ± 0.04	9.07 ± 0.20	-169.3 ± 0.3	1352 ± 2
12DBC2	-0.673 ± 0.021	8.59 ± 0.14	-0.56 ± 0.02	8.33 ± 0.15	-159.2 ± 0.2	1259 ± 1
2UBC3					-274.0 ± 1.2	1301 ± 5
1BC3	-0.842 ± 0.027	8.50 ± 0.16	-0.95 ± 0.02	8.40 ± 0.11	-189.1 ± 0.4	1269 ± 2
5BC3	-1.004 ± 0.034	8.86 ± 0.19	-1.31 ± 0.05	9.28 ± 0.25	-237.1 ± 0.5	1269 ± 2
11BC3	-0.728 ± 0.031	8.84 ± 0.20	-0.70 ± 0.04	8.64 ± 0.24	-168.2 ± 0.3	1237 ± 2
14BC3	-0.975 ± 0.024	8.19 ± 0.13	-0.60 ± 0.03	8.58 ± 0.19	-162.8 ± 0.3	1239 ± 1
2DBC3	1.072 ± 0.026	9.93 ± 0.21	1.22 ± 0.04	9.95 ± 0.39	-171.4 ± 0.3	1264 ± 1
2ACC7	-0.498 ± 0.014	9.85 ± 0.12	-0.52 ± 0.01	8.67 ± 0.08	-58.6 ± 0.2	1269 ± 1
10ACC7					-282.0 ± 1.5	1280 ± 6
2.1BYP	-0.468 ± 0.023	8.19 ± 0.17	-0.81 ± 0.03	8.37 ± 0.15	-194.1 ± 0.4	1256 ± 2
2.2BYP	-0.274 ± 0.014	8.66 ± 0.13	-0.68 ± 0.02	8.59 ± 0.16	-158.6 ± 0.3	1283 ± 1
15.1BYP	-0.592 ± 0.032	8.54 ± 0.23	-0.63 ± 0.03	9.04 ± 0.23	-122.0 ± 0.2	1246 ± 1
15.2BYP	-0.932 ± 0.028	8.75 ± 0.16	-0.96 ± 0.02	8.52 ± 0.13	-185.4 ± 0.3	1266 ± 2
36BYP	-1.049 ± 0.029	8.55 ± 0.15	-1.19 ± 0.04	8.88 ± 0.20	-202.2 ± 0.4	1258 ± 2
59BYP	-0.300 ± 0.017	8.50 ± 0.16	-0.05 ± 0.02	8.62 ± 0.10	-151.9 ± 0.3	1250 ± 1
74BYP	0.409 ± 0.024	7.42 ± 0.15	0.05 ± 0.02	8.75 ± 0.15	-195.5 ± 0.3	1259 ± 2
92BYP	-0.116 ± 0.009	8.52 ± 0.06	-0.20 ± 0.01	8.24 ± 0.13	-141.2 ± 0.3	1236 ± 1
2.1TCOL	-0.444 ± 0.008	8.66 ± 0.06	-0.59 ± 0.01	8.89 ± 0.06	-155.9 ± 0.3	1265 ± 1
2.2TCOL	-0.736 ± 0.015	8.53 ± 0.10	-0.72 ± 0.01	8.17 ± 0.09	-144.9 ± 0.3	1238 ± 1
8.1TCOL	-0.590 ± 0.013	8.49 ± 0.09	-0.46 ± 0.01	9.22 ± 0.11	-167.1 ± 0.3	1249 ± 1
8.2TCOL	-0.710 ± 0.010	8.26 ± 0.06	-0.76 ± 0.01	8.74 ± 0.06	-151.9 ± 0.3	1252 ± 1
2.1ECOL	-0.503 ± 0.011	8.45 ± 0.09	-0.85 ± 0.01	8.48 ± 0.08	-132.1 ± 0.3	1241 ± 1
2.2ECOL					-278.8 ± 1.1	1279 ± 4
3.1ECOL	-0.507 ± 0.007	8.66 ± 0.05	-0.71 ± 0.01	8.71 ± 0.08	-163.9 ± 0.3	1234 ± 2
3.2ECOL	-0.522 ± 0.011	9.26 ± 0.09	-0.70 ± 0.01	8.60 ± 0.09	-89.5 ± 0.2	1281 ± 1
7MATCH	-0.930 ± 0.007	8.48 ± 0.04	-1.09 ± 0.01	9.68 ± 0.05	-188.6 ± 0.3	1249 ± 2
5SUND2	-0.842 ± 0.010	8.88 ± 0.06	-0.72 ± 0.01	8.56 ± 0.07	-140.9 ± 0.3	1233 ± 1
3SEED			-0.95 ± 0.01	8.61 ± 0.03	-203.6 ± 0.4	1273 ± 2
17SEED	-0.970 ± 0.008	8.55 ± 0.04	-1.15 ± 0.01	8.58 ± 0.03	-181.1 ± 0.3	1272 ± 2
2EXP					-266.4 ± 0.6	1411 ± 3
10EXP	-0.784 ± 0.037	7.79 ± 0.21	-0.67 ± 0.03	7.98 ± 0.20	-201.2 ± 0.4	1262 ± 2
1.1DUMP	-0.218 ± 0.018	8.43 ± 0.18	-0.14 ± 0.02	8.86 ± 0.11	-130.6 ± 0.2	1232 ± 1
1.2DUMP	-0.057 ± 0.024	9.11 ± 0.16	0.02 ± 0.02	8.87 ± 0.12	-170.5 ± 0.3	1275 ± 2
6DUMP	0.807 ± 0.021	7.58 ± 0.19	0.85 ± 0.02	7.52 ± 0.19	-153.2 ± 0.3	1241 ± 2
9DUMP	-0.597 ± 0.019	8.39 ± 0.13	-1.19 ± 0.03	8.84 ± 0.16	-150.9 ± 0.3	1299 ± 1
13.1DUMP	0.574 ± 0.027	7.54 ± 0.19	0.72 ± 0.02	7.18 ± 0.18	-174.6 ± 0.3	1239 ± 1
13.2DUMP	9.600 ± 0.259	101.97 ± 6.94	10.00 ± 0.26	72.13 ± 4.94	-215.5 ± 1.5	1272 ± 7
22L.SEED	-1.215 ± 0.018	10.00 ± 0.10	-1.09 ± 0.02	8.76 ± 0.08	-168.0 ± 0.3	1250 ± 1
22R.SEED					-213.4 ± 0.5	1259 ± 2
1L.UND1	-0.936 ± 0.008	9.35 ± 0.05	-1.22 ± 0.01	9.43 ± 0.06	-150.8 ± 0.3	1245 ± 1
1R.UND1			-1.05 ± 0.01	8.53 ± 0.07	-182.0 ± 0.3	1277 ± 2

Table B.3: Electrical calibration data of the FLASH BLMs: Threshold calibration

BLM name	τ^0_{single}	$\frac{\mathrm{d}\tau_{\mathrm{single}}}{\mathrm{d}x_{\mathrm{offs\_in}}}$ $\cdot 10^{-3}$	τ^0_{multi}	$\frac{\mathrm{d}\tau_{\mathrm{multi}}}{\mathrm{d}x_{\mathrm{offs\_in}}}$ $\cdot 10^{-3}$	τ^0_{int}	$\frac{\mathrm{d}\tau_{\mathrm{int}}}{\mathrm{d}x_{\mathrm{offs\_in}}}$ $\cdot 10^{-3}$
3L.UND1	-0.712 ± 0.014	8.18 ± 0.09	-0.93 ± 0.01	8.25 ± 0.07	-149.4 ± 0.3	1248 ± 1
3R.UND1	-1.171 ± 0.005	8.52 ± 0.03			-228.0 ± 0.5	1325 ± 2
5L.UND1	-0.986 ± 0.010	8.21 ± 0.05	-1.03 ± 0.01	8.56 ± 0.06	-186.8 ± 0.4	1238 ± 2
5R.UND1	0.036 ± 0.024	8.67 ± 0.17	0.28 ± 0.03	7.90 ± 0.19	-185.6 ± 0.4	1298 ± 2
1L.UND2	-1.066 ± 0.013	8.34 ± 0.07	-0.82 ± 0.01	8.82 ± 0.09	-198.6 ± 0.3	1276 ± 2
1R.UND2	-0.854 ± 0.005	8.41 ± 0.03	-0.71 ± 0.01	8.26 ± 0.04	-145.5 ± 0.3	1224 ± 2
3L.UND2	-0.955 ± 0.008	8.45 ± 0.05	-0.60 ± 0.01	8.41 ± 0.06	-149.1 ± 0.3	1227 ± 1
3R.UND2	-0.909 ± 0.006	9.51 ± 0.04	-0.96 ± 0.01	8.52 ± 0.04	-183.9 ± 0.4	1281 ± 2
5L.UND2	-0.894 ± 0.055	9.46 ± 0.35	-0.49 ± 0.03	7.31 ± 0.20	-197.9 ± 0.4	1255 ± 2
5R.UND2	-0.503 ± 0.041	8.95 ± 0.33	-0.53 ± 0.03	7.24 ± 0.18	-163.2 ± 0.3	1256 ± 2
1L.UND3			-0.89 ± 0.03	8.50 ± 0.19	-170.0 ± 0.3	1268 ± 1
1R.UND3	-1.066 ± 0.006	8.38 ± 0.03			-185.5 ± 0.4	1282 ± 2
3L.UND3					-230.4 ± 0.5	1290 ± 2
3R.UND3	-0.294 ± 0.002	8.48 ± 0.02	-0.78 ± 0.01	8.50 ± 0.04	-140.1 ± 0.3	1236 ± 1
5L.UND3					-220.2 ± 0.4	1248 ± 2
5R.UND3	-0.479 ± 0.040	8.71 ± 0.32	-0.46 ± 0.03	8.38 ± 0.27	-182.7 ± 0.3	1259 ± 2
1L.UND4	-1.039 ± 0.030	8.30 ± 0.16			-187.7 ± 0.3	1256 ± 2
1R.UND4	-1.032 ± 0.005	8.47 ± 0.03	-1.10 ± 0.01	8.53 ± 0.03	-207.3 ± 0.5	1251 ± 2
3L.UND4	-0.282 ± 0.010	8.62 ± 0.10	-0.81 ± 0.02	8.52 ± 0.14	-171.4 ± 0.3	1259 ± 2
3R.UND4	-0.711 ± 0.004	8.45 ± 0.03	-0.74 ± 0.00	8.54 ± 0.03	-139.6 ± 0.3	1307 ± 2
5L.UND4					-206.8 ± 0.4	1266 ± 2
5R.UND4	-0.449 ± 0.029	7.74 ± 0.21	-0.64 ± 0.04	8.51 ± 0.25	-187.7 ± 0.3	1263 ± 2
1L.UND5					-210.5 ± 0.5	1282 ± 3
1R.UND5	-1.127 ± 0.005	8.35 ± 0.03	-1.15 ± 0.00	8.55 ± 0.02	-161.6 ± 0.3	1246 ± 2
3L.UND5			-1.45 ± 0.08	9.75 ± 0.44	-248.3 ± 0.6	1288 ± 3
3R.UND5	-0.782 ± 0.004	8.83 ± 0.03	-0.87 ± 0.01	8.67 ± 0.03	-59.7 ± 0.2	1299 ± 1
5L.UND5	-1.046 ± 0.054	9.83 ± 0.33	-0.85 ± 0.04	8.59 ± 0.27	-154.5 ± 0.3	1277 ± 2
5R.UND5	-0.775 ± 0.035	8.28 ± 0.21			-213.3 ± 0.4	1255 ± 2
1L.UND6	-0.263 ± 0.025	9.32 ± 0.17	-0.59 ± 0.04	9.26 ± 0.28	-133.9 ± 0.3	1240 ± 1
1R.UND6	-0.573 ± 0.004	8.44 ± 0.03	-0.61 ± 0.00	8.60 ± 0.03	-131.1 ± 0.3	1078 ± 1
3L.UND6	-0.390 ± 0.016	8.74 ± 0.13	0.15 ± 0.02	8.02 ± 0.13	-158.9 ± 0.3	1251 ± 1
3R.UND6			-0.87 ± 0.01	8.49 ± 0.04	-150.0 ± 0.3	1035 ± 2
5L.UND6	-0.948 ± 0.012	8.47 ± 0.06	-0.89 ± 0.01	8.48 ± 0.07	-185.1 ± 0.4	1267 ± 2
5R.UND6	-0.702 ± 0.025	7.37 ± 0.14	-0.62 ± 0.02	6.94 ± 0.13	-154.3 ± 0.3	1248 ± 2
average	-0.46 ± 1.35	10.03 ± 11.59	-0.52 ± 1.38	9.50 ± 7.73	-177 ± 41	1257 ± 50

Table B.3: Electrical calibration data of the FLASH BLMs: Threshold calibration (continued)

B.2 Superconducting cavities

B.2.1 Conversion between accelerating gradient and field amplitude

An important figure of merit for a cavity is the maximum energy ΔW it can transfer to a passing beam particle. In practice, this value is often normalized to the elementary charge and the length of the cavity L and quoted in the form of an *accelerating gradient:*

$$G = \frac{\Delta W}{eL} \qquad [G] = \mathrm{V/m} \tag{B.1}$$

The field amplitude can only be deduced from this value in a precise way if the field profile of the cavity is known. A good approximation, however, can be made by assuming a perfect N-cell cavity.

Because the length of each cell is half the wavelength of the rf wave, a standing electromagnetic wave builds up in the resonator:

$$E(z,t) = \hat{E}\sin(kz)\sin(\omega t)$$

The temporal and spatial periodicities are related by $\omega = kc$:

$$E(z,t) = \hat{E}\sin(kz)\sin(kct)$$

For a particle entering the first cell at $t = z = 0$ with a velocity very close to c, the time of flight is given by $t = z/c$, and the electric field becomes a function of only the longitudinal position:

$$E(z) = \hat{E}\sin^2(kz)$$

It is now straightforward to calculate the energy gain of this particle by integrating the electric field over the length of the cavity, i.e., over N half-cells:

$$\begin{aligned}
\Delta W &= e\int_0^{N\frac{\lambda}{2}} E(z)\mathrm{d}z \\
&= e\int_0^{N\frac{\lambda}{2}} \hat{E}\sin^2(kz)\mathrm{d}z \\
&= e\hat{E}\left[\frac{kz - \cos(kz)\sin(kz)}{2k}\right]_0^{N\frac{\lambda}{2}}
\end{aligned}$$

Inserting $k = 2\pi/\lambda$, this reduces to

$$\Delta W = \frac{N}{4}e\hat{E}\lambda.$$

For a 9-cell TESLA cavity with its rf frequency of 1.3 GHz, the following values are obtained:

$$\begin{aligned}
&\Delta W \approx e\hat{E}\cdot 0.52\,\mathrm{m} \\
\Leftrightarrow \quad &\hat{E} \approx \frac{\Delta W}{e}\cdot 1.93\,\mathrm{m}^{-1}.
\end{aligned}$$

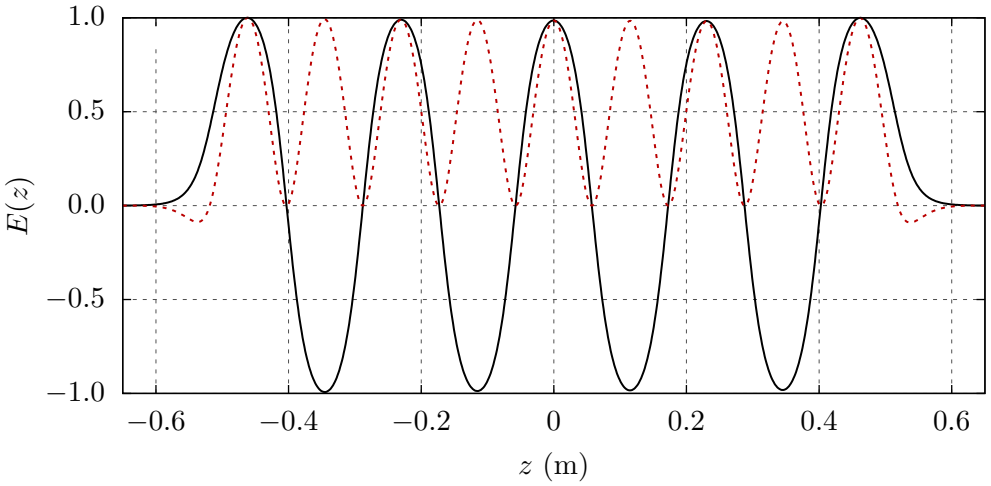

Figure B.2: Field profile of a TESLA cavity. The amplitude of the electric field has been normalized to unity. The dotted curve indicates the accelerating field experienced by a particle passing the cavity with a velocity of c under the optimal acceleration phase.

This calculation is of course oversimplified. For the field profile used in all Astra simulations, shown in Fig. B.2, a numerical integration yields the following conversion factors:

$$\begin{aligned} \Delta W &= e\hat{E} \cdot 0.5401\,\mathrm{m} \\ \Leftrightarrow \quad \hat{E} &= \frac{\Delta W}{e} \cdot 1.8515\,\mathrm{m}^{-1}. \end{aligned}$$

Conversion to an accelerating gradient according to (B.1) requires the choice of a value for the cavity length L. Choosing nine half-wavelengths of the 1.3 GHz rf wave yields $L = 1.0377\,\mathrm{m}$, and consequently

$$\begin{aligned} G &= \frac{\Delta W}{e \cdot 1.0377\,\mathrm{m}} \\ G &= 0.5205\hat{E} \\ \Leftrightarrow \quad \hat{E} &= 1.9214G. \end{aligned}$$

B.2.2 Distribution of cavity gradients at FLASH

A TTF- or XFEL-type cryomodule contains eight 9-cell niobium cavities. Table B.4 shows a typical distribution of accelerating gradients for a beam energy of about 980 MeV.

	Cavity gradients (MV/m)					
	ACC1	ACC2	ACC3	ACC4	ACC5	ACC6
C1	12.0	16.1	24.6	21.5	21.5	21.5
C2	12.0	12.7	24.6	21.5	21.5	21.5
C3	12.0	18.3	24.6	21.5	21.5	21.5
C4	12.0	18.9	24.6	21.5	21.5	21.5
C5	17.7	18.1	24.6	21.5	21.5	21.5
C6	17.7	15.1	24.6	21.5	21.5	21.5
C7	17.7	18.6	24.6	21.5	21.5	21.5
C8	17.7	19.2	24.6	21.5	21.5	21.5

Table B.4: Cavity gradients used in the simulation of the FLASH linac

B.3 Vacuum chambers

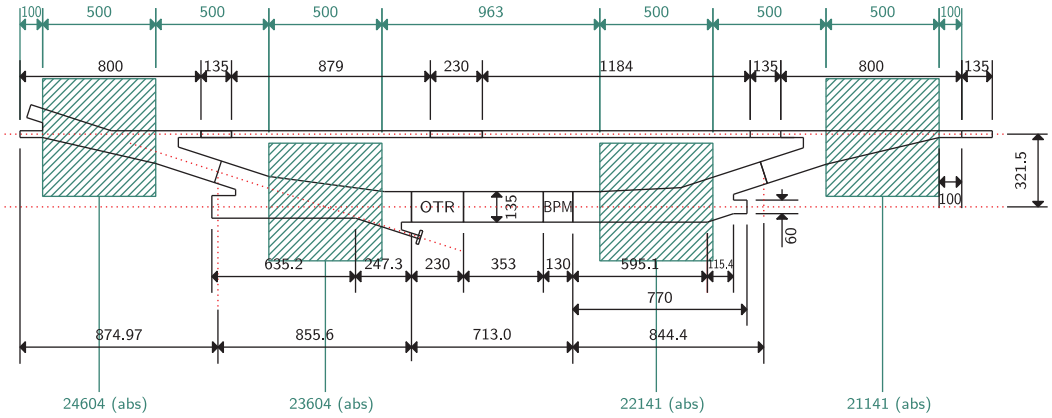

Figure B.3: Drawing of the BC2 vacuum chamber. The hatched areas indicate the position of the dipole magnet yokes. The beam enters the chamber from the right. All measures in mm.

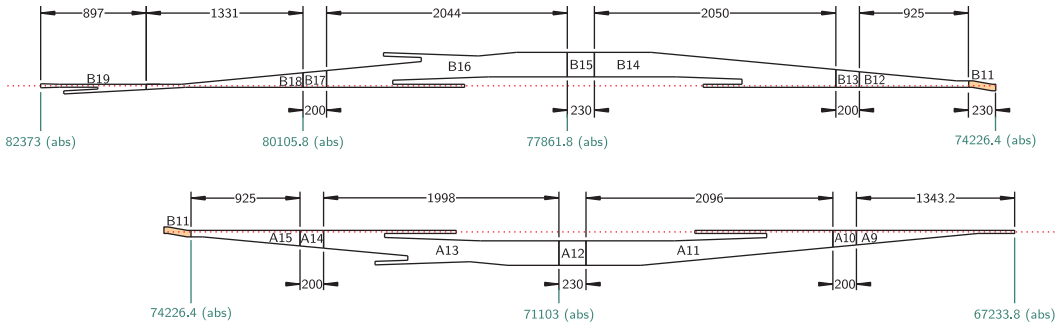

Figure B.4: Drawing of the BC3 vacuum chamber. The beam enters the chamber from the right. All measures in mm.

C Linear dispersion generated by a dipole

In a magnetic dipole field B, an electron of momentum p follows a circular trajectory with a bending radius of

$$R = \frac{p}{eB}.$$

For a rectangular dipole magnet as sketched in Fig. C.1, the deflection angle can be calculated from the bending radius R and the magnet length L:

$$\alpha = \arcsin\frac{L}{R} \approx \frac{L}{R}$$

Combining the equations, the dependence of the deflection angle on the particle momentum is obtained:

$$\alpha = \arcsin\frac{eBL}{p} \approx \frac{eBL}{p} \tag{C.1}$$

Thus, the bending angle is inversely proportional to the particle momentum.

It is straightforward to calculate the total transverse deviation from the axis of incidence at a given distance Δs from the dipole exit:

$$D = R(1-\cos\alpha) + \Delta s \sin\alpha \tag{C.2}$$

The first term describes the offset accumulated within the dipole, the second term is the offset accumulated in the adjacent drift space. If (C.2) is evaluated for

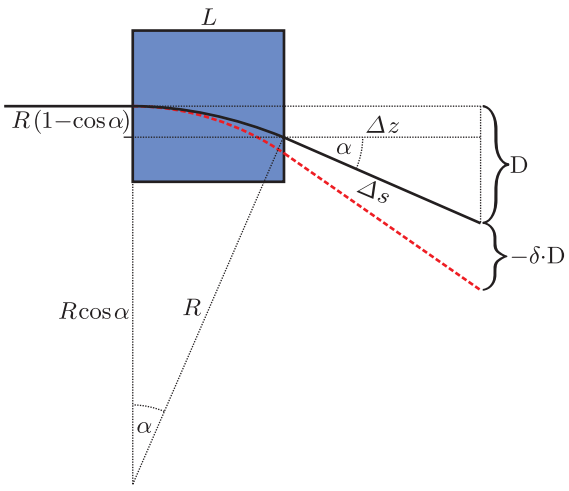

Figure C.1: Dispersion generated by a dipole magnet. The black solid curve indicates the design trajectory of a particle with momentum p_0, the red dotted curve is the trajectory of a particle with momentum $p_0 + \delta p_0$.

the design bending angle (i.e. for a particle of design momentum p_0), D is also the *linear dispersion* generated by the dipole. Considering a particle with a small momentum difference $\delta = (p - p_0)/p$, its transverse offset from the design orbit can then be approximated as $-\delta D$. This is of course only valid if the incident beam is free of dispersion.

Dispersion within the dipole

If ϕ_d is the deflection angle at an intermediate point of the particle trajectory within the dipole, the dispersion is given by

$$D(\phi_\mathrm{d}) = R(1 - \cos\phi_\mathrm{d}).$$

By introducing s_d as the arc length within the dipole, the equation can be rewritten as

$$D(s_\mathrm{d}) = R\left(1 - \cos\frac{s_\mathrm{d}}{R}\right). \tag{C.3}$$

At the exit of the dipole area, $s_\mathrm{d} = \alpha R$, and (C.3) is equivalent to (C.2) with $\Delta s = 0$.

Longitudinal parametrization

Equations (C.2) and (C.3) give the linear dispersion as a function of the drift path length s_d and Δs. It is sometimes useful to reparametrize them in longitudinal coordinates; the projections of the drift path lengths to the axis of incidence to the dipole can be expressed as follows:

$$z_\mathrm{d} = R\sin\frac{s_\mathrm{d}}{R}$$
$$\Delta z = \Delta s\cos\alpha$$

Hence the dispersion is given by

$$D = \begin{cases} 0 & \text{in front of the magnet,} \\ R\left(1 - \cos\left(\arcsin\dfrac{z_\mathrm{d}}{R}\right)\right) & \text{within the magnet,} \\ R(1-\cos\alpha) + \Delta z\tan\alpha & \text{behind the magnet,} \end{cases}$$

or, using the length L of the magnet,

$$D = \begin{cases} 0 & \text{in front of the magnet,} \\ L\cdot\dfrac{1-\cos\left(\arcsin\frac{z_\mathrm{d}}{R}\right)}{\sin\alpha} & \text{within the magnet,} \\ L\cdot\dfrac{1-\cos\alpha}{\sin\alpha} + \Delta z\tan\alpha & \text{behind the magnet.} \end{cases} \tag{C.4}$$

Symbols and constants

The numerical values of fundamental constants specified here are the "2006 CODATA recommended values", unless stated otherwise. Absolute uncertainties are given in concise form: 12.345(67) reads as 12.345 ± 0.067.

Symbol	Description	Value	Units
α	Fine-structure constant	$7.297\,352\,5376(50) \cdot 10^{-3}$	
A	Atomic mass number		
A^*	Molar mass, average molecular weight		kg/mol
β	Particle speed in units of c		
c	Speed of light in vacuum	$2.997\,924\,58 \cdot 10^{8}$	m/s
e	Elementary charge	$1.602\,176\,487(40) \cdot 10^{-19}$	C
γ	Relativistic Lorentz factor: $\gamma = 1/\sqrt{1-\beta^2}$		
h	Planck constant	$6.626\,068\,96(33) \cdot 10^{-34}$	J s
L_{rad}	Radiation length		m
m_e	Electron mass	$9.109\,382\,15(45) \cdot 10^{-31}$	kg
N_a	Avogadro constant	$6.022\,141\,79(30) \cdot 10^{23}$	mol^{-1}
ρ	Material density		$\mathrm{kg/m}^3$
r_e	Classical electron radius	$2.817\,940\,2894(58) \cdot 10^{-15}$	m
X_0	Density-normalized radiation length		$\mathrm{kg/m}^2$
Z	Atomic number		

References

[Ack07] W. Ackermann, et al., "Operation of a free-electron laser from the extreme ultraviolet to the water window." *Nat. Photonics* 1, 336–342, 2007.

[Agh07] A. Aghababyan, et al., *The European X-Ray Free-Electron Laser technical design report*. Report DESY 2006-097, DESY, Hamburg, Germany, 2007.

[Ago03] S. Agostinelli, et al., "Geant4—A simulation toolkit." *Nucl. Instr. and Meth. A* 506, 250–303, 2003.

[All06] J. Allison, et al., "Geant4 developments and applications." *IEEE T. Nucl. Sci.* 53(1), 270–278, 2006.

[Amd67] G. Amdahl, "Validity of the single processor approach to achieving large-scale computing capabilities." *Proc. AFIPS'67*, 30, pp. 483–485, Atlantic City, USA, 1967.

[And00] J. Andruszkow, et al., "First observation of self-amplified spontaneous emission in a free-electron laser at 109 nm wavelength." *Phys. Rev. Lett.* 85(18), 3825–3829, 2000.

[And01] J. Andruszkow, P. Jurkiewicz, and F. Tonisch, *8-channel Fast-ADC with 14 bit resolution—reference manual, version 1.1*. DESY, Zeuthen, Germany, 2001.

[And07] S. Anderson, et al., "Fast neutron radioactivity and damage studies on materials." *Proc. PAC'07*, pp. 581–583, Albuquerque, USA, 2007.

[And49] P. A. Anderson, "The work function of copper." *Phys. Rev.* 76(3), 388–390, 1949.

[Art02] J. Arthur, et al., *LCLS conceptual design report*. Report SLAC-R-593, SLAC, Menlo Park, USA, 2002.

[Art08] V. Artisyuk, et al., "Radiation dose as a barrier against proliferation for advanced fuel compositions." *J. Nucl. Sci. Technol.* 45, 1009–1015, 2008.

[Atk89] K. Atkinson, *An introduction to numerical analysis*. 2nd edn., Wiley, 1989.

[Ayv06] V. Ayvazyan, et al., "First operation of a free-electron laser generating GW power radiation at 32 nm wavelength." *Eur. Phys. J. D* 37, 297–303, 2006.

[Bae03] J. Bähr, et al., *Behavior of the TTF2 rf gun with long pulses and high repetition rates*. Report TESLA 2003-33, DESY, Hamburg, Germany, 2003.

[Bal03] V. Balandin, et al., *Studies of the collimator system for the TTF phase 2*. Report TESLA 2003-17, DESY, Hamburg, Germany, 2003.

[Bal08] V. Balandin and N. Golubeva, *Matlab functions for calculations of the linear beam optics of FLASH linac, version 1.4*. DESY, Hamburg, Germany, 2008.

[Bal08a] K. Balewski et al., "Status of Petra III." *Proc. EPAC'08*, pp. 1977–1979, Genoa, Italy, 2008.

[Ben07] S. Benson, et al., "The 4th generation light source at Jefferson Lab." *Nucl. Instr. and Meth. A* 582, 14–17, 2007.

[Ber05] M. Berger, et al., *ESTAR, PSTAR, and ASTAR: Computer programs for calculating stopping-power and range tables for electrons, protons, and helium ions (version 1.2.3)*. Online database [2009-02-24], National Institute of Standards and Technology, Gaithersburg, USA, 2005, URL http://physics.nist.gov/Star.

[Ber63] C. Bernardini, et al., "Lifetime and beam size in a storage ring." *Phys. Rev. Lett.* 10, 407–409, 1963.

[Ber64] M. J. Berger and S. M. Seltzer, *Tables of energy losses and ranges of electrons and positrons*. Report NASA-SP-3012, National Aeronautics and Space Administration, Washington, USA, 1964.

[Ber75] B. L. Berman and S. C. Fultz, "Measurements of the giant dipole resonance with monoenergetic photons." *Rev. Mod. Phys.* 47, 713–761, 1975.

[Bet30] H. Bethe, "Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie." *Ann. Phys.* 397(3), 325–400, 1930.

[Blo33] F. Bloch, "Zur Bremsung rasch bewegter Teilchen beim Durchgang durch Materie." *Ann. Phys.* 408(3), 285–320, 1933.

[Boh03] I. Bohnet, et al., "Dark current measurements at the PITZ rf gun." *Proc. DIPAC'03*, pp. 242–244, 2003.

[Bon84] R. Bonifacio, C. Pellegrini, and L. M. Narducci, "Collective instabilities and high-gain regime free electron laser." *Opt. Commun.* 50(6), 373–378, 1984.

[Bra07] A. Brandt, *Development of a finite state machine for the automated operation of the LLRF control at FLASH*. Ph.D. thesis, University of Hamburg, Hamburg, Germany, 2007.

[Bro01] P. Brown, et al., "Ultimate performance of the LEP RF system." *Proc. PAC'01*, pp. 1059–1061, Chicago, USA, 2001.

[Bro04] D. Broggio, et al., "Degradation of the scintillation yield of anthracene under high-fluence carbon ion beams." *Radiat. Meas.* 39(3), 283–287, 2004.

[CAE00] CAEN S.p.A., Viareggio, Italy, *SY527 universal multichannel power supply system technical information manual*, 2000.

[CAE98] CAEN S.p.A., Viareggio, Italy, *SY527 high voltage boards technical information manual rev. n. 0*, 1998.

[Cam05] P. Cameron, *Differential current measurement in the BNL energy recovery linac test facility*. Report C-A/AP/#203, Brookhaven National Laboratory, Upton, USA, 2005.

[CAR04] A. Peters, H. Schmickler, and K. Wittenburg (eds.), *Proc. Workshop on "DC Current Transformers and Beam-Lifetime Evaluations"*, CARE-Note-2004-023-HHH, Lyon, France: Coordinated Accelerator Research in Europe, CARE-N3 networking for high energy, high intensity hadron beams, 2004.

[Cas03] P. Castro, *Beam trajectory calculations in bunch compressors of TTF2*. Technical note 2003-01, DESY, Hamburg, Germany, 2003.

[Cho07] A. Cho et al., "Large Hadron Collider: Having a blast, wish you were here." *Science* 315, 1652–1656, 2007.

[CRC03] *CRC handbook of chemistry and physics*. 84th edn., Boca Raton, USA: CRC Press, 2003.

[Dav54] H. Davies, H. A. Bethe, and L. C. Maximon, "Theory of bremsstrahlung and pair production. II. Integral cross section for pair production." *Phys. Rev.* 93(4), 788–795, 1954.

[DET06] DETEC-Europe, Vannes, France, *Scintillating products*, 2006.

[Doh04] M. Dohlus, et al., "Start-to-end simulations of SASE FEL at the TESLA Test Facility, phase 1." *Nucl. Instr. and Meth. A* 530, 217–233, 2004.

[Doh08] M. Dohlus, J. Rossbach, and P. Schmüser, *Ultraviolet and soft X-ray free-electron lasers*. Springer, 2008.

[Eli76] L. R. Elias, et al., "Observation of stimulated emission of radiation by relativistic electrons in a spatially periodic transverse magnetic field." *Phys. Rev. Lett.* 36(13), 717–720, 1976.

[ET01] Electron Tubes Limited, Ruislip, UK, *38 mm photomultiplier 9972B series data sheet*, 2001.

[Fas03] A. Fassò, et al., "The physics models of FLUKA: status and recent developments." *Proc. CHEP'03*, La Jolla, USA, 2003. Paper MOMT005.

[Fei04] J. Feikes, et al., "The BESSY low alpha optics and the generation of coherent synchrotron radiation." C. Biscari and W. Chou (eds.), *ICFA Beam Dynamics Newsletter*, 35, pp. 82–95, International Committee for Future Accelerators, 2004.

[Fer05] A. Ferrari, et al., *Fluka: a multi-particle transport code*. CERN, 2005.

[Fis67] M. Fishman and D. Reagan, "The SLAC long ion chamber system for machine protection." *IEEE T. Nucl. Sci.* 14(3), 1096–1098, 1967.

[Flo00] K. Flöttmann, *ASTRA—A Space Charge Tracking Algorithm*. DESY, Hamburg, Germany, URL http://www.desy.de/~mpyflo/. Software and manual available online [2009-04-08].

[Flo03a] K. Flöttmann, et al., "Tolerances of TTF-2 first bunch compressor." *Proc. PAC'03*, pp. 965–967, Portland, USA, 2003.

[Fow28] R. H. Fowler and L. Nordheim, "Electron emission in intense electric fields." *Proc. R. Soc. A* 119(781), 173–181, 1928.

[Fro06a] L. Fröhlich, et al., "First operation of the FLASH machine protection system with long bunch trains." *Proc. LINAC'06*, pp. 262–264, Knoxville, USA, 2006.

[Fro07a] L. Fröhlich and S. Meykopff, "Improvements of the tracking code Astra for dark current studies at FLASH." *Proc. FEL'07*, 2007.

[Fro08] L. Fröhlich, *A new Astra for parallel computing*. DESY, Hamburg, Germany, URL http://tesla.desy.de/~lfroehli/astra/. Software and manual available online [2009-03-10].

[Goe03] P. Göttlicher, *Test of the prototype of the analog part "alarm-integrator for the TTF beam loss monitor"*. Internal document, group FEB, DESY, Hamburg, Germany, 2003.

[Goe04] P. Göttlicher, *Manual for a differential test pulse driver*. DESY, Hamburg, Germany, 2004.

[Goe05] W. Goettmann, et al., "Beam loss position monitor using Cerenkov radiation in optical fibers." *Proc. DIPAC'05*, pp. 301–303, Lyon, France, 2005.

[Gol95] S. V. Golovkin, et al., "Radiation damage studies on new liquid scintillators and liquid-core scintillating fibers." *Nucl. Instr. and Meth. A* 362, 283–291, 1995.

[Gor95] P. Gorodetzky, et al., "Quartz fiber calorimetry." *Nucl. Instr. and Meth. A* 361, 161–179, 1995.

[Hah00] U. Hahn, et al., "Design and performance of the vacuum chambers for the undulator of the VUV FEL at the TESLA test facility at DESY." *Nucl. Instr. and Meth. A* 445, 442–447, 2000.

[Ham04] A. Hamdi, M. Luong, and M. Werner, "VHDL design and simulation of a fast beam loss interlock for TTF2." *Proc. EPAC'04*, pp. 2020–2022, Lucerne, Switzerland, 2004.

[Ham04a] A. Hamdi, *Conception et réalisation d'un système numérique rapide pour la protection d'un accélérateur linéaire*. Diploma thesis, Conservatoire National des Arts et Métiers, Paris, France, 2004. Also published as report DAPNIA-04-156.

[Ham07] A. Hamdi, et al., "Toroid Protection System for FLASH." *Proc. DIPAC'07*, Venice, Italy, 2007.

[Ham99] M. M. Hamada, et al., "Radiation damage studies on the optical and mechanical properties of plastic scintillators." *Nucl. Instr. and Meth. A* 422, 148–154, 1999.

[Han05] J.-H. Han, *Dynamics of electron beam and dark current in photocathode RF guns*. Ph.D. thesis, University of Hamburg, Hamburg, Germany, 2005.

[Hen03] H. Henschel, et al., "Optical fibre dosimeter for SASE FEL undulators." *Proc. DIPAC'03*, pp. 248–250, Mainz, Germany, 2003.

[Hen05] S. Henderson, *SNS parameters list*. Report SNS-100000000-PL0001-R13, Spallation Neutron Source, 2005.

[Hof01] G. H. Hoffstätter, "The Luminosity Upgrade of HERA." K. Hirata, et al. (eds.), *ICFA Beam Dynamics Newsletter*, 24, pp. 8–18, International Committee for Future Accelerators, 2001.

[Hof08] G. Hoffstätter. Personal communication.

[ICR93] *ICRU publication 51—Quantities and units in radiation protection dosimetry*. ICRU Publications, 1993.

[ICR97] *ICRP publication 74—Conversion coefficients for use in radiological protection against external radiation*, *Annals of the ICRP*, vol. 26/3. Elsevier, 1997.

[ILC07] *International Linear Collider reference design report – Volume 1: Executive summary*. Report, ILC global design effort and world wide study, 2007.

[ITU96] *Electrical characteristics for balanced double-current interchange circuits operating at data signalling rates up to 10 Mbit/s*. ITU-T recommendation V.11/X.27, International Telecommunication Union, 1996.

[Jan00] E. Janata and M. Körfer, *Radiation detection by Cerenkov emission in optical fibers at TTF*. Report TESLA 2000-27, DESY, Hamburg, Germany, 2000.

[Jea79] B. Jean-Marie, V. Lepeltier, and D. L'Hote, "Systematic measurement of electron drift velocity and study of some properties of four gas mixtures: A–CH_4, A–C_2H_4, A–C_2H_4, A–C_3H_8." *Nucl. Instr. and Meth.* 159, 213–219, 1979.

[Jej88] A. Jejcic, et al., "Accelerator simulation on a parallel computer." *Proc. EPAC'88*, pp. 1178–1180, Rome, Italy, 1988.

[Jor08] K. Jordan. Personal communication.

[Ker40] D. W. Kerst, "Acceleration of electrons by magnetic induction." *Phys. Rev.* 58, 841, 1940.

[Ker41] D. W. Kerst and R. Serber, "Electronic orbits in the induction accelerator." *Phys. Rev.* 60, 53–58, 1941.

[Koc59] H. W. Koch and J. W. Motz, "Bremsstrahlung cross-section formulas and related data." *Rev. Mod. Phys.* 31(4), 920–956, 1959.

[Koc79] E. E. Koch, *Synchrotron radiation facilities at DESY, a status report*. Report DESY SR-79/35, DESY, Hamburg, Germany, 1979.

[Kon81] A. M. Kondratenko and E. L. Saldin, "Generation of coherent radiation by a beam of relativistic electrons in an undulator." *Zhurnal Tekhnicheskoi Fiziki* 51, 1633–1642, 1981.

[Kor07] I. Koren and C. M. Krishna, *Fault-tolerant systems*. Morgan Kaufmann, 2007.

[Kra07] D. Kramer, et al., "Secondary electron emission beam loss monitor for LHC." *Proc. DIPAC'07*, pp. 313–315, Venice, Italy, 2007.

[Kra97a] G. A. Krafft, et al., "Electron-beam diagnostics for Jefferson Lab's high power free electron laser." *Proc. PAC'97*, pp. 912–914, Vancouver, Canada, 1997.

[Lan05] K. Langa, et al., "Characterization of 1600 Hamamatsu 16-anode photomultipliers for the MINOS far detector." *Nucl. Instr. and Meth. A* 545, 852–871, 2005.

[Lat81] R. V. Latham, *High voltage vacuum insulation: The physical basis*. Academic Press, 1981.

[Leo94] W. R. Leo, *Techniques for nuclear and particle physics experiments: A how-to approach.* 2nd edn., Berlin: Springer, 1994.

[Leu08] A. Leuschner. Personal communication.

[Li05] Z. Li, et al., "Properties of plastic scintillators after irradiation." *Nucl. Instr. and Meth. A* 552, 449–455, 2005.

[Lil05] L. Lilje, "Experience with the TTF." *Proc. PAC'05*, pp. 1–3, Knoxville, USA, 2005.

[Lun89] H. B. Luna, et al., "Bremsstrahlung radiation effects in rare earth permanent magnets." *Nucl. Instr. and Meth. A* 285, 349–354, 1989.

[Lyn91] G. R. Lynch and O. I. Dahl, "Approximations to multiple Coulomb scattering." *Nucl. Instr. and Meth. B* 58, 6–10, 1991.

[Lyn97] D. W. Lynch, "Tantalus, a 240 MeV dedicated source of synchrotron radiation, 1968-1986." *J. Synchrotron Radiat.* 4, 334–343, 1997.

[Mad71] J. M. J. Madey, "Stimulated emission of bremsstrahlung in a periodic magnetic field." *J. Appl. Phys.* 42(5), 1906–1913, 1971.

[Mas99] M. Maslow, A. Lokhovitskii, and M. Schmitz, *Design of a beam dump for the TTF-FEL phase II (TTF2) project*. Report TESLA-FEL 99-06, DESY, Hamburg, Germany, 1999.

[May02] T. Mayer-Kuckuk, *Kernphysik*. 7th edn., Teubner, 2002.

[Mil00] S. V. Milton, et al., "Observation of self-amplified spontaneous emission and exponential growth at 530 nm." *Phys. Rev. Lett.* 85(5), 988–991, 2000.

[MIL83] *Military handbook—Survivability enhancement, aircraft, nuclear weapon threat, design and evaluation guidelines*. Handbook MIL-HDBK-273(AS), Department of the Navy, Naval Air Systems Command, Washington D.C., USA, 1983.

[MPI09] Argonne National Laboratory, Argonne, USA, *MPICH2, an implementation of the Message Passing Interface standard*. URL http://www.mcs.anl.gov/research/projects/mpich2/. Software package and manuals available online [2009-04-08].

[NAS96] *Single event effect criticality analysis*. Report 431-REF-000273, NASA, Goddard Space Flight Center, Greenbelt, USA, 1996.

[Nei06] G. R. Neil, et al., "The JLab high power ERL light source." *Nucl. Instr. and Meth. A* 557, 9–15, 2006.

[Nel90] W. R. Nelson and Y. Namito, "The EGS4 code system: Solution of gamma-ray and electron transport problems." *Proc. Int. Conf. on Supercomputing in Nuclear Applications*, Mito City, Japan, 1990. Also published as report SLAC-PUB-5193.

[Noe03] D. Nölle, et al., "The beam inhibit system for TTF II." *Proc. DIPAC'03*, pp. 62–64, Mainz, Germany, 2003.

[Obi09] F. Obier. Personal communication.

[Oku94] S. Okuda, K. Ohashi, and N. Kobayashi, "Effects of electron-beam and γ-ray irradiation on the magnetic flux of Nd–Fe–B and Sm–Co permanent magnets." *Nucl. Instr. and Meth. B* 94, 227–230, 1994.

[Pal98] D. T. Palmer, *The next generation photoinjector*. Ph.D. thesis, Stanford University, Stanford, USA, 1998. Also published as report SLAC-R-500.

[Pan63] W. K. H. Panofsky, *The use of a long co-axial ion chamber along the accelerator*. Report SLAC-TN-63-57, SLAC, 1963.

[Pat92] D. R. Patterson, "Preliminary design of the beam loss monitor system for the advanced photon source." *AIP conf. proc.*, vol. 281, pp. 194–203, Accelerator instrumentation fourth annual workshop, 1992.

[PDG06] W.-M. Yao, et al., "Review of particle physics." *J. Phys. G: Nucl. Partic.* 33(1), 1–1232, 2006.

[PDG08] C. Amsler, et al., "Review of particle physics." *Phys. Lett. B* 667, 1, 2008.

[Pel00] M. Pelliccioni, "Overview of fluence-to-effective dose and fluence-to-ambient dose equivalent conversion coefficients for high energy radiation calculated using the Fluka code." *Rad. Prot. Dos.* 88(4), 279–297, 2000.

[Pet61] W. W. Peterson and D. T. Brown, "Cyclic codes for error detection." *Proc. IRE*, vol. 49, pp. 228–235, 1961.

[Pfl03] J. Pflüger, et al., "Radiation Exposure and Magnetic Performance of the Undulator System for the VUV FEL at the TESLA Test Facility Phase-1 after three years of operation." *Nucl. Instr. and Meth. A* 507, 186–190, 2003.

[Pfl03a] J. Pflüger, et al., "Undulator system for the VUV FEL at the TESLA test facility phase-2." *Nucl. Instr. and Meth. A* 507, 228–233, 2003.

[Pfl04] J. Pflüger, et al., "Magnetic properties of SASE 400–900.", 2004. Unpublished measurement data.

[Pow73] R. A. Powell, et al., "Photoemission studies of cesium telluride." *Phys. Rev. B* 8(8), 3987–3995, 1973.

[Ros91] M. C. Ross, "Machine protection schemes for the SLC." *Proc. PAC'91*, vol. 3, pp. 1502–1504, San Francisco, USA, 1991.

[Ryb04] D. K. Rybka, et al., *Irradiation investigations for TESLA and X-FEL experiments at DESY*. Report TESLA 2004-12, DESY, 2004.

[Sai08] Saint-Gobain Crystals, Nemours, France, *BC-400, BC-404, BC-408, BC-412, BC-416 premium plastic scintillators*, 2008.

[Sch01] H. Schlarb, *Collimation system for the VUV free-electron laser at the TESLA Test Facility*. Ph.D. thesis, University of Hamburg, Hamburg, Germany, 2001. Also published as report DESY-THESIS-2001-055.

[Sch02] H. Schlarb, et al., "Expansion of the fast linac protection system for high duty cycle operation at the TESLA Test Facility." *Proc. EPAC'02*, pp. 1966–1968, Paris, France, 2002.

[Sch02b] R. Schurig and P. Michel, "A beam loss monitor with longitudinal resolution." U. Lehnert (ed.), *Radiation Source ELBE—Annual Report 2002*, chap. 1, pp. 13–14, Dresden, Germany: Forschungszentrum Rossendorf, 2002.

[Sch06a] R. Schmidt, et al., "Protection of the CERN Large Hadron Collider." *New J. Phys.* 8(290), 2006.

[Sch06b] S. Schreiber and M. Staack. Personal communication.

[Sch08] S. Schreiber. Personal communication.

[See02] J. Seeman, et al., "PEP-II status and future plans." *Proc. EPAC'02*, pp. 434–436, Paris, France, 2002.

[Sek08] J. Sekutowicz. Personal communication.

[Sel84] S. M. Seltzer and M. J. Berger, "Improved procedure for calculating the collision stopping power of elements and compounds for electrons and positrons." *Int. J. Appl. Radiat. Is.* 35(7), 665–676, 1984.

[Shi08] Y. Shimosaki, et al., "Development of beam loss monitor for the SPring-8 storage ring." *Proc. EPAC'08*, pp. 1284–1286, Genoa, Italy, 2008.

[Sie08] M. Siemens. Personal communication.

[Sku08] J. Skupin, et al., "Undulator demagnetization due to radiation losses at FLASH." *Proc. EPAC'08*, pp. 2308–2310, Genoa, Italy, 2008.

[Spe03] J. Spencer and J. Volk, "Permanent magnets for radiation damage studies." *Proc. PAC'03*, pp. 2180–2182, Portland, Oregon, 2003.

[Sri99] T. Srinivasan-Rao, et al., "Simulation, generation, and characterization of high brightness electron source at 1 GV/m gradient." *Proc. PAC'99*, pp. 75–77, New York, USA, 1999.

[Sta08] M. Staack. Personal communication.

[Ste04b] L. Steffen and M. Reinecke, *Beam loss monitor – DOOCS software*. DESY, Hamburg, Germany, 2004.

[Sto07] M. Stockner, et al., "Classification of the LHC BLM Ionization Chamber." *Proc. DIPAC'07*, pp. 328–330, Venice, Italy, 2007. Also published as report CERN-AB-2007-029.

[Swa79] W. P. Swanson, *Radiological safety aspects of the operation of electron linear accelerators*. International atomic energy agency, 1979.

[TTF02] *SASE FEL at the TESLA facility, phase 2*. Report TESLA-FEL 2002-01, DESY, Hamburg, Germany, 2002.

[TTF08] "Electronic logbook of the TESLA Test Facility.", 2008, URL https://ttfinfo.desy.de/TTFelog-secure/.

[Urs95] R. Ursic, et al., "CEBAF beam loss accounting." *Proc. PAC'95*, 1995.

[Wal93] R. P. Walker, "Interference effects in undulator and wiggler radiation sources." *Nucl. Instr. and Meth. A* 335, 328–337, 1993.

[Wan89] J. W. Wang, *RF properties of periodic accelerating structures for linear colliders*. Ph.D. thesis, Stanford University, Stanford, USA, 1989. Also published as report SLAC-R-339.

[Wei04] H. Weise, "The TESLA XFEL project." *Proc. EPAC'04*, pp. 11–15, Lucerne, Switzerland, 2004.

[Wer08] M. Werner. Personal communication.

[Wie03] H. Wiedemann, *Particle accelerator physics I, Basic principles and linear beam dynamics*. 2nd edn., Springer, 2003.

[Wie03a] H. Wiedemann, *Particle accelerator physics II, Nonlinear and higher-order beam dynamics*. 2nd edn., Springer, 2003.

[Wie77] H. Wiedemann, *Present and future colliding beam facilities at SLAC*. Report SLAC-PUB-1971, SLAC, Stanford, USA, 1977.

[Wit00] K. Wittenburg, "The PIN-diode beam loss monitor system at HERA." *AIP Conf. Proc.*, vol. 546, pp. 3–22, BIW'00, Boston, USA, 2000. Also published as report DESY-HERA-00-03.

[Wit02] R. L. Witkover and D. Gassner, "Design of an improved ion chamber for the SNS." *AIP Conf. Proc.*, vol. 648, pp. 337–344, BIW'02, Upton, USA, 2002. Also published as report BNL-69185.

[Wit94] K. Wittenburg, "Preservation of beam loss induced quenches, beam lifetime and beam loss measurements with the HERAp beam-loss-monitor system." *Nucl. Instr. and Meth. A* 345, 226–229, 1994.

[Yu92] L. H. Yu, et al., "Effect of wiggler errors on free-electron laser gain." *Phys. Rev. A* 45(2), 1163–1176, 1992.

[Zha03] X. Zhang, et al., "Experimental studies of beam-beam effects in the Tevatron." *Proc. PAC'03*, pp. 1757–1759, Portland, USA, 2003.

Acknowledgements

> He held up his hand and smiled a smile that stretched not merely from ear to ear, but seemed to extend some way beyond the mere confines of his face.
> "Thank you, ladies and gentlemen!" he cried. "Thank you very much. Thank you so much."
>
> *(Max Quordlepleen, in Douglas Adams' The Restaurant at the End of the Universe)*

Enumerating all persons who supported me and my work during the last years is an insurmountable task. I can only mention a few, and beg forgiveness of everyone I leave out. You are surely not forgotten.

I would first like to thank my advisors, Jörg Roßbach and Dirk Nölle, for their support during the completion of this work, and for lots of helpful criticism. I thank Rüdiger Schmidt for taking the time out of his busy schedule to review this thesis. I am also grateful to Bernhard Schmidt for participating on my committee.

More than anyone else, I have to thank Martin Staack. Work at FLASH is sometimes chaotic, and that the MPS keeps functioning against all odds is entirely Martin's fault. Thanks for the great teamwork and for teaching me the subtleties of Murphy's law.

I am deeply indebted to Abdallah Hamdi. I cannot count how often he came all the way from France to help us with the TPS—working through the night to the deafening sound of two modulators banging at 5 Hz. *Et je m'excuse pour parler français comme une vache espagnole...*

The commissioning of the FLASH MPS would not have gone as smoothly as it did without the help of Max Görler and Peter Göttlicher in tracking down hardware faults and misconfigurations. Henrik Tiessen and Bastian Michalek deserve special thanks for crawling through the nastiest parts of the accelerator tunnel with me, for help with cabling and mounting of BLMs, and for many great ideas.

I would like to express my gratitude to Sven Karstensen, Timmy Lensch, Frank Schmidt-Föhre, Matthias Werner, and all the others trying to get the XFEL MPS on the right way.

Kay Wittenburg and Gero Kube provided me with a lot of insightful comments on the thesis and were often a source of good advice. Fruitful discussions with Siggi Schreiber, Bart Faatz, and many other people usually found in the control room are also greatly appreciated.

A final nod of appreciation goes to Bolko Beutner, Florian Löhl, Sascha Meykopff, and Eduard Prat for letting me keep them from work with all those discussions on various physical and non-physical topics.