A Unified Distributed DSP-Based
Beam Diagnostics and Global Feedback System
for Ramped Electron Storage Rings:
Development, Construction and Applications

Boris Keil

Contents

Chapter 1

Introduction

DELTA (**D**ortmunder **El**ec**t**ron **A**ccelerator) is a 1.5 GeV synchrotron light source located at the University of Dortmund in Germany. It consists of three main parts: The linear accelerator "Linac", the booster synchrotron "Bodo" (**Bo**oster **Do**rtmund) and the storage ring "Delta"[1] (see fig. 1.1). Electrons of 50 keV energy are generated in the electron gun. The thermal cathode of the gun emits short electron pulses of 2 to 60 ns length. The subsequent Linac increases the energy of the electrons to typically 60 MeV. After passing the first transfer line ("T1"), the electrons are injected into Bodo, a ramped storage ring that accelerates them to a maximum energy of 1.5 GeV within a few seconds. The electrons are extracted from Bodo and injected into the storage ring Delta through the second transfer line ("T2"). By repeating this process, average electron beam currents of more than 100 mA can be accumulated in Delta. A very low pressure in the order of 10^{-8} Pa in the electron beam pipe minimises beam loss due to gas scattering, which allows a beam lifetime of many hours.
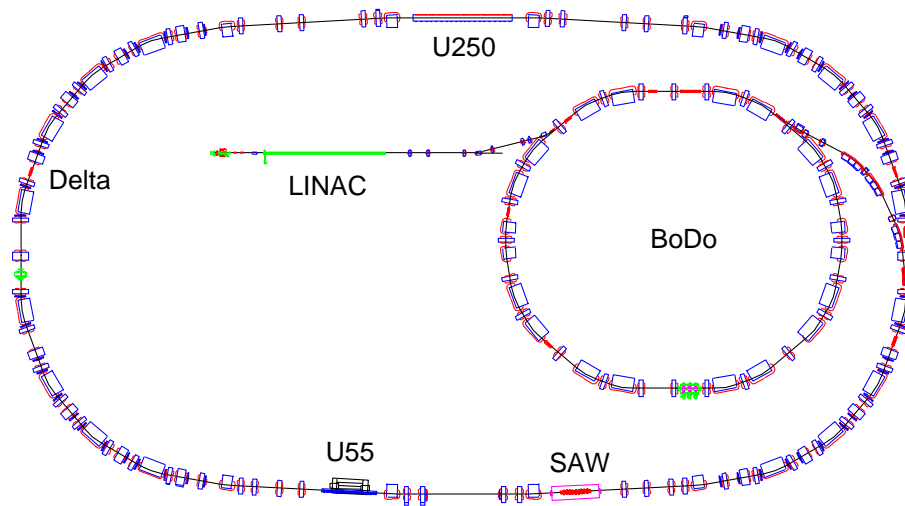
Figure 1.1: *The Electron Accelerator DELTA.*

The main purpose of DELTA is the generation of synchrotron radiation. This electromagnetic radiation is created by the transverse acceleration of high energy electrons either in bending magnets or in dedicated insertion devices like the wiggler (fig. 1.1, "SAW") or undulator ("U55","U250") magnets in the straight sections of Delta. Synchrotron radiation, being emitted tangential to the orbit of the electrons, is extremely bright, has a very small opening angle and a broad energy spectrum ranging from far infrared to X-rays. Therefore, it is perfectly suited to investigate the structure and properties of matter. Applications of synchrotron radiation include solid state physics, chemistry, biology, medicine and lithography for microstructure technology.

---

[1]Both the entire accelerator facility and the storage ring are named DELTA. In order to be non-ambiguous, the storage ring is spelled "Delta", while the name of the accelerator facility is "DELTA".

DELTA does not only serve as a 1.5 GeV synchrotron light source, but also as a test accelerator for accelerator physics research at lower energies up to 300 MeV, e.g. for the development and test of newly developed accelerator components. Furthermore, Delta drives a free electron laser ("FEL") and allows to investigate its interaction with the storage ring.

The booster Bodo is basically a small version of the storage ring Delta. It uses the same magnets, beam pipe, beam position monitor ("BPM") pickups and similar, though faster, magnet power supplies. The electron beam in Bodo can be stored at any energy between 30 MeV and 1.5 GeV, with arbitrary computer-programmable energy ramps and beam optics. Due to its similarity to Delta, Bodo is not only used as an injector, but also serves as a testbed for newly developed storage ring components, hardware and software.

The object of this thesis is the development, construction and application of a distributed DSP[2]-based system that is tailored to the measurement and real-time correction of beam parameters like global orbit and betatron tunes[3] at DELTA, with measurement and correction rates of some 100 Hz to some kHz. This enables the systematic analysis and improvement of beam parameters on a time scale that was not accessible with the previous systems at DELTA, since their control system integration achieved typical measurement and correction rates of a few Hz.

The correction of orbit perturbations and drifting betatron tunes in the booster is important to avoid beam loss during energy ramps and to achieve high beam currents and large charge transfer rates between Linac, Bodo and T2. While boosters of other light sources have injection rates up to some 10 Hz, the Bodo injection rate of less than 0.2 Hz requires much larger beam currents in order to achieve competitive storage ring filling times. This requires fast and precise measurement and correction of different beam parameters, which was either difficult and time-consuming or not possible with the existing diagnostics and control system hardware and software. Since the scope of Delta has shifted from a test machine to a synchrotron light source, a precise real-time correction of beam parameters in the booster increases its potential to replace the storage ring as a test accelerator for accelerator physics research and for the test of accelerator components and monitors.

Orbit perturbations in the storage ring are harmful for synchrotron light source experiments, since they may either change the position, intensity or energy of the synchrotron light that hits a sample, or they may increase the effective beam size or energy spread if the measurement averages over the perturbation. A fast DSP-based orbit feedback would allow the correction of such orbit perturbations that may be caused by mechanical vibrations, 50 Hz power supply noise or magnetic stray fields of booster or pulsed transfer line magnets during injection.

The correction of such orbit perturbations during injection is also important for the envisaged top-up injection mode of the storage ring. Top-up means that the beam current of Delta is kept nearly constant by injecting electrons every few minutes with opened beam shutters [1, 2], which improves the temperature stability of accelerator and beamline components. The necessary very low beam loss rates and stable beam positions at insertion devices also require a precise and fast measurement, control and correction of orbit and betatron tunes in booster and storage ring.

---

[2]A DSP (Digital Signal Processor) is a computer processor that is optimised on real-time processing and transfer of measurement data, e.g. for digital feedback control loops (see section 5.2.2).

[3]The horizontal (vertical) betatron tune is the number of horizontal (vertical) oscillations of an electron around the ideal orbit per turn.

As the storage ring does not provide sufficiently fast steerer magnet power supplies, the DSP system was installed and tested at the booster. In addition to faster power supplies, the booster was available for tests of hardware and software during normal synchrotron light source operation of Delta. This decreased the development time and allowed software tests and frequent reboot of the DSP boards without interference with experiments at the storage ring beamlines.

Chapter two of this thesis contains a summary of accelerator physics theory as far as required for the understanding of the following chapters. Chapter three describes the architecture, components and operation of the booster synchrotron Bodo.

Chapter four introduces the Bodo beam optics that was used for nearly all tests of the DSP-based beam diagnostics and feedback systems in the following chapters. The Bodo optics as well as suitable transfer line optics were developed during the work on this thesis in order to improve the storage ring filling time.

Chapter five describes the architecture of a VMEbus[4] DSP board ("DeltaDSP" board) that was developed and built during the work on this thesis. The DSPs of up to 255 boards can be connected in a ring by a novel fibre optics network for distributed multiprocessing of DSPs. This network ("DeltaNet") transfers measurement data from the internal memory of each DSP to the memories of all other DSPs in real-time. The boards are tailored to the fast acquisition of monitor signals, as well as to the efficient control of the ramped Bodo power supplies and the betatron tune measurement system. The DeltaDSP system is ideally suited for fast global orbit feedbacks, with advantages over existing orbit feedback architectures at other light sources. Chapter five also provides an overview of the system of six DeltaDSP boards that were installed in Bodo in October 2001. Since then, the system controls the ramped power supplies and handles most of the beam diagnostics systems of the booster. Furthermore, the chapter introduces the generic structure of the measurement and feedback DSP software that was developed in this thesis, as well as the integration of the DSP boards into the DELTA control system.

Chapter six describes the architecture and calibration of a DSP-based system that measures and corrects the Bodo betatron tunes in real-time. The system consists of DeltaDSP boards, additional digital and analogue hardware that was developed in this thesis, and parts of the previous all-analogue tune measurement system [4, 5].

Chapter seven introduces the global orbit feedback system that was implemented in Bodo via DeltaDSP boards and describes the BPM data acquisition system that was developed in this thesis. The chapter also includes a characterisation of the commercially available BPM RF front-end electronics. Furthermore, the chapter describes the orbit correction algorithm that was implemented and discusses the influence of the BPM locations on the algorithm.

Chapter eight presents betatron tune measurements at Bodo with operational betatron tune feedback. The chapter evaluates the precision and speed of the correction, and compares the measurement data with betatron tunes that were measured via beam loss monitors.

Chapter nine demonstrates the operation of the global Bodo orbit feedback during ramp cycles. Moreover, the response of a local feedback to an external orbit perturbation is investigated.

After a discussion of possible future applications in chapter ten, chapter 11 concludes the thesis with a summary of the results.

---

[4]VMEbus[3] is an industrial computer bus system that is used at DELTA.

Chapter 2

Theoretical Foundations of Circular Accelerators

This chapter contains definitions and theoretical foundations of accelerator physics that are required for the understanding of the subsequent chapters. In order not to exceed the scope of this work, derivations of formulas will be omitted and only the results will be presented. A more comprehensive treatment of the subjects can be found in accelerator physics literature (e.g. [6, 7, 8, 19, 22]).

2.1 Linear Transverse Motion

This section introduces the basic principles of transverse motion in circular electron accelerators. The description is limited to the case of linear motion for the types of circular accelerators used at DELTA, i.e. electron synchrotrons and storage rings. Electron storage rings are optimised for storing electrons at a certain constant or slowly varying energy, whereas synchrotrons are designed to increase the energy of electrons typically by a factor of 10 to 20 on a timescale of milliseconds to seconds.

All statements made in the following sections and chapters hold true equally for storage rings and synchrotrons, except where stated otherwise. Furthermore, it is assumed that the electrons are relativistic and have an energy of at least some ten MeV, so that approximations can be made that are not valid for non-relativistic electrons. The relation of momentum p and energy E of an electron is [9]

$$E[MeV] = \sqrt{(m_e[MeV/c^2])^2 + (p[MeV/c])^2} \, , \tag{2.1}$$

with $m_e[\text{MeV/c}^2] \approx 0.511$. Since typical electron energies in Bodo and Delta range from 60 resp. 300 MeV to 1.5 GeV, the values of $E[MeV]$ and $p[MeV/c]$ are nearly identical. Therefore, the notions energy dependence and momentum dependence of beam parameters in Bodo and Delta will be used synonymously in the following chapters.

Storage rings for relativistic electrons use magnetic fields to keep the particles on their orbit. These fields determine a closed periodic orbit in a storage ring for electrons with nominal momentum p_0, the so-called ideal orbit. In the following, it is assumed that this ideal orbit is located in a horizontal plane, so that the magnet fields of the bending magnets ("dipoles") have to be perpendicular to this plane. In order to simplify the description of the movement of a certain electron, one can use a coordinate system that is moving with an electron of nominal energy along the ideal orbit [6]. The s-axis unity vector $\vec{e}_s$ is tangent to the ideal orbit, i.e. it has the direction of the velocity vector of an electron moving along it. The z-axis $\vec{e}_z$ is perpendicular to the plane of the ideal orbit (i.e. it points upwards), and the x-axis $\vec{e}_x$ is perpendicular to s- and x-axis ($\vec{e}_x = \vec{e}_z \times \vec{e}_s$). As the electrons in Bodo and Delta circulate in a clockwise direction, the respective x-axis vector points to the outside of the respective accelerator ring.
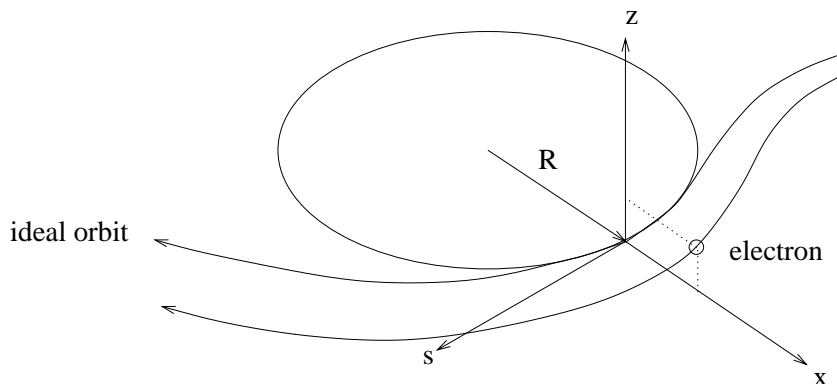
Figure 2.1: *Moving coordinate system.*

In the moving coordinate system, an electron is described by a state vector $(x, x', z, z', \Delta p/p_0)$ that is a function of s, the path length of an electron moving along the ideal orbit with respect to some starting point on the orbit. x' and z' are the derivatives of the coordinates with respect to s. $\Delta p/p_0$ is the relative momentum difference to an electron with nominal momentum p_0. Since $\Delta p/p_0$ is typically in the order of 1% or less, it can be approximated by $\Delta p/p$. By definition, an electron moving along the ideal orbit with nominal momentum has the time- resp. s-independent state vector $(0, 0, 0, 0, 0)$.

In addition to the bending dipoles, the magnet systems of storage rings contain higher order multipole magnets, e.g. focussing magnets to achieve and maintain a small beam size. In the following, the magnet system of an accelerator is called lattice, whereas a specific choice of magnet field strengths of the lattice is called the optics of the accelerator. Furthermore, it is assumed that all magnets have only transverse magnet field components with respect to the moving coordinate system, i.e. $B_s = 0$. Table 2.1 shows the (momentum-normalised) magnet fields for some common accelerator magnet types [7]. The respective constants R, k and m are obtained by a Taylor series expansion of the magnet field at $(x, z) = (0, 0)$. The table contains only terms up to the third order of the series. R is the bending radius of a dipole magnet field, k the focussing strength of a so-called quadrupole magnet field, and m the strength of a sextupole magnet field. By definition, R is positive for clockwise circulating beams.

An accelerator optics that consists only of dipole and "upright" quadrupole fields (see table 2.1) is called linear. The transverse equations of motion of an electron for such a system (with bending magnets only for the horizontal plane) can be approximated by the so-called Hill's equations [6] :

$$x''(s) + \left(\frac{1}{R^2(s)} - k(s) \right) x(s) \;\; = \;\; \frac{1}{R(s)} \frac{\Delta p}{p} \qquad (2.2)$$

$$z''(s) + k(s)z(s) \;\; = \;\; 0 \,. \qquad (2.3)$$

They are only valid for sufficiently relativistic electrons with $x \ll R$, $z \ll R$ and $\Delta p/p_0 \ll 1$. In storage rings, k and R are periodic functions of s. Quadrupole magnets with

Multipole	$\frac{e}{p}B_x$	$\frac{e}{p}B_z$
Dipole (upright)	0	$-\frac{1}{R}$
Quadrupole (upright)	kz	kx
Sextupole (upright)	mxz	$\frac{1}{2!}m(x^2 - z^2)$
...		
Dipole (tilted, 90^o)	$\frac{1}{R}$	0
Quadrupole (tilted, 45^o)	$-kx$	kz
Sextupole (tilted, 30^o)	$-\frac{1}{2!}\tilde{m}(x^2 - z^2)$	$\tilde{m}xz$
...		
Dipole (tilted, ϕ)	$\frac{1}{R}\sin(\phi)$	$-\frac{1}{R}\cos(\phi)$
Quadrupole (tilted, ϕ)	$k(-x\sin(2\phi) + z\cos(2\phi))$	$k(x\cos(2\phi) + z\sin(2\phi))$
...		

Table 2.1: *Magnet multipole fields. The fields of the tilted magnets are obtained by rotating the fields of the "upright" magnets by an angle ϕ around the s-axis ($\vec{e}_x$ towards $\vec{e}_z$, i.e. upright axis towards the inside of the accelerator in case of Bodo or Delta).*

$k < 0$ focus the electron beam horizontally and defocus it vertically, and magnets with $k > 0$ vice versa. The general solution of the equations is [6]:

$$u(s) = \sqrt{\varepsilon_u \beta_u(s)}\ cos(\Psi_u(s) + \phi_u) + D_u(s)\frac{\Delta p}{p_0} \qquad (2.4)$$

$$\Psi_u(s) = \int\limits_0^s \frac{d\tilde{s}}{\beta_u(\tilde{s})}\ , \qquad (2.5)$$

where u denotes either x or z. Electrons with $\Delta p/p_0 = 0$ oscillate transversely to the ideal orbit, with an s-dependent amplitude and phase. The function $\beta_u(s)$ is called betatron function (or short "beta function"), and the oscillation is therefore called betatron oscillation. The integration constant ε_u is called transverse single particle emittance of the electron, as opposed to the average emittance of all electrons in an electron beam. It can also be calculated from the so-called Twiss parameters α, β and γ [6]:

$$\varepsilon_u = \gamma_u u^2 + 2\alpha_u u u' + \beta_u u'^2 \qquad (2.6)$$

$$\alpha_u := -\frac{\beta_u'}{2} \qquad (2.7)$$

$$\gamma_u := \frac{1 + \alpha_u^2}{\beta_u}\ . \qquad (2.8)$$

u denotes either x or z. In storage rings with a static linear lattice, ε_u is a constant of motion as long as the movement of the electron is not disturbed, e.g. by acceleration in a cavity or the emission of synchrotron radiation photons [6].

Electrons with $\Delta p/p_0 \neq 0$ oscillate around a dispersion orbit $u(s) = D(s)\Delta p/p_0$. The dispersion function $D_u(s)$ is a solution of eq. (2.2) for $\Delta p/p_0 = 1$. In storage rings of circumference L_0, both $\beta_u(s)$ and $D_u(s)$ are periodic with the period L_0. $\beta_u(s)$ cannot change its sign [7], it is positive by definition. As eq. (2.3) contains no $\Delta p/p_0$ term and the
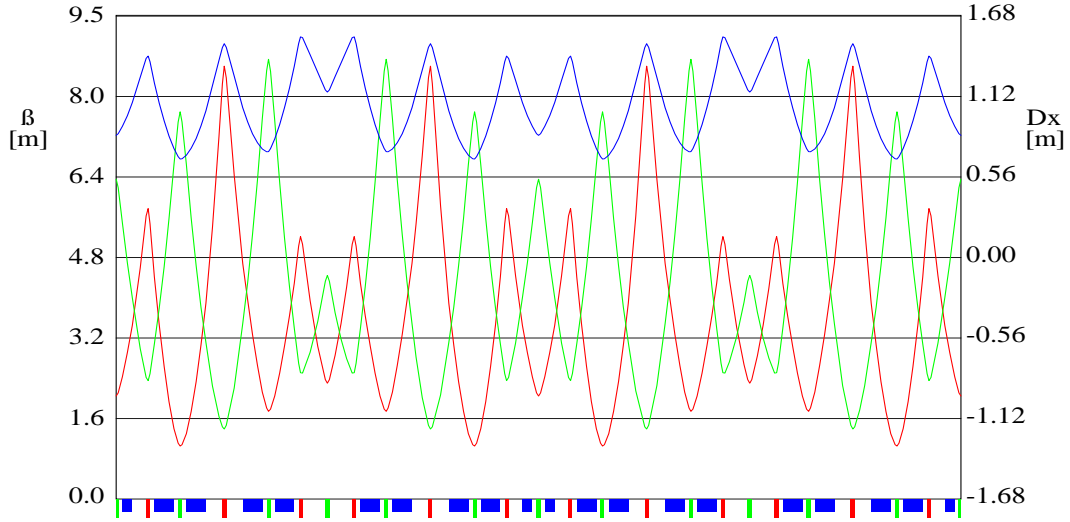
Figure 2.2: *Beta functions (red: x, green: z, left axis) and dispersion (blue, right axis) of the DELTA booster optics "bo-006b". The horizontal axis is the s coordinate (left: s = 0, right: s = $L_{0,Bodo}$ = 50.4 m). Blue boxes on the horizontal axis are dipoles, red and green boxes are horizontally and vertically focussing quadrupoles.*

storage ring is assumed to have only horizontal bending magnets and linear optics, $D_z(s)$ is 0. In practice, magnet field imperfections and misalignments lead to non-zero but small values of $D_z(s)$ ("spurious vertical dispersion"), usually in the order of a few centimeters [10, 11]. In general, a non-zero horizontal dispersion function $D_x(s)$ in storage ring bending magnets causes electrons of different momentum to have different orbit lengths (except for so-called isochronous storage rings [12]). The relative length change per relative momentum variation is called momentum compaction factor α_c:

$$\alpha_c := \frac{\Delta L/L_0}{\Delta p/p_0} = \frac{1}{L_0} \int_0^{L_0} \frac{D(\tilde{s})}{R(\tilde{s})} d\tilde{s} \ . \tag{2.9}$$

Since the beam optics of Bodo and Delta usually have positive or zero dispersion, the respective α_c is positive.

2.2 Longitudinal Motion

The average energy loss per turn of a relativistic electron due to the emission of synchrotron radiation is [6]

$$\Delta E_{SR}[keV] = 88.5 \frac{E^4[GeV^4]}{R[m]} \ , \tag{2.10}$$

with E being the electron energy and R the bending magnet radius. In order to compensate this energy loss, RF cavities in Bodo and Delta generate a longitudinal 500 MHz electric RF field along the beam axis. This field accelerates or decelerates electrons and separates the beam

into so-called bunches that have one RF wavelength distance to each other. The overall average energy change of an electron per turn for an accelerating peak voltage U_0 is

$$\Delta E_L(\Phi) = eU_0 sin(\Phi) - \Delta E_{SR} \ . \tag{2.11}$$

Φ is the RF phase angle that determines if an electron is accelerated or decelerated. U_0 is positive by definition. Electrons with nominal momentum $p = p_0$ and $\Delta E_L(\Phi = \Phi_0) = 0$ have a constant phase angle Φ_0 (nominal phase angle), whereas either $p \neq p_0$ or $\Phi \neq \Phi_0$ will cause so-called synchrotron oscillations of the electrons, i.e. Φ oscillates around Φ_0 (provided that $\Delta \Phi$ and Δp are sufficiently small [6]). This oscillation is caused by a non-zero momentum compaction factor α_c. In case of Bodo and Delta ($\alpha_c > 0$), electrons with nominal RF phase $\Phi = \Phi_0$ and $\Delta p/p_0 > 0$ have an increased revolution time, so that Φ, $\Delta E_L(\Phi)$ and therefore $\Delta p/p_0$ decrease from turn to turn, and vice versa for electrons with $\Phi = \Phi_0$ and $\Delta p/p_0 < 0$. The resulting oscillation focuses the beam in the longitudinal phase space, in analogy to transverse betatron oscillations that are caused by focussing magnet fields.

The longitudinal oscillation frequency is called synchrotron frequency $f_s = 1/T_s$. Its value for small $\Delta \Phi$ is [6]

$$f_s = f_0 \sqrt{\frac{h\, e\, U_0\, cos(\Phi_0)}{2\pi \beta^2 E}\left(\alpha_c - \frac{1}{\gamma_e^2}\right)} \tag{2.12}$$

$$U_0 = \sqrt{2 P_{rf} R_s} \tag{2.13}$$

$$sin(\Phi_0) = \frac{\Delta E_{SR}}{e\, U_0} \ . \tag{2.14}$$

h is the integer number of oscillations of the RF wave in the cavity per beam revolution (also called harmonic number). U_0 is the accelerating peak voltage, Φ_0 is the nominal RF phase angle, $\beta = v/c$ and α_c are relative speed and momentum compaction factor, and E is the electron energy. R_s and P_{rf} are shunt impedance of and RF power in the cavity (Bodo: $R_s \approx 9\ M\Omega$, Delta: $R_s \approx 3\ M\Omega$ [13]).

The number of synchrotron oscillations per revolution period in a storage ring of revolution frequency $f_0 = 1/T_0$ is called synchrotron tune Q_s:

$$Q_s = \frac{f_s}{f_0} \ . \tag{2.15}$$

Since the revolution frequency of Bodo is $f_{0,Bodo} \approx 5.95$ MHz and typical synchrotron frequencies are in the order of 10 to some 100 kHz, Q_s is usually in the order of 10^{-3} to 10^{-2}.

2.3 Betatron Tunes

The number of horizontal resp. vertical betatron oscillations per turn is called betatron tune, working point or Q-value of the storage ring:

$$Q_u := \frac{f_u}{f_0} = \frac{\Psi_u(L_0)}{2\pi} \ . \tag{2.16}$$

L_0 is the circumference of the accelerator (i.e. of its ideal orbit), f_u and f_0 are betatron oscillation and beam revolution frequency, and u denotes either x or z. As opposed to Q_s, Q_x and Q_z are normally larger than one. The betatron tunes of a storage ring or synchrotron must be adjusted carefully, since it is not possible to store the electron beam for certain values and value combinations of Q_x and Q_z. The beam will be lost if the betatron tunes drift and reach such values. The reason for this effect are so called optical resonances that are caused by imperfections of the magnet fields (see section 2.4).

Betatron tunes depend on the beam optics resp. quadrupole focussing strength. As both dipole bending radius R and quadrupole magnet strength k depend on the beam energy, the fields of these magnets in a synchrotron must increase proportionally in order to keep optics and betatron tunes constant and to avoid beam loss due to resonance crossing. The betatron tune shift due to a sufficiently small change Δk of the focussing strength of a quadrupole of length l located at $s = s_0$ is [6]

$$\Delta Q_x \approx -\frac{1}{4\pi} \int_{s0}^{s0+l} \beta_x(\tilde{s}) \Delta k(s) d\tilde{s} \qquad (2.17)$$

$$\Delta Q_z \approx +\frac{1}{4\pi} \int_{s0}^{s0+l} \beta_z(\tilde{s}) \Delta k(s) d\tilde{s} . \qquad (2.18)$$

The equation is an approximation valid for $\Delta k/k \ll 1$. By definition, k is negative for horizontally focussing quadrupoles and positive for vertically focussing ones, therefore the signs in the formulas for Q_x and Q_z are different. In case of betatron tune measurement devices that display both tunes and cannot distinguish between them, the sign difference can be used to find out which tune is Q_x and which one is Q_z by changing k. Moreover, the k-dependence of Q_x and Q_z can be used to measure the beta functions.

Depending on the distance of the betatron tunes to the next resonance, the maximum focussing error (also called gradient error) $\Delta k/k_0$ that can be tolerated may be as small as a few 10^{-3}. Due to various physical effects and technical imperfections, the magnet focussing errors in synchrotrons may be much larger, e.g. due to hysteresis and saturation of magnet fields and different power supply current regulation speed for dipoles and quadrupoles. Therefore, betatron tunes of synchrotrons and ramped storage rings usually have to be stabilised by some kind of dynamic magnet field resp. power supply current correction (see chapter 6).

The measurement of the beta function $\beta(s)$ by changing k and measuring Q requires $\Delta k = k - k_0$ to be sufficiently small. This follows from formulas (2.17) which assume that the variation $\Delta \beta(s) = \beta(s) - \beta_0(s)$ due to the change of k can be neglected. The variation of β with k is [7]:

$$\frac{\Delta \beta_u(s)}{\beta_{u,0}(s)} = \pm \frac{1}{2sin(2\pi Q_{u,0})} \oint \beta_{u,0}(\tilde{s}) \Delta k(\tilde{s}) \cos[2Q_{u,0}(\Psi_{u,0}(s) - \Psi_{u,0}(\tilde{s}) + \pi)] d\tilde{s} . \qquad (2.19)$$

The sign of the right hand term is negative for $u = x$ and positive for $u = z$. The Fourier series of the right hand side is [7]:

$$\frac{\Delta \beta_u(s)}{\beta_{u,0}(s)} = \pm \frac{Q_{u,0}}{2} \sum_{n=-\infty}^{\infty} \frac{F_n e^{in\Phi}}{Q_{u,0}^2 - (n/2)^2} . \qquad (2.20)$$

Eq. (2.20) shows that even small variations of k will cause large oscillations of $\beta(s)$ with a period of $2Q$ along the orbit ("beta-beat") if the respective betatron tune is close to an integer or half integer resonance. If k is varied in a single short quadrupole of length l located at $s = s_0$ and the quadrupole can approximated by a thin magnetic lens, the relative betatron function variation at the location of the quadrupole is

$$\frac{\Delta\beta_u(s_0)}{\beta_{u,0}(s_0)} = \pm\frac{l\beta_{u,0}(s_0)\Delta k}{2\ \tan(2\pi Q_{u,0})}\ . \tag{2.21}$$

The sign of the right hand term is positive for $u = x$ and negative for $u = z$. Equations (2.17) are only valid as long as the right hand side of eq. (2.21) is much smaller than one. Otherwise the variation of β has to be included:

$$\Delta Q_x \approx -\frac{1}{4\pi}l\bar{\beta}_x(s_0)\Delta k\left[1+\frac{l\bar{\beta}_x(s_0)\Delta k}{2\ \tan(2\pi Q_{x,0})}\right] \tag{2.22}$$

$$\Delta Q_z \approx +\frac{1}{4\pi}l\bar{\beta}_z(s_0)\Delta k\left[1-\frac{l\bar{\beta}_z(s_0)\Delta k}{2\ \tan(2\pi Q_{z,0})}\right]\ . \tag{2.23}$$

2.4   Optical Resonances

The general condition for an optical resonance is [6, 14]

$$n_x Q_x + n_z Q_z + n_s Q_s = n\ , \tag{2.24}$$

with n_x, n_z, n_s and n being integers whose greatest common divisor is one. Q_x, Q_z and Q_s are betatron tunes and synchrotron tune. The sum

$$o_{res} = |n_x| + |n_z| + |n_s| \tag{2.25}$$

is defined as the order of the resonance and determines its strength resp. impact on the beam, which decreases with increasing order o_{res}.

Resonances with $Q_u = n$ are driven by magnetic field imperfections of dipoles. If e.g. Q_x is an integer n, a dipole magnet imperfection will lead to a transverse kick of the electron orbit each time the electron passes the respective magnet. Due to the phase advance of $n \cdot 360^o$ per turn, the kicks add up every turn and therefore lead to increasing horizontal oscillation amplitudes until the electron hits the beam pipe. Similarly, quadrupole field imperfections drive resonances with $Q_u = n$ and $2Q_u = n$, sextupole fields drive resonances with $3Q_u = n$ etc. [7]. The impact of a resonance is therefore not only determined by n, but also by the quality and imperfections of the accelerator magnets and their power supplies.

The strength of a resonance is also determined by the symmetry or periodicity of the lattice. If a lattice consists of n identical periods ("superperiodicity n"), the impact of certain resonances is usually reduced significantly [7], provided that the periodicity is not distorted e.g. by large magnet field or power supply current errors in some of the periods. The lattice of Bodo has a superperiodicity of $n = 2$, consisting of 4 quadrants that are mirror-symmetrical with respect to two orthogonal axes. The Delta lattice has only one such mirror axis and therefore $n = 1$.

The tunes Q_x and Q_z of Bodo and Delta must have a sufficient distance to transverse resonances ($Q_u = n$) of low order, since they usually cause complete ($n \leq 2$) or partial (n small, but $n > 2$) beam loss. Furthermore, even tunes that are sufficiently close to a resonance may cause beam loss, since resonances have a certain "thickness" which is called stop band width [7]. The stop band width depends on the alignment precision of the magnets, on the magnetic field errors, on the order and type of the resonance and on the magnet lattice. The avoidance of resonances is especially important for synchrotrons, since the rapidly changing magnet fields often lead to drifting betatron tunes and therefore beam loss when tunes cross a destructive resonance. Since the number of lost electrons depends on how long the tune is close enough to a resonance, resonances are more problematic for ramped storage rings like Bodo, since its typical ramp cycle frequency is about 50 times lower than the ramp cycle frequencies of other light source boosters [15, 16] (see section 3.2).

Resonances with $n_s = 0$ and $n_x \neq 0 \neq n_z$ are called transverse resonances ("sum resonances" or "difference resonances" for same or different sign of n_x and n_z). They are caused by magnet fields that result in a coupling of horizontal and vertical betatron oscillations, e.g. by tilted quadrupole fields or sextupole fields (see table 2.1). The transverse coupling resonances with $|n_x| = |n_z| = 1$ are also called linear coupling resonances. The linear sum resonance can be destructive (i.e. cause beam loss), because it preserves only the difference between horizontal and vertical emittance ($\varepsilon_x - \varepsilon_z = \text{const.}$), while the individual values and therefore the beam size may diverge. In case of the linear difference resonance, the sum $\varepsilon_x + \varepsilon_z$ of the emittances is constant, which only leads to a redistribution of the transverse emittances between x and z plane, while the overall beam size remains finite [8, 17].

Resonances with $n_s \neq 0$ and $n_x \neq 0$ or $n_z \neq 0$ are called synchro-betatron resonances. They are caused by a coupling of transverse and longitudinal motion of the electrons, e.g. by to non-zero dispersion in a cavity or by fast periodic transverse kicks that depend on the longitudinal coordinate [14, 18].

2.5 Chromaticity

The chromaticities ξ_x and ξ_z of a storage ring optics are defined as betatron tune variation per relative momentum variation:

$$\xi_x := \frac{\Delta Q_x}{\Delta p/p_0} = +\frac{1}{4\pi} \oint [k(s) - m(s)D_x(s)]\beta_x(s)ds \qquad (2.26)$$

$$\xi_z := \frac{\Delta Q_z}{\Delta p/p_0} = -\frac{1}{4\pi} \oint [k(s) - m(s)D_x(s)]\beta_z(s)ds . \qquad (2.27)$$

The integrals are calculated once along the complete orbit. The sign difference for x and z results from the sign definition of k (see section 2.1). In case of zero sextupole fields (i.e. $m(s) = 0$), the resulting chromaticity is called natural chromaticity $\xi_{u,0}$ of the beam optics. Both $\xi_{x,0}$ and $\xi_{z,0}$ are usually negative. $\xi_{z,0}$ can be changed to zero or positive values by installing sextupole magnets with $m = m_z > 0$ at locations where $D_x(s)$ and β_z are large and β_x is small ("vertical sextupoles"). Respectively, ξ_x can be increased by installing sextupole magnets with $m = m_x < 0$ at locations where $D_x(s)$ and β_x are large and β_z is small ("horizontal sextupoles"). Large differences of β_x and β_z in the sextupoles guarantee that horizontal

sextupoles change mainly ξ_x and vertical sextupoles change mainly ξ_z , so that both ξ_x and ξ_z can be changed to zero or positive values simultaneously with minimal field strengths m_x and m_z. Positive chromaticities are usually necessary in storage rings to avoid the so-called head-tail instability which leads to beam loss at high single bunch beam currents and negative chromaticity [8]. Furthermore, large absolute chromaticity values may limit the momentum acceptance of the accelerator, because momentum oscillations would lead to large betatron tune oscillations that cause beam loss if the nominal tune is sufficiently close to a destructive optical resonance.

2.6   Beam Dimensions in Storage Rings

When an electron of nominal energy moves along the ideal orbit and emits a synchrotron radiation photon at a location of non-zero dispersion, the momentum loss Δp will cause the electron to perform betatron oscillations around a dispersion orbit $D(s)\Delta p/p$. On the other hand, the emission of synchrotron radiation photons at locations of small or zero dispersion reduces the transverse momentum of electrons with large betatron oscillation amplitudes, whereas the RF cavity can only restore the longitudinal component of the momentum vector. Therefore, synchrotron radiation leads to damping of transverse and also longitudinal oscillations with large amplitudes, but simultaneously excites transverse and longitudinal oscillations for electrons with zero or small oscillation amplitudes.

This results in an equilibrium between excitation and damping both in transverse and longitudinal directions, leading to a Gaussian distribution of transverse and longitudinal oscillation amplitudes [19]. As the RF systems of Bodo and Delta have a frequency of approx. 500 MHz, the RF phase space focussing (see section 2.2) leads to a longitudinal charge distribution that consists of max. 84 (Bodo) resp. 192 (Delta) equally spaced electron bunches, with typical bunch lengths in the order of a few centimeters. The charge density ρ of a bunch of N electrons located at (x_0, z_0, s_0) can be approximated by [6]

$$\rho(x,z,s) = \frac{-eN}{(2\pi)^{3/2}\sigma_x\sigma_z\sigma_s} \, \exp\left[-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(z-z_0)^2}{2\sigma_z^2} - \frac{(s-s_0)^2}{2\sigma_s^2}\right] \, . \qquad (2.28)$$

The transverse beam dimensions $\sigma_u(s)$ and divergences $\sigma_{u'}(s)$ are given by [7]

$$\sigma_u(s) = \sqrt{\varepsilon_u\beta_u(s) + (D_u(s)\,\Delta E/E)^2} \qquad (2.29)$$

$$\sigma_{u'}(s) = \sqrt{\varepsilon_u\gamma_u(s) + (D'_u(s)\,\Delta E/E)^2} \, . \qquad (2.30)$$

u denotes either x or z, ΔE is the standard deviation of the particle energy, and ε_u is the equilibrium emittance. The longitudinal beam dimension σ_s is [8]

$$\sigma_s = v_e \frac{|\alpha_c - \frac{1}{\gamma_e^2}|}{2\pi f_s}\frac{\Delta E}{E_0} \qquad (2.31)$$

$$\approx c\,\frac{|\alpha_c|}{2\pi f_s}\frac{\Delta E}{E_0} \, . \qquad (2.32)$$

v_e is the speed of the electron, α_c and γ_e are momentum compaction factor and relativistic gamma factor, $f_s = 1/T_s$ is the synchrotron oscillation frequency, and ΔE is the beam

energy spread. Eq. (2.32) is an approximation that is valid for the electron energy range of Bodo and Delta.

The horizontal equilibrium emittance $\varepsilon_{x,0}$ of an electron storage ring with ideal linear optics and constant energy is [8]

$$\varepsilon_{x,0} \;=\; C_q \gamma_e^2 \frac{\langle H/|R^3|\rangle_s}{J_x \langle 1/R^2\rangle_s}\,, \tag{2.33}$$

where J_x, ϑ, H and C_q are defined as

$$J_x \;:=\; 1 - \vartheta \tag{2.34}$$

$$\vartheta \;:=\; \frac{\oint \left[\frac{D}{R}\left(-2k + \frac{1}{R^2}\right)\right] ds}{\oint \frac{1}{R^2} ds} \tag{2.35}$$

$$H \;:=\; \gamma D^2 + 2\alpha D D' + \beta D'^2 \tag{2.36}$$

$$C_q \;:=\; \frac{55 hc}{64\sqrt{3}\,\pi mc^2} = 3.84 \cdot 10^{-13} m\,. \tag{2.37}$$

The brackets $\langle\rangle_s$ denote the average of the respective term along the complete orbit. In an ideal linear machine, the synchrotron radiation also leads to an equilibrium vertical emittance of [8]

$$\varepsilon_{z,0} \;=\; C_q \frac{\langle \beta_z/|R^3|\rangle_s}{J_z \langle 1/R^2\rangle_s} \tag{2.38}$$

$$J_z \;:=\; 1\,. \tag{2.39}$$

It should be noted that the usually extremely small vertical emittance value $\varepsilon_{z,0}$ of equation (2.38) is only valid for an ideal flat storage ring with linear optics. In practice, spurious vertical dispersion, coupling between horizontal and vertical betatron oscillations, mutual interactions of the electrons within a bunch (intra beam scattering) and interactions of the electrons with the residual gas in the beam pipe usually lead to much higher values of ε_z. Therefore, $\varepsilon_{z,0}$ is usually neglected, and the horizontal (uncoupled) emittance $\varepsilon_{x,0}$ of a storage ring beam optics is called *the* emittance ε_0 of the optics. The coupling constant κ between horizontal and vertical emittance is defined as:

$$\kappa = \frac{\varepsilon_z}{\varepsilon_x}\,. \tag{2.40}$$

Therefore, transverse emittances of a storage ring as a function of coupling factor κ and uncoupled horizontal emittance ε_0 are

$$\varepsilon_x = \frac{1}{1+\kappa}\varepsilon_0 \tag{2.41}$$

$$\varepsilon_z = \frac{\kappa}{1+\kappa}\varepsilon_0\,. \tag{2.42}$$

Typical values of κ are in the order of 10^{-1} to 10^{-3}.

The natural energy spread of a storage ring due to an equilibrium of synchrotron radiation damping and excitation is [8]

$$\frac{\Delta E}{E_0} \;=\; \gamma_e \sqrt{C_q \frac{\langle 1/|R^3|\rangle_s}{J_E \langle 1/R^2\rangle_s}} \tag{2.43}$$

$$J_E \;:=\; 2 + \vartheta\,. \tag{2.44}$$

Equation (2.43) is only valid in the zero-current limit, i.e. as long as the charge density is below the threshold for the so-called turbulent bunch lengthening instability (also called bunched beam microwave instability). The threshold beam current of this self-stabilising instability increases with $\Delta E/E_0$. Therefore turbulent bunch lengthening does not cause beam loss, but just increases the bunch energy spread until it is just large enough to avoid the instability. The threshold is ("Boussard criterion") [8]:

$$\frac{\Delta E}{E_0} \geq \left(\frac{e I_0 f_s |Z_\parallel/n|}{F \sqrt{2\pi} \; \beta^2 E_0 f_0 |\alpha_c - \frac{1}{\gamma_e^2}|^2} \right)^{1/3} \; .$$ (2.45)

I_0 is the average circulating single beam current per bunch, $|Z_\parallel/n|$ is the normalised longitudinal broadband impedance, and F is one for Gaussian charge distributions [22]. For a given beam current in a ramped storage ring, the natural energy spread $\Delta E/E_0$ in eq. (2.43) scales with E, but the turbulent energy spread in eq. (2.45) scales with $1/E^{1/3}$ for a constant synchrotron frequency (i.e. RF cavity power $P_{rf} \sim E^2$) or $1/E^{1/2}$ for $P_{rf} = \text{const.}$. Therefore, the equilibrium bunch length at low booster beam energies and high single bunch currents is usually determined by turbulent bunch lengthening, whereas the bunch length at the maximum (i.e. extraction) energy is usually determined by the natural energy spread.

When electrons are injected into a storage ring (e.g. from Bodo to Delta), their emittance is usually higher than the equilibrium emittance of the storage ring. Assuming an ideal linear lattice and constant beam energy, synchrotron radiation damping will cause the emittances to converge to the equilibrium emittances. The damping times for the exponential damping of the transverse and longitudinal oscillations are [19]

$$\tau_x \;\; = \;\; \frac{2T_0 E}{W_0} \frac{1}{J_x}$$ (2.46)

$$\tau_z \;\; = \;\; \frac{2T_0 E}{W_0} \frac{1}{J_z}$$ (2.47)

$$\tau_E \;\; = \;\; \frac{2T_0 E}{W_0} \frac{1}{J_E} \; .$$ (2.48)

W_0 is the energy loss due to synchrotron radiation per turn and T_0 the revolution time. The characteristic damping time $\tau_{SR} := 2T_0 E/W_0$ for a storage ring with constant bending magnet radius R and revolution frequency f_0 is

$$\tau_{SR} \;\; \approx \;\; 22.6 \; ms \; \frac{R[m]}{f_0[MHz] \; E^3[GeV^3]} \; .$$ (2.49)

As the sum of the so-called damping partition numbers $J_x + J_z + J_E \;\; = \;\; 4$ is constant [6] ("Robinson's Theorem"), the sum of the damping factors $1/\tau \;\; = \;\; 1/\tau_x + 1/\tau_z + 1/\tau_E$ of a storage ring lattice is independent of the beam optics. Since Bodo and Delta have separated function lattices (i.e. separated dipole and quadrupole magnets), $J_x \approx 1$ [6]. However, if the beam is moved onto a dispersion orbit by increasing the RF frequency (and thus decreasing the circumference and momentum), the resulting increase of J_x (see eq. (2.34) and (2.35)) can lead to a significant reduction of the horizontal emittance for a positive momentum compaction factor [20], while simultaneously reducing J_E and therefore increasing energy spread and bunch length. Since the longitudinal damping partition J_E should not become negative (to avoid antidamping and beam loss), the maximum attainable emittance reduction factor is limited by the value at which J_E changes its sign.

2.7 Beam Dimensions in Synchrotrons

In order to achieve short storage ring filling times and good electron transfer rates from booster to storage ring, the beam dimensions of the booster electron beam must be small enough to guarantee a loss-free transfer of the electrons from the booster through the transfer line into the storage ring. The equations for the beam dimensions in section 2.6 are only valid if the energy of the electron beam is either constant or changing on a timescale that is much longer that the damping times τ_x, τ_z and τ_E. This holds true for the storage ring Delta, but not for the ramped storage ring Bodo, since its energy is changing periodically, typically between 60 MeV and 1.5 GeV in a period of 6-10 seconds. Typical synchrotron radiation damping times in Bodo are in the order of 30 to 60 seconds at 60 MeV and 2 to 4 milliseconds at 1.5 GeV. Therefore, synchrotron radiation does not change the beam dimensions significantly during the low energy part of the energy ramp. However, the low-energy emittance is influenced by so-called adiabatic damping that changes the emittance inversely proportional to the beam energy. The overall emittance change $d\varepsilon_x/dt$ due to synchrotron radiation and adiabatic damping in synchrotrons with no emittance coupling and constant bending magnet radius R is [7, 19]:

$$\frac{d\varepsilon}{dt} = -\varepsilon \left[\frac{dE/dt}{E} + 2\alpha_x \right] + \frac{55}{48\sqrt{3}} \frac{\Delta E_{SR} E_c \langle H \rangle_s}{T_0 E^2} \tag{2.50}$$

$$E_c := \frac{3hc\gamma_e^3}{4\pi R} . \tag{2.51}$$

α_x is the horizontal damping constant, E_c is the critical photon energy of the bending magnets, i.e. the energy that splits the photon power spectrum in two halves of equal power [6]. $\langle H \rangle_s$ is the average of the function $H(s)$ (see eq. (2.36)) in the bending magnets of the synchrotron, ΔE_{SR} is the energy loss per turn, and T_0 the revolution time. The dE/dt term causes the adiabatic damping, $2\alpha_x$ the radiation damping and the right hand term the radiation excitation.

When the Linac beam is injected into Bodo with an energy of 60 MeV and an emittance of about 1 μm rad [21], the emittance decreases due to adiabatic damping $\sim 1/E$ until reaching a local minimum at some 100 MeV. Then synchrotron radiation starts to dominate the emittance change, and the E^2 dependence of the equilibrium emittance (eq. (2.33)) causes the emittance to increase until the beam is extracted at 1.5 GeV. In case the beam is extracted at a lower energies e.g. for the operation of the free electron laser of the DELTA storage ring, the emittance might not yet have reached its equilibrium value. The actual value then depends on the ramp period and the time dependence of the energy ramp.

Furthermore, interactions of the electrons with each other (e.g. so-called intra beam scattering at high beam currents) or with residual gas atoms or ions in the beam pipe may increase the transverse and longitudinal beam dimensions, depending on various parameters like beam energy, energy ramp speed, beam current and longitudinal filling pattern (i.e. number of bunches) [8].

2.8 Orbit

In an ideal storage ring with ideal magnet fields, the center of charge of the electron bunches in the ring moves along the ideal closed orbit that leads through the centers of all quadrupole magnets.

If an ideal orbit is perturbed by a horizontally or vertically deflecting dipole magnet field ("orbit kick") with the kick angle $\Delta u'$ (u denotes x or z) at $s = s_0$, this results in a new closed orbit that is different from the ideal orbit. The resulting horizontal orbit deviation $u(s_0)$ and orbit angle $u'_{\varepsilon \to +0}(s_0 + \varepsilon)$ is [6]

$$u(s_0) \;=\; \Delta u' \frac{\beta_u(s_0)}{2 \, \tan(\pi Q_u)} \tag{2.52}$$

$$u'_{\varepsilon \to +0}(s_0 + \varepsilon) \;=\; \frac{\Delta u'}{2}\left(1 - \frac{\alpha_u(s_0)}{\tan(\pi Q_u)}\right) \;. \tag{2.53}$$

The formulas are an approximation that is valid for linear optics and short dipole magnets where the overall kick can be approximated by a kick in the center of the magnet. Since the equations of motion (2.2) and (2.3) of the electrons describe a harmonic oscillation with an envelope and phase that depend on the longitudinal coordinate s, the orbit perturbation that results from one dipole kick has a sine-like shape, with a phase jump at $s = s_0$.

Equation (2.52) shows that the amplitude of the perturbed orbit is minimal for half-integer betatron tunes Q_u. It diverges for integer betatron tunes, because the phase advance per turn is a multiple of 2π, and the dipole kicks add up coherently at every turn of the electron beam, which finally leads to beam loss. Furthermore, the orbit perturbation for a given kick angle $\Delta u'$ is larger if the kick is applied at locations with large values of β_u, since the growth of the single particle emittance of an electron at $s = s_0$ is larger if the phase ellipse of the electron motion in the 2-dimensional (u, u') phase space is "flat" (i.e. u values of the ellipse are large and u' values are small with respect to other locations s in the ring, see eq. (2.6)).

Since it is not possible to avoid orbit perturbations completely, most storage rings have horizontal and vertical steerer electromagnets (i.e. additional small dipole electromagnets that kick the beam in the respective plane). By measuring the orbit with beam position monitors ("BPMs") and applying appropriate orbit corrections with the steerer magnets, orbit perturbations are usually reduced to a fraction of the initial amplitude.

Singular Value Decomposition

A common method to calculate the steerer magnet kicks that are required to correct the orbit perturbations at a given number of BPMs is the so-called SVD (**S**ingular **V**alue **D**ecomposition) method [25]. To apply this method, one needs to know the so-called beam response matrix B_{ij} that is defined as the orbit change Δu_i at BPM no. i divided by the angle kick Δk_j of steerer magnet no. j that generated the orbit change:

$$B_{ij} = \frac{\Delta u_i}{\Delta k_j} \;. \tag{2.54}$$

B can be obtained either from a theoretical model of the accelerator optics or by measuring it directly. Assuming linear beam optics, B_{ij} is independent of the size of the kick, and the overall orbit change Δu_i that is generated by a number of steerer magnet kicks Δk_j is a superposition of the individual orbit perturbations:

$$u_i = B_{ij} k_j \;. \tag{2.55}$$

In eq. (2.55), the Einstein sum rule applies.

The $m \times n$ beam response matrix B can be expressed as the product $U\Sigma V^T$ of three matrices, where U and V are orthogonal $m \times m$ and $n \times n$ matrices:

$$B = U\Sigma V^T \, . \tag{2.56}$$

Σ is a $m \times n$ matrix with zero off-diagonal elements and non-negative diagonal elements that appear in non-increasing order. The product of the three matrices is called singular value decomposition ("SVD") of B [25]. The SVD orbit correction method calculates the pseudo-inverse M of the matrix B:

$$M = VDU^T \, . \tag{2.57}$$

D is a $n \times m$ matrix with zero off-diagonal elements and diagonal elements that are obtained by inverting the diagonal elements of Σ if these elements are larger than a certain cutoff value. If an element of Σ is smaller than the cutoff value, the respective diagonal element of D is set to zero. If B is quadratic and its column vectors are linearly independent, M is the inverse of B (if the cutoff value is zero). This means that the number of BPMs per plane is equal to the number of steerer magnets for that plane, and that the orbit at each BPM can be corrected independently. If the number of steerer magnets is smaller than the number of BPMs, the SVD method may not be able to eliminate the orbit perturbation completely, but it will minimise the RMS value of the perturbation. If the number of steerers is larger than the number of BPMs, the SVD method corrects the orbit perturbation completely and minimises the RMS value of the required steerer kicks (since there are more steerers available than required). Therefore, the SVD method is one of the most common methods to correct the orbit of circular accelerators. Synchrotron light sources that use the SVD method for their global orbit feedback include ESRF [46], APS [49], SLS [58] and Bessy II [57].

However, if there are too few BPMs, the SVD method can cause large orbit kicks that result in large orbit deviations at locations in the accelerator that are not visible at the available BPMs. An example for this effect is a subsequent series of three steerer magnets with no BPM in between. Appropriate kicks of the steerers can create a local orbit bump that is nearly invisible at the BPMs outside the bump. This is equivalent to the occurrence of diagonal elements of the beam response matrix Σ that are zero or very small compared to the largest diagonal element. Therefore, the cutoff value for the inversion of the diagonal elements must be large enough to suppress such effects. Another solution to this problem is either to reduce the number of steerer magnets that are used for the orbit correction (i.e. for the SVD calculation) or to install additional BPMs. Furthermore, BPM noise and nonlinear effects (in case of measured beam response matrices instead of calculated ones) can cause small but non-zero diagonal elements that should also be suppressed by a sufficiently large cutoff value.

Chapter 3

The Ramped Storage Ring Bodo

This chapter describes the architecture and components of the ramped storage ring Bodo, as far as required for the understanding of the subsequent chapters.

3.1 Magnet Lattice

Bodo has a so-called missing dipole FODO lattice. The layout of the lattice is shown in fig. 3.1. A regular FODO lattice consists of an alternating series of equally spaced horizontally focussing and defocussing quadrupole magnets with dipole magnets of equal bending angle in between. In Bodo, the dipoles next to two of the quadrupole magnets QD1 (see fig. 3.1) are missing in order to obtain the required space for the RF cavity and the extraction kicker magnet that is required to transfer the electrons from the booster to the storage ring (see section 3.4).

Unlike many other booster synchrotrons of existing or planned light sources, Bodo has six quadrupole magnet families (i.e. sets of quadrupoles with different focussing strength) instead of two. The resulting flexibility of the beam optics allows varying distances of the quadrupole
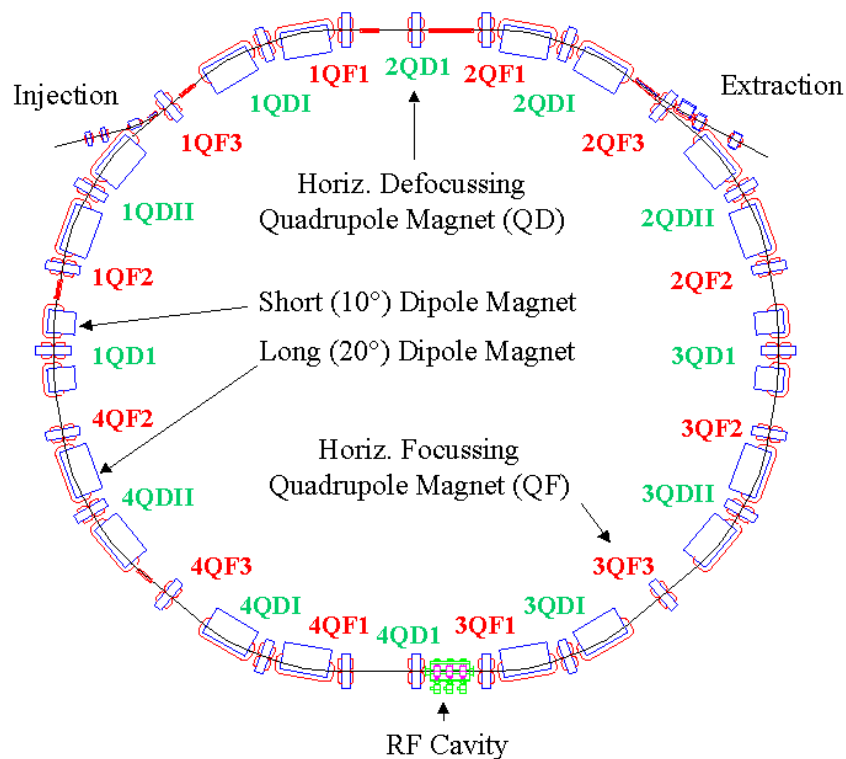
Figure 3.1: *Dipole and quadrupole magnets of Bodo.*

magnets, so that one can obtain additional space for kicker and septum magnets (see section 3.3) while keeping the lattice compact. Moreover, Bodo has two kinds of dipole magnets: Sixteen magnets with 20 degrees deflection angle, and four with 10 degrees deflection angle. While the flexibility of the Bodo beam optics is a desirable feature for a test machine, the optimisation of beam parameters like emittance, dispersion and beta functions is not as simple as for regular FODO lattices, since the beam optics parameter space (i.e. the focussing strengths of the quadrupole families) has six dimensions instead of two.

3.2 Magnets and Power Supplies

As opposed to many other synchrotron light sources that use a capacitor-inductor resonant circuit ("White Circuit" [15, 16]) of fixed frequency (typically 10 to 50 Hz) for their booster synchrotron focussing and bending magnets, the magnets of Bodo have DC power supplies that allow the booster to be operated both as a storage ring and as a synchrotron (i.e. ramped storage ring). The shape of the power supply current ramps is programmable and arbitrary.

The advantage of the DC power supply solution is the significant cost reduction compared to White Circuits and the flexibility to use the booster as a test storage ring for newly developed machine component or beam monitors, machine modelling etc. while the storage ring serves as a synchrotron radiation source. The drawback is a storage ring injection frequency of typically 0.1 to 0.2 Hz compared to 10-50 Hz for typical fast cycling synchrotrons. The booster to storage ring charge transfer efficiency must be about two orders of magnitude higher compared to fast cycling synchrotrons in order to achieve the same storage ring filling time. Therefore, beam loss due to variations of orbit or betatron tunes during the ramp cycle that are usually tolerated in normal synchrotrons have to be avoided in Bodo. This requires precise control and correction of all beam parameters and therefore magnet power supply currents.

Dipole Magnets

All 20 rectangular laminated iron dipole magnets (16 long and 4 short ones, see fig. 3.1) of Bodo are connected in series. The resulting overall inductance is $L_B = 555\ mH$, and the overall resistance is $R_B = 320\ m\Omega$. Their time constant is therefore $\tau_B = L_B/R_B = 1.734\ s$. The maximum current of the dipole magnet power supply is $I_{B,max} = 995\ A$, with a maximum voltage of $U_{B,max} = 356\ V$ and a minimum voltage of $U_{B,min} = -310\ V$ [28]. The current regulation of the power supply is unipolar, i.e. the current is always $\geq 0\ A$, but the output voltage of the power supply can be both positive and negative to allow both a fast increase and decrease of the booster beam energy. The control system interface of the power supply has three analogue signal inputs for the set current and its first and second derivative.

The minimum current ramp time for a power supply with maximum resp. minimum output voltage U_{limit} connected to a resistance R in series with an inductance L is

$$t_{ramp} = -\frac{L}{R} \cdot \ln\left(1 - \frac{I_{stop} - I_{start}}{U_{limit}/R - I_{start}}\right) ,$$

where I_{start} and I_{stop} are the currents at the start and the end of the ramp. Assuming $U_{limit} = +336\ V$ for up and $U_{limit} = -290\ V$ for down ramps and start and end currents

of 37.7 A (60 MeV) and 980 A (1488 MeV), the resulting minimum ramp times of the dipole power supply for up and down ramps are

$$
\begin{aligned}
T_{B,60\to1488} &= 4.632 \ s \\
T_{B,1488\to60} &= 1.2 \ s \ .
\end{aligned}
$$

The voltages and currents are slightly lower than the maximum values to account for possible magnet resistance changes due to the temperature of the water that cools the copper coils of the magnets. Bodo ramp curves usually have an injection plateau of $0.2 \ s$ length, a $0.2 \ s$ flat top for extraction, and four smoothing segments of $0.1 \ s$ each to connect injection plateau and flat top with the rising and falling edges of the ramp. The resulting minimum ramp time is

$$
T_{B,60\leftrightarrow1488} = 6.63 \ s
$$
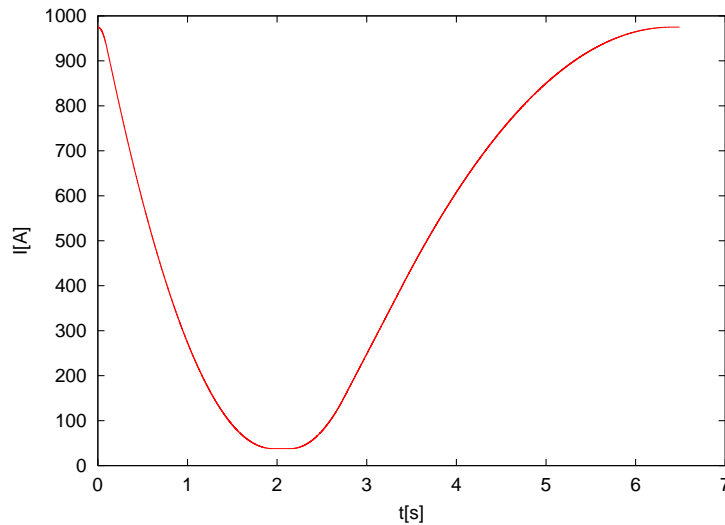
Figure 3.2: *Typical Bodo dipole current ramp, with 62 MeV minimum and 1483 MeV maximum beam energy.*

Quadrupole magnets

The four laminated iron quadrupole magnets of each of the six quadrupole families are connected in series, with an overall inductance of $L_Q = 1.6 \ H$ and a resistance of $R_{Q,up} = 1.66 \ \Omega$. The resulting time constant is $\tau_{Q,up} = L_Q/R_{Q,up} = 0.964 \ s$. The maximum power supply current is $I_{Q,max} = 60 \ A$, and the minimum and maximum voltages are $U_{Q,min} = 0 \ V$ and $U_{Q,max} = 110 \ V$.

For cost reasons, the booster quadrupoles have the same commercial switch mode power supplies as the storage ring, with some modifications to account for the required ramp speed of the booster. The power supplies chop a DC input voltage at a frequency of 14 kHz to charge an output capacitor that is connected to the load (i.e. the magnets). If the semiconductor switch

("MOSFET") that connects DC input voltage and capacitor (via some passive circuit elements) is non-conducting, the magnet current flows through a free-wheeling diode and decays with the time constant of the L-R load circuit, provided that the resistance and inductance in the internal freewheeling current flow path of the power supply are negligible. Since the power supplies cannot generate a negative DC output voltage, the original power supply design would have resulted in very long ramp periods, since the ramp down time from 1.5 GeV to 60 MeV would have been determined by the time constant $\tau_{Q,up}$ of the magnets:

$$T_{Q,1500\rightarrow60,slow} = -\tau_{Q,up} \cdot \ln(60MeV/1500MeV) = 3.10s.$$

In order to obtain shorter ramp cycle periods, the power supplies have a modified power circuit and faster control loop in their current regulation circuits. The modification includes the insertion of a resistor of $R_{Q,Diode} = 7.5\ \Omega$ in series with the free-wheeling diode. This results in different time constants for up and down current ramps, with an effective resistance of $R_{Q,down} = R_{Q,Diode} + R_{Q,up} = 9.16\ \Omega$ for down ramps. The respective time constant is $\tau_{Q,down} = L_Q/R_{Q,down} = 0.175\ s$. The corresponding minimum ramp down time is

$$T_{Q,1500\rightarrow60} = -\tau_{Q,down} \cdot \ln(60MeV/1500MeV) = 0.57s.$$

The minimum ramp up time for $I_{start} = 1.94\ A$ and $I_{stop} = 57.8\ A$ (corresponding to a ramp from 60 MeV to 1500 MeV with a quadrupole focussing strength of $k = 3.9\ m^{-2}$) is

$$T_{Q,60\rightarrow1500} = 1.955s.$$

Therefore, the minimum repetition rate for 1500 MeV injections from booster into storage ring is limited by the maximum and minimum voltages of the dipole magnet power supply of the booster.

Steerer and Sextupole Magnets

Steerer magnets are dipole magnets with small magnetic fields that are used for orbit correction. In Bodo, the quadrupole magnets are also used to generate steerer (i.e. dipole) and sextupole magnet fields by appropriate additional coils on and inside the magnets.

Figures 3.3 show Bodo quadrupole magnets with additional coil windings to obtain a horizontally (left picture) or vertically (right picture) deflecting magnetic field in the center of the magnet. The orbit change Δu generated by a dipole field is proportional to the beta function β_u at the location of the kick. Therefore, horizontally deflecting steerer coils are installed at local maxima of β_x (i.e. at horizontally focussing quadrupoles, see chapter 2), and vertically deflecting steerer coils at local maxima of β_z (i.e. at vertically focussing quadrupoles).

Sextupole fields are generated by superposition of the magnetic fields of two sets of coils: one set of windings on the magnet jokes ("outer coil set", exactly like the horizontal steerers) and another set on aluminum carriers directly above and below the beam pipe resp. the ideal orbit ("inner coil set") [29]. By connecting the two sets in series and choosing an appropriate geometry and number of coil windings, the dipole field of the outer coils is eliminated by the inner coils, and the remaining magnetic field is a sextupole field. Higher multipole contributions that can be neglected except for orbit perturbations in excess of $\pm13\ mm$ [29]. All horizontally focussing quadrupoles of Bodo have such additional sextupole windings that are connected in
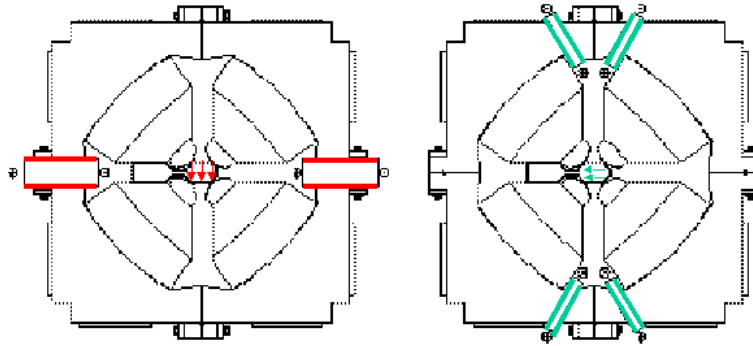
Figure 3.3: *Vertical ("VK") and horizontal ("HK") steerer magnets (i.e. orbit correctors) of Bodo. The red and green lines indicate the magnet coils that generate the additional dipole magnet field in the quadrupole center. The red and green arrows in the center of the magnets indicate the direction of the magnetic field.*

series to a single unipolar power supply ("horizontal sextupoles"), and all vertically focussing quadrupoles except the QD1 quadrupoles (see fig. 3.1) also have sextupole windings that are connected in series to another unipolar power supply ("vertical sextupoles").

Since Bodo is a ramped storage ring, the resulting time-dependent magnetic fields of dipole and quadrupole magnets generate eddy current in the beam pipe, which consists of stainless steel of 3 mm thickness (see section 3.5). Due to the long ramp cycle periods compared to booster synchrotrons that use White Circuits, the resulting magnetic fields (e.g. sextupole magnet fields generated by eddy currents in the dipole magnet beam pipe) can be neglected [31].

As mentioned in section 3.4, Bodo has three steerer magnet coils ("DC extraction coils") named HK-DC1, HK-DC2 and HK-DC3 (see fig. 3.4) that are mounted onto 20 degree bending magnets. These magnets allow to generate a local orbit bump ("DC extraction bump") which is required for the extraction of the beam. Orbit measurements showed that the 10 degree dipoles required additional steerer magnet coils (named "HK-DCB") to compensate orbit perturbations caused by these magnets (see section 9.1.1). While the normal steerer magnets that are integrated into the quadrupoles have bipolar power supplies, all HK-DC coils only have unipolar power supplies.

3.3 Injection

The beam is injected from the Linac and the transfer line T1 into Bodo by an eddy current injection septum magnet and two injection kicker magnets. The injected beam moves along one side of the copper septum blade of 1 mm thickness, where the magnetic field of the pulsed septum is high and the beam is bent in the direction of the closed Bodo orbit. At the other side of the septum blade (towards the closed orbit), the magnetic field is nearly zero due to eddy currents in the blade. Therefore, the stored Bodo beam is not disturbed by the septum magnet pulse of 46 microseconds length when it passes the septum during the turns after injection. When electrons are injected into Bodo, the septum bends the electron beam nearly parallel to the closed orbit, injection kicker 2 kicks the beam so that it crosses the closed orbit at the location of injection kicker 3, and this kicker bends the beam onto the closed orbit (see fig. 3.5). This on-orbit injection (with kicker 1 not being used) is the standard injection mode for Bodo.
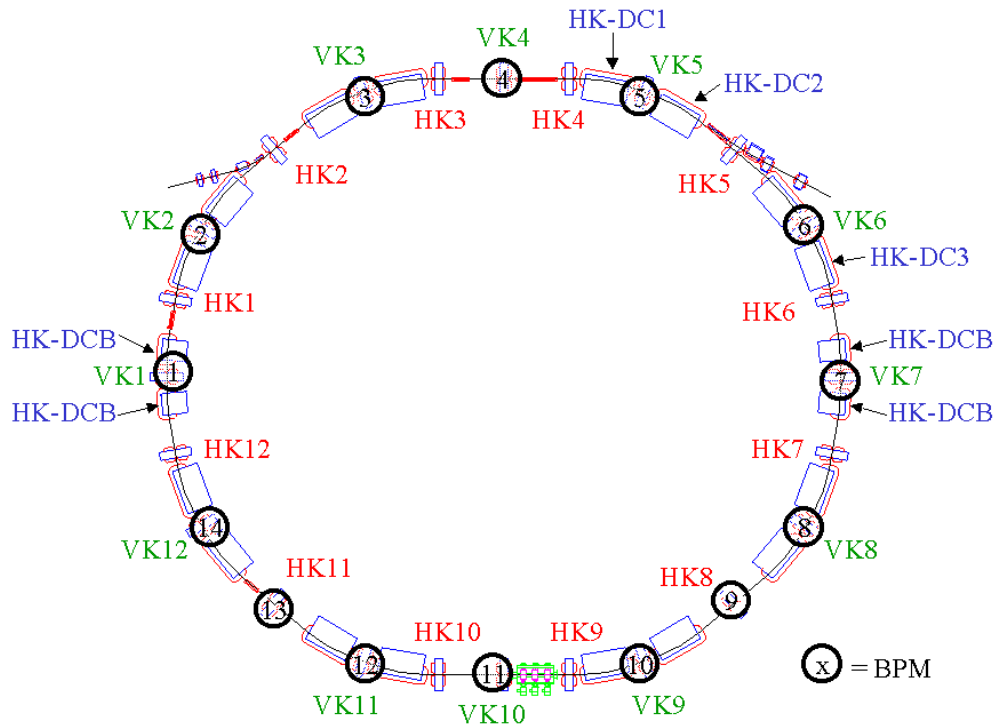
Figure 3.4: *Beam position monitors, horizontal (red, "HK") and vertical (green, "VK") steerer magnets (i.e. orbit correctors) of Bodo. The blue magnet names ("HK-DC") indicate additional orbit correctors that are integrated into some of the 10 and 20 degree dipole magnets by additional coils on the dipole magnet yokes, with dedicated power supplies for these coils. The "HK-DCB" orbit correctors are additional steerer coils windings on the four 10 degree bending magnets, with the coils of all four magnets connected in series to a single power supply.*

The magnetic fields of the kicker magnets have short rise and fall times of only one Bodo turn (i.e. 168 ns). Therefore, the kicker timing system can be adjusted so that all kickers have their full amplitude when a sufficiently short electron bunch train is injected and have zero amplitude on the following turns, so that the beam orbit is not disturbed on subsequent turns.

3.4   Extraction

The beam is extracted from Bodo by kicking it with the extraction kicker magnet so that it moves through the slit of the extraction septum that deflects it into the transfer line T2. The kicker has a rise time of only one Bodo turn so that a sufficiently short bunch train is kicked with the full amplitude at the turn of extraction without being perturbed during the previous turns. To keep the required voltage for the capacitors of the extraction kicker within reasonable limits, three steerer magnet coils ("DC extraction coils") named HK-DC1, HK-DC2 and HK-DC3 (see fig. 3.4 at the end of the following section) that are mounted onto 20 degree bending magnets generate a local orbit bump ("extraction bump") that moves the stored beam close to the blade of the extraction septum. Therefore, the extraction kicker amplitude has just to be large enough to shift it across the blade. Fine-tuning of orbit angle and amplitude at the extraction septum can be achieved by an additional local orbit bump using four horizontal steerer dipole magnets (see section 3.2) that can control orbit angle and amplitude at the septum independently.
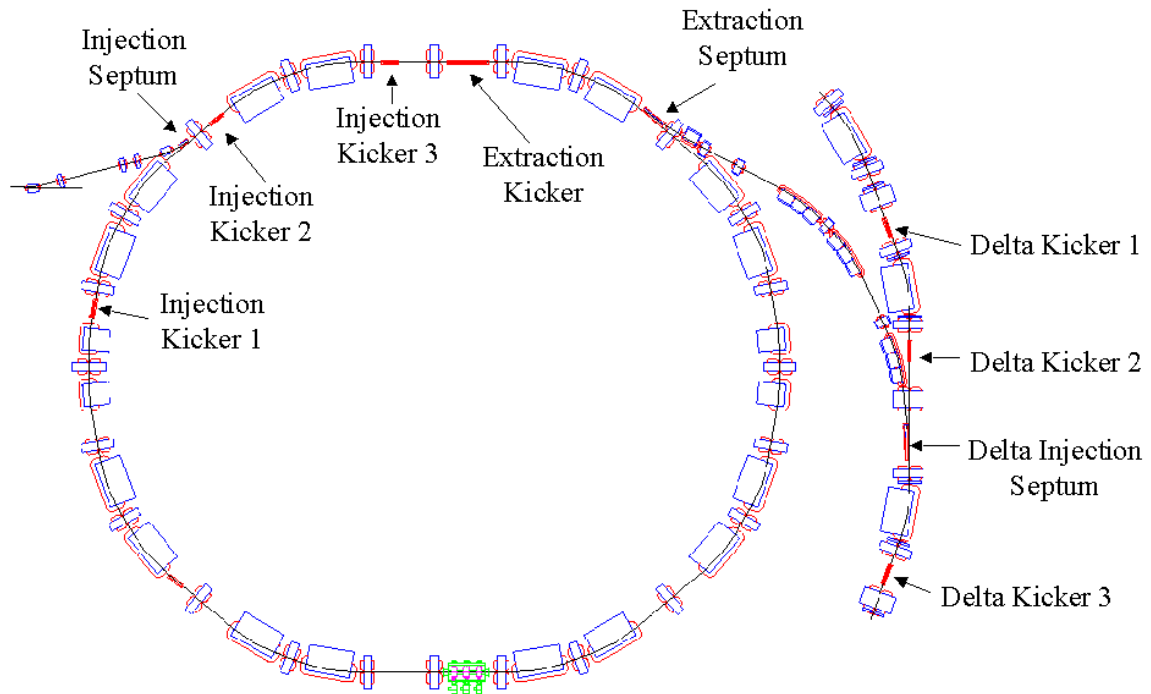
Figure 3.5: *Injection and extraction magnets of Bodo.*

After passing the transfer line, the beam is injected into the storage ring Delta via the Delta injection septum and a local injection kicker bump in the storage ring (see fig. 3.5). At 1.5 GeV, the Delta injection kicker orbit bump also has to be supported by a static local orbit bump generated by steerer dipole magnets in order to move the stored beam sufficiently close to the Delta injection septum blade.

3.5 Beam Position Monitor Pickups and Beam Pipe

Fig. 3.6 shows a cross section of a capacitive button beam position monitor ("BPM") that is used both in Bodo and Delta to measure the beam orbit. The monitors are integrated in the beam pipe that normally consists of stainless steel (material type "sst 1.4429") of 3 mm thickness. The pipe has a keyhole cross section and consists of two tubes that are connected by a so-called pumping slit of 8 mm height and 35 mm width. The electron beam moves through the right tube at the outside of the storage ring. The left tube at the inner side of the storage ring contains integrated vacuum pump modules [30]. The distance from the ideal orbit in the center of the right tube to the walls of the beam pipe is 20 mm in the vertical and 37 mm in the horizontal direction. At the locations of the BPMs, the beam pipe is thicker than the usual 3 mm to allow the installation of the BPM pickup buttons. Furthermore, the outer shape of this thick pipe fits exactly into the quadrupole magnets and allows a movement of ± 1.8 mm for the so-called variable-type and ± 0.07 mm for the so-called fixed-type BPMs. These two BPM types, having slightly different beam pipe diameters, are meant to avoid magnet movements due to bending of the beam pipe, which may be caused by temperature variations of inner and outer side of the pipe. The different diameters allow the beam pipe (which lies in the quadrupole magnets) to move at certain locations while being fixed at other locations.
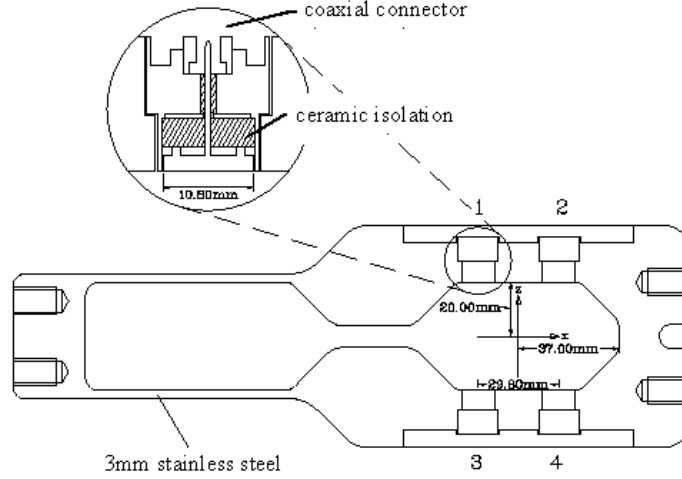
Figure 3.6: *Cross section of the DELTA beam pipe at the location of a beam position monitor (BPM). One of the four capacitive button pickups is magnified. The pickup buttons have a diameter of 10.8 mm and are isolated from the beam pipe by* Al_2O_3 *ceramic.*

When the electron beam passes the four capacitive button BPM pickups, it induces a voltage at each button that depends on the beam position, on the beam current and on the longitudinal charge distribution. In the vicinity of the ideal orbit, the horizontal and vertical beam position x and z can be obtained from these voltages [24]:

$$x = a_x(x,z)\frac{U_2(x,z) - U_1(x,z) + U_4(x,z) - U_3(x,z)}{U_1(x,z) + U_2(x,z) + U_3(x,z) + U_4(x,z)} \qquad (3.1)$$

$$z = a_z(x,z)\frac{U_1(x,z) - U_3(x,z) + U_2(x,z) - U_4(x,z)}{U_1(x,z) + U_2(x,z) + U_3(x,z) + U_4(x,z)} \qquad (3.2)$$

The so-called monitor coefficients $a_x(x,z)$ and $a_z(x,z)$ are functions that depend on the geometry of the beam pipe and the locations of the pickup buttons. They are nearly constant at small values of x and z, but change up to some 10 percent at large values of x and z. At present, the orbit measurement systems used for Bodo and Delta use $a_x(0,0)$ and $a_z(0,0)$ to calculate the orbit from the button voltages. This causes a systematic difference (i.e. error) between calculated and real orbit for large values of x and z. Figures 3.7 and 3.8 show the difference between the beam positions $x_{sig,0}$ and $z_{sig,0}$ that were calculated using the monitor coefficients $a_x(0,0)$ and $a_z(0,0)$ at the ideal orbit and the real beam positions x_{real} and z_{real} Figures 3.9 and 3.10 show the respective relative differences. The BPM pickup voltages and the required sums and differences that were used for the figures originate from a numerical simulation that was made for the design of the DELTA BPM system [24].

As long as the orbit perturbations are smaller than $\Delta x = \pm 5$ mm and $\Delta z = \pm 4$ mm, the absolute difference is smaller than 200 μm or 5%, which is in the order of the mechanical positioning error of the BPM beam pipe [24]. Furthermore, the most important application of BPMs in Bodo and Delta is the reproduction of certain reference beam orbits, i.e. to measure and avoid deviations from these reference orbits. This requires a high stability and reproducibility of the measurement system, whereas smaller reproducible differences between measured and real orbit are mostly of little practical importance. Therefore, $a_x(x,z)$ and $a_z(x,z)$ can be substituted by $a_x(0,0)$ and $a_z(0,0)$ as long as orbit perturbations are not too large.
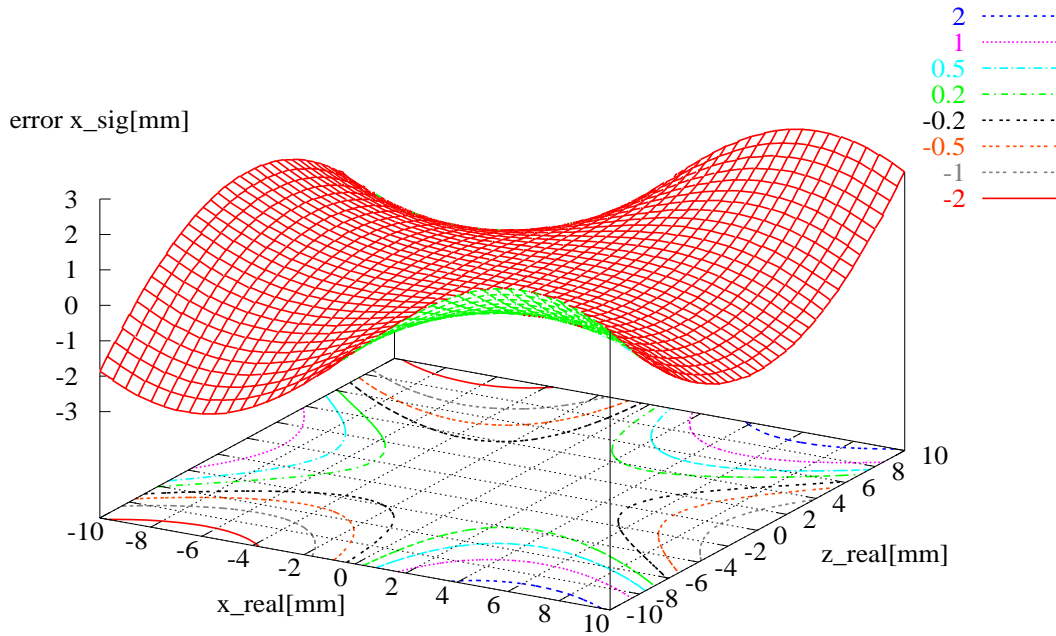
Figure 3.7: *Absolute difference between the horizontal beam position $x_{sig,0}$ calculated with the monitor coefficient $a_x(0,0)$ and the real beam position x_{real}, for different real beam positions x_{real} and z_{real} [24]. The coloured lines in the (x,z)-plane indicate locations of same difference (ranging from -2 mm to $+2$ mm) between calculated and real orbit.*
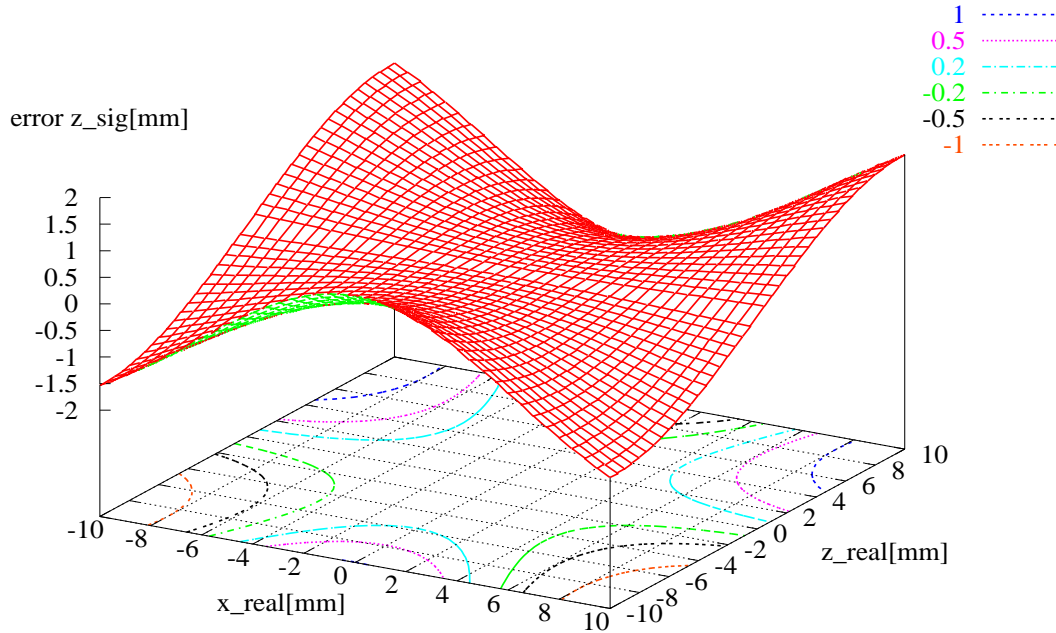
Figure 3.8: *Absolute difference between the vertical beam position $z_{sig,0}$ calculated with the monitor coefficient $a_z(0,0)$ and the real beam position z_{real} [24].*
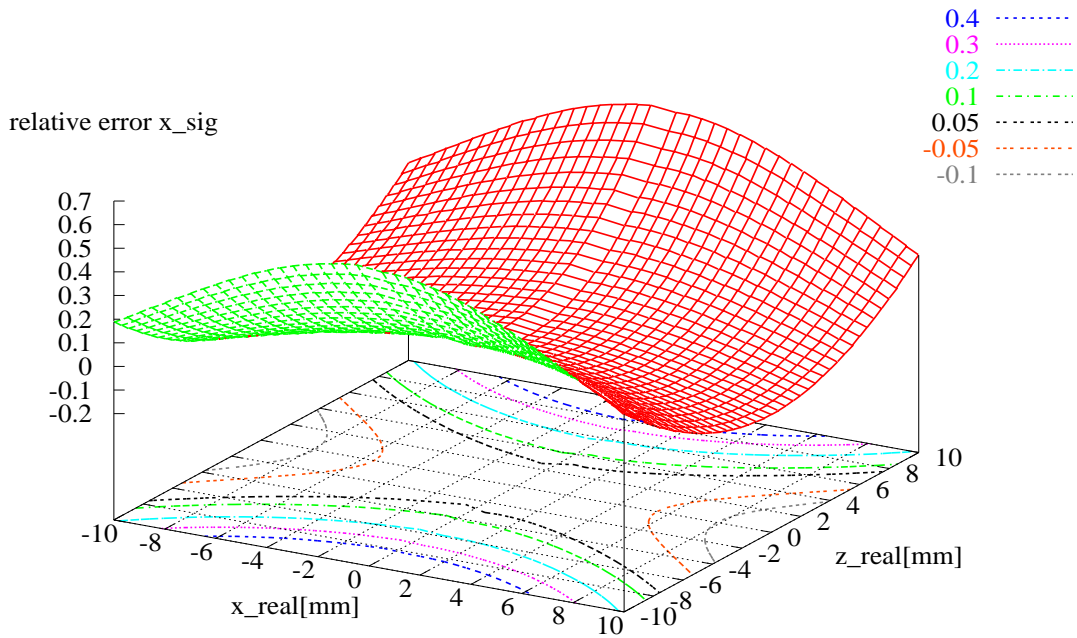
Figure 3.9: *Relative difference between the horizontal beam position $x_{sig,0}$ calculated with the monitor coefficient $a_x(0,0)$ and the real beam position x_{real}, for different real beam positions x_{real} and z_{real} [24]. The coloured lines in the (x,z)-plane indicate locations of same relative difference (ranging from -10% to $+40\%$) between calculated and real orbit.*
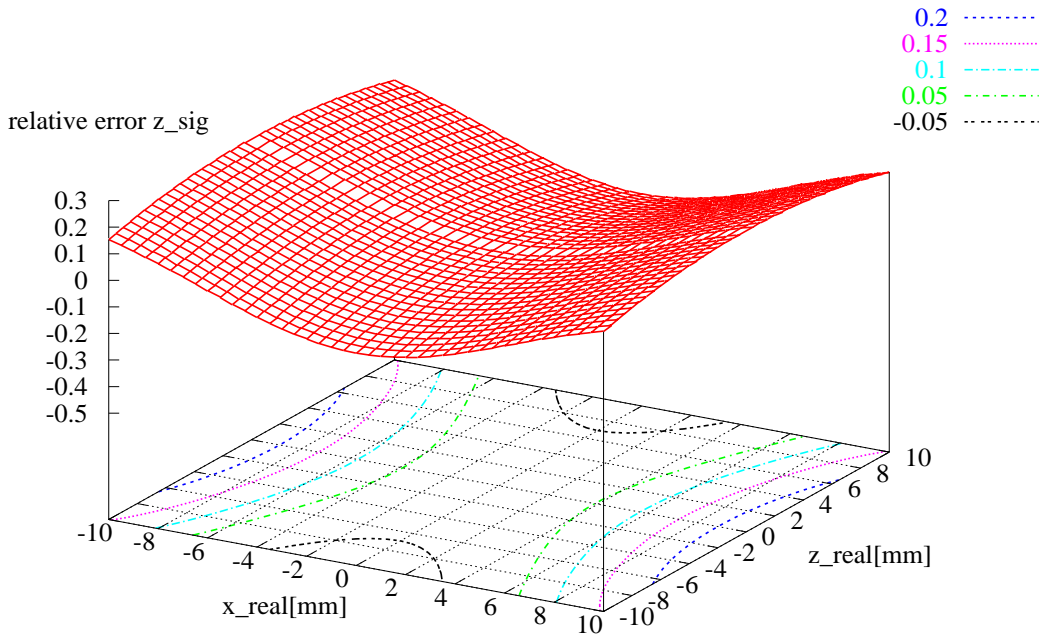
Figure 3.10: *Relative difference between the vertical beam position $z_{sig,0}$ calculated with the monitor coefficient $a_z(0,0)$ and the real beam position z_{real} [24].*

The frequency ("Fourier") spectrum of the BPM button voltages depends on the longitudinal beam charge distribution and is dominated by integer multiples of the beam revolution frequency. The n-th multiple of the revolution frequency is called n-th revolution harmonic. The Fourier amplitudes of these revolution harmonics depend on the number of electrons in the different RF buckets (see fig. 6.4 on page 88). Most Fourier amplitudes can be very small or zero for certain patterns. If, for example, all bunches in a storage ring have the same charge and length, the longitudinal charge distribution is a periodic function that has a period equal to the RF cavity oscillation period. In this case, the amplitudes of all revolution harmonics are zero except for integer multiples of the so-called harmonic number h of the storage ring that is equal to the number of RF buckets (i.e. h is the RF cavity frequency divided by the revolution frequency). Therefore, the electronic systems that measure the beam orbit in Bodo and Delta use only this frequency component of the BPM pickup signals, which guarantees that the BPMs work for any filling pattern. A detailed description of the beam position measurement and feedback system that was developed in this thesis will be given in chapters 7 and 9.

Fig. 3.4 shows the locations of the BPMs and steerer magnets of Bodo. 12 of the BPMs are located in vertically focussing quadrupoles ("QDs", that are local maxima of the vertical betafunction β_z and minima of the horizontal betafunction β_x, see chapter 4), but only two BPMs are located in horizontally focussing quadrupoles ("QFs", with local maxima of β_x and minima of β_z). Therefore, the horizontal orbit perturbations measured by the BPMs in QDs are typically two to four times smaller than the perturbations in adjacent QFs, depending on the beam optics (see chapter 4).

3.6   RF System

The 500 MHz RF cavity that accelerates the electrons in Bodo consists of three cells, with an overall shunt impedance of $R_s = 9$ $M\Omega$. The maximum RF forward power of the transmitter that powers the cavity is 30 kW. During the work on this thesis the power was limited to 18 kW in order to reduce the heat load and therefore the risk of damage of the ceramic window that separates the ultra high vacuum in the cavity from the air in the RF waveguide that powers the cavity. The transmitters of the cavities in Bodo and Delta are connected to the same RF frequency generator. Since the length of the closed beam orbit is inversely proportional to the RF frequency, the RF generator is usually tuned to the frequency at which the horizontal beam orbit of the storage ring Delta is centered (typically 499.820 MHz). This frequency is not constant, but may change due to temperature changes of the DELTA machine hall and of its concrete floor, which results in magnet movements on a time scale of days to months, e.g. due to daily and seasonal temperature changes of the environment. The RF frequency is one of many parameters that is changed by the Delta orbit correction system in order to optimise the orbit [32]. The orbit correction system applies corrections to the Delta orbit corrector magnets and RF frequency at a repetition rate of 2 Hz. Therefore, the booster RF frequency cannot be chosen independently, and it must be possible to operate the booster at an RF frequency that may change from cycle to cycle.

The RF power of the booster cavity is ramped synchronously with the booster magnet currents to provide low RF power at the beam injection energy (60 MeV) and high RF power at the extraction energy (up to 1.5 GeV). The mechanical frequency tuning system of the cavity consists of three motor-driven plungers. It is only activated for a short period at the top of the

cavity power ramp just before extraction. This guarantees that the cavity is resonant at 1.5 GeV, i.e. when the highest accelerating voltage is required.

At 1.5 GeV, the energy loss of an electron per turn is 134.1 keV [27]. Assuming a maximum cavity power of 14 kW, the maximum accelerating voltage in the zero beam current limit is 0.5 MV, and the resulting RF phase angle $\Phi_{1500,min}$ (see eq. (2.11)) at 1.5 GeV is

$$\Phi_{1500,min} = 15.5^o \qquad (3.3)$$

In order to obtain a constant longitudinal tune Q_s and synchrotron frequency f_s in the zero current limit, the cavity power P_{RF} has to increase approximately with the square of the beam energy (see eq. 2.12), provided that the cosine of the cavity phase angle is close to unity. A constant synchrotron frequency has the advantage that the beam does not have to cross synchro-betatron resonances (see section 2.4) during the booster energy ramp if the betatron tune frequencies are also kept constant. However, the booster RF power ramp that is used for energy ramps between 60 MeV and 1.5 GeV is not proportional to the square of the beam energy ramp, but it was modified empirically with respect to such a ramp both to maximise the booster beam current that can be ramped up to 1.5 GeV and to achieve stable operation of the plunger regulation circuit and the cavity cooling system.

Chapter 4

Optimisation of Booster and Transfer Line Optics

As explained in the introductory chapter, the object of this thesis is the development and application of a unified DSP-based beam diagnostics and feedback system that can measure and minimise deviations of booster and storage ring beam parameters from desired "optimal" values. The first applications of the system include the measurement and real-time correction of the beam orbit and betatron tunes of Bodo (see chapters 6 to 9).

As far as the orbit of Bodo is concerned, "optimal" usually means that the orbit passes through the magnetic centers of the quadrupoles, except for well-defined local orbit bumps that may be required for beam extraction or injection. However, the definition of "optimal" with respect to betatron tunes and beam optics is less obvious, since the six quadrupole families of Bodo allow a variety of different beam optics and betatron tunes.

When being used as an injector for Delta, an "optimal" booster optics should maximise the beam charge that can be injected into Delta while minimising beam loss and radiation levels during the charge transfer. During the first years of operation, typical charge transfer rates of 10 to 20 %, average Bodo beam currents of 3-6 mA and ramp periods of 7.3 s resulted in typical storage ring filling times of 30 to 80 minutes in order to accumulate 100 mA. Thus, the long time required for refilling the storage ring after hardware failures with subsequent beam loss led to significant reduction of available beam time for synchrotron radiation users.

Beam parameter feedbacks for Bodo and Delta cannot achieve optimal machine performance if the reason for suboptimal performance is not only the deviation from the desired optics and orbit, but also the desired optics itself. Charge transfer rates of up to 100 % at lower Delta energies indicated that the charge transfer rate from Bodo to Delta might be improved by booster optics with lower equilibrium emittance than the "bo-006b" Bodo optics that was used so far, since the emittance scales with the square of the energy. Therefore, new Bodo optics with lower emittance than the "old" "bo-006b" optics were developed during the work on this thesis, as well as suitable optics for both transfer lines. One of the new Bodo optics ("bo-101c") was used for nearly all tests of the DSP-based feedback systems in the subsequent chapters, and it is also used for regular 1.5 GeV operation of Bodo since December 2001, together with new transfer line optics.

Since the main subject of this thesis is the DSP-based measurement and optimisation of beam parameters, this chapter will only motivate the choice of the "bo-101c" optics and discuss some theoretical aspects and properties of old and new optics that are relevant for DSP-based diagnostics and feedback applications in the following chapters. A more detailed overview of the development and simulation of the new Bodo and transfer line optics can be found in ref. [52]. An experimental in-depth investigation of the optics and of related beam dynamics during injection and extraction is beyond the scope of this thesis, but is envisaged in the context of two other ongoing Ph.D. theses [33, 34]. Furthermore, such investigations would have required beam diagnostics systems like transfer line BPMs, energy spectrometers, optical diagnostics for

phase space tomography or booster BPMs with turn-by-turn resolution that were not available during the work on this thesis for cost and manpower reasons.

Section 4.1 summarises the advantages of a small Bodo emittance both for normal operation and for the envisaged top-up operation of Delta. Sections 4.2.1 and 4.2.2 describe the development and properties of the new optics, as well as simulations of their sensitivity to magnetic field and alignment errors. Sections 4.2.3 and 4.2.4 give an overview of the transfer line optics and the Delta injection process, and section 4.3 concludes the chapter with a summary and discusses the improvement of the charge transfer between Bodo and Delta that was achieved with the new optics.

4.1 Advantages of a Low Booster Beam Emittance

In order to achieve a high charge transfer efficiency from booster to storage ring, the transverse beam size of the booster must be sufficiently small to fit into the available transverse phase space of the transfer line and the storage ring. The phase space is limited either by the mechanical aperture of the beam pipe or by the so-called dynamic aperture which is the maximum storage ring aperture in which particles may circulate without getting lost by chaotic particle movement due to nonlinear (e.g. sextupole) magnetic field components [6]. If the emittance of the booster beam is too large, many of the electrons that are extracted from the booster will hit either the beam pipe of the transfer line T2 and the storage ring Delta (see fig. 1.1) or the septum blades of booster extraction or storage ring injection septum and get lost.

Even if the emittance of a synchrotron light source booster is just small enough to allow a theoretical charge transfer rate of 100%, alignment tolerances, magnetic field errors, stray fields and suboptimal settings or drifts of magnet currents may lead to charge transfer rates that are much lower. In the case of DELTA, the transfer line T2 from Bodo to Delta contains two dipole and two septum magnets that are pulsed by discharging a voltage-controlled capacitor bank, which is equivalent to a voltage-regulated power supply. Therefore, temperature drifts of the magnet coils change their current, which changes the orbit of the injected beam and leads to a low charge transfer efficiency and beam loss in T2 and Delta. Furthermore, the transfer line and storage ring do not have enough beam position monitors that could be used by machine operators to measure the beam orbit during the first injection turn precisely enough to correct these magnetic field drifts. Moreover, the storage ring itself has too few beam position monitors to measure the orbit and angle of the stored beam at the location of the injection septum blade precisely, while the uncorrected storage ring orbit may drift up to several mm per hour due to temperature-related magnet movements and field drifts of a superconducting wiggler magnet [32]. However, the orbits of injected and stored beam are essential for a high charge transfer efficiency, since they determine the distance between injected and stored beam in the transverse phase space. These problems can also be alleviated by a decrease of the booster beam emittance, since a smaller beam size makes the charge transfer efficiency less sensitive to orbit drifts in transfer line and storage ring.

4.1.1 Booster Emittance and Top-Up Injection

Beam loss due to large booster emittance is also a problem for the so-called top-up operation of synchrotron light sources. Top-up means that the storage ring is not filled just every few hours

with closed beam shutters of the synchrotron light beamlines during injection. Instead, electrons are injected quasi-continuously every few seconds or minutes with opened beam shutters in order to maintain a nearly constant beam current in the storage ring [1, 2].

If new electrons are injected only every few hours with a typical beam current decrease of 30 to 50 percent in between, this leads to a change of the synchrotron light heat load both on the beam pipe of the accelerator and on the beamline experiments. The resulting movement of the Delta beam pipe leads to movement of magnets and therefore of the electron beam [36]. The movement of the synchrotron light often leads to suboptimal results of beamline experiments, because it can change the intensity or energy of the monochromatic part of the synchrotron light that is used for an experiment, move the synchrotron light beam away from a sample etc. Furthermore, the changing heat load on components in a synchrotron light source beamline e.g. on mirrors or monochromator crystals can limit the resolution and performance of the beamline, e.g. by deformation of mirrors or monochromator crystals.

Due to its quasi-constant beam current, top-up mode operation of a storage ring avoids these problems by a constant heat load on beam pipe and beamlines. This can significantly improve the quality and reproducibility of the results of synchrotron radiation experiments. Furthermore, top-up mode leads to a higher average beam current and a quasi-infinite beam lifetime that is nearly independent of vacuum pressure, emittance coupling, or small vertical aperture due to mini-gap insertion devices.

In order to use top-up mode injection, beam loss during injection must be so small that the beam shutters which protect the beamlines from radioactive radiation can be left open without exceeding the permitted radiation levels in the beamline experimental stations. This requires a sufficiently small booster emittance. Furthermore, high beam loss rates might be dangerous for low gap permanent magnet undulators, since beam loss can lead to partial demagnetisation of these magnets [39].

Due to the numerous disadvantages of a large booster emittance, booster synchrotrons of modern state-of-the-art synchrotron light sources are designed to have a very small beam emittance at the extraction energy of the beam. The booster synchrotron of the Swiss Light Source (SLS) which uses top-up operation [2] has a nominal emittance of 9 nm rad at 2.4 GeV, which corresponds to 1.56 nm rad at 1 GeV [37].

Compared to this, the beam optics ("bo-006b") that was used for the booster synchrotron Bodo until December 2001 had an emittance of 212 nm rad at 1 GeV, which is 136 times larger and results in a transverse beam size that is nearly 12 times larger if the beta functions have the same size. This difference results from the different size and magnet lattices of the two boosters. The large circumference of the SLS booster (270 m) allows the installation of 93 combined function bending magnets with bending radii of 50.7 m and 11.2 m. Bodo has only 20 bending magnets with a bending radius of 3.35 m due to its comparatively small circumference of 50.4 m. Furthermore, the lattice of the SLS booster was optimised for a small beam emittance, resulting in an average dispersion of about 0.3 m, whereas the average dispersion of the Bodo "bo-006b" optics is about 1 m. According to eq. (2.33) and (2.36), the equilibrium emittance decreases with decreasing dispersion (resp. "H-function", eq.(2.36)) in the dipole magnets and also decreases with increasing dipole magnet bending radius, which explains the much smaller emittance of the SLS booster.

4.2 Development of Low Emittance Bodo Optics

4.2.1 Emittance Optimisation

In case of regular FODO lattices (consisting of equally spaced focussing and defocussing quadrupoles with equally spaced dipoles of the same type in between) the emittance can be derived and minimised analytically. The minimum emittance is reached for a betatron phase of approx. 135^o per FODO cell [7]. However, Bodo consists of an irregular FODO lattice with missing bending magnets, different kinds of bending magnets, varying distances between quadrupoles and six instead of two quadrupole families. Furthermore, boundary conditions (e.g. power supply current limits) have to be accounted for.

Therefore, a numerical method was used to find beam optics for Bodo that have a lower emittance than the "bo-006b" optics that was used so far. An iterative gradient descent method (using the PC program "Optics" [27]) was applied to the quadrupole focussing strengths ("k-values") of the six Bodo quadrupole families, starting with the k-values of existing optics, e.g of the the "bo-006b" optics.

Since the quadrupole magnets of Bodo have a maximum absolute focussing strength of $|k|_{Q,max,1.5} = 3.98 \ m^{-2}$ at 1.5 GeV due to the maximum quadrupole power supply current of 60 A, the iteration algorithm should not generate beam optics with larger k-values. Furthermore, the beta functions should not exceed a certain limit, since large beta functions lead to large orbit perturbations caused by quadrupole magnet misalignments (see eq. (2.52)). Therefore, an additional boundary condition is required to avoid excessive beta functions. In the case of Bodo, calculations show that optics with beta functions larger than 20 m often required k-value variations as small as a few percent until the optics program could not find a solution for the optics (i.e. there was no stable optics for that k-value setting, which practically results in beam loss). Therefore, the iteration method should not generate optics with beta functions much larger than 20 m. Furthermore, large beta functions result in a large beam size in Bodo at 60 MeV when the Linac beam of fixed emittance is injected, which reduces the available transverse phase space and therefore can lead to beam loss and reduced beam currents.

The algorithm that was developed (see fig. 4.1) varies the focussing strength ("k-value") of each of the 6 quadrupole families iteratively in small steps, typically $0.0005 \ m^{-2}$. Each quadrupole family is varied until a local emittance minimum is reached, the k-value is at the current limit, or the beta functions become too large. In the latter case, the algorithm tries to find another quadrupole family that reduces the beta functions below the limit again while still achieving an overall emittance reduction. If this reduction can be done with several families, the algorithm will select the family that achieves the lowest emittance when reducing the beta functions to the limit value. Thus, the algorithm reduces the emittance iteratively. It stops either if the desired emittance is reached or if the search takes too long ("timeout").

4.2.2 Properties of Previous and New Bodo Optics

The emittance optimisation method resulted in several new booster optics, three of which are presented in table 4.1. The theoretical equilibrium emittances of the optics are 101 nm rad ("bo-101c"), 88.1 nm rad ("bo-101e") and 65.0 nm rad ("bo-103b") at 1 GeV beam energy, compared to 212 nm rad for the "old" optics ("bo-006b"). The "bo-101e" optics has the smallest
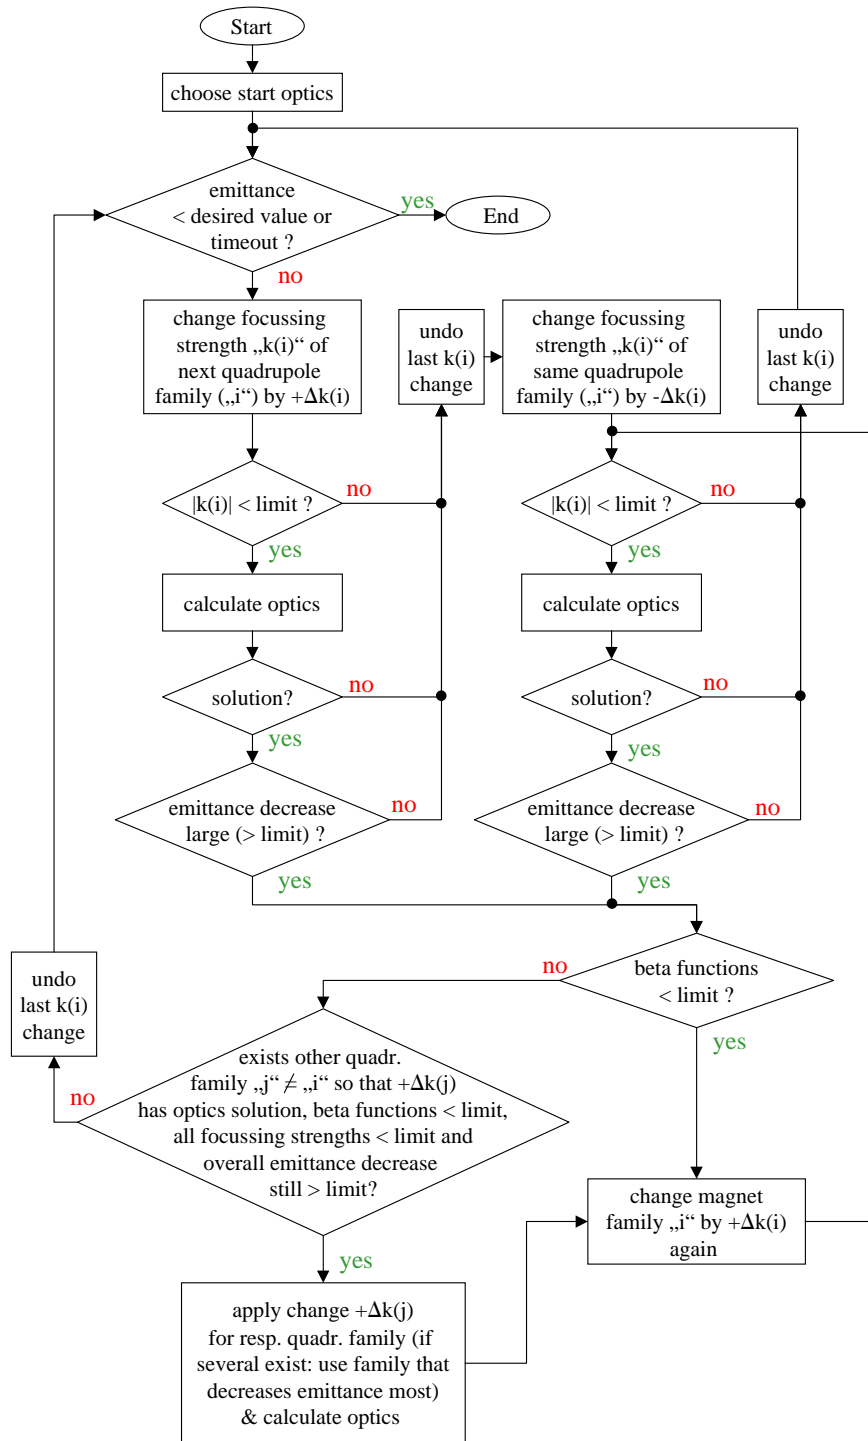
Figure 4.1: *Emittance optimisation algorithm.*

emittance that was obtained with the algorithm for a maximum absolute k-value of $3.95 m^{-2}$, while providing moderate maximum beta functions of about 10 m and a betatron tune with sufficient distance from destructive resonances. In case of the "bo-103b" optics, the limitation for the maximum k-value was ignored, and a further reduction of the emittance was achieved. Therefore, this optics requires a very strong focussing that cannot be achieved with the available quadrupole power supplies (see below). Since the equilibrium emittance scales with the square of the beam energy, all emittances were calculated for an energy of 1 GeV to allow easier comparison and scaling.

The "bo-101c" optics was actually not obtained directly from the emittance optimisation algorithm, but experimentally by varying the quadrupole focussing strengths, betatron tunes and orbit corrector kicks of the "bo-101e" optics until the average beam current of the booster reached a local maximum, while simultaneously adjusting the transfer line and injection kicker magnets appropriately after each variation ("transfer line matching" [6]). As will be shown in section 4.2.2.1, the "bo-101c" optics has a very low sensitivity to magnet alignment and field errors. This low sensitivity and the small number of available orbit correctors during first tests of the optics explain why the obtained beam current was higher than for the "bo-101e" optics.

The parameters in table 4.1 were calculated with the PC program "Optics", except for the chromaticities which were also calculated with the program "MAD" [26] because both programs use a different calculation method and thus obtain different values. While "Optics" uses the plain Hill's equations (2.2) and (2.3) and only scales the quadrupole focussing inversely to the electron energy, "MAD" accounts for additional effects, e.g. the change of the optics due to the energy-dependent length change of the electron orbit which cannot be neglected for accelerators with a large momentum compaction factor. Therefore, the chromaticities calculated with "MAD" are different from (and more realistic than) the chromaticities calculated with "Optics". However, the results of "MAD" and "Optics" e.g. for the betatron tunes are identical,

Parameter / Optics	bo-006b	bo-101c	bo-101e	bo-103b
Emittance [$nm\ rad$], 1 GeV	212	101.1	88.1	65.0
horizontal tune	2.811	3.6244	3.7877	4.8563
vertical tune	2.277	2.5586	2.5415	2.2205
horiz. chromat. ("Optics")	-3.032	-4.989	-5.786	-10.278
horiz. chromat. ("MAD")	-3.466	-4.920	-5.559	-9.9181
vertical chromat. ("Optics")	-2.684	-3.529	-3.493	-4.0764
vertical chromat. ("MAD")	-0.715	-2.117	-2.3249	-3.5192
momentum compaction factor	0.117	0.0754	0.0687	0.04545
damping partition Jx	0.7188	0.819	0.835	0.8909
k(QF1) [m^{-2}]	-2.477	-2.9559	-3.0700	-4.6140
k(QF2) [m^{-2}]	-2.907	-3.9395	-3.95	-4.1359
k(QF3) [m^{-2}]	-2.529	-3.037	-3.1916	-3.6588
k(QD1) [m^{-2}]	1.97	2.8573	2.8854	2.5736
k(QDI) [m^{-2}]	1.97	2.1864	2.1335	2.2123
k(QDII) [m^{-2}]	1.97	2.1538	2.2398	2.2992

Table 4.1: *Calculated parameters of different previous ("bo-006b") and newly developed ("bo-1....") Bodo beam optics.*

Figure 4.2: *Calculated beta functions (**red:** $\beta_{\mathbf{x}}$**, green:** β_z*, left axis) and dispersion (**blue:** D_x*, right axis) of the Bodo optics* **"bo-006b"**. *The horizontal axis is the s coordinate (left: s = 0, right: $s = L_{0,Bodo} = 50.4$ m). Blue boxes on the horizontal axis are dipoles, red and green boxes are horizontally and vertically focussing quadrupoles. The emittance is* **212 nm rad** *at 1 GeV.*

Figure 4.3: *Calculated beta functions (red: $\beta_{\mathbf{x}}$, green: β_z, left axis) and dispersion (blue: D_x, right axis) of the Bodo optics* **"bo-101c"**. *The emittance is* **101.1 nm rad** *at 1 GeV. The maximum absolute quadrupole focussing strength is* $3.94\ m^{-2}$ *(for quadrupole family QF2), which is slightly below the power supply limit of* $3.98\ m^{-2}$ *at 1.5 GeV.*

Figure 4.4: *Calculated beta functions (**red:** $\beta_{\mathbf{x}}$**, green:** β_z, left axis) and dispersion (**blue:** D_x, right axis) of the Bodo optics **"bo-101e"**. The emittance is **88.1 nm rad** at 1 GeV. The maximum absolute quadrupole focussing strength is* $3.95\ m^{-2}$ *(for quadrupole family QF2), which is slightly below the power supply limit of* $3.98\ m^{-2}$ *at 1.5 GeV.*

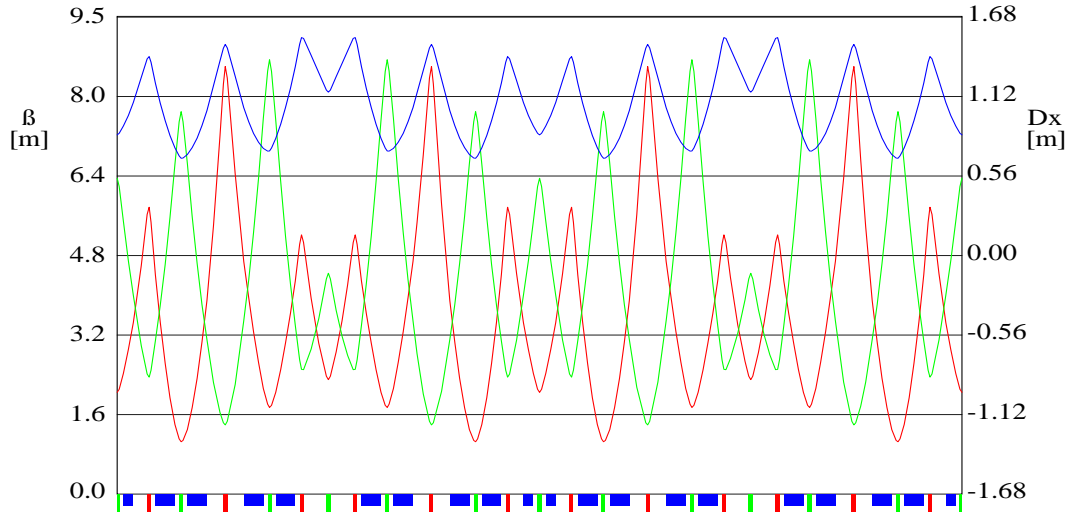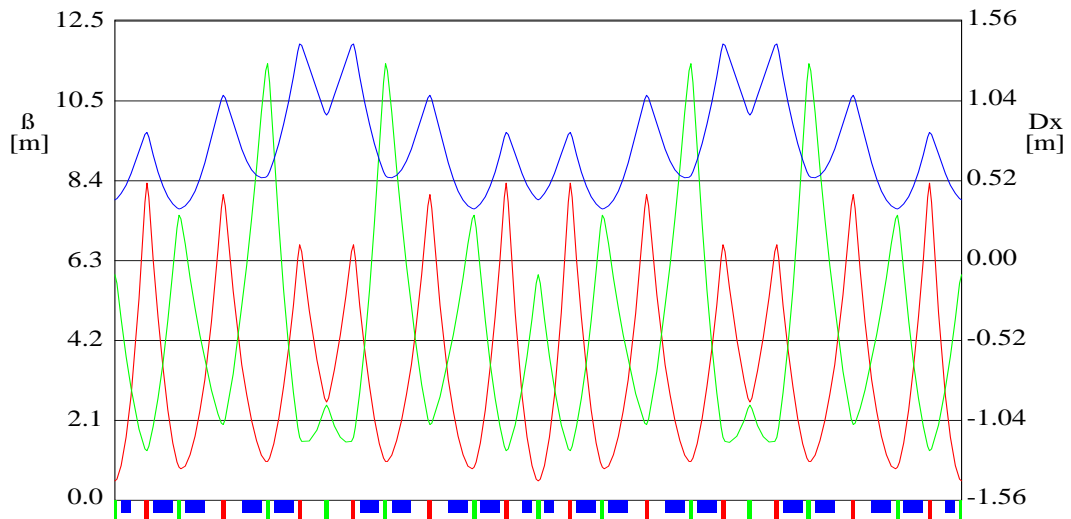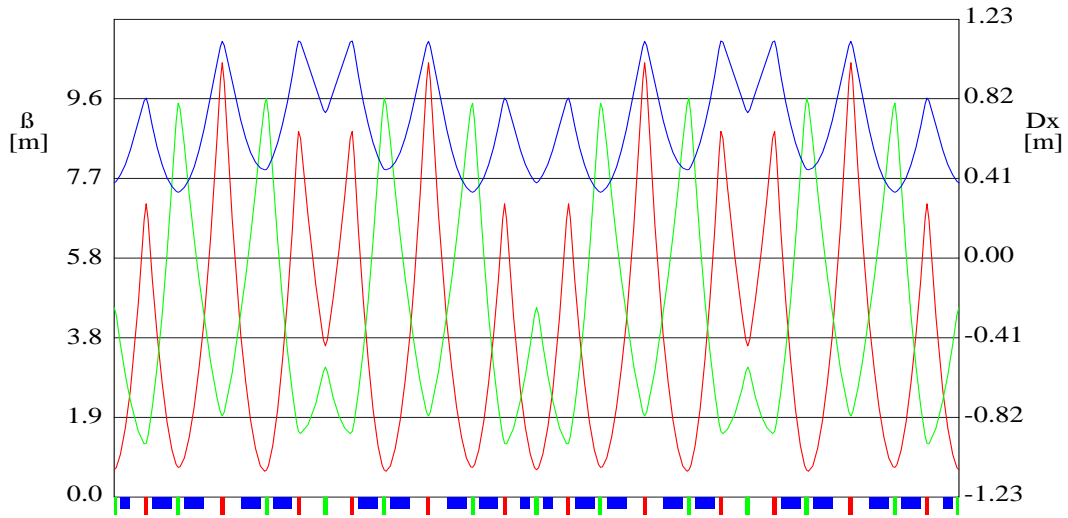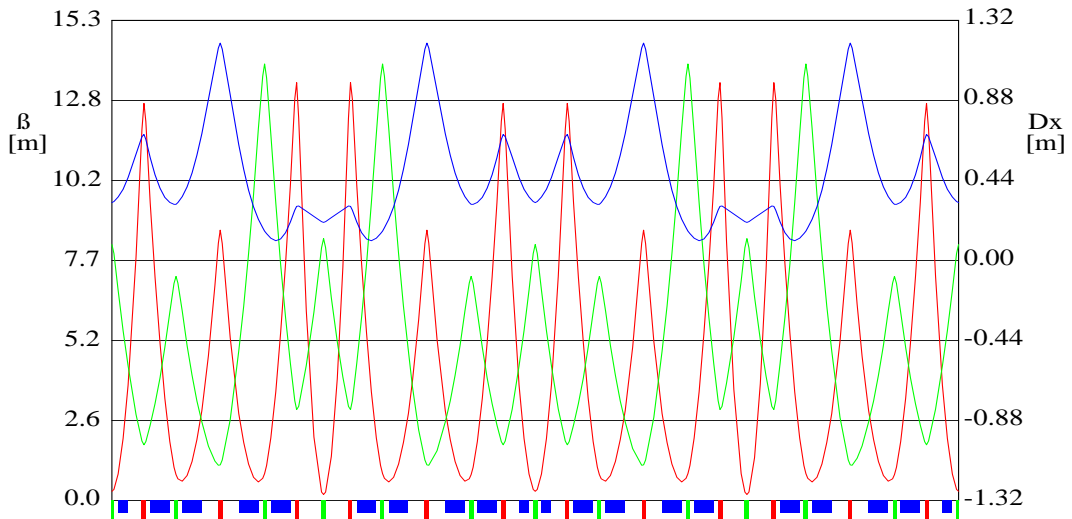Figure 4.5: *Calculated beta functions (red:* $\beta_{\mathbf{x}}$*, green:* β_z*, left axis) and dispersion (blue:* D_x*, right axis) of the Bodo optics **"bo-103b"**. The emittance is **65.0 nm rad** at 1 GeV, but the required quadrupole focussing strength (*$4.614\ m^{-2}$ *for family QF1 and* $4.136\ m^{-2}$ *for family QF2) exceeds the power supply current limits at 1.5 GeV.*

with relative differences below 10^{-4}.

The "bo-101c" and "bo-101e" optics can be used up to 1.5 GeV with the available quadrupole power supplies, as their maximum quadrupole focussing strength is 3.95 m^{-2} and the power supply limit is 3.98 m^{-2}. Since the actual extraction energy is typically 1.488 GeV instead of 1.5 GeV and the magnet current is not proportional to the magnetic field at high energy and high quadrupole strengths (due to saturation of the magnet iron), the required magnet current of 58.3 A has a small but sufficient safety margin to the maximum current of 60 A. This does not restrict smaller betatron tune adjustments before beam extraction.

By allowing magnet currents above the power supply limit, the "bo-103b" optics could achieve an emittance of 65.0 nm rad at 1 GeV. However, this requires new power supplies with currents of at least 73.25 A at 1.5 GeV for the quadrupole family QF1 and 64.2 A for QF2.

The emittance optimisation method was also applied to a Bodo magnet lattice with inverted polarities for the quadrupole magnet families, i.e. horizontally focussing and defocussing quadrupoles were exchanged. However, the smallest emittance that was obtained with the algorithm for such a lattice was 190 nm rad, compared to 65.0 nm rad for the "bo-103b" optics. The asymmetry, which does not exist for regular FODO lattices, can be explained by the irregular spacing of focussing and defocussing magnets, as well as by the different kinds of dipole magnets in Bodo. Polarity changes of single quadrupole families were also tested, which results in optics that have no FODO structure. The best result was obtained for a polarity inversion of all quadrupoles except the QF3 family, resulting in undesired large vertical beta functions of 30 m and and a suboptimal emittance of 168 nm rad.

It should be noted that the emittance optimisation algorithm in fig. 4.1 was developed empirically and might not reach the absolute emittance minimum for a given set of boundary conditions, despite the above attempts to vary the initial conditions and quadrupole polarities. However, the improvements of charge transfer rates and beam currents that were obtained with one of the newly developed optics ("bo-101c") and with the respective transfer line optics were already sufficient to achieve the desired reduction of storage ring filling times. Therefore, more advanced algorithms (e.g. combinations of gradient descent and Monte-Carlo methods or genetic algorithms) or analytic methods were not investigated.

4.2.2.1   Sensitivity to Magnetic Field and Alignment Errors

When the Bodo optics was to be changed from the "bo-006b" high emittance optics one of the low emittance optics ("bo-101c" or "bo-101e"), only some of the horizontal steerer magnets and no vertical steerer magnets were equipped with power supplies for cost reasons. Therefore, the sensitivity of the orbit to misalignments and magnetic field errors was simulated the for different optics, since such imperfections might lead to large orbit perturbations that cannot be corrected sufficiently with the few available steerer magnets. Large orbit perturbations reduce the transverse phase space that is available for the injected Linac beam, which may result in low or zero beam currents in Bodo.

The simulations were performed with the PC program "Optics" [27]. Technical details and histograms of the simulations are included in appendix A. The results are summarised in table 4.2. Row 1 and 2 of the table show the average horizontal and vertical peak orbit perturbations for transverse magnet misalignments, with random alignment errors in the range of ± 0.3 mm in

Plane, Error / Optics	bo-006b	bo-101c	bo-101e	bo-103b
x, random position offset ±0.3 mm max.	2.5 ± 0.2	2.2 ± 0.2	3.5 ± 0.2	8.5 ± 0.5
z, random position offset ±0.3 mm max.	1.7 ± 0.2	1.7 ± 0.2	1.8 ± 0.2	3.5 ± 0.2
x, random dipole field error ±0.2 % max.	6.7 ± 0.3	4.8 ± 0.2	6.8 ± 0.3	12.5 ± 0.5
z, random dipole tilt ±0.5 mrad max.	1.7 ± 0.2	1.5 ± 0.2	1.4 ± 0.2	2.7 ± 0.3

Table 4.2: *Simulated average peak orbit perturbations for random alignment errors and magnetic field errors of all Bodo magnets.*

both planes for all Bodo magnets. Row 3 shows the average horizontal peak orbit perturbations for random magnetic field errors in the range of ±0.2 % for all Bodo dipoles. Finally, row 4 shows the average vertical peak orbit perturbations for random dipole tilt misalignments of all dipoles, with random rotations of ±0.5 mrad around the beam axis.

The magnetic field of a misaligned quadrupole magnet is a superposition of a quadrupole and a dipole field. Therefore, the magnet deflects ("kicks") the orbit with a kick angle that is proportional to the quadrupole focussing strength. Dipole field errors cause undesired horizontal kicks, while dipole tilt errors result in vertical kicks. The sensitivity of an optics to alignment and magnetic field errors does not only depend on the kick angles, but also on the betatron phase between the magnets that cause the undesired orbit kicks, on the betatron tune and on the size of the beta functions at the locations of the kick [7]. If the beta functions at the locations of orbit kicks are large or the betatron tune is close to an integer resonance or the phase between different orbit kicks leads to coherent superposition and amplification of the orbit perturbations, this will result in a very high sensitivity to misalignments (see eq. 2.52), whereas optics without these properties are less sensitive.

Because of its very low emittance, the "bo-103b" optics has the highest quadrupole strengths, and its tunes have a distance of 0.14 (Q_x) and 0.28 (Q_z) to integer resonances. Therefore, the optics is most sensitive to misalignments and field errors. Although the "bo-101c" optics has larger quadrupole strengths than the "bo-006b" optics, it has the smallest sensitivity to alignment and field errors, because its tunes have a larger distance to integers. The larger perturbations of the "bo-101e" optics in the horizontal plane compared to the "bo-101c" optics are also caused by the horizontal tune, while the focussing strengths of both optics are similar.

Due to the limited number of available steerer magnet power supplies and the suboptimal placement of BPMs (see section 7.3) in Bodo, the "bo-101c" optics was considered to be the best compromise between low emittance and low sensitivity to magnetic field and alignment errors.

4.2.3 T1 Transfer Line Optics

In order to maximise the injected beam current for the "bo-101c" optics, a suitable optics for the T1 transfer line (named "t1-101c-5") was developed during the work on this thesis. Ideally, a transfer line for on-axis injection should allow maximum charge transfer from Linac to booster while preserving the beam emittance of the Linac. Emittance preservation is achieved by adjusting the focussing strengths of the six quadrupole magnets in the T1 transfer line so
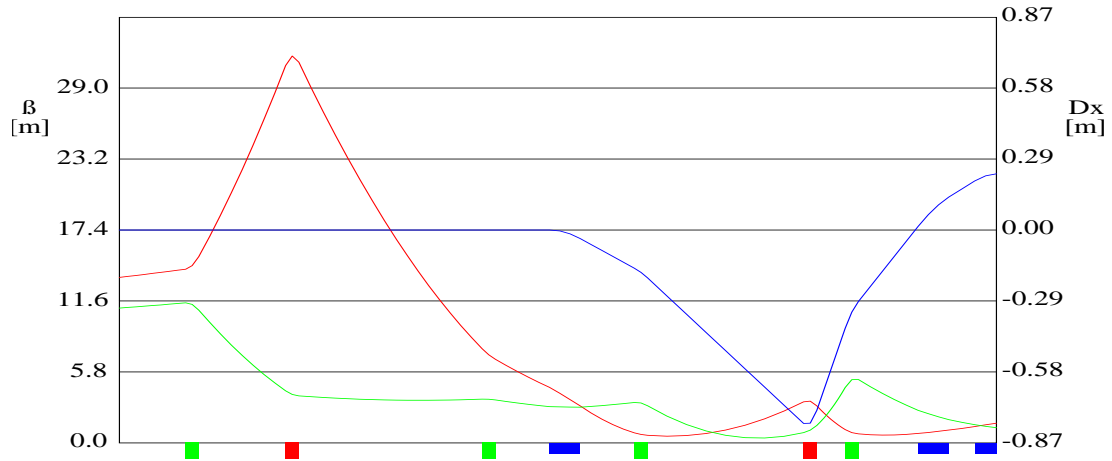
Figure 4.6: *Calculated optical functions of the* **T1 transfer line optics "t1-101c-5"** *(red: β_x,* **green:** β_z**, blue: dispersion** D_x*). The electrons move from left to right through horizontally defocussing (green) and focussing (red) quadrupole magnets, two bending magnets (blue) and the Bodo injection septum (blue, outer right side).*

that the optical functions β_x, α_x, β_z, α_z, D_x and D'_x of the T1 optics match the values of the periodic solution of the Bodo optics at the injection septum blade ("transfer line matching" [42]). A mismatch of β or α leads to an effective increase of the emittance in the respective plane ("emittance blow-up") for all injected electrons. A mismatch of D_x or D'_x results in an additional horizontal emittance blow-up for injected off-energy electrons. As shown in ref. [52], a matched T1 optics will cause beam loss of more than 50 % of the injected beam when it enters Bodo through a slit in the injection septum, and the injected beam will only use a small fraction of the available transverse phase space in Bodo. The septum slit is a rectangular hole of 10 mm height and 7 mm width, which is very small compared to the beam cross section of a matched optics at this point. On the other hand, T1 optics with very small values of beta functions and dispersion at the septum slit will avoid beam loss there, at the expense of emittance blow-up and more beam loss in Bodo.

In order to optimise the charge transfer from Linac to Bodo for a given Bodo optics, the T1 optics should have an "optimal mismatch" so that the number of beam cross sections that fit through the septum slit is both maximal and identical to the number of beam cross sections that fit into the aperture of Bodo. Figure 4.6 shows the T1 optics "t1-101c-5" that was optimised according to these criteria via numerical simulation of the injected beam with the PC program Optics. Technical details of the simulation and an overview of aperture limitations and beam envelopes for T1 and Bodo optics at 60 MeV are contained in ref. [52]. The "t1-101c-5" optics provides a relative aperture of more than two horizontal standard deviations for the injected beam both at the septum slit and during subsequent turns in Bodo, which corresponds to a maximum charge transfer of more than 90 % for a Linac beam with $\Delta p/p = 0.005$ and $\varepsilon_x = \varepsilon_z = 1000$ nm rad at 60 MeV [21].

It should be noted that the above charge transfer rates are theoretical maximum values for ideal conditions, e.g. for a sufficiently large dynamic aperture and a negligible mismatch of the

injected beam in the longitudinal phase space. As shown in ref. [52], the chromaticity of both old ("bo-006b") and new ("bo-101c") Bodo optics can be compensated without reducing the dynamic aperture significantly below the mechanical one. However, the injected beam is not matched in the longitudinal phase space because the Linac has no 500 MHz RF system in order to bunch the electrons synchronously to the Bodo RF phase before injection. Nevertheless, experimental tests at Bodo showed that T1 optics with optimised mismatch still achieve an improvement of the Bodo beam current by typically a factor of two ("bo-101c" optics) or three ("bo-006b" optics) with respect to a matched T1 optics [52], with maximum average Bodo beam currents of more than 10 mA at 1.5 GeV after DSP-based optimisation of betatron tunes and orbit (see following chapters). A 500 MHz bunching system for the Linac may lead to further improvements of the beam current, but this may require the installation of a system that compensates the beam loading in the Bodo cavity.

4.2.4 T2 Transfer Line Optics

In order to achieve a high charge transfer from booster to storage ring, new T2 transfer line optics were developed in this thesis for the "bo-006b", "bo-101c" and "bo-103b" Bodo optics using the PC program "Optics". The properties of these T2 optics are summarised in table 4.3. The table also contains the optics "t2-006b-old" that was used for the T2 transfer line until December 2001. The T2 optics are named similar to the respective Bodo optics, i.e. "t2-101c" belongs to "bo-101c" etc.. The optical functions of one of the T2 optics ("t2-101c") are

Parameter / Optics	t2-006b-old	t2-006b	t2-101c	t2-103b	del-008a
k(Quad1) $[m^{-2}]$	0.45	3.128	4.245	2.283	-
k(Quad2) $[m^{-2}]$	-2.577	-3.761	-4.133	-3.611	-
k(Quad3) $[m^{-2}]$	2.417	2.622	2.790	2.634	-
k(Quad4) $[m^{-2}]$	-1.790	-3.170	-3.290	-3.203	-
$\sigma_{x,max}[mm]$	3.02	4.44	3.54	2.836	-
$\sigma_{z,max}[mm]$	1.19	1.04	0.841	0.781	-
$\beta_{x,DSept}[m]$	8.97	1	1	1	1.96
$\alpha_{x,DSept}$	0.457	0.2	0.2	0.2	1.295
$\beta_{z,DSept}[m]$	5.673	7	7	7	4.969
$\alpha_{z,DSept}$	-2.177	-1.963	-2.29	-2.609	-2.652
$D_{x,DSept}[m]$	0.835	0.497	0.497	0.497	0.497
$D'_{x,DSept}$	0.436	0.312	0.186	0.209	-0.433

Table 4.3: *Calculated parameters of previous ("t2-006b-old") and newly developed ("t2-006b", "t2-101c", "t2-103b") T2 transfer line beam optics at 1.5 GeV. The quadrupole "Quad1" is the first one which is passed by the beam, "Quad4" the last one. $\sigma_{x,max}$ ($\sigma_{z,max}$) is the maximum horizontal (vertical) beam standard deviation of the beam in the transfer line, assuming 10 % emittance coupling and the energy spread and emittance of the respective Bodo optics at 1.5 GeV. The aperture of the T2 beam pipe is not constant, but always larger than $\Delta x = \pm 15mm$ horizontally and $\Delta z = \pm 8$ mm vertically (assuming that alignment errors of the pipe are smaller than 1 mm). The right column contains the optical functions of the "del-008a" Delta optics at the Delta injection septum slit.*

shown in fig. 4.7. An overview of aperture limitations of Bodo and T2 that are relevant for beam extraction are included in ref. [52]. The following discussion assumes a beam energy of 1.5 GeV.

Phase Space Matching

Altogether, there are six optical functions of the T2 optics that must have certain (optimal) values at the Delta injection septum. However, the T2 contains only four quadrupole magnets that can be varied to obtain these optimal values, therefore the maximum number of optical functions that can have arbitrary desired values is four. Furthermore, not all combinations for a given set of four optical functions can be achieved [7].

Therefore, only β_x, α_x and D_x in table 4.3 have nearly optimal values for the newly developed optics. The values for β_x and α_x at the end of the transfer line are optimal when the horizontal phase space ellipse of the T2 beam fits best into the acceptance ellipse of the Delta beam (see blue ellipses in fig. 4.8, page 44). This leads to optimal values for β_x and α_x that are different from the values of the periodic solution of the "del-008" optics. Since matching of D_x' results in an excessive mismatch in the vertical plane, the values of D_x', β_z and α_z for the newly developed optics in table 4.3 are a compromise between horizontal emittance blow-up due to mismatch of D_x' and vertical emittance blow-up due to mismatch of β_z and α_z.

As will be shown below, the resulting mismatch of D_x' for the newly developed optics in table 4.3 is tolerable. The mismatch of β_z and α_z is also not critical because it allows to store six vertical standard deviations of the injected beam for all newly developed optics. Since the emittance of the injected beam is damped down to the emittance of the stored beam within a few milliseconds (at 1.5 GeV) by synchrotron radiation, the particle loss due to this mismatch should be negligible.

Aperture Limitations in Bodo and T2

According to table 4.3, more than 3 horizontal and 5 vertical standard deviations of the beam cross section fit into the T2 beam pipe for each of the optics except for the septum slits. The main aperture limitation for the extracted beam is the extraction septum slit of 6 mm height and 8 mm width. Only 1.76 horizontal standard deviations or 92 % of the beam charge fit through the slit in case of the "bo-006b" optics, in contrast to at least 2.73 standard deviations or 99.4 % for the three newly developed optics (see table B.1, page 160).

Phase Space at the Delta Injection Septum

Figure 4.8 shows the calculated horizontal phase space ellipses of stored and injected beam at the Delta injection septum for an injected beam with nominal energy and zero energy spread. The Delta beam is moved close to the septum blade by a local DC orbit bump. An additional local kicker orbit bump shifts the acceptance ellipse of the stored beam across the septum blade for a single turn when the Bodo beam is injected.

The Delta acceptance ellipse in the figure corresponds to a horizontal mechanical aperture of 34 mm instead of the ideal value of 37 mm. The safety margin of 3 mm is necessary to
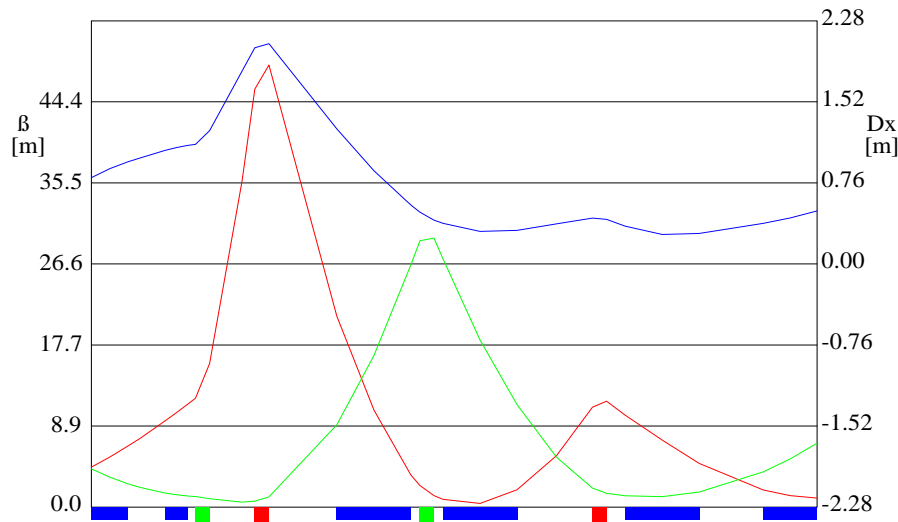
Figure 4.7: *Calculated beta functions (**red:** β_x**, green:** β_z, left axis) and dispersion (**blue:** D_x, right axis) of the **T2 transfer line optics "t2-101c"**. The beam moves from left to right along the horizontal s-axis, through the Bodo extraction septum bending magnet and first ("short") dipole magnet (blue), horizontally defocussing (green) and focussing (red) quadrupole, through the second ("long") dipole magnet (actually consisting of three dipole magnets connected in series, with two quadrupoles in between) and the Delta injection septum (blue, right side). The length of the transfer line is 11.327 m. All dipole magnets currents in the transfer line are pulsed. The septa have pulse lengths of $88\mu s$ and $83\mu s$ (half sine wave), the two dipoles have pulse lengths of 100 ms (one full sine wave). However, the beam passes the transfer line within 38 ns.*

account for local orbit bumps of this magnitude that are usually required for the "U250" undulator beamline (see fig. 1.1) in the "upper" (north) straight section of Delta, where the absolute maximum of the horizontal beta function limits the acceptance for the injected beam. Furthermore, magnetic stray fields of T2 transfer line magnets generate orbit perturbations of more than 1 mm during injection that also reduce the available aperture.

The amplitudes of the local orbit bumps in fig. 4.8 were chosen in such a way that the DC bump does not limit the acceptance for electrons with $\Delta p/p = 0$, and that the kicker bump leaves a safety margin of 1 mm to the septum blade for a kicked stored beam of 5 standard deviations and $\pm 0.175\%$ energy offset with respect to the nominal energy of Delta. The safety margin allows smaller orbit perturbations by suboptimal kicker timing and stray fields of pulsed T2 magnets without risk of beam loss for the stored beam during injection. Furthermore, beam movements in the order of 1 mm can be caused by the Delta orbit feedback system that may change the Delta RF frequency by several kHz within minutes to hours in order to correct the orbit. The safety margin guarantees that such frequency changes do not move the stored Delta beam too close to the septum (that has no BPM to measure the position of the stored beam at its blade), which would result in partial beam loss.

According to fig. 4.8, the phase space ellipses of the "bo-101c" and "bo-103b" optics fit into the Delta acceptance, while a larger part of the beam of the "bo-006b" optics gets lost, espe-
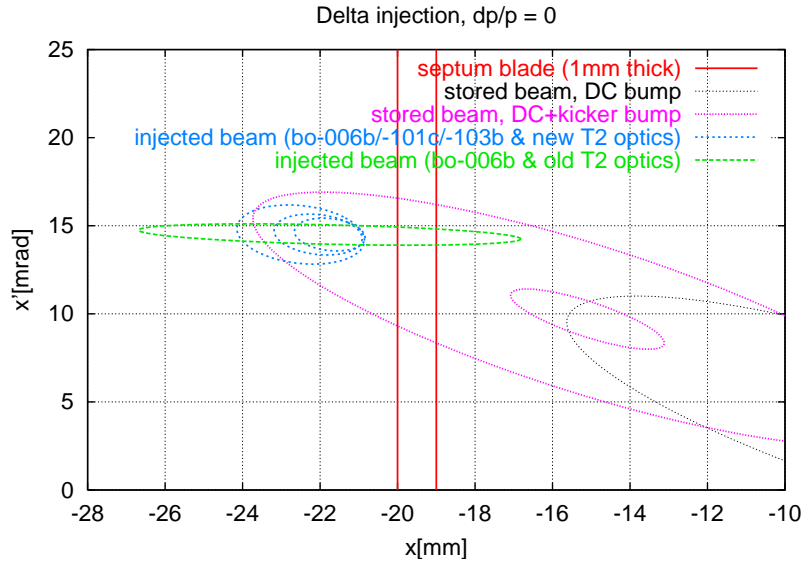
Figure 4.8: *Calculated horizontal phase space of the injected 1.5 GeV Bodo beam and stored Delta beam at the Delta injection septum slit (for the del-008a Delta optics, emittance* 21.78 *nm rad ,* 1% *coupling). Stored beam (right side) and injected beam (left side) are separated by the copper septum blade of 1 mm thickness (indicated by red lines). The black ellipse section on the outer right side encloses the stored Delta beam which was moved towards the septum by a local static orbit bump. The two pink ellipses on the right enclose the Delta beam, kicked towards the septum by an additional injection kicker orbit bump for 1 turn. The black ellipse and the outer pink ellipse ("acceptance ellipse") have a size of 42 standard deviations (= mechanical aperture limit for* $\Delta p/p = 0$*), the inner pink ellipse has a size of 5 standard deviations (for* $\Delta p/p = 6.82 \cdot 10^{-4}$*, i.e. its size includes the beam enlargement due to dispersion and natural energy spread). The three blue ellipses and the green ellipse on the left show 2.5 standard deviations of the injected 1.5 GeV Bodo beam for 10 % coupling and* $\Delta p/p = 0$ *(outer blue ellipse: "bo-006b"/"t2-006b" optics (* $\varepsilon = 477$ *nm rad), middle blue ellipse: "bo-101c"/"t2-101c" optics (* $\varepsilon = 277.5$ *nm rad), smallest blue ellipse: "bo-103b"/"t2-103b" optics (* $\varepsilon = 146$ *nm rad), flat green ellipse: "bo-006b"/"t2-006b-old" optics).*

cially when using the "t2-006b-old" T2 optics instead of the newly developed optics. However, the figure is only valid for an injected beam with zero energy spread, while the relative energy spread of the Bodo beam at 1.5 GeV is in the order of $\Delta p/p = 7 \cdot 10^{-4}$ (see table B.1, page 160).

Therefore, figure 4.9 shows the phase space ellipses for an injected beam with zero energy spread and energy offsets of $\Delta p/p = \pm 0.175$ %. This corresponds to an energy mismatch of about ± 2.5 standard deviations with respect to the Bodo energy spread at 1.5 GeV for any optics, since the energy spread of all optics is nearly identical. The ellipses of injected beam and acceptance are shifted with respect to fig. 4.8 due to non-zero dispersion D_x (and its derivative D'_x) both of T2 and Delta optics at the septum. The matching of D_x leads to identical "horizontal" (i.e. x) shifts of injected and stored dispersive beam ellipses in the figure, while the mismatch of D'_x results in different "vertical" (i.e. x') shifts of the ellipses. Despite the D'_x mismatch, the energy-mismatched phase space ellipses of the "bo-101c" and "bo-103b" optics still fit into the Delta acceptance ellipse. This means that the beam loss of an injected
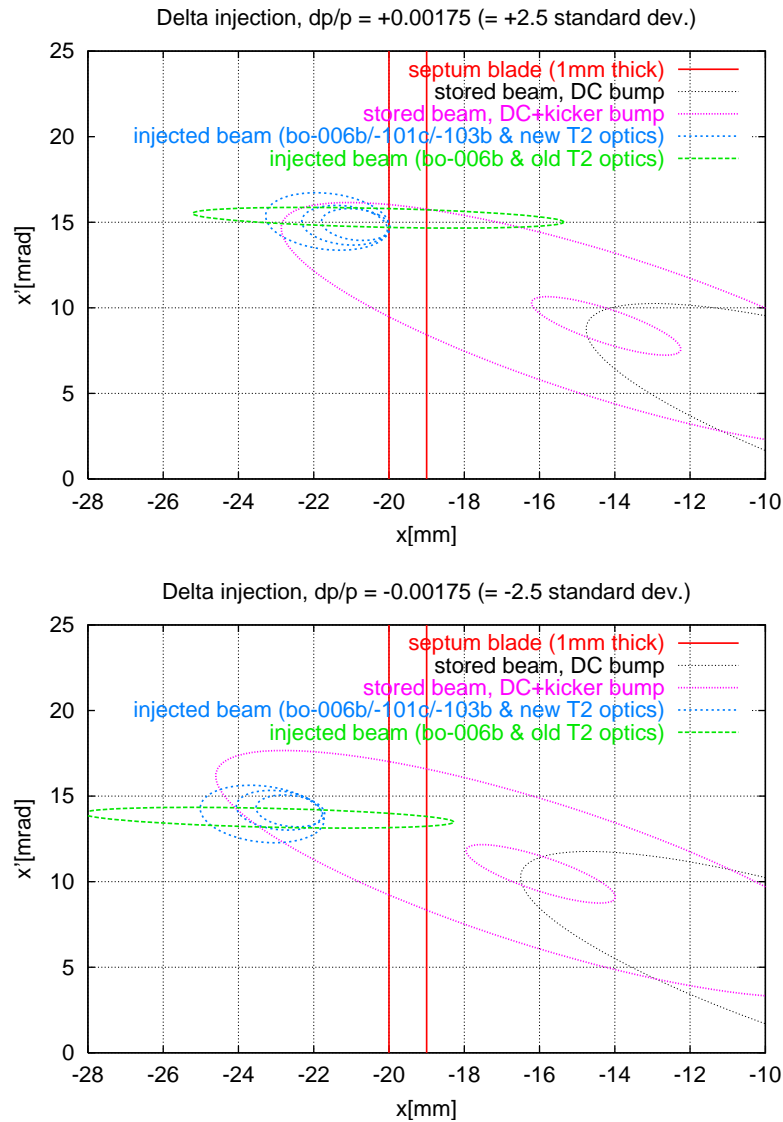
Figure 4.9: *Calculated horizontal phase space of the injected 1.5 GeV Bodo beam at the Delta injection septum slit. Left side (blue ellipses and green ellipse): 2.5 horizontal standard deviations for Bodo/T2 optics "bo-006b"/"t2-006b" (largest blue ellipse), "bo-101c"/"t2-101c" (medium size blue ellipse), "bo-103b"/"t2-103b" (smallest blue ellipse) and "bo-006b"/"t2-006b-old" (flat green ellipse). All parameters are identical to fig. 4.8 except that the injected beam has an energy that is +0.175 % = +2.5 Bodo standard deviations larger (upper figure) or −0.175 % = -2.5 Bodo standard deviations smaller (lower figure) than the nominal energy of the Delta beam. Therefore, the injected beam will oscillate around a dispersion orbit in Delta with a betatron frequency of typically 23.8 MHz until the synchrotron oscillation (with a frequency of typically 15 kHz) changes the energy of the injected electrons. The outer pink ellipse on the right side is the acceptance ellipse of this dispersion orbit.*

beam *without* energy mismatch but with natural energy spread is negligible for these optics. In contrast, part of the injected beam for the "bo-006b" optics will be lost, and the loss increases with growing energy spread.

4.3   Summary and Experimental Results

Section 4.2.1 showed that the Bodo quadrupole power supplies allow the reduction of the emittance by a factor of 2.4 with respect to the "bo-006b" optics that was used so far. The installation of two power supplies with higher current limits would allow a larger reduction by a factor of 3.3. Because of the limited number of steerer magnet power supplies and the suboptimal positions of the Bodo BPMs (see section 7.3), the "bo-101c" optics was considered as most suitable optics for Delta injection at 1.5 GeV, despite its slightly smaller emittance reduction by a factor of 2.1 instead of 2.4. Although the "bo-101c" optics requires stronger sextupole magnet fields than the "bo-006b" optics because of higher absolute values of the natural chromaticities, the chromaticity can be compensated without reducing the dynamic aperture significantly below the mechanical one [52]. This should allow to increase the charge transfer rate from Linac to Bodo by overcoming the aperture limitation of the injection septum slit with an "optimised mismatch" of the T1 optics, which causes a horizontal emittance blow-up in Bodo and therefore requires a large Bodo aperture. Despite the insufficient number of quadrupoles in the T2 transfer line, section 4.2.4 showed that a compromise between dispersive and vertical mismatch should result in a substantial increase of the charge transfer rate from Bodo to Delta when using the "bo-101c" instead of the previous Bodo and T2 optics.

As mentioned in the introduction, the goal of this chapter was the development of a Bodo beam optics that will improve the performance of Bodo as 1.5 GeV injector when being combined with the DSP-based beam diagnostics and feedback systems that were developed in this thesis. Due to the reasons given above, the "bo-101c" optics was chosen for all tests of the diagnostics and feedback systems in the following chapters of this thesis. Because of the lack of suitable beam diagnostics systems for Linac and both transfer lines (e.g. BPMs, energy spectrometer, longitudinal and transverse phase space diagnostics etc.) and for Bodo and Delta (e.g. BPMs with turn-by-turn time resolution, beam cross section measurement after injection etc.), the new optics could not be commissioned and optimised by a detailed experimental investigation of the charge transfer between Linac, Bodo and Delta. Nevertheless, the optics were commissioned successfully with the few available booster and storage ring monitors for beam current, average beam position and betatron tunes, by iterative variation of unknown parameters (e.g. transfer line orbit corrector currents) and optimisation of the resulting Bodo and Delta beam currents.

Figure 4.10 shows the measured charge transfer rates for the Delta injection during first tests of the "bo-101c" optics. The rates were calculated from the average Bodo beam current before extraction and from the average Delta beam current before and some 100 ms after injection. Typically, only 10-20% of the Bodo beam charge could be accumulated in Delta with the "bo-006b" Bodo optics during 18 hours of machine operation by different operators. In contrast, the "bo-101c" optics achieved rates up to 65% after optimisation of T2 dipoles, septa and orbit correctors. The drifts and variations of the transfer rates result from trigger jitter and temperature drifts of pulsed voltage-controlled transfer line magnets and of drifts of the Delta orbit. Due to the lack of suitable monitors, the machine operators have to optimise the transfer
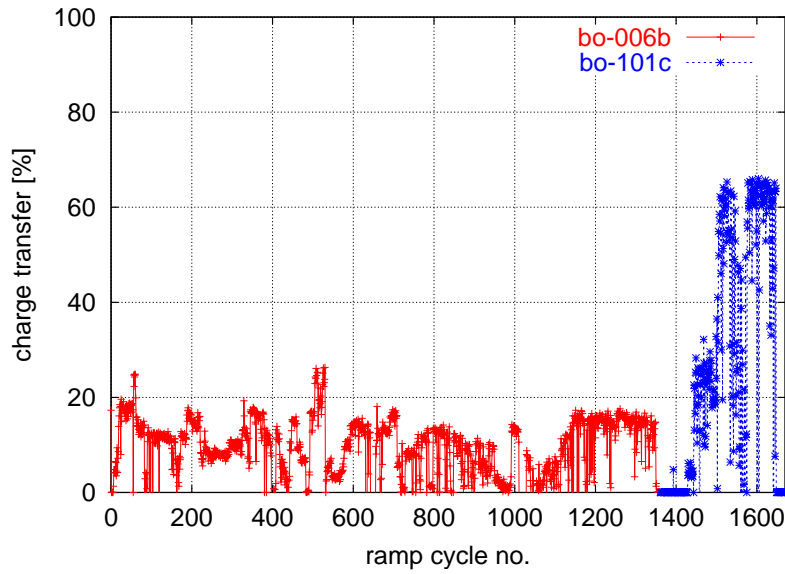
Figure 4.10: *Measured charge transfer rate from Bodo to Delta for the bo-006b optics (left side, ramp cycles 0-1360) and for the bo-101c optics (right side, ramp cycles 1361-1650). The measurements were taken over an 21 hour period at 1.5 GeV Delta energy, with a switch to the "bo-101c" optics after 18 hours.*
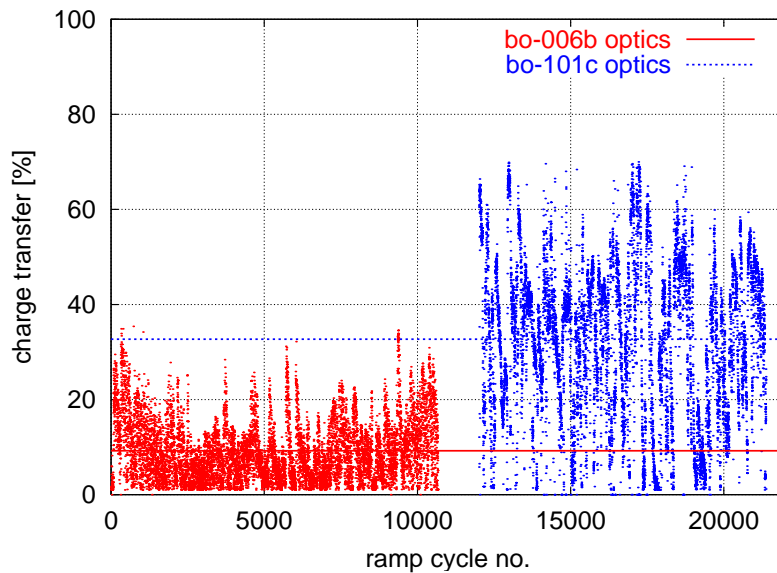
Figure 4.11: *Measured charge transfer rates from Bodo to Delta. The measurements were taken over a period of one month of regular machine operation at 1.5 GeV with the "bo-006b" optics (ramp cycles 0 to 10660, red dots, one dot per ramp cycle) and another month with the "bo-101c" optics (ramp cycles 12000 to 21367, blue dots). The horizontal lines indicate the average transfer rates of 9.3% ("bo-006b optics", red line) and 32.7% ("bo-101c" optics, blue line).*

rate by iterative variation of magnet currents and voltages in T2 and Delta, with the beam current as the figure of merit. Depending on the skills and experience of the operator, this results in varying transfer rates. In order to estimate the typical improvement of the transfer rates for operation with different operators, the rates were measured for both optics over a period of one month each. Figure 4.11 shows that the transfer rates for the month with the "bo-101c" optics are about 3 times larger than for the month with the "bo-006b" optics.

Although the measurements were taken for the same Delta optics, the improvement may not only result from the different Bodo and transfer line Optics, but also from other factors that could not be measured due to lack of monitors, e.g. the beam position in the transfer line or the position of stored and injected beam at the Delta septum slit. Furthermore, the question remains why the maximum transfer rates are still much lower than 100 %. While the achieved rates are sufficient for normal 1.5 GeV operation of Delta, the envisaged top-up injection mode might require even higher transfer rates in order to decrease beam loss and radiation levels. The envisaged installation of BPMs and beam cross section monitors in the T2 transfer line in the context of another Ph.D. thesis should allow to analyse these problems by comparison of theoretical and measured T2 optics and orbit. Moreover, it is recommended to install BPM systems with turn-by-turn time resolution at critical locations in Delta, in addition to an envisaged additional BPM at the Delta septum blade. The improved beam diagnostics would allow to distinguish between the different factors that contributed to the overall improvement of the change transfer rate by a comparison of the theoretical beam positions in fig. 4.8 with the real beam positions. These might be different from the theoretical ones, e.g. due to magnetic stray fields of injection septum, pulsed T2 dipole magnets and imperfections of the kicker orbit bump. In addition to an improvement of beam diagnostics systems, a future implementation of the top-up injection mode should include the analysis of the available mechanical and dynamic aperture in Delta, since the improvement of suboptimal apertures might lead to a further increase of the charge transfer rates.

During first tests, the Bodo beam current for the "bo-101c" optics was typically 1-2 mA,
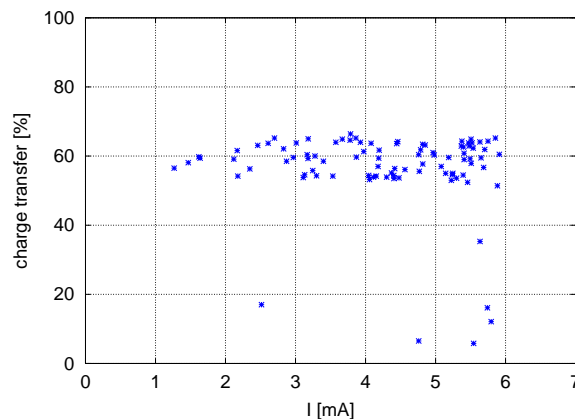
Figure 4.12: *Plot of the measured charge transfer rate from Bodo to Delta as a function of the Bodo beam current before extraction at 1.5 GeV. The measurements were taken during 100 successive ramp cycles with the "bo-101c" optics. There is no significant correlation between transfer rate and beam current. The occasional lower transfer rates are caused by timing jitter of the Bodo extraction kicker.*

compared to 3-6 mA for the "bo-006b" optics. Figure 4.12 shows the transfer rates and currents for the "bo-101c" optics after first DSP-based improvements of the beam parameters. The figure proves that the increased charge transfer rates were not caused by the smaller initial beam current, because the transfer rates are still in the order of 60 % and not correlated with the improved beam current of 3-6 mA. After further DSP-based optimisation and stabilisation of optics and orbit and after stability improvements of Linac and T1 components, Bodo beam currents of typically 8 mA were achieved with the "bo-101c" optics (see fig. 4.13).
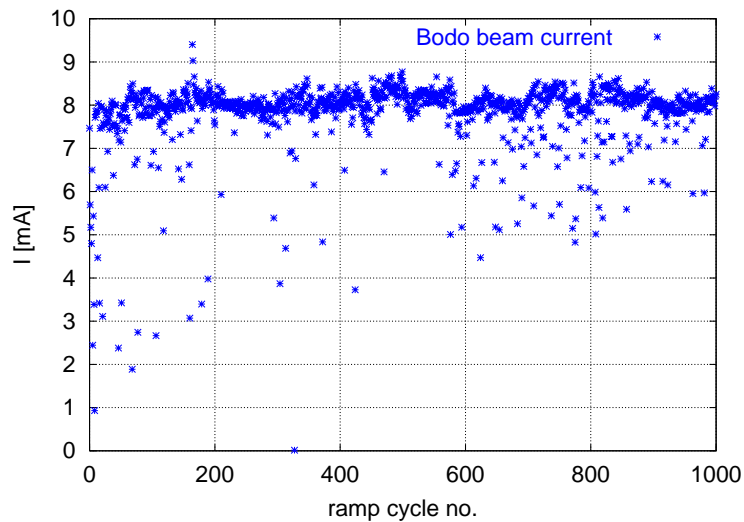
Figure 4.13: *Plot of the measured Bodo beam current before extraction at 1.5 GeV during 1000 successive ramp cycles with the "bo-101c" optics.*

The remainder of this thesis will describe the development, test and operation of a DSP-based beam diagnostics and feedback system that was used to measure and optimise the beam parameters of Bodo. The system is not only designed to improve the performance of Bodo as injector for 1.5 GeV synchrotron light source operation of Delta, but it was also designed for a future use of Bodo as a test machine for newly developed accelerator components, hardware and software. As will be explained in the next chapter, the use of the DSP system is not limited to Bodo, but it may also be used for a variety of other applications, e.g. for a fast orbit feedback in Delta.

Chapter 5

Development of a Unified Beam Diagnostics and Global Feedback System

5.1 Motivation

Despite the effort to eliminate sources of orbit perturbations in the construction phase of boosters and storage rings, various imperfections of the machine usually still lead to constant and time-dependent deviations of the closed orbit from the ideal orbit that can range from micrometers to tens of millimeters.

Common sources of constant orbit perturbations are field errors of dipole magnets and position adjustment errors of quadrupole magnets. Time-dependent orbit perturbations can be caused e.g. by magnet stray fields from external sources, by magnets of adjacent accelerators, or by magnet movements due to temperature changes or ground vibrations. Since ground vibrations can be resonantly amplified by magnet girders and since movements of a quadrupole magnet cause beam movements that are typically 5-20 times larger, even very small vibrations of a large number of magnets can add up to significant orbit oscillations, typically at frequencies of 5 to 60 Hz [46] . In general, the time scale of common time-dependent perturbations ranges from milliseconds or seconds (for stray fields or mechanical vibrations) to hours or days (for magnet movements due to temperature changes of the environment).

Orbit perturbations have a variety of undesired effects. Large static orbit deviations can increase the vertical beam emittance [24] due to increased coupling of transverse betatron oscillations, decrease the beam lifetime, or it may be impossible to store the beam at all. Time-dependent orbit perturbations can lead to time-dependent variations of the number or the energy of the photons that hit the sample of a synchrotron light experiment, leading to suboptimal results of the experiment. Therefore, the orbit at synchrotron light source insertion devices is usually required to have a stability better than 5-10 percent of the cross section and angular spread of the electron beam at the source point. Typical beam stability requirements range from some 10 micrometers for the horizontal plane to a few micrometers for the vertical plane, depending on the emittance coupling. Maximum orbit drifts that can be tolerated by certain experiments at synchrotron radiation beamlines may be less than one micrometer [23].

In booster synchrotrons, orbit perturbations may be one reason for low beam currents, beam loss at injection or during beam acceleration, and suboptimal booster to storage ring charge transfer rates. Furthermore, orbit perturbations in the booster may cause partial loss of the injected beam in the storage ring, which is undesired especially for machines that use the so-called top-up injection mode (see section 4.1.1) which requires very low beam loss rates because injection takes place with opened beam shutters of the synchrotron radiation beamlines.

In order to eliminate orbit perturbations, present operational state-of-the-art orbit feedback systems in synchrotron light source storage rings apply orbit corrections at a rate up to 1.5 kHz

for 160 BPMs and 38 steerer magnets per plane [49] and 4.4 kHz for 16 correctors and 16 BPMs only in the vertical plane [46]. These feedback systems consist of analogue-to-digital converters (ADCs) that sample analogue beam position signals, one or more DSPs that calculate the required steerer magnet corrections, and digital-to-analogue converters (DACs) that control a number of sufficiently fast steerer magnet power supplies.

If the feedback algorithm is global and therefore needs all BPM data to calculate the correction current for each steerer magnet, the feedback system has to be distributed all over the accelerator, with a real-time information flow from each BPM to each steerer magnet. In order to minimise noise, state-of-the-art light sources minimise the length of analogue cables in this information flow chain and therefore use digital communication and synchronisation links for real-time data transfer of typically several Mbyte/s over distances up to some 100 m [45, 46]. The mainly digital electronic systems for these feedbacks are often partial or complete in-house developments (sometimes in cooperation with industrial suppliers) that are tailored to the respective I/O interfaces, BPM and orbit corrector systems of the machines [45, 46, 49]. Orbit feedbacks that consist mainly of commercially available components usually require a very large number of electronics systems (e.g. more than 100 VMEbus boards for a 1.5 kHz system with 160 BPMs and 38 correctors per plane [49]), because the individual components are not optimised for this application. This results in several times higher hardware costs compared to an optimised technical solution.

In this thesis, a digital global feedback system has been developed that can be used to calculate and apply corrections for 64 BPMs and 32 steerer magnets in each plane of a storage ring (e.g. Delta) simultaneously at a rate of more than 4 kHz, with only eight of the newly developed VMEbus ("DeltaDSP") boards. The number of DeltaDSP boards in the system is variable (between 1 and 255) in order to adapt the system performance and computing power to the number of BPMs and steerers and to the desired bandwidth of the feedback. The novel design of the boards allows higher performance at a fraction of the price of comparable commercially available equipment.

The DeltaDSP boards are interconnected by a novel fibre optics real-time network ("DeltaNet") that was also developed in this thesis. DeltaNet has a ring topology and distributes real-time measurement data, trigger and synchronisation signals between the DeltaDSP boards.

Since the present steerer magnet power supplies of Delta allow only orbit corrections at a rate of about 2 Hz [32], the newly developed DSP-based orbit feedback system could not yet be tested at Delta, but at the booster Bodo that has steerer magnet power supplies with a bandwidth of more than 100 Hz. Furthermore, the BPM readout system of Delta is based on CANbus fieldbus modules that allow BPM readout rates of a few Hz with 12 bit or $4.9\ \mu m$ ADC resolution [32], whereas the ADCs of the Bodo BPM readout system that was developed in this thesis have a maximum readout frequency of 20 kHz and 16 bit or $0.3\ \mu m$ ADC resolution. In order to implement a fast global orbit feedback at Delta, BPM system and steerer magnet power supplies need an upgrade in analogy to the Bodo system.

The use of the newly developed feedback architecture is not limited to orbit feedbacks. The DeltaDSP boards were designed to allow the integration of different kinds of monitors and actuators (i.e. devices that control the electron beam, e.g. power supplies) into the system.

About 85 percent of the circuit design of the DeltaDSP boards is not visible on the printed circuit boards ("PCBs"), since it was implemented in several so-called field programmable

gate arrays ("FPGAs"). FPGAs are ICs that contain a large number of small digital circuit elements (e.g. logical gates, flip-flops, memory). The connections between these elements are (re-)programmable by one of the two DSPs on each DeltaDSP board. This allows an FPGA to perform nearly every desired task, limited only by its size and speed. Therefore, 85 percent of the DeltaDSP boards can be rewired within seconds by remote control, and the boards can be adapted to a variety of different tasks and interfaces in the accelerator without any physical hardware modifications or redesigns of the printed circuit boards.

After several months of successful burn-in tests, the DeltaDSP boards were installed in Bodo in October 2001, together with newly developed software for the various tasks of the DSPs and for the generic control system integration of the boards (see section 5.4). Since then, the DeltaDSP boards control dipole, quadrupole, sextupole and steerer magnet power supplies of Bodo, as well as the RF power of the cavity. Furthermore, DeltaDSP boards handle nearly all diagnostics of Bodo, e.g. beam loss monitors ("BLMs"), beam position monitors ("BPMs"), beam current and lifetime measurement, as well as power supply current and voltage measurement. The system works very reliable, with no hardware failures and only two hours of machine down-time in the first year of operation due to one software bug.

The respective data acquisition system that was developed in this thesis consists of external standalone ADC boards that sample the beam diagnostics and power supply signals with 16 bit resolution and sampling rates between 4 and 100 kHz. The ADCs are connected to DeltaDSP boards that synchronise the sampling triggers for all ADCs in the whole accelerator via DeltaNet, with less than 300 ns conversion trigger jitter between any two ADCs in Bodo. This data acquisition system allows the correlated analysis of different beam diagnostics and power supply signals during the booster energy ramp, with sufficient speed and precision to detect reasons for bad performance of the accelerator. This was not possible with the previous data acquisition system that had data sampling frequencies up to 10 Hz at 12 bit resolution and a typical time stamp precision of 100 ms.

Moreover, DeltaDSP boards were used to build a system that measures and corrects the betatron tunes of Bodo during the acceleration cycles in real-time. The newly developed system is described and characterised in chapter 6. It replaces an all-analogue betatron tune measurement system [4, 61] that was not capable of an automatic tune correction and required continuous manual adjustment in order to obtain usable measurements.

Every DeltaDSP board can receive nearly all beam diagnostics data of a machine in real-time via DeltaNet, typically in periods of some 10 microseconds. Moreover, the DSP boards can handle nearly all devices (e.g. magnets power supplies) that control and change the beam parameters. Therefore, the newly developed DSP system can be used to automate the operation of a complete accelerator facility (i.e. of booster, storage ring, transfer lines etc.) by controlling and optimising all kinds of beam parameters automatically in real-time, e.g. orbit, beam loss, lifetime, betatron tunes, beam optics, beam size, emittance coupling, charge transfer rates etc.. This allows the construction of an "intelligent" accelerator ("cybernetic machine") that requires a minimum of man-power or none at all to operate it.

The following sections introduce the general concepts and architecture of the newly developed beam diagnostics and feedback system, as well as a brief introduction to digital signal processors. A more detailed description of applications (e.g. a real-time betatron tune feedback and a global orbit feedback in Bodo) and the analysis of beam parameter measurements and corrections that were made with the system is deferred until the following chapters.

5.2 Digital Feedback Systems

5.2.1 Accelerator Control Systems

From the control system point of view, an electron accelerator consists of sensors (called "monitors") that measure the properties either of the electron beam (e.g. the orbit) or of devices that control the electron beam (e.g. power supply currents), and of actuators that determine and change the properties of the electron beam (e.g. orbit corrector magnet power supplies). In case of the sensors, the task of an accelerator control system is to record the sensor signals fast and precise enough to allow a sufficiently precise comparison of desired and actual sensor value on the required time scale. Since accelerators usually have a large number of sensors (typically several hundred), the sensor data has to be recorded synchronously (e.g with "time stamps" for analogue to digital conversions, using a global clock) in order to allow the analysis of beam perturbations that occurred at a certain time.

Furthermore, the control system has to control the actuators with sufficient speed and precision to allow adjustments of beam parameters on a sufficiently small time scale, e.g. in order to compensate undesired orbit movements using orbit corrector magnets whose magnetic field can be changed sufficiently fast. In many cases, time-dependent changes of actuators also have to be synchronised, e.g. in case of the focussing magnets of the booster synchrotron that must change their magnetic fields proportionally to the beam energy which itself depends on the dipole magnet current.

If beam parameters need too be corrected periodically on a time scale larger than one second, they can be adjusted manually by operators in the accelerator control room by a graphical user interface that visualises the sensor data and that allows manual changes of the actuator parameters. If the corrections must take place on a time scale much smaller than one second, they have to be calculated and applied automatically.

If the perturbations of the beam parameters are periodic and reproducible, the sensor data can be recorded once, then the time-dependent actuator values (correction curve) that are required for the correction can be calculated, and afterwards the actuators just have to be changed periodically ("feed-forward"), using the constant correction curve that must be properly synchronised to the beam parameter perturbation. If the beam parameter perturbations are not periodic and reproducible, the processor must acquire the sensor data and calculate and apply the actuator corrections in real-time, i.e. fast enough to compensate the perturbations ("feedback"). It is also possible to combine feedback and feed-forward if beam parameters have perturbations that are mainly periodic, but drift slightly from period to period.

Analogue and Digital Control Loops

The parameters of the electron beam of a synchrotron light source like DELTA are controlled by hundreds of actuators and monitored by hundreds of sensors. Set value changes of one actuator usually affect a large number of sensor signals in a nonlinear way. The manual optimisation of beam parameters by machine operators in this high-dimensional parameter space is very time-consuming and often hardly possible. Furthermore, the compensation of beam parameter perturbations may require simultaneous set value changes for tens or even hundreds of actuators at rates of many kHz, e.g. in order to correct time-dependent orbit perturbations.

The automated optimisation of beam parameters at the required time scale can be achieved by electronic feedback systems. Control loops for beam parameter feedback systems (e.g. orbit feedback systems) can be implemented either by analogue or digital electronics. Digital feedback systems usually consist of an analogue front-end that measures the beam properties and converts them to a voltage signal, of ADCs (analogue to digital converters) that digitise these signals, of one or several processors that calculate the required actuator settings, and of DACs (digital to analogue converters) that generate output voltages that control the actuators (e.g. power supply currents).

Compared to analogue systems, digital feedback loops have several advantages, especially in case of feedbacks that use a large number of sensors and actuators (e.g. a global beam orbit feedback). While the properties and feedback characteristics (e.g. speed and frequency response) of analogue systems are usually modified either by turning potentiometers or by replacing components on printed circuit boards, digital control loops can be tuned and modified by either changing input parameters of the feedback software or the software itself. Therefore, large digital feedback systems in an accelerator can be tuned, supervised and modified on a timescale of seconds by a graphical user interface during synchrotron light source operation. In case of analogue feedback systems, tuning, maintenance and modification takes orders of magnitude longer, since these tasks require individual changes of a large number of subsystems, e.g. tuning of many potentiometers or modification of a large number of printed circuit boards. Furthermore, digital feedbacks allow fast self-tests and automatic calibration, and the behaviour of a large number of feedback loops (e.g. unstable feedbacks that result in beam loss) can be visualised, analysed and modified on graphical user interfaces in real-time. Therefore, digital feedback systems are the preferred solution at modern synchrotron light sources [44, 45, 46, 47].

5.2.2   Digital Signal Processors

Processors that are optimised for the processing of measurement data in real-time are called Digital Signal Processors ("DSPs"). As opposed to normal computer processors in desktop computers, DSPs usually have a large number of input/output ("I/O") ports, high speed internal memory and a direct memory access ("DMA") unit that can transmit and receive e.g. measurement data or set values of actuators through the I/O ports at a high guaranteed data rate without usage of the processor core, which is the part of the processor that performs the calculations.

Furthermore, DSPs have command sets that are optimised for common signal processing tasks, e.g. matrix-vector multiplications for SVD orbit correction algorithms, fast Fourier transforms ("FFTs") for the frequency analysis of orbit perturbations, and digital filters to reduce the noise of BPM readings. Many DSPs can perform each command in a single processor clock cycle, whereas many desktop processors require twenty or more cycles to fetch the operands of a command, perform the calculation and store the result. Moreover, most DSPs allow very fast and deterministic interrupts of the processor core operation without the overhead of most desktop processors. Therefore, very high interrupt frequencies can be achieved without affecting the processor performance. Furthermore, most DSPs can handle a large number of interrupt sources without the software overhead that is common to other processors. Efficient interrupt handling is essential for real-time environments (e.g. feedback systems) that require fast response to
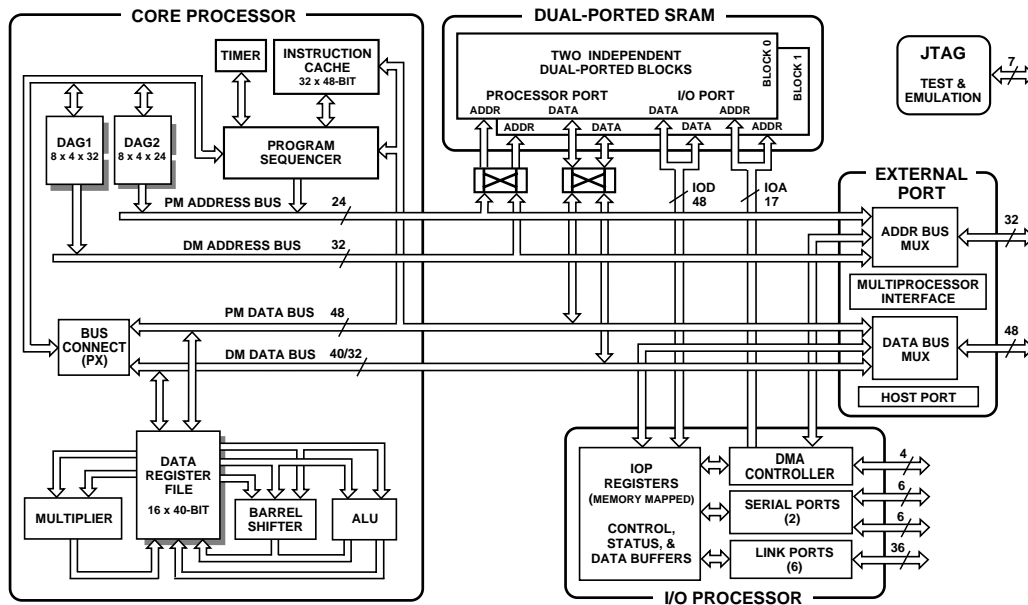
Figure 5.1: *Simplified architecture of the DSP "ADSP-21062" [50] that was used for the DeltaDSP board which was developed in this thesis. The left side ("core") contains the part of the DSP that performs feedback calculations. The lower right part shows the I/O processor ("DMA unit") that can receive and transfer data (e.g. BPM readings) via six so-called link ports (up to 32 Mbyte/s continuous data rate each) and four serial ports (32 Mbit/s each) simultaneously without needing the core. The upper right part shows two internal dual-ported memory banks that can be accessed by the core and the DMA unit simultaneously. The external port on the right side has a 48-bit data bus that is used to access external data and program memory.*

external events or trigger signals while maintaining a guaranteed speed of the processor calculations. Finally, many DSPs can easily be combined to large multiprocessor systems of tens or hundreds of DSPs in order to adapt the available computing power and data transfer rates to the respective real-time application, e.g. to the size of an accelerator and to the number of BPMs and orbit correctors of a global feedback system.

Figure 5.1 shows the architecture of the DSP that was used for the system developed in this thesis. In addition to the features described above, the DSP has two internal dual-ported memory banks (upper right side of fig. 5.1) that can be accessed simultaneously by the processor core (e.g for feedback calculations) and the DMA unit (e.g. for measurement data transfer). The DSP is optimised for floating point calculations up to 40 bit precision, which is necessary for fast scientific calculations. It has six so-called link ports that can be used for data transfer with several other DSPs, with data transmission rates of up to 32 Mbyte/s per port. The link ports allow the design of very large multiprocessor farms that can have a variety of different topologies (e.g. with DSPs and link port connections being arranged in lines, rings, matrices or cubes). The choice of the optimal topology is determined by the data flow of the multiprocessing application.

5.3   The DeltaDSP System

5.3.1   Hardware Architecture

As mentioned above, the DeltaDSP feedback system was developed for different beam diagnostics and feedback applications. Since it should be possible to adapt the system to different requirements concerning computing power and data transfer rates, it was designed as a distributed real-time multiprocessing system that consists of a variable number of VMEbus DSP boards (up to 255 boards with 2 DSPs per board). In case of orbit feedbacks, this concept has the advantage that the number of BPMs and steerer magnets that are used for the global feedback may be increased in the future with only moderate degradation of the feedback bandwidth. Therefore, additional monitors (e.g. photon BPMs of planned synchrotron radiation beamlines) can be easily added to an existing DeltaDSP feedback system, without the rapid decrease of the feedback bandwidth that is inherent to conventional orbit feedback systems that use a centralised architecture with a single DSP board [46]. The bandwidth of such systems is usually dominated by the speed of the feedback algorithm, because the computation time e.g. for the SVD algorithm (see section 2.8) is proportional to the number of BPMs and the number of orbit correctors that are used for the feedback. The algorithm is well suited for distributed multiprocessor systems, since it basically consists of a matrix-vector multiplication, with one matrix row for each corrector magnet. If each processor controls the set currents of a number of correctors, it must only perform the multiplication for the respective rows of the matrix. Therefore the speed of the algorithm can be increased by increasing the number of processors.

DeltaNet has a ring topology where every DeltaDSP board is connected to its next two neighbours in the ring by a fibre optics cable. This allows the boards to be installed close to the BPMs and magnet power supplies that are usually distributed all around the storage ring, while minimising the DeltaNet fibre optics cable length and delays from DSP board to DSP board and maximising the data transmission rates. Data, trigger and synchronisation signals travel unidirectionally around the ring. The overall cable length of such a DeltaNet ring is in the order of the circumference of the storage ring and remains nearly constant when more DeltaDSP boards are added, which may be necessary in order to integrate additional monitors and actuators into the system. In contrast, a star topology with a centralised data concentrator ("hub") would cause the overall cable length to increase proportionally to the number of DSP boards.

Master DSP and Universal Accelerator Device Interface

Fig. 5.2 shows the data flow on a DeltaDSP board. Each board contains two DSPs. One DSP ("master") is designated for local measurement data acquisition, data analysis, real-time feedback algorithms and magnet power supply control. The second DSP ("slave") handles the DeltaNet data transfer, as well as some administration tasks. The master DSP is connected to a multi-purpose I/O piggyback module that can be plugged onto the DSP mainboard. This allows easy adaption to new accelerator components by developing a new piggyback without having to modify the DSP mainboard that has a much higher complexity. An FPGA on the I/O module can control up to eight external DAC boards or similar devices like digital sine generators for the Bodo tune measurement system (see chapter 6). The DAC boards (with one 16-bit DAC each) are connected in parallel to a differential 16-bit ribbon cable bus (up to 10 m
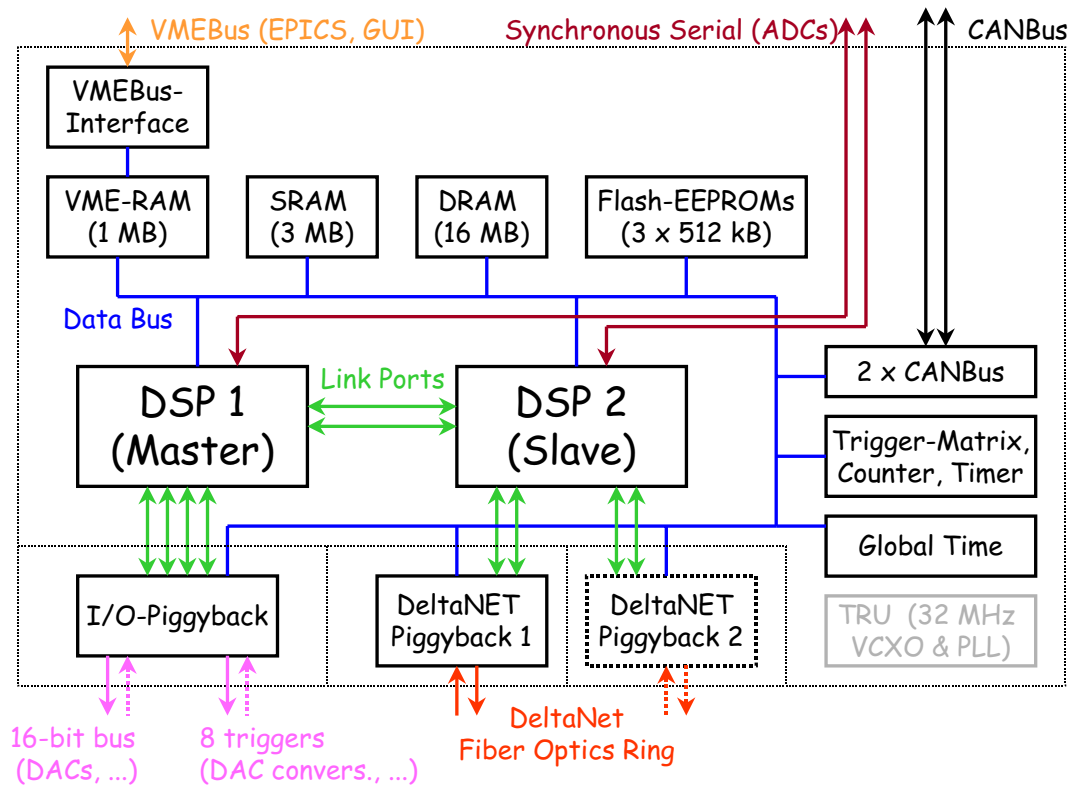
Figure 5.2: *Simplified architecture and data flow of a DeltaDSP VMEbus board. Each DeltaDSP board can be equipped with an optional 2nd DeltaNet interface (indicated by dotted lines) in order to build redundant networks, double the transmission rates or couple feedbacks that use different DeltaNet rings by one DeltaDSP board that is part of both rings. Up to 8 external 16-bit DAC boards can be connected in parallel to the 16-bit parallel data bus of the I/O piggyback module. Since all connected DACs see the same data on the bus, individual conversion triggers determine which of the DACs uses the 16-bit value on the bus for a digital-to-analogue conversion.*

length, with termination resistors) that is connected to the I/O module of the DeltaDSP board. The I/O module provides additional trigger outputs for each of the eight DACs. If a DAC board receives such a trigger signal, it will sample the data on the bus and make a digital-to-analogue conversion. Each such DAC write access takes 125 ns, so the master DSP can write data e.g. to one DAC at a rate of 8 MHz, or to 8 DACs at a rate of 1 MHz per DAC. The master DSP and the FPGA of the I/O module are connected by four link ports that allow a guaranteed and sustained overall data transfer rate of 128 Mbyte/s. Furthermore, the master DSP can access internal FPGA registers and thus DACs directly by the 48-bit main data bus of the processor.

Control of External DACs

The FPGA on the I/O module can operate in two different modes: direct mode and ramp mode. In direct mode, the DSP core (i.e. a software command) writes new values to a DAC by setting one of eight DAC value registers in the I/O module FPGA to the desired value by using the normal processor data bus. In ramp mode, the DSP stores up to eight DAC ramp curves (e.g.

for Bodo magnet power supplies) in its internal memory and transfers them automatically to the FPGA and thus to the DACs by the four link ports (one port for two DACs) using the DMA (direct memory access) unit of the DSP. The DMA unit just has to know where the ramp curve begins and how long it is. Once started, it transfers the data automatically from the internal memory of the DSP to the DACs without any further software commands, so that the DSP core can perform e.g. feedback algorithms while the data is being sent to the DACs. In case of the ramp mode, the master DSP can determine the DAC update frequency and thus the temporal length of a given DAC ramp curve by programming one of several timers of another FPGA ("trigger-matrix/counter/timer"-FPGA, see fig. 5.2). The timer issues a programmable number of triggers at a programmable frequency, and the FPGA on the I/O module writes new set values to all DACs each time it receives a trigger.

In order to modify a DAC ramp in ramp mode e.g. for feedback applications, the master DSP can modify the respective data words of a DAC ramp curve in its internal memory just before it is transferred to the FPGA on the I/O module (and thus to the DACs). Alternatively, the DSP can write a desired (signed) DAC offset value to one of eight dedicated DAC offset registers in the FPGA of the I/O module. The FPGA adds the contents of each offset register to the respective DAC value of the original ramp curve just before it is written to the DAC. In case of the Bodo betatron tune feedback, the DSP can store a quadrupole current ramp curve in its internal memory and make real-time corrections by using these FPGA offset registers. The FPGA provides an optional overflow protection in order to avoid the 16-bit DAC set value to wrap around (e.g. from 65535 to 0) when the offset is added or subtracted. Such undesired wraparounds could cause set current jumps (e.g. of 60 A in case of quadrupole magnets) that would lead to beam loss and power supply failures.

The 16-bit differential external DAC bus and the DAC trigger connectors of the DeltaDSP board can not only be used to control devices, but also to read data from external devices, since the FPGA of the I/O module can switch all outputs to inputs by a software command. Furthermore, the FPGA can be reconfigured (i.e. internally rewired) to perform other tasks than the control of DACs, e.g. digital power supply current regulation for fast betatron tune feedbacks or detection of digital pulses of the beam loss monitors ("BLMs") that are used at Bodo and Delta [60] (with up to 16 BLMs per DSP board).

Slave DSP and DeltaNet Data Transfer

In addition to the construction of very large multiprocessing systems via link port connections, the DSPs that are used for the DeltaDSP board can be used to build an additional local multi-processing system (consisting of up to 6 DSPs) that uses the normal 48-bit data bus of the DSPs and some additional signals (for bus arbitration etc.) for inter-DSP communication [51]. Each DSP can access the internal registers and memory (256 Kbyte) of up to 5 other local DSPs via memory-mapping. Therefore, the master DSP of a DeltaDSP board can transmit and receive measurement data via DeltaNet data simply by accessing the data transmission and reception buffers of the second DSP on the DeltaDSP board that is actually connected to DeltaNet. The slave DSP is responsible for administrative tasks and also organises the measurement data transfer via DeltaNet. It receives and transmits DeltaNet data via link ports that are connected to an FPGA on a DeltaNet piggyback module that can be plugged onto the main board. The DeltaNet module has a fibre optics interface with one input and one output plug on the front panel of the DSP board, connected to the next neighbours in the DeltaNet ring.

Each DeltaDSP board can be equipped with two DeltaNet modules so that one DSP board can be part of two DeltaNet rings. This feature is not yet used at DELTA, but will allow the construction of a hierarchy of interconnected DeltaNet rings that can control even large accelerator facilities which consist of a number of smaller accelerators or accelerator subsystems (e.g. Linac, booster, storage ring, cooler rings, different interconnected feedbacks in the same storage ring, etc.). Furthermore, the second interface can be used to implement fail-safe DeltaNet rings that switch the data transfer automatically from first to second ring if the first ring fails (e.g. due to a broken fibre optics cable). This feature is of special interest in case of possible future applications of DeltaNet in orbit feedbacks at storage rings or linear accelerators that have very high particle energies. The particle beam in such accelerators usually contains so much energy that failures of the feedback data transmission may cause beam loss that leads to damages of the beam pipe or to quenches of superconducting magnets or RF cavities.

In-System Reprogramming

The DSPs can be programmed either in the programming language C [73] (with some DSP-specific language extensions) or in DSP assembler that can be embedded in the C code. Since the C-compiler for the DSPs is optimised for the generation of fast and therefore short programs, the program code for master and slave DSP fits into one of three 512 Kbyte flash memory EEPROMs (electrically erasable/programmable memory) on the DeltaDSP board. The EEPROM also contains the data that is required to program ("wire") the five FPGAs on the DSP board and its piggyback modules. The internal wiring of the FPGAs is volatile, therefore they have to be (re-)programmed after power-on. The flash EEPROM has a socket and can be programmed externally, but it is usually programmed in-system during normal operation by the master DSP, using a graphical user interface (GUI) in the DELTA control room. If the DeltaDSP board is switched on or rebooted, the master DSP loads its program code from the flash EEPROM and programs the FPGAs if they are not yet programmed. Then it downloads boot code to the slave DSP. The process of compiling new DSP code ("firmware"), reprogramming the EEPROM and rebooting the DSP board only takes a few minutes, which allows fast feedback software development cycles during normal accelerator operation. The external DACs keep their set values when the DeltaDSP board is booted, and the DeltaNet interface also operates continuously during reboot. Therefore individual DeltaDSP boards can be booted without losing a stored electron beam or disturbing the feedback network.

Memory and Accelerator Control System Interface

The DeltaDSP board is equipped with 3 Mbyte fast SRAM (static random access memory) of 48 bit width that can be accessed in a single processor clock cycle. It contains DSP program code that can be executed directly by the DSPs, as well as data that is accessed frequently and therefore requires fast access time. 16 Mbyte of DRAM (dynamic random access memory) of 32 bit width have a longer access time and are mainly used to store large amounts of measurement data (e.g. BPM data during a Bodo ramp). The communication with the EPICS[1] control system was implemented by 1 Mbyte of dual-ported memory (see fig. 5.2, "VME-RAM", left upper side). It can be accessed "from two sides", i.e. both from the VMEbus by EPICS and

---

[1]EPICS (**E**xperimental **P**hysics and **I**ndustrial **C**ontrol **S**ystem) is a universal accelerator control system that is used at a large number of accelerator facilities [71, 72]

from the local DSP data bus by the DSPs. The EPICS control system sends commands and data to the DSP board by writing it to this memory via VMEbus. A VMEbus write access to some dedicated locations in this memory causes a DSP interrupt and tells the DSP to execute a certain task, e.g. to start a measurement, to change DAC values or to send measurement data back to the control system. The DSPs can also write data to the dual-ported memory and generate VMEbus interrupts by writing to dedicated memory locations. The VMEbus interrupts tell the EPICS control system e.g. that it should read new measurement data from the dual-ported memory (for GUI visualisation, archiving etc.).

CANbus Interfaces

The DeltaDSP board also has two CANbus interfaces. CANbus is a fault-tolerant industrial fieldbus with length-dependent data rates up to 1 Mbaud. It is used to control about half of the devices in the DELTA accelerator. At present, most CANbus devices are connected to VMEbus CAN boards that are controlled by the main CPU board in the VMEbus rack. The CANbus interface of the DeltaDSP board allows the seamless integration of all CANbus-controlled accelerator devices into future DeltaDSP feedback systems, e.g. in the context of a cybernetic machine that controls and optimises the whole accelerator autonomously.

ADCs and Beam Diagnostics Data Acquisition

The data acquisition ADC system for beam and device diagnostics data consists of stand-alone ADC boards with two DC-isolated 16-bit ADCs on each board. The boards have short analogue connections to the analogue beam diagnostics front-ends (e.g. BPM electronics). The data is transferred to DeltaDSP boards by a proprietary synchronous serial bus over distances up to 50 m at a speed of 10 Mbit/s (using differential transmission via shielded CAT5 Fast Ethernet cable). Up to 32 ADCs can be connected to one serial bus in a daisy-chain fashion, with a DeltaDSP board at the end of the chain. Each DeltaDSP board can handle four ADC buses, i.e. 128 ADC boards or 256 ADCs. The ADC boards sample beam diagnostics data synchronously at frequencies between 4 and 100 kHz. The serial data is transferred directly into the internal memory of the DSPs by their DMA unit through the synchronous serial interfaces of the DSPs. DeltaNet synchronises the sampling triggers and data acquisition for all ADCs in the whole accelerator, with typically less than 300 ns time stamp error between any two ADCs in the machine. This synchronisation is essential for a high bandwidth of distributed beam parameter feedbacks.

Compatibility to the Previous System

The external DAC interface of the DeltaDSP boards is compatible to the previous VMEbus power supply control boards [56] ("FGs", German abbreviation for "Führungsgrößengeneratoren") that controlled the ramped magnet power supplies and RF power of Bodo before the DeltaDSP boards were installed. The FGs are synchronised digital function generators with a 16-bit differential output that can control only one 16-bit DAC. An FG consists of a 16 Kbyte first-in-first-out ("FIFO") memory (containing the DAC set value changes e.g. for magnet ramp curves), a timer (to determine the ramp speed) and two synchronisation chain connectors (an input and an output) that allow synchronously ramping chains of

FGs. Such FG chains were used in Bodo and are still used in Delta to ramp the magnet power supply currents. While the FGs are not "intelligent" in the sense that they cannot change the ramp curves dynamically (e.g. for feedback applications), the DeltaDSP boards allow the DAC set values to be calculated and modified in real-time.

Furthermore, one DeltaDSP board can replace up to eight FGs, and the DeltaDSP boards provide FG-compatible DAC and FG chain synchronisation signals so that FGs and DeltaDSP boards can be mixed and FG systems can be slowly upgraded to DeltaDSP systems. Due to this compatibility the FGs would have been available as a fast fallback solution if technical problems with the first DeltaDSP prototypes had occurred. The availability of a quick fallback solution was important since the migration from FGs to DeltaDSP boards in Bodo was made during a period of normal synchrotron light source operation of the accelerator. However, the fallback solution turned out to be not necessary, since the first DeltaDSP prototype is identical to the final board version, and all DeltaDSP boards work continuously without any hardware failures since March 2001. After several months of software development and burn-in tests, the FGs in Bodo were completely replaced by DeltaDSP boards in October 2001.

5.3.2 DeltaNet

5.3.2.1 Global Measurement Data Distribution

DeltaNet is a novel inter-DSP communication system that was developed in this thesis in order to combine a large number of DSPs with link ports to a distributed global real-time multiprocessing system for accelerator feedbacks. Existing fast global orbit feedbacks at synchrotron light sources that use DSPs with link ports also have arranged them in a ring around the storage ring, but the DSPs can only exchange data directly with their next neighbours in the ring via direct point-to-point connections of the DSP link ports [45]. Therefore, the DSPs in such a system have only local information about the beam orbit, and the feedback system is only a series of local feedbacks and not a truly global feedback. Global orbit correction with such systems is nevertheless possible, but only with certain correction methods that use a superposition of more or less local corrections or orbit bumps (e.g. SVD) and only if the storage ring was designed so that BPMs, steerers and beam optics lead to a sufficiently diagonal inverted beam response matrix.

However, an effective global orbit feedback at Delta should use nearly all BPMs to calculate optimal corrections for each steerer magnet [32]. Therefore, DeltaNet was designed to send the beam diagnostics data of every DeltaDSP board in a DeltaNet ring to every other DeltaDSP board and not just to its neighbour boards. This allows truly global feedbacks and does not impose any restrictions on the type of the feedback algorithm, beam optics and the number and positions of BPMs and steerers that participate in the feedback. Furthermore, the DSP boards can be arranged in any order, and additional DSP boards, ADCs and DACs can be inserted anywhere into the DeltaNet ring e.g. in order to integrate additional photon BPMs or steerers or to increase the overall computing power. In case of the local "next-neighbour" feedback communication structure described above, this would lead to problems since each DSP board must have a neighbour that is connected to certain BPMs. Furthermore, DeltaNet was designed not only as an orbit feedback network, but as a unified system that makes all beam diagnostics data of the machine available to all DeltaDSP boards in order to allow many different kinds of feedbacks and automatic DSP-based beam parameter optimisations. This also

requires a flexible global data distribution concept that can hardly be achieved with local "next-neighbour" communication schemes. In principle, the local "next-neighbour" communication schemes described above can also distribute the data globally, but only via repetitive point-to-point transfer from DSP to DSP by software. Compared to DeltaNet (that is implemented in hardware), this would take at least an order of magnitude more time, especially if so-called 32-bit cyclic redundancy checksums ("CRC checksums") are used [59] that are required for reliable detection of transmission errors.

5.3.2.2   DeltaNet Speed: Reflective Internal DSP Memory

In order to achieve substantial damping of a periodic orbit oscillation, the overall feedback loop delay of a digital orbit feedback should be less than 10% of the oscillation period [46]. The overall loop delay $\Delta\, t_{loop}$ is the sum of the BPM electronics response time, of the delay $\Delta\, t_{ADC \rightarrow DAC}$ in the digital part of the loop and of the delays due to power supply response time and eddy currents in corrector magnets and beam pipe. The delay $\Delta\, t_{ADC \rightarrow DAC}$ in the digital part of the loop results from the BPM ADC conversion time, from the speed of the processors that use the ADC data to calculate the power supply DAC set values, from the DAC settling time and from all data transmission delays between ADCs, processors and DACs. If the delays are split equally between analogue and digital parts of the loop, orbit corrections up to a cutoff frequency of 200 Hz require $\Delta\, t_{ADC \rightarrow DAC} = 250\ \mu$s, resulting in 4 kHz correction rate. A global orbit feedback in Delta with 64 BPMs and 32 steerers in each plane and 4 kHz correction rate requires global real-time network data transfer rates of several Mbyte/s. In order to avoid the network to become a bottleneck that limits the feedback bandwidth, the network should be able to distribute all BPM data to all DSPs in an overall time of about $0.3 \cdot \Delta\, t_{ADC \rightarrow DAC} = 75\ \mu$s.

Furthermore, a real-time data transmission system for DELTA should use fibre optics cable (up to some 100 m length) to avoid transmission errors due to the noise generated by pulsed high voltage components in the Delta ring.

The required guaranteed real-time performance could neither be achieved with the network and fieldbus systems available at DELTA (e.g. Ethernet, CANbus, GPIB, RS232), nor with other conventional commercial network or fieldbus systems like IBM token ring, IEEE1394 or Profibus. Furthermore, none of these networks and fieldbuses supports direct connectivity to DSP link ports, or provides the global synchronisation and trigger distribution features of DeltaNet that are described in the following sections. Some global orbit feedback systems use switched 100-Mbit Ethernet for data transmission [64], but this results in network transmission delays about an order of magnitude larger than DeltaNet (at similar baud rates), which limits the feedback bandwidth to a fraction of the respective DeltaNet value. Other accelerator feedbacks that were developed in parallel to this thesis use commercially available "reflective memory" VMEbus boards and separate VMEbus DSP boards to achieve global data distribution with short delays and high data transmission rates [48]. The data transfer between DSP board and reflective memory board in such systems has to be done via DSP software commands that access the VMEbus, which results in significant software overhead. Furthermore, the VMEbus causes undesired delays in the feedback loop in case several boards try to access the VMEbus simultaneously.

DeltaNet avoids such problems by the concept of "global reflective internal DSP memory". The DeltaNet interface was completely implemented in hardware, mainly one FPGA. It transmits ("reflects") the measurement data periodically from the internal memory of one slave DSP

to the internal memories of all other slave DSPs in the network in real-time. To achieve this, an FPGA on the DeltaNet piggyback module ("DeltaNet FPGA") is connected to two link ports of the slave DSP (see fig. 5.2), one port for transmission and one for reception of measurement data packages. The FPGA is also connected to the actual DeltaNet ring. If a slave DSP wants to transmit a data from its internal memory to the internal memory of all other DSPs in the network, it just has to start the DMA unit that automatically transfers the data to the FPGA. The FPGA puts this data into a DeltaNet data package and transfers it to the DeltaNet FPGAs and the slave DSPs of all other DeltaDSP boards in the ring.

To achieve this, FPGA is connected to a parallel-to-serial converter which sends a serial 160 Mbaud bitstream to a fibre optics transmitter. The DeltaNet package travels along a fibre optics cable "downstream" to the fibre optics receiver of the next DSP board in the ring. The serial bitstream of the receiver is connected to a serial-to-parallel converter that sends the data in parallel to the DeltaNet FPGA of the receiving DSP board. Therefore, DeltaNet is actually a ring that physically consists of point-to-point data connections. This allows much higher transmission rates than buses [59]. If a DeltaDSP board sends a data package that originates from the DSP on the board ("own data package"), the DeltaNet FPGAs of all other DeltaDSP boards in the ring receive it and re-transmit it until the package has travelled once around the ring. Then the FPGA that first sent the data package removes its own package from the ring so that it does not circulate forever. Furthermore, each FPGA sends all received DeltaNet packages (except for its own packages) also to the slave DSP. The DSP uses its DMA unit to store the incoming link port data automatically in its internal memory.

Thus, the data transmission and reception requires only the DMA unit of the slave DSP and the DeltaNet interface hardware, but not the DSP core (i.e. no software commands) except for the short configuration of the DMA unit. Since the internal DSP memory is dual-ported, the DMA unit can access the memory for DeltaNet data transmission and reception without slowing down the DSP. Thus, the DSP can access the internal memory simultaneously without "wait states". This results in less software overhead and in faster as well as in more deterministic transfer rates than commercially available solutions of existing synchrotron light source orbit feedbacks. Furthermore, the DeltaNet interface basically consists of three ICs and a fibre optics transceiver. This resulted in costs of one DeltaNet piggyback module that were about 20 times lower than the price of commercially available reflective memory VMEbus boards.

Network Access Arbitration

In order to avoid simultaneous access of several DeltaDSP boards to the network (which would result in data collisions), DeltaNet uses the so-called token passing scheme [59] for network access arbitration. The "token" is the permission to write to the network, and it is passed unidirectionally from node to node. Each DeltaNet piggyback has a unique identifier ("DeltaNet ID") that is programmed at power-up by the EPICS control system. After power-up or disconnection of cables, all DeltaNet FPGAs in the ring wait until all point-to-point connections in the ring are functional, i.e. until all DeltaDSP boards have power, all DeltaNet FPGAs are initialised and configured, and all fibre optics cables are connected. Then the DeltaNet nodes negotiate which one has the highest ID, and this one gets the token. The DeltaNet node with the token sends its first data package. It may be empty in case the DSP did not provide data to be sent. The package contains the token, therefore the nodes loses the permission to send another package. The next DeltaDSP board in the ring receives this data package and sends it

to the next DeltaNet node, but removes the token from the package before sending it. Then it sends its own (possibly empty) data package (with the token), the next DeltaNet node receives it, and so on.

The length of a data package is limited to 252 bytes plus some additional bytes for package header, checksums, token flag etc.. Thus, the token passing scheme guarantees that each DeltaNet node gets access to the network at regular deterministic intervals. Moreover, the DeltaNet package size can be limited to less than 252 bytes. This provides a guaranteed minimum correction frequency for beam parameter feedbacks in accelerators, in contrast to many commercially available networking systems with similar average data throughput rates that can be slowed down massively by very large data packages of individual network nodes [59]. The resulting increased delays in the feedback systems can result in an unstable feedback loop with subsequent beam loss.

5.3.2.3 Accelerator-Wide Feedback Synchronisation

DeltaNet is not only used for data transmission, but it also synchronises all DSP main board clocks in a DeltaNet ring (or even a hierarchy of interconnected DeltaNet rings) so that their clock frequency is identical, with a typical jitter and clock edge drift of less than 10 ns between different boards, even on long timescales. One DeltaDSP board in a DeltaNet ring (or a hierarchy of rings) is the clock master. It obtains its clock frequency either from an on-board oscillator or optionally from an external clock input. The master clock signal is multiplexed onto the same fibre optics cable that is used for data transmission. This is achieved by the parallel-to-serial converter on the DeltaNet piggyback module that converts 8 parallel bits (at 16 MHz byte rate) into 10 serial bits (at 160 MHz bit rate), which allows the insertion of a guaranteed amount of clock edges into the serial bitstream. Even if no data is transmitted to the parallel-to-serial converter it will continue to send 10-bit characters (so-called SYNC characters). This allows the serial-to-parallel converter of the next DeltaNet node in the ring to recover the master clock signal from this bitstream even if the upstream DeltaDSP board sends no measurement data, e.g. because the board was just rebooted.

Since the clock signal that is recovered from the serial bitstream might contain missing edges due to transmission errors, it is fed into a so-called timing recovery unit ("TRU") that consists of a voltage-controlled crystal oscillator ("VCO") with a phase-locked loop ("PLL") that tracks the clock edges of the input signal clock and smoothes the clock frequency. The TRU provides output frequencies of 16 and 32 MHz, with a guaranteed maximum frequency deviation of of $\Delta f_{rel,TRU} = \pm 100 \; ppm$ (part per million). The 32 MHz signal of the TRU is used as the main board clock for all ICs (e.g. DSPs, FPGAs) on the DeltaDSP board. The guaranteed maximum deviation of the nominal frequency is essential for the operation of the DSP board, since many ICs would not tolerate large deviations from the frequency or missing clock pulses due to DeltaNet transmission errors. As the main board clock is also used for the parallel-to-serial converters on the DeltaNet piggyback modules, the clock signal travels from the clock master around the complete DeltaNet ring.

The global clock synchronisation has the advantage that all hardware and even all software commands of all DSP boards are automatically synchronised. Therefore, magnet current ramps in Bodo are automatically synchronised due to synchronous timer clocks in the FPGA that determines the DAC update frequency. Furthermore, DeltaNet can distribute trigger signals, one
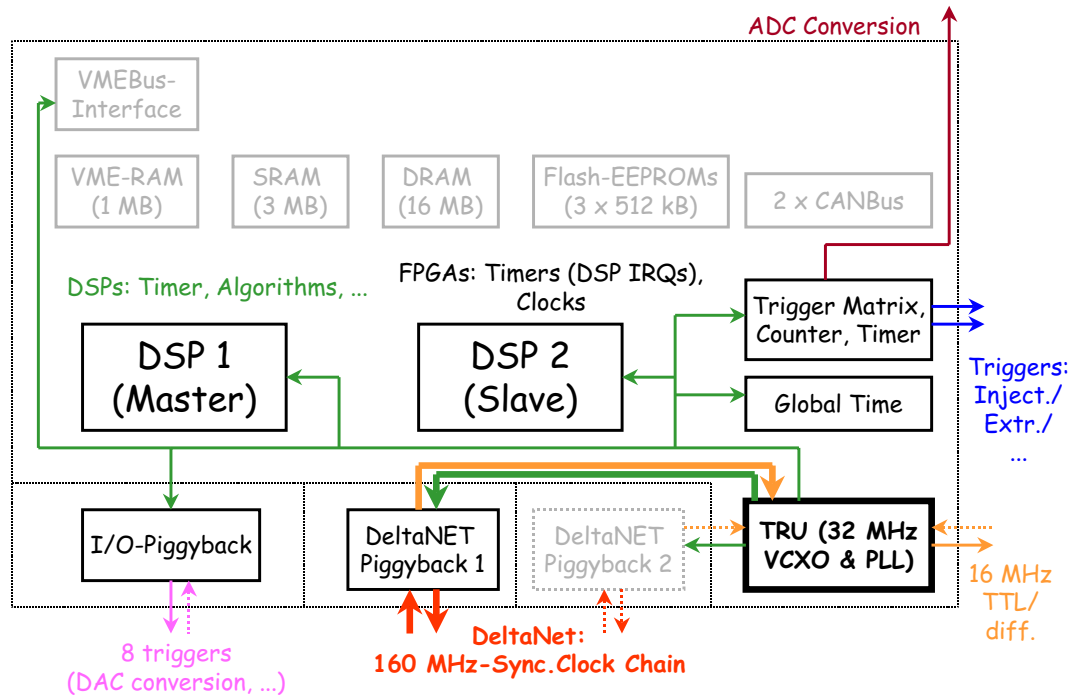
Figure 5.3: *Global DeltaDSP main board clock synchronisation via DeltaNet. The synchronisation provides identical main board clocks on all DeltaDSP boards connected to DeltaNet, with a typical jitter resp. clock edge drift of less than 10 ns between different boards, even on long timescales.*

of which is used to reset a time stamp clock on all DSP boards. Therefore, DeltaNet does not only guarantee the synchronous sampling of ADCs and DACs in the complete DeltaNet ring, but it also provides global time stamps with typically 300 ns precision for all beam diagnostics data. This allows the correlated analysis of the data of different beam monitors (e.g. BPMs, beam loss monitors, beam current monitors, betatron tune monitors etc.) connected to different DeltaDSP boards on a time scale that is about 5 orders of magnitude below the time stamp precision of the previously used CANbus data acquisition system.

The trigger-matrix FPGA can be configured to generate an interrupt in case the PLL of the TRU cannot track the incoming clock signal any more. This may occur if the master clock for the DeltaNet ring is generated externally (which is possible) and the external clock is beyond the $\pm 100\,ppm$ tracking range of the TRU. Loss of synchronisation would result in asynchronously ramping Bodo magnets and thus undesired betatron tune drifts, as well as asynchronous data acquisition and imprecision of time stamps for diagnostics data acquisition. The TRU-generated DSP interrupt can be used to indicate such states to the operator in the Delta control room.

It should be noted that DeltaNet will also work if the DeltaDSP boards in the ring are not synchronised and the individual main board clock frequencies have arbitrary values between $32\,MHz - 100\,ppm$ and $32\,MHz + 100\,ppm$. DeltaNet allows a maximum number of 255 nodes (i.e. DeltaDSP boards) in the ring. This number and the maximum number of bytes per DeltaNet data package are chosen in such a way that 254 nodes that send data packages of maximum size at a master clock frequency of $32\,MHz + 100\,ppm$ will not cause an overflow

in the reception buffer of a subsequent node that has a master frequency of $32\,MHz - 100\,ppm$ (and therefore receives data 200 ppm slower than it is being transmitted by the 254 stations). In order to avoid buffer overflows, all packages on the DeltaNet ring are usually separated by three SYNC characters. Depending on the filling state of the DeltaNet FPGA input buffer, the FPGA varies the number of these inter-package SYNC characters when received data packages are re-transmitted to the next node. The variation scheme that was implemented guarantees that buffer overflows are impossible and that all data packages are separated by at least two SYNC characters.

Transmission Error Detection

The use of fibre optics cable for the DeltaNet transmission system results in better noise immunity and less transmission errors compared to copper cables. The resulting number of transmission errors is small, but according to fibre optics transceiver data sheets not zero. The relative bit error rate depends on the input power of the fibre optics receiver and can increase from 10^{-13} to as much as 10^{-3} when the optical input power changes by only $-5\,dB$ [65]. Whilst the number of transmission errors in new fibre optics systems is usually very small, it can increase significantly due to ageing and therefore darker transmitter LEDs, dirty connectors, damaged or heavily bent cables etc. Without an error detection mechanism, the resulting data errors e.g. in BPM data of a feedback system could cause beam loss once in a while, and the reason for this beam loss might be hard to detect. Therefore, DeltaNet data packages contain so-called 32-bit cyclic redundancy checksums ("CRC checksums") that provide an extremely low probability of not detecting a transmission error [59]. This probability would be much higher in case of more primitive checksums that could be generated e.g. by calculating the sum of all bytes in the package.

DeltaNet data packages consist of a package header with a 32-bit CRC checksum at its end, and the actual data (e.g. BPM positions), with another 32-bit CRC checksum at the end. The package header contains the identifier ("ID") of the node that generated the package. In order to avoid a data package to circulate forever on the DeltaNet ring, each station removes its own data packages after they circulated once around the ring. However, there is a small but non-zero probability that some bits in a package might flip due to transmission errors and the resulting data package has valid checksums and a node ID that is different from all existing node IDs in the ring. In order to avoid such packages to circulate forever, every package header contains an package age counter that is set to the number of nodes in the ring when the package is generated. This age counter is decreased by one each time the package is received and re-transmitted by the other DeltaNet FPGAs while the package travels around the ring. The header CRC is checked and re-calculated by the DeltaNet FPGAs each time the age counter is decremented (before re-transmission of the package), therefore changes of the age counter due to transmission errors can also be detected. When the package arrives at its sender, the age counter reaches zero, and such packages are also removed from the ring, no matter what the package ID is. Since the age counter has a maximum value of 255, it is impossible that packages travel around the ring indefinitely, even in the improbable case of transmission errors that generate packages with valid CRCs and unknown sender ID. Furthermore, all packages with wrong header or data CRCs or with an invalid package format are also removed from the DeltaNet ring. This guarantees a very high data integrity and reduces the probability of beam loss due to transmission errors in DeltaNet feedback systems to nearly zero.

In case a package with the DeltaNet token gets lost, the DeltaNet node with the highest ID will issue a new package (with a token) after a programmable timeout ("master token time-out"). The node with the highest ID ("master node") is determined after power-up or after reconnection of disconnected fibre optics cables by automatic (hardwired) negotiation between the DeltaNet FPGAs. In case a non-master ("slave") DeltaNet node did not receive a token after several master token timeout intervals, it resets all DeltaNet nodes automatically. The nodes wait until all point-to-point connections are working properly again. Then they negotiate which node has the highest ID, and after that they resume normal operation of the ring.

Each DeltaNet FPGA has an internal state machine that can be in the states "bad transmission quality", "searching for master ID", "trying to close ring" and "normal". All state machines of all FPGAs in the ring are synchronised, therefore every DeltaDSP board knows if the ring is operating normally or not. As soon as all fibre optics connections in a ring are operating without transmission errors for at least 1 second, the DeltaNet state changes from "bad transmission quality" to "searching for master ID", and the FPGAs start to negotiate which one has the highest DeltaNet ID ("master node"). If one node finds out that it is the master, it changes to the "trying to close ring" state and sends a DeltaNet package with a flag that tells all other nodes that a master has been found and which ID it has. If the respective data package has travelled once around the ring without any transmission error, all nodes change to the state "normal", the master node sends its first regular data package (with the token), and the DeltaNet ring commences normal operation.

The number of transmission errors, the state of the ring and the quality of the individual point-to-point connections is visualised in the DELTA control room, which simplifies maintenance e.g. in case of transmission errors due to cable imperfections. Furthermore, the state of the DeltaNet ring is also visualised by a two-color LED on the DeltaDSP front panel, which allows fast detection of faulty cables without access to the control system.

It should be noted that all DeltaNet features described above (power-up, negotiation, time-outs, error handling, access arbitration) are implemented in hardware, therefore the network will continue to operate even if the software on individual DSP boards in the ring crashes. This allows software development on dedicated DeltaDSP testbed boards during normal synchrotron light source operation of the machine without affecting the operation and data transfer of other boards in the same ring that are vital for the operation of the accelerator.

5.3.2.4 Global Trigger Distribution

In order to avoid cycle-to-cycle betatron tune shifts and orbit drifts in Bodo, all magnet power supplies not just have to ramp synchronously, but they also have to start their ramps at the same time. Therefore, DeltaNet can not only transfer data and synchronise the main board clocks, but it can also distribute trigger signals e.g. to start the Bodo ramp. The synchronicity and start time jitter between any two DACs must be smaller than 6 μs in order to keep the resulting relative betatron tune shift below 10^{-4}. As mentioned above, DeltaNet achieves a clock main board (and therefore DAC output) synchronisation with a jitter that is much smaller, therefore the ramps will be sufficiently synchronous as long as they are started synchronously.

DeltaNet triggers are implemented by trigger flag bits in the headers of regular DeltaNet data packages. Therefore, the exact time when a trigger signal that is generated e.g. by a timer on a DeltaDSP board arrives at the DeltaNet nodes in the ring varies typically by some ten
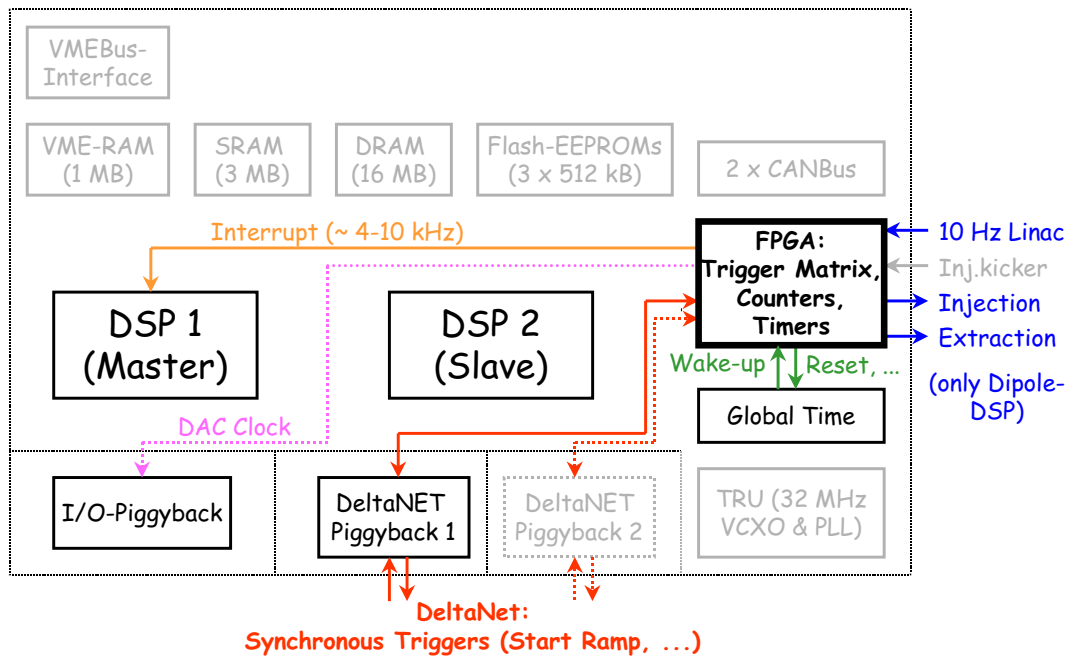
Figure 5.4: *Trigger and interrupt signal routing capability of a DeltaDSP board. The "trigger-matrix/counter/timer"-FPGA on the DSP board contains a 31x25 matrix with 31 trigger inputs and 25 trigger outputs. The DSP can route trigger and interrupt signals from any input to many different outputs of the matrix by connecting the respective matrix "crossing points" via corresponding bits in internal FPGA registers. This allows trigger and interrupt signals of various sources (VMEbus, DeltaNet, front panel plugs, timers etc.) to be routed to various destinations (e.g. VMEbus interrupts, DeltaNet triggers, front panel trigger outputs, timer start triggers, internal counters, DAC conversion, DSP interrupt).*

microseconds, because the DeltaNet node has to wait for the token before it can transmit its package with the triggers. However, this is no problem, since it is not important if the Bodo ramp is started ten microseconds earlier or later, as long as all power supplies start their ramp at the same time. Therefore, DeltaNet provides a trigger distribution mechanism that guarantees that all DeltaDSP boards in Bodo start the ramp simultaneously, with an asynchronicity of less than $300\ ns$.

The ramp start trigger is generated by a timer in the trigger-matrix FPGA on the DeltaDSP board that controls the dipole power supply. The timer is clocked by an external 10 Hz trigger of the DELTA timing system that triggers the Linac RF system (see fig. 5.4, right side). The 10 Hz triggers are generated by dividing the 50 Hz frequency of the 220 V power line, which guarantees that 50 Hz ripple e.g. on Linac magnet power supplies does not cause variations of the beam position or energy at the end of the Linac, because the electrons always see the same phase of a given magnet power supply current ripple [62].

DeltaNet Ramp Trigger Equalisation Scheme

As mentioned above, the age counter in the DeltaNet package header is decremented at every DeltaNet node and therefore the 32-bit CRC header checksum has to be checked and re-

calculated at every node. This causes an overall delay of 1.5 μs between reception and re-transmission of a circulating DeltaNet package at every single DeltaNet node. Since the Bodo DeltaNet ring consists of 7 DeltaDSP boards, the resulting overall delay in the ring is larger than 10 μs, which does not include the fibre optics cable delay of about 5 ns/m resp. 0.75 μs for 150 m. As this delay exceeds the desired limit of 6 μs, the DeltaNet FPGAs are able to delay the received trigger signals by a programmable counter in the FPGA so that all FPGAs in the DeltaNet ring receive DeltaNet triggers at different times, but send the triggers to the DSP mainboard (e.g. to start a magnet ramp) nearly at the same time. During the initialisation of the DeltaNet nodes, all nodes measure the time that a data package requires to travel once around the DeltaNet ring. Furthermore, the nodes detect how many other DeltaNet nodes the ring contains and what IDs and order these nodes have. If a DeltaDSP board is configured by the EPICS control system to accept DeltaNet triggers from a certain DSP board, it calculates how long the trigger will take to travel from node to node. Then it programs its trigger delay counter so that the trigger will be synchronous to the last node in the ring (which has a delay counter value of 0). This last node is not the DeltaDSP board just upstream of the trigger generator board, but the trigger generator board itself, i.e. the dipole DSP board sends a ramp start trigger onto the network and starts to ramp when it receives its own trigger that has travelled once around the ring. This trigger equalisation scheme reduces the asynchronicity of DeltaNet triggers typically by a factor of 30, which is sufficient for all desired applications of the system.

Beam Injection and Extraction Triggers

Since the injection and extraction in Bodo must be synchronised to the energy ramp, the dipole DeltaDSP board also generates the required injection and extraction triggers for the DELTA timing system. However, these triggers are not sent onto the DeltaNet ring, but just to differential trigger outputs of the DSP board that are connected directly to the timing system. The trigger delays with respect to the start of the ramp can be adjusted by a graphical user interface (GUI) in the DELTA control room. The trigger is generated by software, but due to the deterministic interrupt behaviour of the DSPs the jitter is typically ± 2 μs, which results in a negligible relative extraction energy error that is below 10^{-5} in case of a worst-case extraction at 500 MeV at an energy change of 1 GeV/s.

5.3.3   Implementation at Bodo

Fig. 5.5 shows an overview of the DeltaDSP system in Bodo. The blue boxes in the center of the figure are the DeltaDSP boards, the blue arrows indicate the DeltaNet ring. The DSP boards control one dipole, six quadrupole and two sextupole magnet circuits, as well as orbit correctors for the DC extraction bump and for the bending angle correction of the 10 degree dipole magnets. Furthermore, the DSP boards control 16 regular steerer magnets (integrated into quadrupoles) and the RF power of the cavity. Moreover, the DSP boards read the signals of beam loss monitors ("BLMs"), beam current monitor and beam position monitors ("BPMs"). One DSP board ("bo-tune") is used for betatron tune measurement. It controls DDS ("direct digital synthesis") sine generators that determine the beam excitation and beam oscillation analysis frequencies, as well as the amplitude of the beam excitation kicker magnet. The DSP board also samples an analogue signal of the betatron tune RF front-end. A detailed description of the DSP-based betatron tune measurement and feedback system is given in chapter 6.
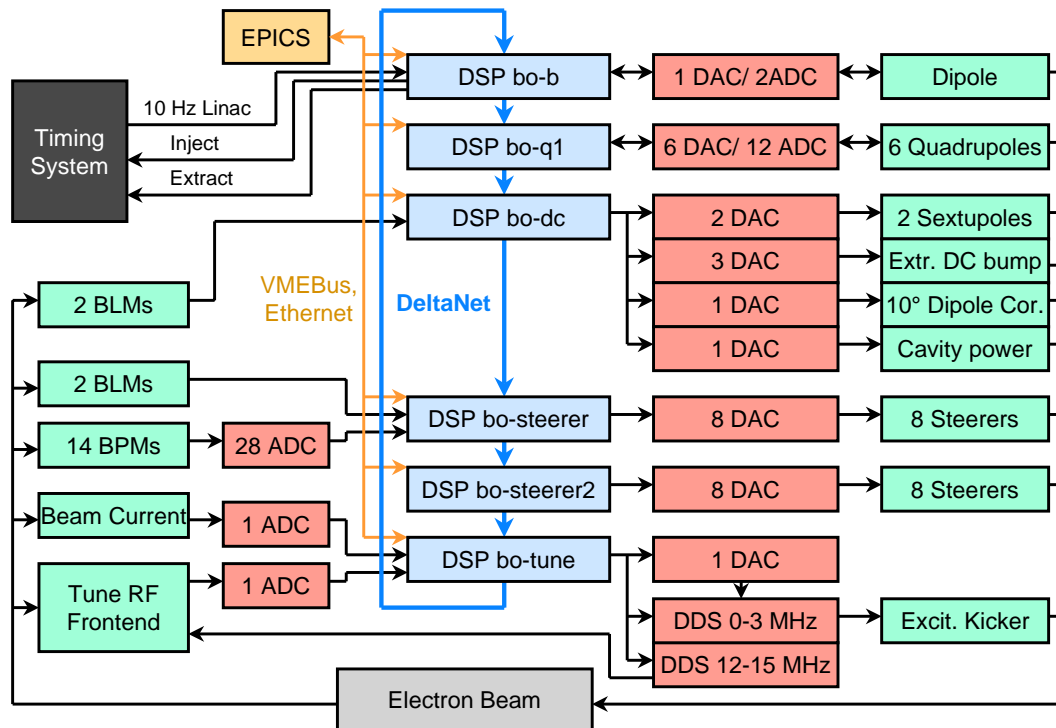
Figure 5.5: *Overview of the distributed DSP multiprocessing system of Bodo, including the actuators (green boxes, right side) and sensors (green boxes, left side) that are handled by the system. Blue boxes are DeltaDSP boards, red boxes are ADCs, DACs and DDS (direct digital synthesis) function generators, green boxes are beam diagnostics systems, magnets with power supplies, and the RF cavity with its transmitter.*

5.4   Software Architecture of the DeltaDSP System

This thesis does not only comprise the development of hardware for a novel DSP-based unified beam diagnostics and feedback system, but also the development of all software that was necessary to integrate the system into the DELTA control system. This includes all beam diagnostics and real-time feedback software on the DSP boards itself, as well as device drivers for the communication with the EPICS control system (via VMEbus) and all graphical user interfaces. The following sections give an introduction to the concepts and architecture of the newly developed software and to the DELTA control system in general.

5.4.1   The DELTA Control System

During the work on this thesis, the DELTA control system was changed from a system than was developed in-house to EPICS (**E**xperimental **P**hysics and **I**ndustrial **C**ontrol **S**ystem) [71, 72]. EPICS is a multi-purpose control system software architecture that is used at numerous accelerators and other large experimental facilities (e.g. telescopes). The "old" DELTA control system did not have all features that are necessary for the efficient operation of modern synchrotron light sources. While the available man-power did hardly allow to implement the required features into the old control system, EPICS now provides these features. Furthermore, EPICS is
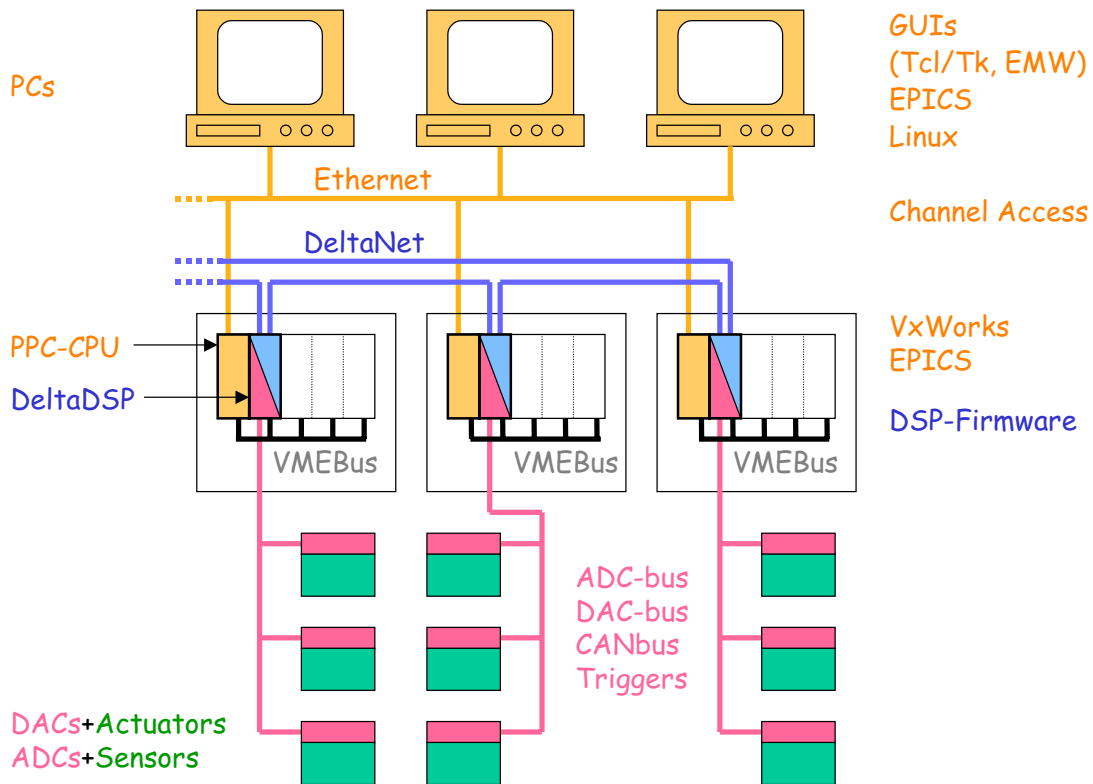
Figure 5.6: *Hardware Layers of the Bodo Control System.*

an open system which is maintained and continuously improved by a world-wide community, which reduces the required in-house man-power for the long-term maintenance of the system. One design principle of EPICS is the extensive use of configuration files in order to allow the adaption of the system to different accelerator facilities without having to modify software. Despite these concepts, EPICS still did not provide all software (e.g. drivers for FGs or DeltaDSP boards) that was required to run the complete hardware of DELTA. EPICS rather consists of a software toolbox that was used to create those parts of the "new" DELTA control system which are not machine-specific.

This thesis includes the development of all software that was required to integrate the novel DeltaDSP beam diagnostics and feedback system into the EPICS control system. Furthermore, the previously used control system hardware (FGs) and digital I/O boards that control the Bodo magnet power supply interlock systems were also integrated into EPICS in order to have a quick hardware fallback solution during first test of the new DSP hardware without having to change the complete control system. While the FGs are not used in Bodo anymore, they still control most power supplies in the storage ring Delta, therefore the EPICS integration of the FGs during the work on this thesis was necessary for both rings. Both new DSP hardware, software and EPICS control system were installed in Bodo in October 2001. Since then, no hardware failures occurred, and the down time of the machine during synchrotron light source operation due to newly installed software of the DSP system was 2 hours in one year.

Graphical User Interfaces

Figure 5.6 shows the hardware layers (text on left side) and software layers (text on right side) of the DSP-based Bodo control system. The accelerator is controlled and monitored by graphical user interfaces ("GUIs") on Linux PCs in the Delta control room. The GUIs were mainly developed with the graphical user interface builder EMW ("Epics MegaWidgets") [66] that was also developed during the work on this thesis in cooperation with a company that built the turnkey Linac for the Swiss Light Source and used EMW to control it [67]. EMW is based on an object-oriented extension of the Tcl/Tk programming language [68] called "IncrTcl/IncrTk" [69]. EMW is not only used for all Bodo GUIs, but also for most GUIs of the other parts of DELTA. It allows to develop new GUIs very quickly without writing software, simply by mouse- and keyboard-based placement and configuration of a variety of graphical elements (so-called "EMW widgets", e.g. buttons, sliders, 2D-graphs etc.) in a main window. The widgets have connectivity to EPICS, and the placement and configuration of the widgets are saved in a configuration file that is read by EMW each time the GUI is started (in order to generate the GUI). While the GUIs of the previous DELTA control system consisted of a large number of programs (one for each GUI) from different authors, all EMW-based GUIs use the same software (namely EMW) and consist of one configuration file per GUI. This simplifies both the maintenance of the software and the operation of the accelerator, since all GUIs have a uniform "look and feel" (e.g. color schemes for alarms, device errors etc.). Furthermore, EMW GUIs can be developed by people with no knowledge of programming languages, whereas experienced users can easily program additional EMW widgets where necessary (that can afterwards be used by all EMW GUI designers). The use of the programming language IncrTcl/IncrTk allows rapid prototyping of new widgets, whereas the extension of common EPICS GUI builders [70] that are based on the C programming language and the so-called "Motif" widget library usually takes several times longer.

Channel Access

The GUIs on the Linux PCs communicate via Ethernet with the a CPU VMEbus board ("VME-CPU", mainly of type "MVME 2306" from Motorola, Inc.) in the first slot of a VMEbus rack. The remaining VMEbus boards (e.g. FGs, digital I/O boards or DeltaDSP boards) in the rack are connected to the actuators and sensors of the accelerator. VMEbus is a robust industrial computer bus system which is used by most accelerator facilities that use EPICS. The VME-CPU board uses the real-time control system VxWorks [75] which executes EPICS software that is required to communicate both with the PCs and with the actuators or monitors of the accelerator. The VME-CPU exchanges data with the GUIs on the PCs by the EPICS-specific "channel access" ("CA") protocol (via Ethernet). Among other features, CA provides a fast and efficient platform-independent exchange of data (e.g. floating point numbers, waveforms, integers, strings, etc.).

EPICS Records

The key elements of the EPICS control system are the so-called records. A record consists of several named fields that describe the properties of an actuator or sensor (or of a part of an actuator or sensor). EPICS provides records of different data types (analogue, waveform, string,

binary etc.) in order to set actuator values ("output records") and to read sensor values ("input records"). The number of fields of a record depends on the record type, as well as the names and types of the fields. The records are located on the VME-CPU. They are created from a configuration file (when the VME-CPU is booted) that determines the names of the records and the initial values of the record fields. Each record must have a unique name that should describe its function. The values of most record fields can be read and set by GUIs on the Linux PCs. In order to access a record field, the operator (or designer of a GUI) has to provide a string that consists of the record name and the field name, separated by a dot. In case of a record (named e.g. "bo-qf1-i:set") that is used e.g. to control the power supply set current of the quadrupole "qf1" in Bodo, there are fields for the desired set current ("VAL"), for the engineering unit ("EGU"), for minimum an maximum current values, alarm limits, readout periods, etc. . By convention, records names at Delta usually consist of a dash-separated concatenation of the machine section (e.g. "bo" for Bodo), the name of the machine component ("qf1") and the property of the machine component ("i" for the current) that is handled by the record.

In order to set an actuator value or to read a sensor value, a record must be "processed". This can be done periodically (by specifying a period in the "SCAN" record field) or triggered either by a GUI (e.g. by writing a power supply set current to the "VAL" field of an analogue output record) or by the interrupt of a VMEbus board (e.g. to tell an analogue input record to transfer new orbit data from a DeltaDSP board to the "VAL" field of the analogue input record). When the value of a record field is displayed by an EMW GUI, the so-called "monitor" mechanism (that is part of the EPICS communication system between PCs and VME-CPUs) automatically updates the GUI when the record field changes its value.

5.4.2 EPICS Integration of the DeltaDSP System

Figure 5.7 shows the different software layers that were developed during the work on this thesis (black, green, blue and red boxes) in order to integrate the DeltaDSP system into the EPICS control system. The yellow boxes indicate the software layers that were already provided by EPICS, the white boxes on the left indicate the operating systems of the respective computers. The grey boxes on the righthand side of fig. 5.7 indicate measurement data files and configuration files that are accessed either directly (if hard disk and software are located on the same computer) or by NFS (network file system [59]) via Ethernet. The extensive use of configuration files allows to adapt the system to different requirements without major software modifications.

The DeltaDSP boards in Bodo are used for a variety of different applications, e.g. global orbit feedback, betatron tune measurement, betatron tune feedback, digital magnet power supply control, beam loss measurement etc. . Therefore, the software on the different DSP boards is not identical, but differs from board to board (red box in fig. 5.7, "user algorithms"). However, the DeltaDSP system is meant to be a universal feedback system, and it should be possible to use it for different (present and future) beam diagnostics and feedback applications with a minimum of additional software development. Therefore, the DeltaDSP software concept separates "user software" on the DeltaDSP board from the software that is common to all DSP boards ("system software", blue boxes in fig. 5.7). In order to allow the EPICS control system e.g. to control magnet power supplies and to read beam monitor values, every DeltaDSP board provides so-called "DSP variables" (for scalar values) and "DSP arrays" (for a set of monitor values). Variables and arrays are accessed via VMEbus by specifying their number. They are

Software Layers: Files:
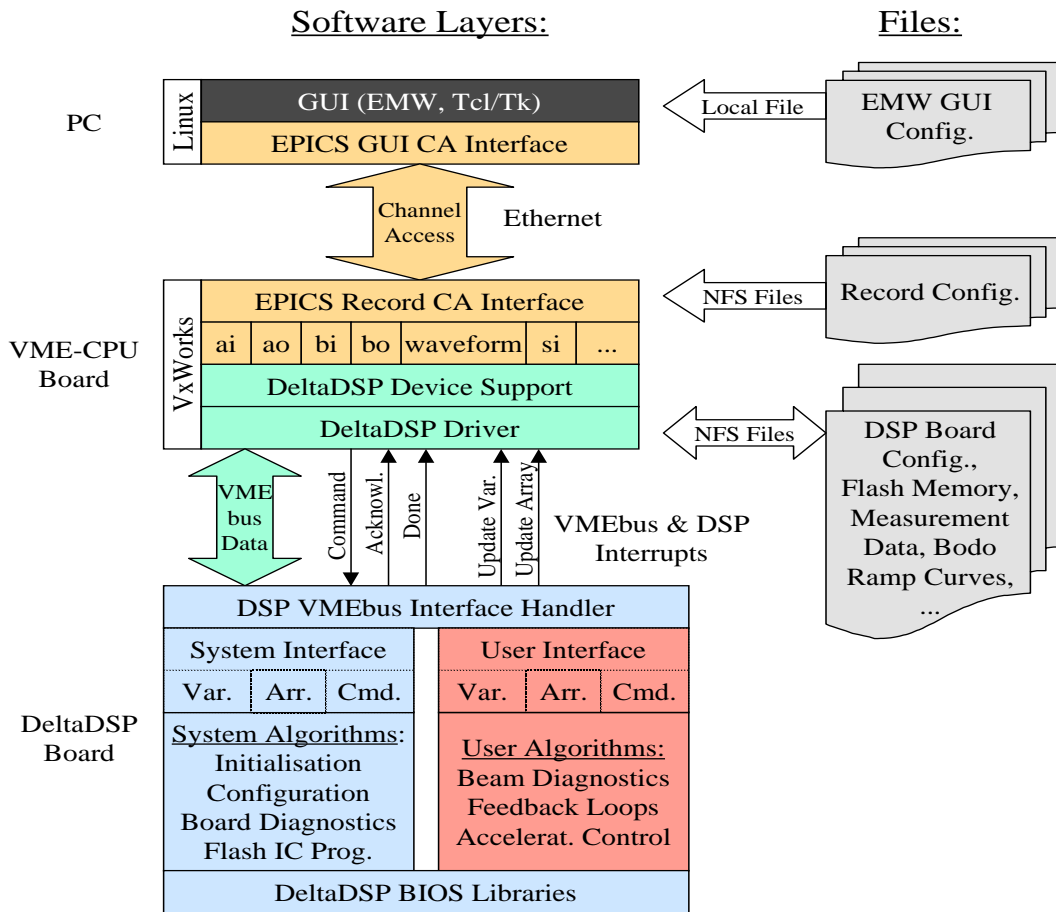
Figure 5.7: *Generic integration of the DeltaDSP feedback system into EPICS.*

separated into those that are common to all DSP boards ("system variables", "system arrays") and those that are specific for the individual user software of the DSP board ("user variables", "user arrays").

An example for a system variable is the number of DeltaNet transmission errors (that is counted by the DeltaNet FPGA in order to monitor the quality of the fibre optics connections), or the number of ADCs that is connected to each serial ADC interface line of the DSP board. The number of ADCs is monitored in order to stop a feedback loop and thus avoid beam loss if the ADC lines are disconnected. An example for a user variable is the beam position at a certain BPM, whereas user arrays are used e.g. to read the position of one or several BPMs as a function of time (i.e. a larger number of successive BPM ADC samples).

Furthermore, DeltaDSP boards provide "system commands" that are common to all DSP boards (e.g. in order to program the DSP software and FPGA "wiring" into the on-board flash memory), and "user commands" that are specific for certain boards. Each command has a number, some optional arguments and a return value that is generated by the DeltaDSP board after the command was executed.

The DeltaDSP boards in a VMEbus rack are controlled by the DeltaDSP driver software of the VME-CPU. The DeltaDSP driver (fig. 5.7, lower green box) is an EPICS-independent C++

[74] class that provides full access to all features of every DeltaDSP board (including all commands, variables and arrays that are provided by the DSP board). The driver sends commands to the DSP board by writing the number and arguments of the command to certain locations in the dual-ported VMEbus memory. Then the driver increments the contents of a certain location ("DSP command counter") in the DSP VMEbus memory by one (with wraparound when the maximum value is reached), which causes a master DSP interrupt on the DeltaDSP board (fig. 5.7, black "Command" arrow). Since the interrupt is shared with other interrupt sources on the board, the VMEbus interface handler of the interrupted DSP board checks the value of the DSP command counter. If the value has changed since the last interrupt, the DSP reads the number and arguments of the command and stores them in a queue for later execution. Then it generates a VMEbus interrupt (fig. 5.7, black "Acknowledge" arrow) which tells the DeltaDSP driver on the VME-CPU that the DSP has received but not yet executed the command. The commands are executed later in a loop outside the interrupt context, and the master DSP generates a VMEbus interrupt for each executed command (fig. 5.7, black "Done" arrow) in order to tell the DeltaDSP driver that the command is finished. Due to this asynchronous handshake mechanism, the DeltaDSP driver can issue a large number of commands without having to wait for the execution of each command (since the execution could take a while and would therefore slow down the VME-CPU).

Read and write access to DSP arrays and variables is also implemented by commands, therefore the handshake mechanism is the same. Whilst DSP variable values are directly returned to the DeltaDSP driver with the "Done" interrupt, DSP arrays are read and set by writing part or all of the array data to a buffer in the dual-ported VMEbus memory of the DSP board. Address and length of the buffer are specified by the DeltaDSP driver in the array read or set command. Since some arrays contain several MByte of data and might not fit completely into the buffer, the DeltaDSP driver and the DSP software provide a mechanism to read or set large arrays in several steps. As the channel access network protocol of EPICS is not able to transfer very large data arrays efficiently and fast from the DSP board to the Linux PCs, the DeltaDSP driver provides EPICS-indepentent functions to write data files from Linux PCs to DSP arrays and vice versa (via NFS). These functions are used to transfer e.g. several MByte of BPM measurement data from the 16 MByte DRAM memory of the DeltaDSP board to a file, or to write desired DAC settings of magnet ramp curves from a file to arrays in the DSP memory.

Furthermore, the generic part of the software on the DeltaDSP board provides software libraries that were developed during the work on this thesis in order to have comfortable access to all hardware features of FPGAs and DSPs on the board. These libaries include e.g. functions to program FPGA timers, configure the DeltaNet FPGA, program the on-board flash memory, write ramp curves to magnet power supplies, etc. . Furthermore, the libraries provide simple access to the global shared memory system of the DeltaNet network, which allows to read measurement data from other DeltaDSP boards and transfer data to other DeltaDSP boards with minimal additional programming effort.

In order to transfer data between EPICS records and a DeltaDSP board, appropriate so-called EPICS device support software was developed. This software is the interface between the hardware-independent EPICS records and the EPICS-independent DeltaDSP driver that handles all hardware access to the DeltaDSP boards (see fig. 5.7). The device support software reads a configuration string in the "DTYP" field of an EPICS record after the VME-CPU is booted. Each time a record is processed, the device support of the record initiates a user or system DSP command or a read or write access to a user or system DSP variable. Which of

these actions is performed depends on the contents of the "DTYP" field of the record. This field also determines which DSP board in a VMEbus crate is accessed and how record field values or data files are read from or written to the DSP board, as well as the type, scaling, offset and bit width of the transferred variables etc. . Supported record types include analogue, binary, multibit-binary, short and long integer and string input and output records, as well as waveform records. The latter can be used to visualise DSP arrays. String records are used mainly for data transfer between DSP arrays and data files on Linux PCs (simply by writing the filename to the "VAL" field of the record).

In addition to periodic processing of records (e.g. in order to read BPM values every 100 ms), the device support and DeltaDSP driver support record processing triggered by the DeltaDSP boards via dedicated VMEbus interrupts generated by the master DSP (see fig. 5.7, "Update..." arrows). This feature can be used to update records automatically each time new beam monitor data is available on a DeltaDSP board. However, at present the typical the data aquisition frequency of the DeltaDSP boards ranges from kHz to tens of kHz, whereas update rates beyond some ten Hz for record fields and GUIs make little sense. Therefore, it is usually sufficient to update records periodically by polling the DeltaDSP boards at the desired frequency. Since the polling mechanism uses an asynchronous interrupt-driven handshake scheme as described above, it has no performance drawbacks with respect to purely DSP-triggered record processing.

Visualisation and file storage of beam monitor data at sampling rates up to many kHz is presently achieved in an oscilloscope-like fashion by storing measurement data in RAM on the DeltaDSP boards. The data (e.g. beam positions as a function of time) can then be visualised as 2D-plots on EMW GUIs (using DSP arrays and waveform records), with zoom and scroll features of the time axis of the plot. Furthermore, the data can be stored in files via string output records, by writing the file name to the "VAL" field of the record.

Due to the generic software concept, all software that was developed for the EPICS integration of the DeltaDSP boards is identical for all boards, except for the "user software" on the DeltaDSP boards. Therefore, the newly developed DeltaDSP system allows rapid prototyping and quick development of new DSP-based beam diagnostics and feedback applications at DELTA, as well as at other accelerator facilities. Nearly all parts of the existing software can be adapted to new applications by configuration files, and only the actual feedback loop on the DSP board has to be implemented in software.

Chapter 6

Development of a DSP-Based Real-Time Betatron Tune Feedback System

This chapter describes a DSP-based betatron tune measurement and real-time feedback system for the booster Bodo that was developed in this thesis. The tune measurement system is based on the previous all-analogue tune measurement system of Bodo that was developed in a diploma thesis [4]. While the RF front-end (see following sections) of the previous system could also be used for the DSP-based system, most of the remaining analogue subgroups and the data acquisition system [61, 63] were replaced by DSP-controlled digital hardware that is based on the universal DeltaDSP beam diagnostics and feedback hardware and software architecture described in chapter 5. The DSP-based system allows 10 times faster tune measurement at higher resolution compared to the previous system. Furthermore, all parameters of the DSP-based system are controlled, monitored and optimised automatically in real-time by DSPs and the EPICS control system via graphical user interfaces, whereas the previous all-analogue hardware required continuous manual adjustment during booster ramps in order to obtain reasonable measurements and avoid beam loss. The increased measurement speed and the DeltaNet integration of the measurement system allowed the implementation of a tune feedback that corrects and optimises the betatron tunes of Bodo in real-time during arbitrary energy ramps of the booster. Section 6.1 gives a brief introduction to tune measurement methods for circular relativistic electron accelerators. The remaining sections describe the architecture of the Bodo tune measurement and feedback system. Furthermore, tests and calibration measurements of the system (without electron beam) are presented. Measurements and a characterisation of the tune measurement and feedback system with the Bodo beam are described in chapter 8.

6.1 Betatron Tune Measurement in Circular Accelerators

6.1.1 Coherent and Incoherent Tune

As defined in section 2.3, the number of horizontal resp. vertical betatron oscillations per turn is called betatron tune, working point or Q-value of a particle storage ring. In order to measure the tunes, one can kick the beam transversely for a single turn and measure the frequency of the resulting coherent oscillations of the electrons (i.e. oscillations of the center of charge of the beam) during the following turns [90]. It should be noted that the resulting tune is called coherent betatron tune, in contrast to the incoherent tune, which is defined as the number of betatron oscillations of an individual electron in a bunch which is not performing coherent oscillations. Self-fields of electrons in circular accelerators with very flat beam pipes, high single bunch charge densities and low energies can lead to significant differences between coherent and incoherent tunes [77]. The following discussion is limited to measurement techniques for the coherent tunes of circular accelerators.

6.1.2   Excitation of Coherent Betatron Oscillations

Coherent transverse beam oscillations can be generated by transverse magnetic or electric fields that kick the beam either once (on a timescale of one turn) or periodically (i.e. every turn) [90]. In case of Bodo and Delta, the beam is excited by a slotted-pipe in-vacuum kicker magnet that generates a diagonal transverse field [78]. A single strong kick of this magnet would result in simultaneous coherent oscillations in both planes, whereas periodic excitation with a small periodic sine-wave magnetic field will only lead to significant coherent oscillations in one plane if the kicker frequency matches the betatron frequency of that plane since the x and z plane of the beam are basically two weakly coupled oscillators of high quality factor. Instead of exciting the beam with a mono-frequency sine wave, some accelerator facilities also use a superposition of many frequencies, i.e. white noise [88]. However, this would require a much higher power for the amplifier that drives the kicker magnet in order to achieve a sufficient amplitude of the frequency component that matches a betatron frequency. In order to allow independent excitation of both vertical and horizontal oscillations with moderate amplifier power, the DSP-based tune measurement and feedback system that was developed in this thesis uses periodic beam excitation with a digitally generated sine wave that has a DSP-controlled frequency, phase and amplitude.

6.1.3   Detection of Coherent Betatron Oscillations

This subsection discusses some common methods for the detection of coherent betatron oscillations.

Turn-By-Turn Beam Orbit Reconstruction

A finite energy spread and different oscillation amplitudes of individual electrons in a circular accelerator cause a finite tune spread due to non-zero chromaticity and nonlinear optics. In case of a single short transverse kick of the beam, this leads to a growing decoherence of individual electron oscillations and thus to a decreasing center-of-charge oscillation amplitude of the initially coherent bunch oscillation [22]. After a single kick, one can sample the beam positions at one or all BPMs turn-by-turn. If all BPMs are sampled and the phase advance between the BPMs is not too large, one can calculate the differences to the beam orbit before the kick, normalize these differences to the square root of the theoretical beta functions and plot the resulting points as a function of the phase advance. Since the resulting plot is an exponentially decaying sine oscillation, the tune frequencies can be obtained by appropriate numerical analysis of the data. While simple fast Fourier transform ("FFT") methods have the advantage of fast computation, refined methods achieve higher precision (up to tune errors below $2 \cdot 10^{-4}$ [79]) at the expense of increased computation time.

Single BPM Turn-By-Turn Sampling

Instead of sampling all BPMs, one can also sample the turn-by-turn beam positions at only one BPM and perform the same computations (e.g. FFT) as described above. However, this will only deliver the so-called "fractional" betatron tune, i.e. the absolute value of its distance to

the next integer, since the method cannot distinguish between tunes of e.g. 1.1, 1.9, 2.1, 2.9 etc.. The distinction between tunes "above" (e.g. 1.1) and "below" (e.g. 1.9) an integer can be made by a small variation of one or several quadrupole magnet fields and a comparison of measured and expected fractional tune frequency shift. In case of the usual Bodo optics, an increase of the focussing strength of one or all horizontally focussing quadrupoles ("QFs") will increase the horizontal tune significantly and decrease the vertical tune slightly, and vice versa for the horizontally defocussing quadrupoles ("QDs"). The integer part of the tune can be obtained either from a theoretical optics model or by deflecting the beam statically with a single horizontal or vertical steerer magnet and analysing the resulting constant orbit perturbation (that is a static betatron oscillation with a phase jump at the location of the kick, see section 2.8, eq. 2.52). As the tunes of the Bodo optics in chapter 4 are below the value of 5 and the tune error of theoretical optics models is usually much smaller than 10 percent, the integer part of the tune in Bodo can be safely predicted by an optics model.

Single BPM Frequency Analysis

Instead of kicking the beam once and sampling BPM positions turn-by-turn, one can also excite the beam with a transverse sine magnetic field of a sufficiently "fast" in-vacuum kicker magnet (in order to avoid beam pipe eddy currents). When the kicker excitation frequency f_{kick} is swept from 0 to $f_0/2$ (with f_0 being the beam revolution frequency, e.g. 5.95 MHz for Bodo), the kicker will excite coherent beam oscillations when the frequency crosses the fractional vertical or horizontal tune frequencies. Instead of using turn-by-turn BPM position sampling, one can also detect these coherent oscillations by analysing the frequency spectrum of the direct BPM pickup signals, e.g. with a spectrum analyser. If f_{kick} is equal to the horizontal fractional tune frequency and thus generates coherent oscillations, the frequency spectrum of a BPM pickup will contain respective sidebands of the n-th revolution harmonics, i.e. peaks in the spectrum at the frequencies $n \cdot f_0 \pm f_{kick}$ (with n being an integer) [8]. Therefore, tunes can be measured by sweeping the kicker frequency from 0 to $f_0/2$, analysing the spectrum simultaneously at the frequency $n \cdot f_0 + f_{kick}$ or $n \cdot f_0 - f_{kick}$ and detecting the two tune frequency peaks in the spectrum (provided the kicker is able to excite oscillations in both transverse planes). This tune measurement scheme is called swept-frequency method [81].

Phase-Locked Loop

Phase-locked loop ("PLL") tune measurement systems excite the beam at its betatron frequency with an oscillator-driven kicker. The system tracks the betatron frequency by measuring the phase difference between excitation and beam oscillation and correcting the excitation frequency so that the phase difference remains 90°. Thus, the kicker frequency is always identical to the usually drifting betatron frequency. Since such closed-loop regulation system are very sensitive to frequency-dependent phase delays and drifts in the signal chain, they are mainly used in very large circular accelerators that have lengths in the order of some kilometres and revolution or fractional betatron tune frequencies in the order of 100 kHz or less [84, 85, 86]. In the case of Bodo with its revolution frequency of 5.95 MHz and betatron tunes between 2 and 3 MHz, the sensitivity to phase delays and drifts in the signal chain would be much larger compared to existing PLL-based tune measurement systems, since a given time delay results in a phase delay that is about 60 times larger compared to a storage ring with 100 kHz revolution

frequency. Furthermore, PLL-based tune measurement systems require a significant amount of time in order to find the tunes and thus to "lock" the PLL. Therefore they are mainly used for storage rings that are ramped either very slowly or not at all so that there is sufficient time (in the order of seconds) for the locking procedure [85]. As will be shown in chapter 8, the Bodo tune feedback requires tune measurement at a rate of at least some 100 Hz immediately after beam injection in order to achieve reliable real-time correction of tune drifts. Due to the high betatron frequencies and the required measurement speed, PLL-based tune measurement is not suitable for Bodo.

Choice of the Detection Method

The "old" analogue Bodo tune measurement system was based on the so-called swept-frequency method, whereas the DSP-based system achieves more than 10 times faster tune measurement by combining the swept-frequency and the "chirp" method [81]. Common implementations of the "chirp" method excite the beam by sweeping the kicker frequency in a small window around the betatron tunes, which would sound like a chirping bird if the revolution and tune frequencies of the machine were sufficiently low. Then wavelet or FFT analysis of turn-by-turn BPM position samples is used to obtain the tunes [82, 83] . In contrast, the DSP-based "chirp" method that was used for Bodo is based on frequency domain analysis (see below).

Since the commercially available RF front-end electronics of the beam position monitors in Bodo and Delta are not capable of turn-by-turn beam position measurements [80], the respective tune measurement schemes described above (that need turn-by-turn beam position sampling) were not used. However, the universal I/O piggyback module of the DeltaDSP board and its reprogrammable FPGA allow the integration of future turn-by-turn BPM electronics systems into the DeltaDSP system with minor or no hardware modifications of the piggyback and without modifications of the DeltaDSP mainboard.

6.2 Motivation

6.2.1 Use of Bodo as a Testbed

At the time when Bodo was designed (in the early 1990s), booster synchrotrons of existing and planned synchrotron light sources usually used capacitor-inductor resonant circuits ("White Circuit") of fixed frequency (typically 10 to 50 Hz) for their focussing and bending magnets [15, 16]. In contrast, the magnets of Bodo have DC power supplies. They allow the booster to be operated both as storage ring and synchrotron (i.e. ramped storage ring) with arbitrary energy and magnet focussing ramps (within the current and voltage limits of the power supplies). Therefore, beam can be stored in Bodo at any energy between 30 MeV and 1.5 GeV. The vacuum and RF system allows beam lifetimes of at least some ten minutes at high beam energies. Furthermore, Bodo uses the same magnets and beam pipe as the storage ring. The similar design allows the booster to serve as an ideal testbed for the storage ring, e.g. for newly developed accelerator components, monitors, hardware and software. Furthermore, the flexible beam optics of Bodo and the wide energy range of 30 MeV to 1.5 GeV allows machine physics studies even at very low energies that are not feasible in Delta (that has an energy range of

about 300 MeV to 1.5 GeV). The availability of Bodo as a test machine becomes even more important since the scope of Delta has changed in the last years from a test accelerator to a synchrotron light source. The growing need of beam time for synchrotron radiation users and the required availability, uptime and stability of machine hardware and software imposes severe limitations on tests and modifications of machine components, hardware and software in the storage ring. Since the booster is only required for refills of the storage ring, the remaining time is fully available for testbed applications during normal synchrotron light source operation of Delta.

Bodo as a test machine requires precise control and correction of beam optics and betatron tunes, e.g. for machine modelling or for the characterisation of newly developed monitors. Precise tune control allows e.g. to generate well-defined beam loss rates by resonance crossing for tests of newly developed beam loss monitor systems [76]. Furthermore, measurement techniques for other beam parameters like beam optics and chromaticity are based on betatron tune measurement, and the respective measurement errors depend directly on the tune measurement precision.

6.2.2 Storage Ring Filling Time

The drawback of the magnet and power supply design of Bodo is the limitation of the Delta injection frequencies to typically 0.1 to 0.2 Hz, compared to 10 to 50 Hz of White Circuit synchrotrons (see section 3.2). Therefore, Bodo must deliver a beam charge that is two orders of magnitude larger in order to achieve the same storage ring filling times as conventional White Circuit boosters. Furthermore, drifting Bodo betatron tunes that touch or cross a partially destructive resonance spend 100 times more time on or in the vicinity of the resonance, leading to beam loss that is up to 100 times larger than in White Circuit boosters (assuming the same relative beam loss rates per time). Consequently, beam loss and betatron tune drifts that may be tolerated in conventional boosters cannot be accepted in Bodo. High charge transfer rates between Linac, T1, Bodo, T2 and Delta require well-defined beam optics and betatron tunes in Bodo, both during the ramp to avoid beam loss and at the time of injection and extraction to guarantee optimal charge transfer. The new "low-emittance" beam optics that were developed for Bodo in chapter 4 have stronger focussing and larger tunes than the previously used optics, which causes increased absolute tune drifts and optics changes for a given relative error in dipole or quadrupole power supply currents (see fig. 6.1). The use of one of these optics for routine machine operation also requires precise measurement and correction of betatron tunes.

6.2.3 Tune Drift in Bodo

In an ideal booster synchrotron, orbit, quadrupole field strengths and thus betatron tunes are constant during energy ramps. Ideally, this only requires to increase the quadrupole magnet currents proportionally to the beam energy (assuming no field saturation). In reality, a variety of different sources can result in differences between ideal and actual tune in Bodo:

- Desired and undesired orbit perturbations that change the betatron tunes due to sextupole magnet fields

- Different current regulation systems for dipole and quadrupole magnet power supplies
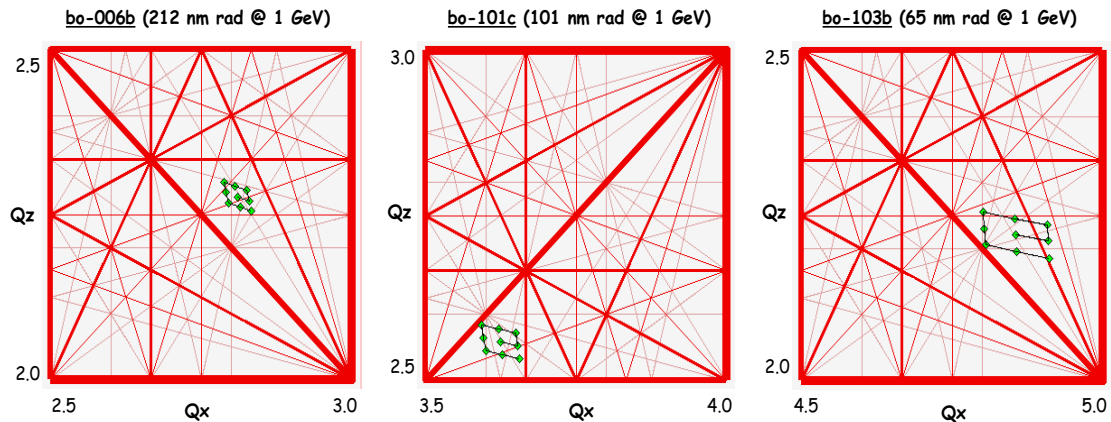
Figure 6.1: *The figures show the theoretical betatron tune diagrams of the "bo-006b", "bo-101c" and "bo-103b" Bodo optics. The red lines indicate transverse resonances, with decreasing line thickness for increasing resonance order. The central of the nine green markers in each diagram indicates the nominal tune $(Q_{x,0}, Q_{z,0})$. The parallelogram of the surrounding eight green markers indicates the betatron tunes for a $\pm 0.5\%$ variation of the focussing strengths for the horizontally (QF) and/or vertically (QD) focussing quadrupole families. While one-dimensional resonances up to third order usually cause fast beam loss, the low-order coupling resonances are only destructive if their lines have negative slope ("difference resonances", see section 2.4).*

- Replacement of faulty power supplies without proper calibration

- Power supply aging

- Temperature-, optics- and energy-dependence of current regulation systems

- Different magnetic hysteresis and saturation for different optics, injection and extraction energies and ramp lengths

- Coherent tune shift due to self-interaction of electrons at high beam currents

- Residual gas ions

In order to achieve constant energy-independent quadrupole focussing strengths, injector rings can use quadrupole power supplies with special current regulation systems that "track" the dipole power supply currents [87]. In contrast, Bodo uses simple stand-alone switched mode power supplies that are actually modified Delta power supplies. While this has the advantage of being much more cost-effective compared to complicated and expensive current tracking systems, the resulting differences in dipole and quadrupole current regulation systems and thus magnetic fields may lead to tune changes in the order of 0.05 with subsequent beam loss on a time scale of 50 ms (see chapter 8). Furthermore, nonlinearities in the current regulation circuits cause an energy- and time-dependence of the time constants of the current regulation loop. This leads to different monotonic or even oscillatory magnet currents and tune drifts for the same relative energy change in case of different optics or energies. Power supply aging and replacement of power supplies without proper calibration are also common sources of tune drifts.

Hysteresis and saturation of magnetic fields that are different for dipoles and quadrupoles also require appropriate corrections of the quadrupole currents up to several percent that depend on energy, optics, shape and length of the magnet ramps.

In case of modifications of the vacuum system, increased residual gas pressure and the resulting large amounts of residual gas ions can also cause a shift of the fractional betatron tune that depends on beam current, filling pattern, optics, pressure and residual gas composition [43]. Direct or indirect self-interaction of electrons at high beam currents and low beam pipe cross sections can also lead to betatron tune shifts [79].

Finally, time-dependent orbit perturbations during ramps cause time-dependent betatron tune shifts due to fields of sextupole magnets (see chapter 2). Desired sextupole fields are generated by the integrated sextupole magnets of Bodo that are used to compensate the natural chromaticity and increase the chromaticity-related momentum acceptance during injection. The integrated steerer magnets that are required for orbit correction do not generate an ideal dipole magnetic field, but they also have a non-vanishing sextupole field component [29]. During booster ramps, the orbit can change by many millimeters e.g. due to energy-dependent magnetic field errors or because of the orbit bump that is required to extract the beam. These orbit perturbations cause significant changes of the betatron tunes.

6.2.4 Previous Tune Measurement and Correction Scheme

Before the installation of the DSP tune feedback system, the betatron tunes of Bodo ramps were corrected manually by iterative editing of quadrupole magnet current offsets in a magnet ramp data file. The current ramps for QD and QF quadrupoles had to be corrected every 50 or 100 ms, with linear spline interpolation functions between correction values at different points of the magnet current ramp curves. For each correction point, the tune was measured with the swept-frequency method at the respective point of the ramp, with a time resolution between 20 and 30 ms (depending on the sweep speed) [4, 61, 63]. Then the currents were changed in such a way that the tune at that point should match the desired tune. After downloading the ramps to the synchronised digital function generators ("FGs") that controlled the power supplies, the tunes were measured again. In case of uncorrected ramps, these steps were repeated iteratively (beginning at the time of injection) until the tunes were sufficiently close to the desired tunes during the whole ramp cycle (that has a length of typically 6500 to 10000 ms). In order to measure the tunes at high energies, the amplitude of the tune excitation kicker had to be adjusted manually during the ramp from a low amplitude at low energies to avoid beam loss to a high amplitude at high energies in order to get sufficient beam current and tune signal quality. Since the system provided six different frequency intervals with fixed limits for the measurement of fractional tune frequencies, the intervals had to be changed (by a manual switch) after each correction in case vertical and horizontal tune frequencies required different frequency windows. Furthermore, the relatively slow measurement period of 20 to 30 ms was too large to track fast tune shifts that lead to beam loss within a single measurement period.

6.3 The DSP-Based Tune Measurement and Real-Time Feedback System

6.3.1 Tune Correction Scheme

In addition to normal synchrotron light source operation, both Bodo and Delta are also used as test machines for machine physics and tests of accelerator components, electronics and soft-

ware. As a consequence, Bodo must be operated at different injection and extraction ener-
gies, different beam optics and beam currents, and different lengths and shapes of the ramp
curves. The manual betatron tune correction procedure described in section 6.2.4 usually re-
quired many hours to correct a new ramp and could only be done by sufficiently experienced
accelerator physicists. Whilst DELTA was originally planned mainly as a test machine with
physicists as machine operators, it is now a synchrotron light source that is operational 24
hours per day, mainly with non-experts as operators. Therefore, one goal of this thesis was
the simplification and automation of the booster betatron tune correction for arbitrary energy
ramps, injection and extraction energies, beam currents, optics, orbits and sextupole settings.

Due to the very large number of parameters that have a (usually nonlinear) influence on the
betatron tunes, a model-based prediction of the betatron tune shifts that accounts for all possible
operating and environmental conditions is hardly practicable. Furthermore, purely model-based
correction has the disadvantage that the model requires recalibration of the model parameters
for each change of the underlying hardware or environment. Moreover, all parameters that
affect the model had to be measured and monitored continuously. Therefore, the tune correction
system that was developed in this thesis uses a beam-based approach to correct the betatron
tunes, by a feedback loop that measures the tunes and applies appropriate corrections to the
quadrupole magnet currents in real-time.

Instead of real-time correction, one could also have used an iterative cycle-to-cycle
("CTC") correction scheme where the betatron tunes are measured in one ramp cycle and
the resulting corrections are applied in the following ramp cycles. Since the beam of an un-
corrected ramp (with quadrupole currents that are proportional to the beam energy except for
saturation corrections) usually gets lost soon after injection (see chapter 8) and the calculation
of the corrections may take some time, the CTC method could require a larger number of ramp
cycles until the tunes match the desired values during the complete ramp cycle. Therefore, CTC
correction methods are more suitable for fast-cycling synchrotrons with ramp frequencies of
many Hz, whereas real-time tune feedbacks for such synchrotrons are less suitable due to the
high bandwidth that is required for measurement system and correction magnets and power
supplies.

In case of CTC correction at Bodo (with typical ramp periods between 6 and 10 seconds),
changes of machine parameters that cause tune drifts (e.g. changes of Bodo steerer magnet
settings or of the RF frequency) could only be corrected on a time scale of one minute, which
would slow down the optimisation of beam parameters e.g. for injection optimisation or ma-
chine physics experiments. In contrast, real-time correction of betatron tunes can compensate
tune drifts immediately and allows simultaneous optimisation of the betatron tunes and of other
beam parameters (e.g. the orbit) that normally would change the betatron tunes. Furthermore,
real-time correction can compensate tune drifts that are different from cycle to cycle, e.g. beam-
current dependent drifts due to residual gas ions or charge-dependent self-interaction of elec-
tron bunches. Moreover, real-time tune feedback allows nearly arbitrary non-periodic changes
of energy, beam optics and other beam parameters that change the tunes, which is an attractive
(and sometimes necessary) feature for machine physics experiments and tests of new machine
components and monitors.

Therefore, the DSP system was designed to perform tune correction in real-time, with the
option of additional cycle-to-cycle corrections in order to improve the remaining tune errors
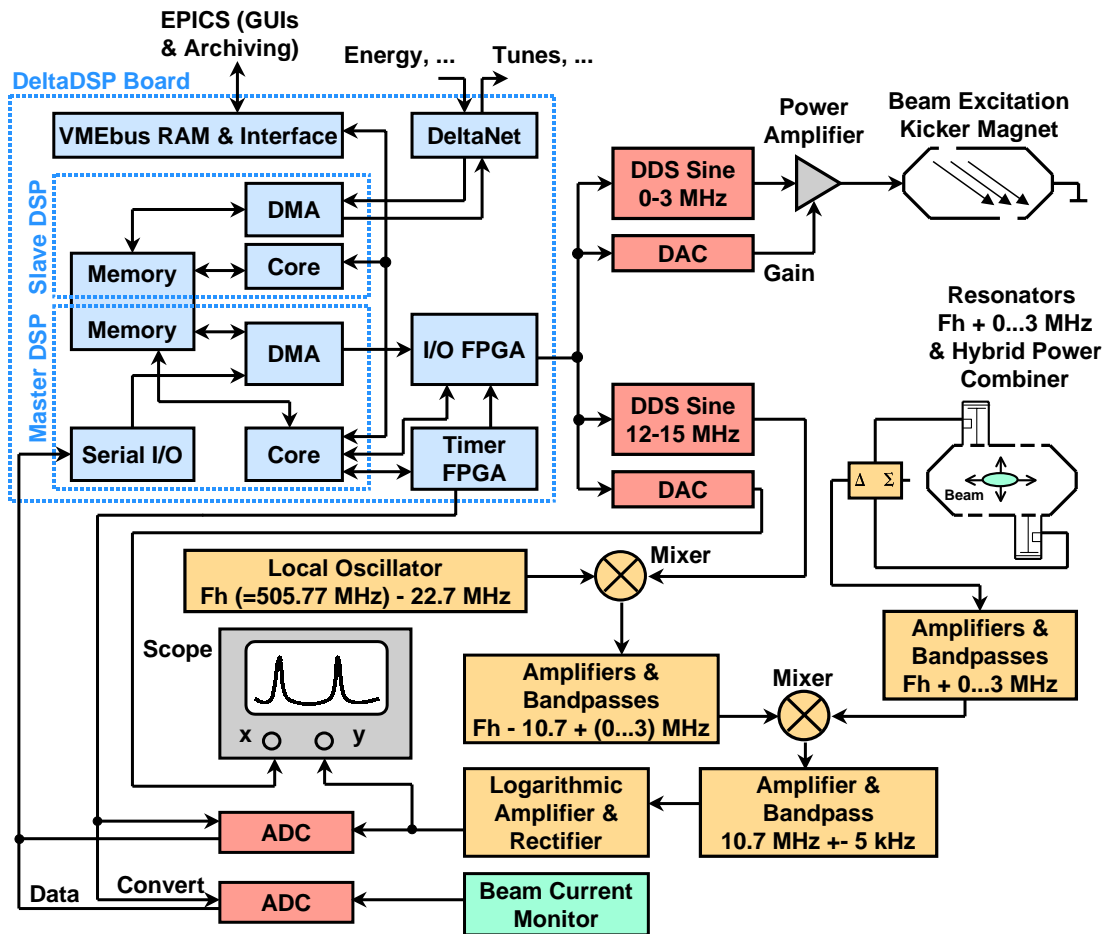for periodic energy ramps.

Figure 6.2: *Architecture of the DSP-based Bodo tune measurement system.*

6.3.2 Tune Measurement System Architecture

Fig. 6.2 shows the architecture of the DSP-based Bodo tune measurement system. The RF front-end (yellow boxes: resonators, RF filters, amplifiers, mixers, rectifier) was developed in diploma theses [4, 5], except for the commercially available local oscillator. The system employs a combination of swept frequency and chirp method and measures the fractional tunes by detecting the respective upper sideband of the 85th revolution harmonic frequency $F_h = 505.770$ MHz (assuming a cavity RF frequency of $f_{RF} = 499.820$ MHz, which is equal to the 84th revolution harmonic). The master DSP of a DeltaDSP board is connected to two external DAC boards and two direct digital synthesis ("DDS") sine function generator boards that were developed within the scope of this thesis. Each DDS board can generate sine wave functions between 0 and 25 MHz. Frequency, phase and amplitude of both DDS generators can be precisely programmed by the DSP at a rate up to 1 MHz. The DDS generator boards use an AD9850 IC of Analog Devices Inc. [91]. This IC contains a DAC that reads digital sine function values from a lookup table at a DAC clock frequency of 100 MHz. The point selection index for the lookup table is changed according to a phase register that is incremented each DAC clock cycle, depending on the frequency that can be programmed in steps of about 1.5 Hz. The discrete voltage steps of the DAC output cause undesired high frequency components

in the DDS generator spectrum that are removed by an appropriate lowpass filter on the DDS generator board.

Excitation of Coherent Beam Oscillations

The first DDS generator and one DAC determine frequency and current amplitude of a slotted pipe diagonal excitation kicker magnet. Since the previous power amplifier for the kicker was not able to generate sufficient current to measure the tunes of the new "bo-101c" low emittance optics at 1.5 GeV (which requires higher excitation frequencies than the previous optics), the amplifier was redesigned in this thesis. The kicker has a very small impedance of 110 nH [78], which is equivalent to the impedance of about 120 mm straight round wire (with 2 mm diameter) 10 mm above a metal mass plane. Therefore, the amplifier design requires extremely short cable paths in order to achieve a relatively flat frequency response and high currents at high frequencies. The impedance-optimised redesign allows up to 16 A peak-to-peak current at 3 MHz, whereas the previous design was limited to 2 A peak-to-peak. As will be shown in chapter 8, this allows reliable betatron tune measurements for fractional tunes between 0 and 3 MHz up to 1.5 GeV beam energy.

Detection of Coherent Beam Oscillations

The beam response (i.e. the amplitude of a coherent beam oscillation) is detected by two diagonal BPM pickup electrodes (see section 3.5, fig. 3.6) that are connected to two coaxial resonators which are directly mounted onto the beam pipe (see fig. 6.2). The resonators act both as (passive) 10 dB amplifiers and as filters for the desired sideband frequency range of $F_h + (0...3)$ MHz [5, 89]. The resonator outputs are connected to a hybrid power combiner that generates the sum and the difference signal of its input signals. The difference signal of the combiner is filtered and amplified before being mixed with the sum frequency $F_h - 10.7 + (0...3)$ MHz of the second DDS generator and a local oscillator. The resulting signal is filtered by a 10.7 MHz bandpass with ± 5 kHz bandwidth. Finally, its signal is amplified logarithmically, rectified and both sampled by an ADC ("tune ADC") and connected to the vertical (y) input of a scope that visualizes the signal in x-y-mode. A second ADC ("beam current ADC") reads the electron beam current that is measured by a commercially available beam current transformer of 4 kHz bandwidth [92]. The horizontal (x) input voltage of the scope is controlled by the master DSP (using a DAC), with the voltage being proportional to the excitation kicker frequency (1 V per MHz, with adjustable offset). Since all tune measurement data can also be displayed on GUIs on the Linux PCs, the scope is actually redundant, but convenient for some measurements (due to the faster display update rate compared to GUIs on Linux PCs).

Measurement of Fractional Betatron Tunes

The DSP measures the tunes by sweeping the kicker DDS generator within the desired frequency range while sweeping the second DDS generator at a frequency that is 12 MHz higher. The frequency difference can be adjusted to be different from 12 MHz to account for RF frequency changes without adjusting the master generator frequency, or to analyse higher harmonics of the kicker excitation signal. Horizontal and vertical tunes are distinguished by the tune shift that is caused by small changes of the Bodo quadrupole magnet focussing strengths.
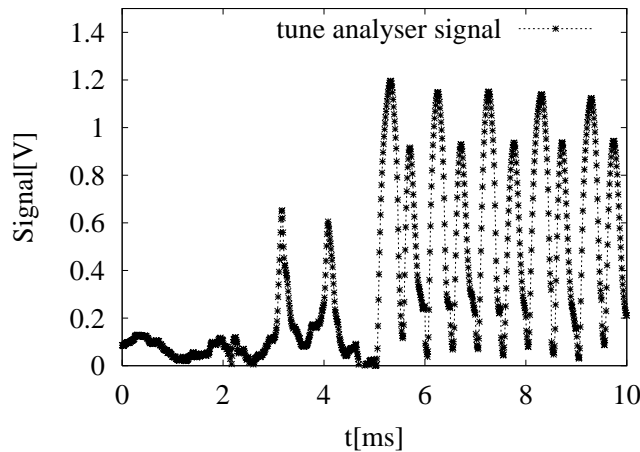
Figure 6.3: *Signal of the tune ADC. After beam injection, the master DSP sweeps the excitation kicker frequency through a user-defined range (e.g. from 1 to 3 MHz at 400 kHz/ms). The DSP measures the fractional betatron tune frequencies by detecting the two corresponding ADC signal peaks (in the left half of the figure, between t=0 and t=5 ms). Then the DSP sweeps periodically around the tune peaks in small frequency windows (e.g. 50 kHz) at a lower sweep rate (e.g. 100 kHz/ms), but with higher measurement precision. The resulting periodic series of x and z tune peaks is visible in the right half of the figure. As a lower sweep speed causes a longer excitation both of the beam (at its fractional betatron tune frequency) and of the resonance filters in the analogue front-end of the measurement electronics, the peaks on the right hand side are larger than those on the left hand side.*

Figure 6.3 shows the tune ADC signal during a tune measurement after beam injection at 60 MeV. The DSP starts the tune measurement when the beam current exceeds a user-defined trigger level (which is typically set to 0.5 mA). At first, the DSP sweeps once through a larger user-defined kicker frequency range (e.g. from 1 to 3 MHz) with a user-defined sweep rate. When the kicker frequency is equal to the fractional betatron tune frequency, the signal of the logarithmic amplifier/rectifier (also called logarithmic detector) in fig. 6.2 reaches a local maximum ("betatron tune peak"). The DSP reads this signal curve by the tune ADC and finds the horizontal and vertical betatron tune peak using a peak search algorithm. The fractional betatron tune is equal to the known excitation kicker frequency at the location of a tune peak. After the initial peak detection, the DSP sweeps continuously in small frequency windows around each of the two peaks (here: 50 kHz window width, 100 kHz/ms sweep rate) and adjusts the windows continuously so that the peaks stay in the middle of each window. In case the distance between two peaks is smaller than the window width, the DSP uses one tracking window for both peaks (with twice the width of the single tracking windows) instead of two windows, and uses a different peak search algorithm that looks for two tune peaks in the window instead of one.

Using this peak tracking method, the time required to measure both tunes was reduced to typically 1.35 ms (compared to 20 ms for the previous tune measurement system), which allows betatron tune correction in real-time.
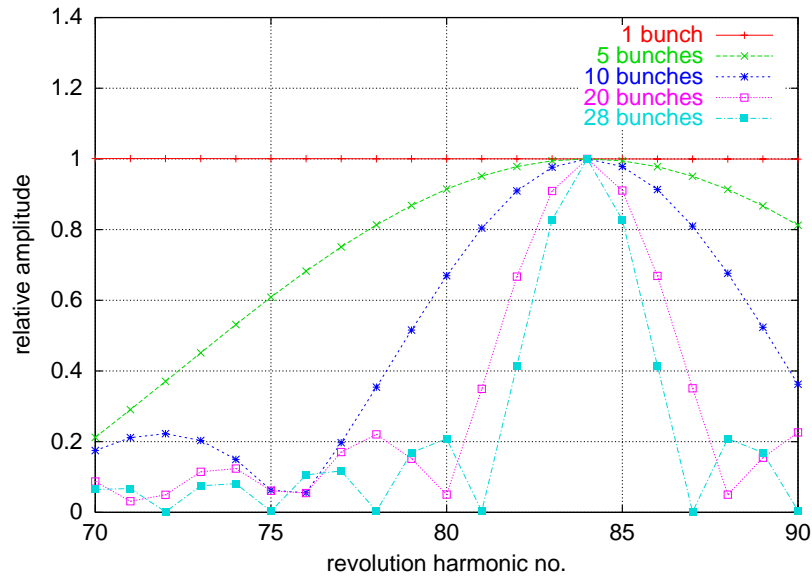
Figure 6.4: *Calculated Fourier transform of the longitudinal charge distribution in Bodo for bunch trains of 1 to 28 subsequent electron bunches (with 3 cm bunch length).*

Choice of the Revolution Harmonic

Figure 6.4 shows the calculated Fourier transform of the longitudinal charge distribution in Bodo for different longitudinal filling patterns. The horizontal axis shows the number of the revolution harmonic, i.e. the frequency of the Fourier component divided by the Bodo revolution frequency (of about 5.95 MHz). The 84th revolution harmonic is the RF cavity frequency (499.820 MHz), therefore it is non-zero for any filling pattern. The relative amplitudes of other revolution harmonics depend on the number of electron bunches. While the 85th revolution harmonic is always larger than 90 % of the 84th harmonic for filling patterns up to 20 subsequent bunches, the calculated amplitude of the 87th harmonic drops to 35 % for 20 bunches and to zero for 28 bunches (assuming the same charge in each bunch).

Before this thesis, the RF front-end of the Bodo tune measurement system had been modified in order to detect the right betatron sidebands of the 87th revolution harmonic instead of the original 85th harmonic, in order to avoid undesired peaks in the spectrum that were caused by revolution harmonic RF signals of the Delta electron beam [61]. Such signals can couple into the metal beam pipe via ceramic gaps and isolations. However, the length of the electron gun pulse for the Linac was increased afterwards from about 12 ns to an adjustable (arbitrary) length. Depending on the adjustment of timing, gun pulse length and RF system phases of the Linac, between one and more than 20 RF buckets in Bodo may now be filled with electrons, compared to about six for the previous gun pulse. In order to make the the dynamic range of the tune measurement system more independent of the filling pattern, the filters and resonators of the RF front-end were re-adjusted during the work on this thesis for a central frequency 1.5 MHz above the 85th harmonic (instead of the 87th). Therefore, the DSP system now detects the right sidebands of the 85th harmonic instead of the 87th. Using the 84th harmonic would provide maximum insensitivity to the filling pattern, but is not suitable due to strong distortions of the tune measurement system by RF leaks in waveguides and cabling of the RF cavity.

The undesired Delta revolution harmonics that were visible in the signals of the Bodo tune RF front-end could be reduced to a negligible level by optimising the suppression of revolution harmonics by the hybrid power combiner in fig. 6.2. The difference output signal of the power combiner should be zero if the beam is in the center of the beam pipe. In order to verify this, a centered beam was simulated by feeding RF signals of equal amplitude and phase into the two unused BPM pickups shown in fig. 6.2. By fine-tuning of resonance frequencies and by using matched attenuators at the combiner inputs, it was possible to achieve a difference signal that was more than 20 dBm smaller than the sum signal for frequencies 0 to 3 MHz above the 85th revolution harmonic. After this adjustment, the undesired Delta revolution harmonics were below the noise level of the Bodo tune signals, which indicates that they had coupled symmetrically into the coaxial resonators, but had not been suppressed properly by the hybrid power combiner due to phase and gain differences of the resonators.

As will be shown below, the DSP system is able to eliminate systematic peaks in the tune frequency spectrum by software. Therefore the DSP-based feedback would also have worked without the elimination of the Delta revolution harmonics as described above, but at the expense of a decreased tune measurement rate.

6.3.3 Digital Signal Processing and Frequency Quantisation

In order to measure the betatron tune frequencies, the DDS generators frequencies must be incremented or decremented in sufficiently small steps. If the kicker frequency is changed in steps that are larger than the fractional tune frequency spread of the beam, this will lead to small or no excitation of coherent oscillations and therefore small or zero tune peak signals. In case of Bodo, frequency steps of 0.5 kHz have proved to be sufficient for a reliable peak detection under all operating conditions. Due to the ± 5 kHz bandwidth of the 10.7 MHz bandpass in fig. 6.2, it is sufficient to sample the tune ADC each time the DDS generator frequency has changed by 1 kHz. However, due to a maximum available ADC sampling rate of 100 ksamples/s, sweep rates above 100 kHz/ms lead to ADC samples just every 2 or 4 kHz (see table 6.1), which limits the frequency resolution of the tune measurement system at high sweep rates.

At 400 kHz/ms sweep rate, two DDS generators and two DACs of the tune measurement system have to be reprogrammed every 1.25 μs, with a jitter that should be much smaller than this value. Since direct write access by the DSP core to the I/O piggyback (that controls the

Sweep Rate [kHz/ms]	DDS Update Period [μs]	DDS Freq. Step [kHz]	ADC Sample Period [μs]	ADC Freq. Step [kHz]
200-400	2.5-1.25	0.5	10	4
100-200	5-2.5	0.5	10	2
0.1-100	5000-5	0.5	10000-10	1

Table 6.1: *Column 2 shows how often the DDS generator frequencies of the tune measurement system are reprogrammed (with 0.5 kHz increments or decrements) to obtain different sweep rates. At sweep rates equal to or lower than 100 kHz/ms, the tune ADC is sampled after two DDS generator updates, i.e. every kHz. Due to a maximum ADC sampling rate of 100 ksamples/s, faster sweep rates lead to ADC samples every 2 or even 4 kHz (columns 4 and 5).*

DACs and DDS generators) at this rate would consume nearly all DSP core processing time, the write accesses are executed automatically by the DMA (direct memory access) unit of the master DSP, while the processor core can continue e.g. to perform calculations or to find tune peaks in the ADC data.

The DMA unit is able to transfer a user-defined number of words from internal DSP memory to the link ports, with user-defined start address and address increment for this transfer. Several identical sub-units of the DMA unit allow independent transfer of several different curves (e.g. one for each DDS generator) via different link ports. The desired set value curves for DDS generators and DACs are stored in the internal memory of the master DSP, and the DMA unit transfers them via four master DSP link ports to the I/O piggyback module FPGA, which writes the values onto the external 16-bit DAC/DDS generator bus. The shapes of the frequency and voltage curves for DDS generators and DACs are programmable and thus arbitrary, but the DDS frequency curve is usually a linear curve from 0 to 3 MHz (with 0.5 kHz increments), and the DAC values for the horizontal input of the scope are proportional to this frequency. Normally, the 2nd DDS generator that determines the detection frequency for coherent beam oscillations requires frequency set values 12 MHz higher than the 1st DDS generator that excites the kicker. Therefore, the linear 0-3 MHz curve is written to both DDS generators, and the FPGA of the I/O piggyback adds a user-programmable offset (usually 12 MHz) to the frequency of the 2nd DDS generator. The start and stop frequency for the peak tracking windows are determined by the memory address where the DMA unit starts to transfer the curve data to the DDS generators and DACs. The direction of a frequency sweep (up or down) and its frequency step width can be chosen by sign and size of the address increment of the DMA unit. The DMA unit works independently of the processor core, which can process measurement data while the DMA unit writes the frequency curves to the DDS generators.

The clock rate at which new values are written via DMA to the DDS generators (and thus the sweep rate) is determined by a programmable timer of the "trigger-matrix/counter/timer"-FPGA (see section 5.3.1). The timer generates a programmable number of update triggers for the I/O module FPGA at a programmable clock rate. A second timer is used to generate ADC conversion triggers at a programmable fraction of the DDS generator update rate. Each time the tune ADC and the beam current ADC get a conversion trigger, they start a new conversion and transmit the ADC data of the previous conversion directly to the internal memory of the master DSP (via DMA and a 10 Mbit synchronous serial bus system that was developed in this thesis, see section 5.3). The synchronicity of ADC and DDS clock rates guarantees that the DSP knows the excitation kicker frequency that belongs to each ADC sample, so that the DSP can determine the fractional betatron tune frequencies for the tune peaks in the ADC data curve.

6.3.4   Static Frequency Resolution

Due to finite bandpass response times in the tune measurement RF front-end, the measured fractional tune frequency of a given constant betatron tune has a systematic error that increases with growing sweep rate. When sweeping from lower to higher frequencies ("up sweep"), the tune peaks as measured by the tune ADC move to higher frequencies ("peak shift"), whereas "down sweeps" result in a peak shift to lower frequencies. This error becomes larger with growing sweep rate.

In order to measure this systematic error as a function of the sweep rate, the electron beam with its betatron sidebands of the 85th revolution harmonic were simulated by the signal of two precision RF generators. The output signals of the generators were added with a power combiner and then amplified with a power amplifier that was connected to one of the two unused BPM pickups of the tune measurement BPM beam pipe. The other two BPM pickups were connected to the coaxial resonators of the tune measurement system as shown in fig. 6.2. The power amplifier had a circulator that avoided undesired standing waves on the cable due to the missing 50 Ω termination at the location of the BPM pickup button.

Figures 6.5 shows the measured peak shifts for up sweeps (four upper curves) and down sweeps (four lower curves) for four simulated fractional betatron frequencies between 500 and 501.5 kHz (with 0.5 kHz increments). The horizontal axis shows the sweep rate, the vertical axis shows the tune peak frequency that was measured by the DSP. The peak frequency values were averaged (over 100 peak detections) in order to reduce the statistical noise-related measurement error. The upper figure shows the sweep speed range up to 100 kHz/ms, the lower figure shows the maximum available sweep speed range up to 400 kHz/ms.

Since the peak frequency shift is a systematic measurement error, the DSP software corrects the tune frequency and the peak tracking windows according to the (smoothed) measured peak frequency shift. The correction can be switched off for testing purposes by a button on a graphical user interface, which is useful e.g. in order to measure the peak frequency shift. The DSP-based peak frequency correction improves the absolute measurement error of the system to typically ± 2 kHz for sweep rates below 100 kHz/ms, while while relative tune frequency changes in the order of 100 kHz or less can be measured with ± 1 kHz resolution.

The peak frequency correction is also important when bidirectional sweeps are used (i.e. alternating sweeps from lower to upper and from upper to lower limit of the sweep frequency window), because the frequency distance between the uncorrected peaks may be larger than half of the sweep window width. In this case, the peaks can not be tracked any more without peak frequency correction, except with a larger window (and thus lower measurement speed).

The peak frequency shift is almost independent of the fractional betatron tune frequency. This was verified by peak shift measurements for various fractional betatron tune frequencies between 100 and 2900 kHz, with RF generators in order to simulate the electron beam as described above. Fig. 6.6 shows the difference of the tune shift for frequencies of 500 and 2500 kHz above the 85th revolution harmonic frequency, averaged over 100 measurements. The differences are within the resolution error of the measurement system. The growing (more or less discrete) differences at high sweep rates in fig. 6.6 are caused by the lower number of ADC samples per DDS generator frequency change and the resulting larger frequency measurement error (see table 6.1). Sweep rates of 400 kHz/ms are only used for the first sweep after injection, in order to find the betatron tune peaks as fast as possible (since the sweep range may be very large, up to 3 MHz), at the expense of a larger measurement error (up to 4 kHz). All further sweeps (i.e. the actual measurement of the tunes during the Bodo ramp) are usually made at sweep rates of 100 kHz/ms or less. Since the sweep windows are quite small, the lower sweep rate still allows both fast measurement and small betatron frequency measurement errors, with negligible tune frequency dependence of the peak shift. As can be seen in fig. 6.6, the DSP system is able to resolve peak frequency changes of 0.5 kHz at sweep rates below 100 kHz/ms, although the ADC just makes a sample each time the excitation kicker frequency has changed by 1 kHz. This higher resolution is achieved by averaging over several measurements. Without
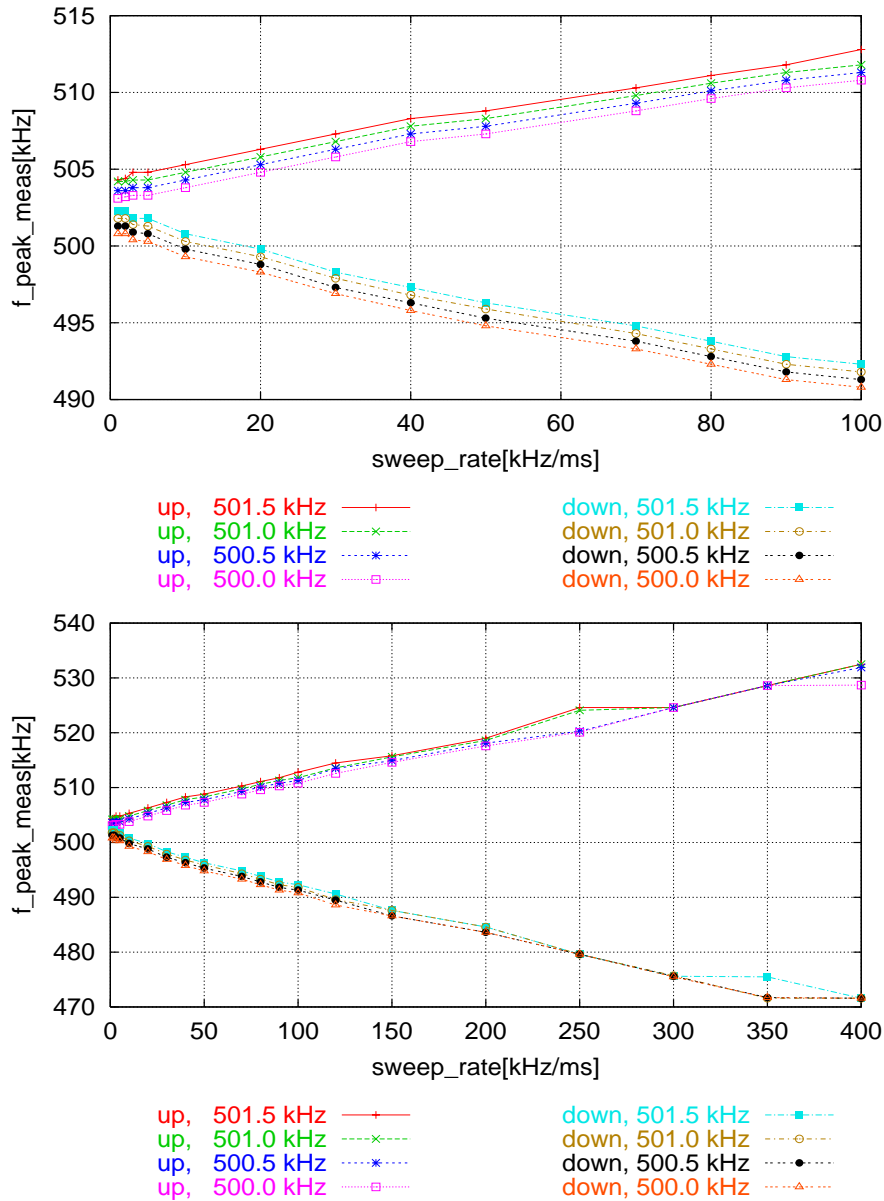
Figure 6.5: *Measured fractional tune frequency as a function of the sweep rate, for up sweeps (curves with positive slope) and down sweeps (negative slope). The plotted frequencies are averaged measurements for fractional tune frequencies between 500 and 501.5 kHz (simulated with a function generator). The plots show a systematic tune measurement error ("tune peak frequency shift") that is caused by finite filter response times of the analogue measurement front-end. The more or less discrete differences at large sweep rates in the lower figure are caused by the frequency quantisation of the DDS generators.*

averaging, the resolution at high sweep rates is still limited by the ADC sampling frequency and the corresponding kicker frequency step width.

In order to investigate the dependence of the peak shift on the amplitude of the betatron sidebands and thus on the beam current and the excitation kicker amplitude, the power of the
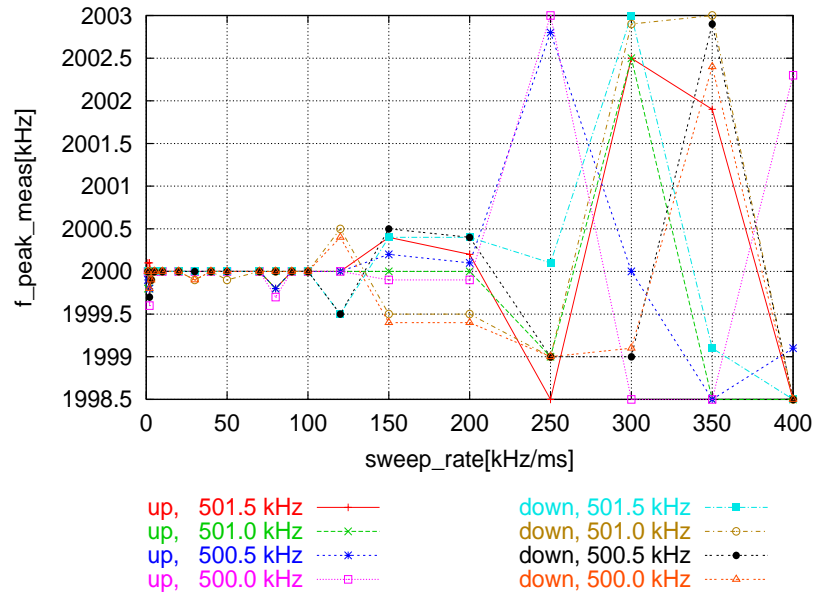
Figure 6.6: *Tune peak frequency shift for a fractional tune frequency of 2500 kHz minus tune peak frequency shift for 500 kHz, plotted against the sweep rate. The peak shift is nearly independent of the tune frequency (within the resolution limits of the system). The more or less discrete differences are caused by the frequency quantisation of the DDS generators and the precision of the frequency readout variable.*

RF generators that simulated the sidebands was varied between -55 and 0 dBm, for simulated fractional betatron tune frequencies of 500 kHz and 2500 kHz. This resulted in tune ADC voltages (of the logarithmic detector, see fig. 6.2) between 0.2 and about 2.6 V. The frequency values were measured via DSP tune peak detection without peak shift correction.

The result of this measurement is shown in figures 6.7 (left side: 500 kHz betatron tune frequency, right side: 2500 kHz, upper plots: frequencies measured by the DSP, lower plots: amplitudes of the logarithmic detector, measured by the tune ADC). The tune peak frequency shift is nearly independent on the height of the betatron sideband for rectifier voltages between 0.5 and 2.2 V (corresponding to RF generator levels between -50 and -10 dBm). The figures show the tune peak amplitudes and fractional frequencies, with several hundred measurements (i.e. tune peak detections) for each RF generator level. Since the measurement was started and stopped manually, the number of measurements is slightly different for different generator output levels. For better visualisation, the measurement numbers were shifted on the horizontal axis, depending on the generator output level.

In order to obtain small errors for the measurement of the real betatron tunes of Bodo, the amplifier gains of the RF front-end of the tune measurement systems and the excitation kicker level are chosen so that the tune peaks both for very small and very large Bodo beam currents are within the range of 0.8 and 2.2 V (without causing significant beam loss or saturating the amplifiers and mixers of the RF front-end by the revolution harmonic signals). While very high sideband powers (with ADC voltages above 2.2 V) lead to a slight increase of the peak frequency error (by about 1 kHz), very low sideband powers (with ADC voltages below 0.8 V) lead to more noise on the tune peak signals, which results in a larger statistical measurement error for the peak frequencies (see upper figures 6.7, lowest RF power level of -55 dBm).
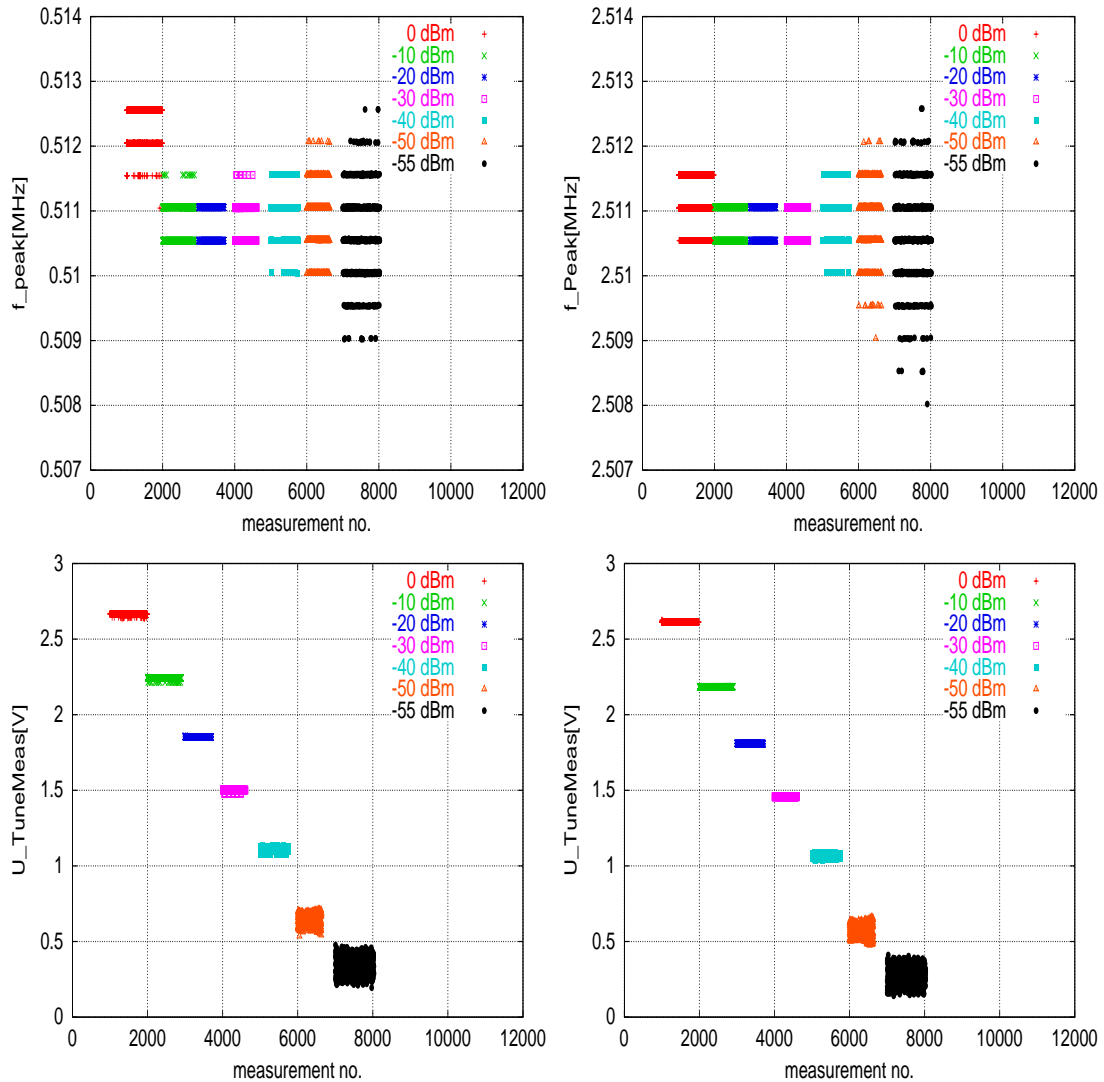
Figure 6.7: *Change of the measured fractional tune frequency (without peak shift correction) with betatron tune sideband power, as measured by the master DSP. The right sidebands of the 85th revolution harmonic were simulated by coupling RF power of the required frequency into one of the two unused tune BPM pickups, therefore the power values have an arbitrary offset with respect to the real electron beam sideband power. The upper figures show that the tune peak frequency shift is basically independent of the sideband power (left plot: 500 kHz simulated fractional betatron tune frequency, right plot: 2500 kHz, up sweeps for both plots). The lower figures show the voltage amplitudes of the logarithmic detector (measured by the tune ADC) of the tune peaks for the respective upper plots. For usual Bodo beam currents, kicker amplitudes and beam energies result in ADC values between 0.8 and 2.2 V.*

Figure 6.8 shows the tune peak voltages of the logarithmic detector (see fig. 6.2) as sampled by the tune ADC for different simulated betatron tune frequencies and output levels of the "simulation" RF generators. The blue curves were recorded for up sweeps (i.e. increasing frequency in the peak tracking window) at generator levels between -55 dBm (lowest curve) and
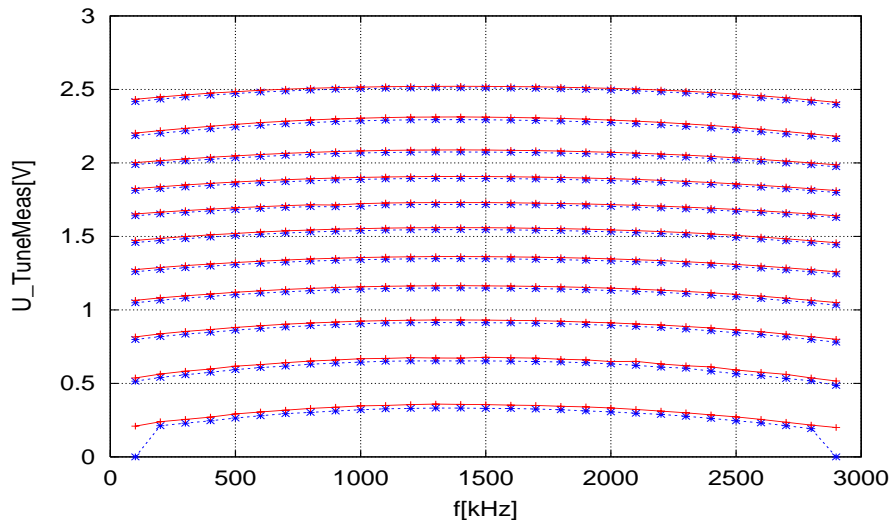
Figure 6.8: *Measurement of the amplitudes of the logarithmic detector by the tune ADC, as a function of the power and frequency of the betatron tune sidebands of the 85th revolution harmonic. The right sidebands of the 85th revolution harmonic were simulated by coupling RF power of the required frequency into one of the two unused tune BPM pickups, therefore the power values have an arbitrary offset with respect to the actual electron beam sideband power. The simulated fractional tune frequencies were incremented in 100 kHz steps (blue: up sweeps, red: down sweeps), the RF generator was incremented from -55 dBm (bottom curve) to -5 dBm (top curve) in 5 dBm steps.*

-5 dBm (highest curve), with steps of 5 dBm. The red curves were recorded for down sweeps. Only the peak frequency of one RF generator was measured while varying it in 100 kHz steps between 100 and 2900 kHz (and the output power in 5 dBm steps). The other RF generator was set to the same output power level. It had a frequency of either 2500 MHz or 500 MHz, whatever value had the largest distance to the frequency of the first RF generator.

The slightly lower voltage for the up sweeps (that is caused by a frequency asymmetry in the RF front-end filter system [4]) is irrelevant, because it does not affect the frequency measurement. Since the filters and resonators in the BPM front-end of the tune measurement system were adjusted for a central frequency that is 1500 kHz above the 85th revolution harmonic, the curves in fig. 6.8 have a slight local maximum at 1500 kHz, with a nearly symmetrical decay to both sides of the maximum. The tune measurement system has an adjustable minimum detection level for the tune peaks in order to suppress noise peaks, therefore the amplitude values at the beginning and end of the lowest curve in fig. 6.8 were not measured (since they were below the noise threshold), but set to zero. The amplitude response of the tune measurement system for a given frequency is not perfectly linear. This is not relevant, since the exact amplitude of a tune peak is not important, only the tune frequency. The measured ADC voltage increases typically by 44 mV per dBm sideband power, with an available dynamic range of more than 50 dB. If necessary, the range can be increased further (e.g. in order to measure tunes for extremely small or large Bodo beam currents) by different attenuators and different amplifier gains in the tune measurement RF front-end.

6.3.5   Dynamic Frequency Resolution

The previous section described the correction of systematic tune peak measurement errors that allowed DSP-based tune measurement with a relative and absolute precision of ± 1 kHz and ± 2 kHz for *constant* fractional betatron tune frequencies and typical Bodo beam currents. However, the betatron tunes in Bodo can drift at rates up to several kHz/ms (see chapter 8). In order to investigate the measurement error of the DSP system for drifting tunes, such tunes were simulated with a similar setup as described in the previous section (two RF frequency generators with combined outputs and a subsequent power amplifier, connected to an unused BPM pickup of the tune measurement BPM block).

However, in contrast to the previous section the frequencies of both RF generators were modulated with the output voltages of two low frequency generators ("LF generators") that generated sine voltage signals with frequencies in the order of 10 to 50 Hz. The modulation voltage of the first RF generator was recorded by the tune measurement system by connecting it to the beam current ADC. The modulation frequencies were calibrated so that the generator frequency could be calculated from the modulation voltage with a precision better than 0.2 kHz. The frequency of the first RF generator was set to the 85th revolution harmonic plus 1500 kHz, the frequency of the second generator to the 85th revolution harmonic plus 2500 kHz. The frequency modulation generated a sine-shaped frequency offset between 0 and about 100 kHz, so that modulator frequency 1 changed periodically between 1500 and approximately 1600 kHz. Figures 6.9 show the actual (set) values the simulated fractional betatron tune frequencies (red curves) and the values that were measured by the DSP (dark blue curves). In order to visualise the measured frequencies together with the measurement error (magenta curve), 1500 kHz were subtracted both from set and measured fractional betatron tune frequency. The measurement was performed with activated peak shift frequency correction by the DSP.

The upper plot in figures 6.9 shows the measured values and the measurement error for the fractional betatron tune frequencies for 10 Hz modulation frequency, 70 kHz width of the peak tracking window, and 100 kHz/ms sweep rate (using only up sweeps). According to the 10 times magnified measurement error (light blue curve), the difference between measured and actual fractional betatron tune frequency is always smaller than ± 1 kHz. The second (middle) plot was made with the same settings, but for a modulation frequency of 50 Hz. While the measurement error for increasing or constant fractional tune frequencies is still good, decreasing frequencies lead to measurement errors up to 10 kHz. According to the lowest plot in figures 6.9, this error can be reduced to about 3 kHz by increasing the width of the peak tracking window from 70 kHz to 100 kHz. The reason for this behaviour is that the tune peak for 70 kHz window width is too close to the lower edge of the window when the tune frequency decreases rapidly, so that the "left" part of the peak is cut off by the lower window edge. Therefore, the filters resp. resonators of the tune measurement system do not see the complete beam excitation, but part of the beam excitation at lower frequencies is missing, which causes an additional tune peak shift in the direction of the sweep.

In principle, this systematic measurement error that depends on the width of the peak tracking frequency window, the sweep rate and the tune change per time could also be measured and corrected by the DSP. However, change rates of fractional betatron tune frequencies in Bodo are usually smaller than 4 kHz/ms without tune feedback, while the second (middle) plot in figures 6.9 has a change rate of nearly 16 kHz/ms. The tune feedback reduces the tune frequency change rates even further, therefore the measurement precision for changing tune frequencies is
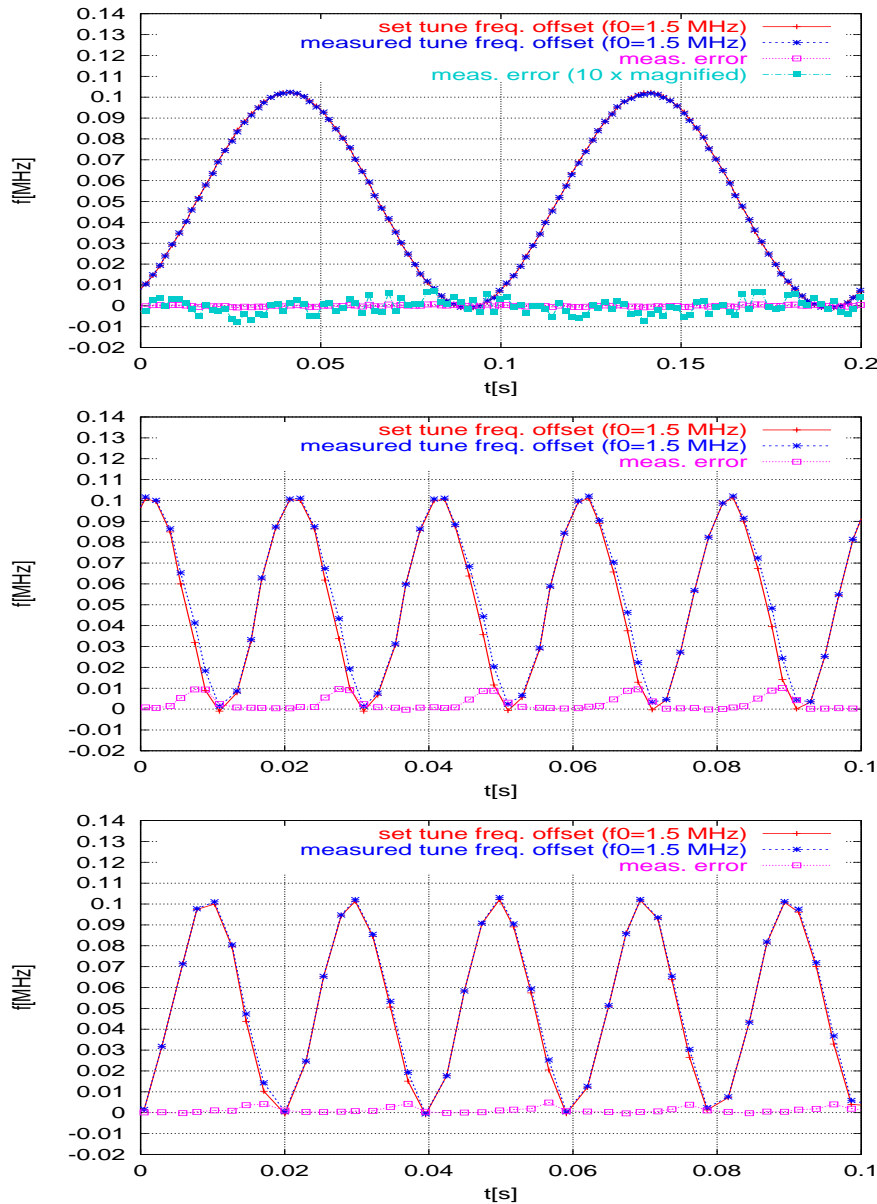
Figure 6.9: *Test of the tune measurement system with RF generators that simulated a fractional betatron tune frequency which changes periodically between 1500 and approximately 1600 kHz. The plots show the set (red curves) and measured frequency values (dark blue curves) minus 1500 kHz, as well as the difference between measured and set values (magenta curves). The measurements were made by the master DSP with activated tune peak frequency correction. Sweep rates of 100 kHz/ms and peak tracking windows of 70 kHz cause negligible measurement errors when the fractional betatron tune frequencies change by 100 kHz at a period of 10 Hz (upper plot). Tune frequency modulations with a period of 50 Hz instead of 10 Hz lead to systematic measurement errors up to 10 kHz (middle plot). The errors can be reduced to 3 kHz by increasing the width of the peak tracking window from 70 to 100 kHz (lower plot), at the expense of lower measurement speed.*

usually as good as for constant tunes, and the additional measurement error is not corrected by the DSP. For tune frequency change rates below 5 kHz/ms and sweep rates up to 100 kHz/ms, peak tracking windows of at least 70 kHz width are sufficient for negligible dynamic measurement errors. The measurement error for faster tune frequency change rates can be improved simply by a larger tracking window.

6.3.6   The Bodo Betatron Tune Feedback System

6.3.6.1   Focussing Magnet Power Supplies

In order to correct fast changes of the betatron tunes in real-time, the quadrupole magnet power supplies must have a sufficiently fast response time to changes of the set current. For budget reasons, Bodo uses the same quadrupole power supplies as Delta, with some modifications of the current regulation time constants and the output driver that allow faster ramping (see section 3.2). Before the installation of the tune feedback systems, betatron tunes had to be corrected manually by iterative modification of quadrupole current correction spline curves in a data file (usually in steps of 50 to 100 ms). In case of rapid changes of the betatron tunes (for uncorrected ramps without correction splines) on the time scale of some 10 ms, it was not possible to correct the tunes with sufficient precision, and the set values for the correction spline often showed an oscillatory pattern and were much larger than the expected current correction. This behaviour showed that the speed of the current regulation system was still too low, which would result either in a low bandwidth (and insufficient tune correction) or instability of a tune feedback system. Therefore, the time constants of the current regulation systems were modified in this thesis and the regulation speed was increased up to the physical limit of the output power stage of the power supply that was actually not designed for fast ramping synchrotrons. The curves in figure 6.10 show the response of the quadrupole power supply current (measured by the ADC of a DeltaDSP board with 5 ksamples/s) to a rapid change of the set current (magenta curves). The modification improved the regulation speed by nearly a factor of two. The resulting slightly larger (but also shorter) overshoots and undershoots cannot be avoided except with a complete redesign of the current regulation unit and the output power stage of the power supply. However, the tune feedback system does not increase the set current instantaneously in a large step, but continuously (with adjustable speed), therefore the overshoots are not critical for the operation of the feedback. The measurements in chapter 8 show that the modified power supply allows sufficiently fast and precise DSP-based real-time corrections of the betatron tunes under all operating conditions of the booster.

Dipole Magnet Power Supply

The Bodo dipole magnet power supply is not regulated by the tune feedback system. It is controlled by a DeltaDSP board that allows set current control either directly by an operator or by predefined periodic current ramps that are synchronised with the current ramps of all other Bodo magnet power supplies. As mentioned in section 3.2, the dipole power supply has control inputs for the set current and its first and second derivative. Before this thesis, only the first input was used, which resulted in time-dependent differences between actual and set current of up to -7% (at the time of beam injection). After installation of the DeltaDSP system, the dipole DeltaDSP board also provides first and second derivative of the set current (via two additional
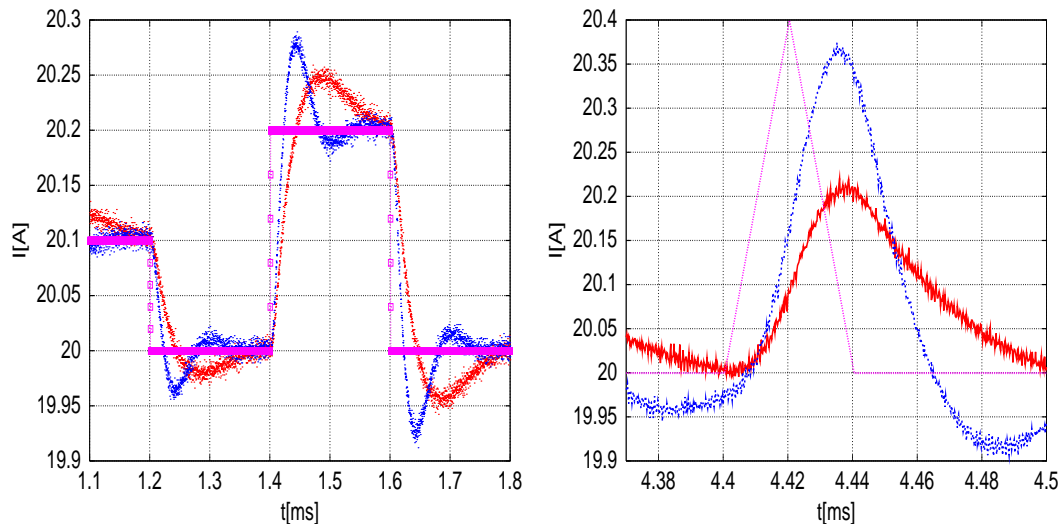
Figure 6.10: *Response of the Bodo quadrupole magnet power supplies to a rectangular (left) and triangular (right) change of the current set value. In order to improve the betatron tune feedback bandwidth, the time constants of the current regulation units were modified. The magenta curves show the set value, the red and blue curves are the quadrupoles currents before (red) and after (blue) the modification. The response time for changes of the current set value was decreased by nearly a factor of 2, which is desired for an efficient and fast betatron tune feedback. The resulting increased (but shorter) undershoots and overshoots are caused by the design of the output stage of the power supply and the design of its analogue current regulation loop.*

DACs), which reduced the maximum difference between set and actual current to typically 0.1% or less (for typical energy ramps).

The dipole DeltaDSP board calculates the beam energy both from set and actual dipole magnet current (using a fitted spline function that is based on magnetic field measurements [93]) and sends the values to all other DeltaDSP boards via DeltaNet. The beam energy is used by several DeltaDSP boards, e.g. by the tune measurement DSP board in order to increase the excitation kicker current amplitude with increasing energy, or by the betatron tune feedback system in order to predict the required change of quadrupole magnet currents and feedback variables with the beam energy (see below).

6.3.6.2 Feedback Loop

Figure 6.11 shows the information flow in the Bodo betatron tune feedback loop. The "bo-tune" DeltaDSP board measures the betatron tunes and sends them to the "bo-q1" DeltaDSP board via DeltaNet. The "bo-q1" DeltaDSP board continuously corrects the quadrupole focussing strengths in order to minimise the difference between desired and measured tunes for arbitrary energy ramps and beam orbits. In addition to the betatron tunes, the feedback system also uses the beam current in order to start and stop the loop if the current is above or below a user-defined beam current trigger level (with some trigger level hysteresis in order to avoid undesired switching due to noise while the beam current is close to the trigger level). Furthermore, the
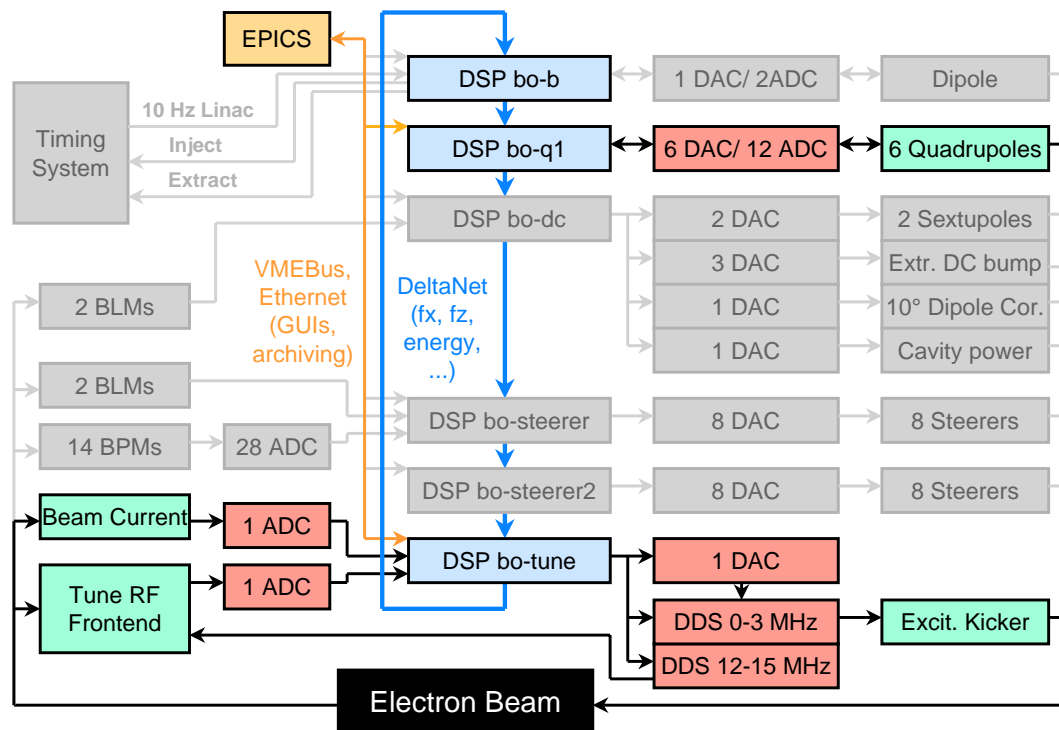
Figure 6.11: *Information flow in the betatron tune feedback loop of Bodo. Light grey boxes and arrows indicate components that are not involved in the loop.*

loop uses the beam energy in order to predict quadrupole current changes and to adapt internal variables of the feedback algorithm. The tune measurement and feedback system is controlled and monitored by the EPICS control system via graphical user interfaces in the DELTA control room.

Depending on the selected sweep speed and the width of the tune peak tracking frequency windows, the betatron tunes are measured typically every 1.4 ms. The DeltaNet data transfer from "bo-tune" DSP board to "bo-q1" DSP board causes a delay of less than 50 μs. The quadrupole DACs are usually updated at an (adjustable) rate of 4 kHz, and the feedback algorithm is executed before each update, causing an additional delay of less than 2 update periods, i.e. of less than 0.5 ms. Therefore, the overall delay between tune measurement and quadrupole DAC update is less than 2 ms. In contrast, the response time of the quadrupole power supply is about 15 ms. Eddy currents in the laminated quadrupole magnets and the 3 mm stainless steel beam pipe at frequencies of 100 Hz or less are not relevant for the tune feedback. The phase delay of the magnetic field in the beam pipe due to beam pipe eddy currents is only 6 degrees at 100 Hz, at an amplitude decrease below 1 percent [94]. Since the tune feedback uses a PID (proportional integral derivative) algorithm [96] with sampling ADCs and DACs, a sufficiently stable feedback loop only improves periodic beam perturbations significantly if they have a frequency below 1/10th of the ADC/DAC sampling frequency [95]. Therefore, the closed-loop bandwidth of the tune feedback is limited mainly by the bandwidth of the power supply current regulation and by the tune measurement frequency. Measurements of the response time of the overall system are presented in chapter 8.

The feedback loop uses two independent PID controllers to correct the tunes. One con-

troller changes the currents of the three horizontally focussing quadrupole families ("QFs"), the other controller changes the currents of the three horizontally defocussing quadrupole families ("QDs"). In principle, the PID controllers could predict the required quadrupole current changes from the theoretical optics model, e.g. using a linear functional approximation of the relative focussing strength changes of QFs and QDs as a function of the fractional betatron tune frequency changes in the vicinity of the nominal tunes. The required relative focussing strength changes could then be calculated by multiplication of a 2x2 matrix with the fractional tune error vector $(\Delta f_x, \Delta f_z)$. Due to the FODO lattice and the properties of the usual Bodo optics, the QF currents have a strong influence on the horizontal fractional tune frequency f_x and a weak influence on the vertical fractional tune frequency f_z, and vice versa for the QD currents (see fig. 6.1). Because of this low transverse coupling, the off-diagonal elements of the above 2x2 matrix were neglected, and the feedback loop uses only the diagonal elements. As will be shown in chapter 8, the influence of the missing off-diagonal elements on the performance on the feedback loop for the usual Bodo beam optics is negligible.

6.3.7 GUI-Based Measurement and Feedback Control

Figures 6.12, 6.13 and 6.14 show graphical user interfaces that are used to configure and adjust the tune measurement and feedback system. The user interfaces were designed with the GUI builder "EMW" [66] (see section 5.4.1) that was developed during the work on this thesis.
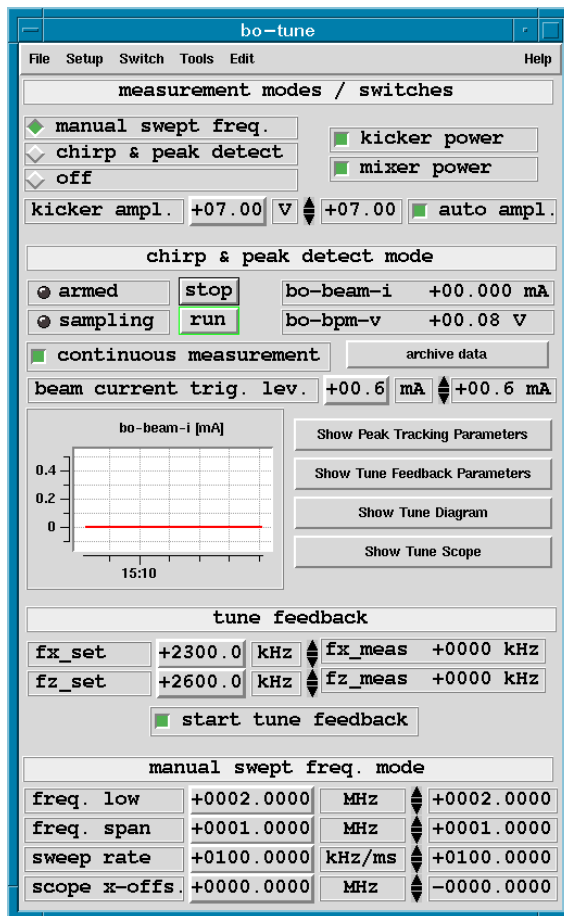

Figure 6.12: *One of the graphical user interfaces ("GUIs") of the Bodo tune measurement and feedback system. The upper part of the window is used to select the measurement mode. The user can choose between swept frequency mode (sweep through one large user-defined frequency window with manual selection of the frequency limits) and chirp mode (with tune peak tracking by the DSP). The excitation kicker amplitude can be adjusted manually or automatically by the DSP. In chirp mode, the tunes can be measured once or continuously, and the measurement data can be archived in data files. The measurement is started when the beam current exceeds a user-defined trigger level. Some buttons are provided to start additional GUIs, e.g. for the configuration of the tune peak search algorithm. The desired frequencies for the fractional tune frequencies can be entered in the lower part of the window. In the lowest part of the window, parameters for the swept frequency mode can be adjusted.*

Different parameter settings of each EMW GUI (i.e. a user-defined subset of the EPICS record values that are used by a GUI) can be saved and re-loaded via so-called setup files. This allowed fast testing of different parameter settings e.g. for the optimisation of the tune peak search algorithm.

The GUI in fig. 6.12 contains only the parameters and switches that are used most frequently by machine operators, while additional "expert" parameters e.g. for the configuration and optimisation of the peak search algorithm are accessed by separate GUIs. In the upper part of the GUI in fig. 6.12, the user can choose between manual tune measurement (via oscilloscope and one user-defined large frequency sweep window, see fig. 6.2) and DSP-based tune measurement via DSP-based tune peak tracking. The manual measurement via oscilloscope was implemented in order to be upward compatible with the previous analogue measurement system, as well as for tests and calibrations of the tune measurement and beam excitation system without electron beam. During normal machine operation, the DSP-based peak tracking mode is used. In this mode, the digital tune measurement can be started and stopped with the "stop" and "run" buttons. The DSP system allows both single shot measurements (e.g in order to measure and record the tunes for a single ramp cycle) and continuous measurement. The measurement is only started when the beam current exceeds a user-defined trigger level (e.g. 0.6 mA in fig. 6.12) when beam is injected. The measurement stops when the beam current is 50 % below the trigger level (e.g. at beam extraction), and starts again when the beam current exceeds the trigger level again (if "continuous measurement" mode is selected). The GUI provides some buttons in order to start other GUIs e.g. for the configuration of the tune peak tracking algorithm. In the lower part of the window, the tune feedback can be switched on and off, and the desired fractional tunes frequencies ("fx_set", "fz_set") can be entered and compared to the measured ones ("fx_meas", "fz_meas"). At present, the set values for f_x and f_z can only be changed manually. However, a time-dependent set value curve (synchronised with the Bodo energy ramp) can be implemented with little effort. This would allow a change of the beam optics during the booster ramp, e.g. for machine physics studies or to use one beam optics that is optimised for beam injection and another optics (with a different betatron tune) that is optimised for beam extraction. The bottom of fig. 6.12 shows the parameters for the manual swept frequency mode of the tune measurement system. The operator can adjust the lower limit and width of the sweep frequency window, the sweep rate, and the x offset for the scope in fig. 6.2.

The GUI in fig. 6.13 is used to modify and optimise parameters of the tune peak tracking algorithm of the master DSP. As mentioned above, the DSP sweeps once through a (usually large) frequency window when the measurement is started and the beam current rises above the trigger level. The lower frequency and the width of this sweep window can be adjusted by the parameters in the left upper part of the GUI, as well as the sweep rate. This initial sweep rate is normally as large as 400 kHz/ms in order to find the fractional tune frequencies after beam injection as fast as possible so that the feedback can correct them as fast as possible. The tune frequencies are then detected by a peak search algorithm. The algorithm looks for all signal peaks in the initial sweep window that are larger than the user-defined peak height in fig. 6.13. The beginning of a peak ("left edge") is defined as a minimum number of successive tune ADC samples ("edge points") with positive sample-to-sample voltage differences ("edge differences") that are each larger than a user-defined positive value. The end of a peak ("right edge") is defined accordingly, but with negative sample-to-sample voltages. In order to find the tune peaks, the algorithm searches alternately for left and right peak edges (starting with the
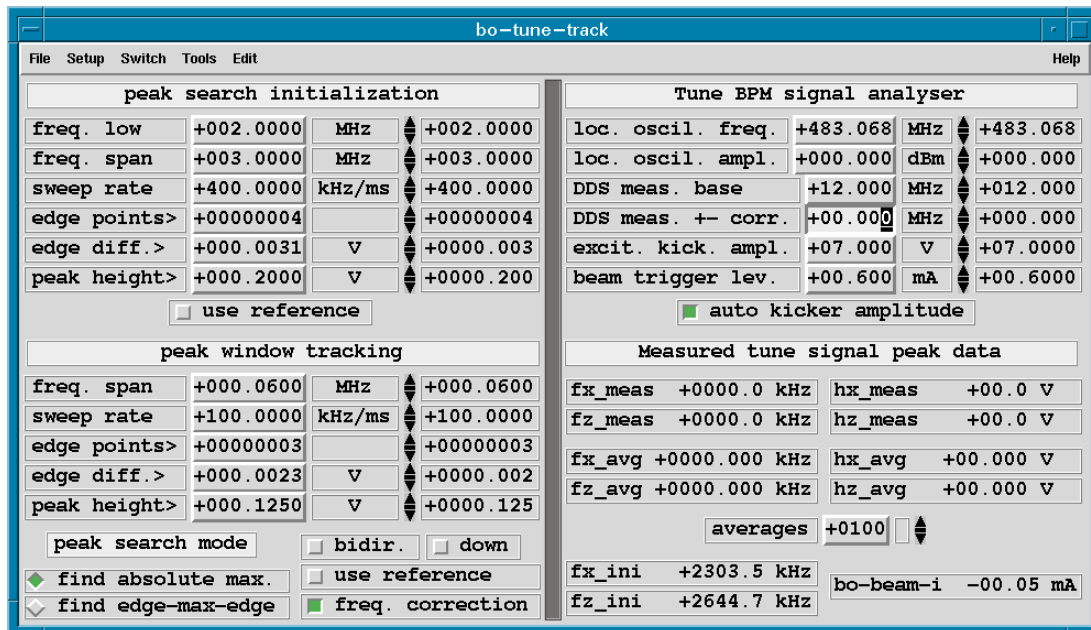
bo−tune−track
File   Setup   Switch   Tools   Edit   Help

peak search initialization

freq. low	+002.0000	MHz	▲	+002.0000
freq. span	+003.0000	MHz	▲	+003.0000
sweep rate	+400.0000	kHz/ms	▲	+400.0000
edge points>	+00000004		▲	+00000004
edge diff.>	+000.0031	V	▲	+0000.003
peak height>	+000.2000	V	▲	+0000.200

☐ use reference

peak window tracking

freq. span	+000.0600	MHz	▲	+000.0600
sweep rate	+100.0000	kHz/ms	▲	+100.0000
edge points>	+00000003		▲	+00000003
edge diff.>	+000.0023	V	▲	+0000.002
peak height>	+000.1250	V	▲	+0000.125

peak search mode   ☐ bidir.   ☐ down
◆ find absolute max.   ☐ use reference
◇ find edge−max−edge   ■ freq. correction

Tune BPM signal analyser

loc. oscil. freq.	+483.068	MHz	▲	+483.068
loc. oscil. ampl.	+000.000	dBm	▲	+000.000
DDS meas. base	+12.000	MHz	▲	+012.000
DDS meas. +− corr.	+00.000	MHz	▲	+000.000
excit. kick. ampl.	+07.000	V	▲	+07.0000
beam trigger lev.	+00.600	mA	▲	+00.6000

■ auto kicker amplitude

Measured tune signal peak data

fx_meas	+0000.0 kHz	hx_meas	+00.0 V
fz_meas	+0000.0 kHz	hz_meas	+00.0 V
fx_avg	+0000.000 kHz	hx_avg	+00.000 V
fz_avg	+0000.000 kHz	hz_avg	+00.000 V

averages +0100 ▲

| fx_ini | +2303.5 kHz | |
| fz_ini | +2644.7 kHz | bo−beam−i −00.05 mA |

Figure 6.13: *Graphical user interface for the configuration of the tune peak tracking system. In the upper left part of the window ("peak search initialisation"), the user can adjust frequency and peak search parameters for the first sweep after the tune measurement is started (e.g. after beam injection), when the DSP sweeps through a user-defined (usually large) frequency window and tries to find the tune peaks in this window. In the left lower part ("peak window tracking"), parameters for the subsequent tracking of the tune peaks (via small moving frequency windows) can be modified. The right upper part of the window allows the adjustment of the local oscillator frequency and output level (see fig. 6.2), of the base frequency of the lower DDS sine generator in fig. 6.2, and of an optional additional small frequency offset for this sine generator that is subtracted for up sweeps and added for down sweeps (e.g. for filter testing purposes). Measured actual and averaged fractional tune frequencies are displayed in the right lower part of the window, as well as the tune frequencies that were measured during peak search initialisation ("fx_ini", "fz_ini"), i.e. at the first sweep after the measurement was started (e.g. after injection).*

ADC data for the lower frequency limit of the window). The frequency than belongs to the ADC sample with the largest value between the left and right edge of a peak is defined as the frequency of the peaks. If the DSP finds at least two peaks in the initial sweep window that are larger than a user-defined peak height (see fig. 6.13), these peaks are assumed to be the tune peaks. Horizontal and vertical tune peaks can be distinguished by small variations of the QF and QD quadrupole currents. The "use reference" button in fig. 6.13 tells the DSP to perform two initial sweeps instead of one: one sweep with beam excitation and one sweep without beam excitation. Then the peak search is performed for the difference of the two resulting ADC data curves, which can be used to suppress peaks in the spectrum that do not originate from the beam, but e.g. from external RF "noise" sources or from revolution harmonics of the Delta beam that may couple into the beam pipe via ceramic gaps.

After the initial single-window sweep and the detection of the two tune peaks, the master DSP starts to sweep in small frequency windows around each tune peak. The left lower half of
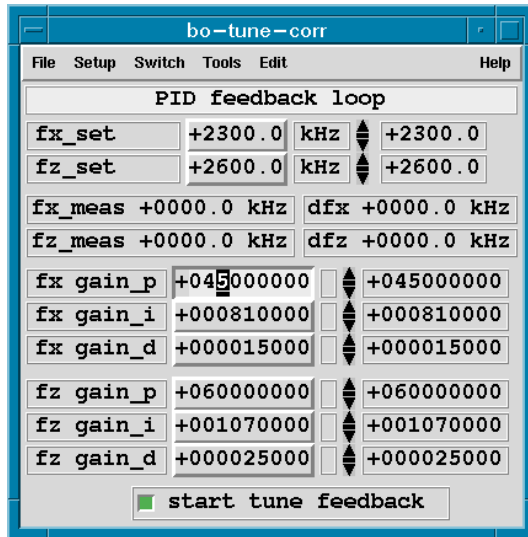
Figure 6.14: *GUI for the tune feedback loop. The upper part is used to enter the desired fractional betatron tune frequencies ("fx_set..." etc.) and to display the measured frequencies and the difference to the desired values. Furthermore, the P, I and D parameters for the PID controller for the horizontal and vertical tune feedback can be adjusted ("fx_gain_ ..." etc.). The feedback can be switched off to allow magnet ramps with user-defined currents.*

fig. 6.13 shows the parameters for this peak tracking algorithm. Width and sweep rate for the frequency windows can be adjusted, as well as the parameters of the peak search algorithm for the frequency windows. The user can choose between two different peak search modes. One works in analogy to the initialisation peak search as described above (fig. 6.13, button "find edge-max-edge"), with one instead of two peaks. The other algorithm just looks for the largest ADC value in the window that has a user-defined minimum distance to the window edges (fig. 6.13, button "find absolute max."). Despite being more primitive, the latter algorithm works reliable as long as the windows are not too large. However, if the distance between the peaks becomes too small, the DSP uses one peak tracking window instead of two. In this case, it automatically uses the "edge-max-edge" algorithm. The "use reference" button has the same functionality as described for the initial peak detection. The GUI in fig. 6.13 allows to choose between peak tracking with up sweeps only, down sweeps only, or alternating up and down sweeps. Furthermore, the DSP can correct the systematic tune peak frequency shift as described in previous sections (button "freq. correction").

The right part of the GUI in fig. 6.13 is used to adjust the local oscillator, the base frequency and an asymmetric frequency offset of the lower DDS generator in fig. 6.2. If the DDS generator is set to 12 MHz, the RF front-end detects betatron sidebands with a distance to the 85th revolution harmonic that is equal to the kicker excitation frequency. Different DDS generator base frequencies can be used to analyse kicker-independent beam excitations or higher harmonics of the kicker driver. The asymmetric DDS generator offset ("DDS +- freq. corr.") in fig. 6.2 is subtracted from the DDS generator frequency for up sweeps and added for down sweeps. Variations of this parameter in the order of some kHz can be used to analyse the duration of the beam excitation (when the kicker has already passed the fractional tune frequency), to improve the shape of the tune peaks, or for filter testing purposes.

Chapter 7

Development of a Distributed DSP-Based Global Orbit Feedback System

This chapter describes the architecture of a distributed DSP-based global orbit feedback system that was developed in this thesis. The system is based on the universal DeltaDSP beam parameter measurement and feedback architecture which is described in chapter 5. It allows global orbit corrections at a rate of several kHz both for Bodo and Delta, which is more than three magnitudes faster than the existing Delta orbit feedback system [32]. As explained in chapter 5, the system is generic in the sense that both hardware and software can be easily adapted to different boosters and storage rings, mainly by modifications in configuration files and with very few modifications of the actual software at some locations in the DSP feedback code where e.g. the maximum number of BPMs was hard-coded for speed reasons. In this thesis, the system was installed and tested at Bodo, since the steerer magnet power supplies in Delta only allow orbit corrections at a rate of about 2 Hz [32], whereas the power supplies of Bodo are more than two orders of magnitude faster. However, Bodo uses the same beam pipe, dipole, quadrupole, combined function steerer magnets and BPM RF front-end electronics as Delta. Due to the generic modular architecture of feedback software and hardware, a future implementation of the feedback system at Delta is possible with minimum effort.

In case of a future implementation at Delta, the DSP-based global orbit feedback system could be used to correct orbit perturbations that are caused e.g. by the ramping booster magnets, by pulsed transfer line dipole magnets or by vibrations of ground, girders and magnets, with typical perturbation frequencies of some Hz to some 10 Hz. Due to its closed-loop bandwidth of a fraction of 1 Hz, the existing Delta orbit feedback system is not able to correct such perturbations. The DeltaDSP system can achieve a closed-loop bandwidth in excess of 100 Hz, which would lead to a significant reduction of the orbit perturbations described above. This is essential both for normal operation and for the envisaged top-up injection mode of the storage ring.

The implementation at Bodo had the advantage that both hardware and software could be tested and developed during synchrotron radiation user operation of the storage ring. As the number of machine physics shifts is much smaller than the number of synchrotron radiation user shifts, the development time for the system could be reduced significantly. Moreover, the use of Bodo instead of Delta for the first tests of the system reduced the risk for synchrotron user shifts, since occasional beam loss in Bodo during tests and optimisation of the system is not critical and may occur even every few minutes, whereas beam loss in Delta is much more critical during user operation.

Moreover, the previous "bo-006b" optics of Bodo did not require orbit corrections by steerer magnets in order to store acceptable beam currents, while the new "bo-101c" optics did not allow to store the beam without orbit corrections by steerer magnets. At first sight, this result may seem unexpected, since the simulations in section 4.2.2.1 show that both optics should have a similar sensitivity to magnetic field errors that cause orbit perturbations. However, the

"bo-006b" optics was obtained empirically by varying magnetic fields of the quadrupoles until the beam current was maximal (without steerer magnets, since Bodo did not have any steerer magnet power supplies at that time), so that the existing magnetic field errors compensated each other in an optimal way by appropriate betatron phase differences between the locations of the magnetic field errors or misaligned magnets. In contrast, the "bo-101c" optics is nearly identical to an optics that was calculated without accounting for magnetic field errors, with much smaller variations of the theoretical quadrupole focussing strengths with respect to the calculated optics in order to obtain higher beam currents while preserving the low emittance of the optics. The installation of the DSP-based orbit measurement and feedback system allowed fast commissioning of the new optics by precise measurement and correction of the beam orbit (see chapter 9).

7.1　Orbit Feedbacks

The necessity of orbit feedbacks was already motivated in section 5.1. In general, orbit feedbacks can be subdivided into local and global feedbacks, as well as into analogue and digital feedbacks. The numerous advantages of digital systems compared to analogue ones were already pointed out in section 5.2.1. Furthermore, orbit feedbacks can be subdivided according to their speed. State-of-the-art synchrotron light sources usually have up to three kinds of orbit feedbacks. "Slow" feedbacks correct thermal drifts, with correction rates in the order of 1 Hz or less [44, 49, 57, 58]. "Fast" orbit feedbacks correct beam movements due to mechanical vibrations and 50 Hz power supply noise, with correction rates in the order of some kHz [45, 46, 48, 64]. RF orbit feedbacks apply transverse corrections kicks at the timescale of one beam revolution or even at the timescale of the bunch-to-bunch distance in order to fight beam instabilities, with correction rates of MHz to hundreds of MHz [97, 98, 99]. Slow orbit feedbacks are usually global, i.e. they use all BPM data to calculate the correction for each steerer, since the required global data transfer on a time scale of seconds can be easily achieved with conventional office networks and computers (e.g. Ethernet and Linux PCs) [32]. Transverse feedbacks in the RF frequency range are in general local and use one BPM and one kicker per plane, since the systems can kick bunches individually on a turn-by-turn basis. Therefore, global systems with many kicks per turn at different locations of the ring usually make little sense. Furthermore, the extreme speed of the system imposes severe limitations on computing time and data transmission speed, so that global systems could hardly reach the desired bandwidth.

7.1.1　Local and Global Orbit Feedbacks

This chapter deals with "fast" feedbacks, i.e. correction rates in the order of kHz. In general, such feedbacks can be implemented locally or globally. Local feedbacks at synchrotron light source storage rings are usually implemented as one feedback per beamline (in one or both transverse planes), where the feedback uses some electron beam BPMs or synchrotron light beam position monitors ("S-BPMs") in order to stabilize the position or angle of the synchrotron light of insertion devices by local orbit bumps at the source point. Since local feedbacks require no global data transfer and need much less computation time, they are easier to design and to implement and are therefore often used when the first one or two beamlines

of a new synchrotron light source are commissioned. However, after some years of operation most existing synchrotron light sources have a very large number of beamlines. In case of one local feedback per beamline, the local orbit bumps of local feedbacks are usually not perfectly closed, therefore each local feedback generates fast *global* orbit perturbations that change the orbit outside the bump. Moreover, each local feedback also tries to correct the global orbit perturbations that are caused by the other feedbacks. This undesired crosstalk between the systems may lead to undesired beam oscillations and unstable operation of the feedbacks, especially if three or even more local feedbacks in a storage ring are operational at the same time. Furthermore, changes of the quadrupole current settings (e.g. in order to correct the betatron tunes) or changes of the orbit itself (in case of nonlinear optics due to sextupole fields) may vary the beta functions. Since the required steerer magnet kicks for a closed orbit bump depend on the beta functions [6], the varying beta functions would cause increasing global orbit perturbations by each local feedback, because previously closed bumps would not be closed any more. This can also result in undesired beam oscillations and unstable feedback operation. In order to avoid such problems, state-of-the-art synchrotron light sources use fast *global* orbit feedback systems [45, 46, 49, 64].

The DeltaDSP system transfers all BPM data to all DeltaDSP boards in the ring via DeltaNet, therefore a future implementation of the system for the Delta storage ring would be global from the beginning, even if the initial system consisted only of a few number of steerer magnets and BPMs around the source points of a few beamlines. Since each DeltaDSP board knows about the orbit at any beamline, the corrections can be calculated globally (e.g. using the SVD algorithm, see section 2.8), so that the steerers stabilise the beam simultaneously at all beamlines. Therefore, the system is inherently global, and the number of beamlines, BPMs and steerer magnets that are used for the feedback may be increased according to the requirements of the synchrotron radiation users, without the growing risk of beam instability and the sensitivity of multiple local feedbacks. Even in case of 64 BPMs and 32 steerer magnets in each plane, a system of 8 DeltaDSP boards in Delta would be capable of calculating and applying global SVD-based orbit corrections at a rate of more than 4 kHz. The following section

7.2 Architecture of the Bodo Orbit Feedback

7.2.1 Beam Position Measurement and Data Acquisition

Figure 7.1 shows the data flow of the Bodo orbit feedback system. The four capacitive pickup buttons of each BPM (see fig. 3.6) are connected to the four RF inputs of commercially available BPM RF front-end electronics [80]. These BPM electronics generate two voltages (x and z voltage) that are proportional to the beam position, averaged over 100 to 400 μs (with maximal $\pm$ 10 V for $\pm$ 10 mm vertical or horizontal beam position, and 0V for a centered beam). The averaging time is determined by an internal 10 kHz clock ("button sample clock") of the BPM electronics. Instead of the internal clock, it is also possible to use an external button sample clock with up to 40 kHz clock rate. The RF signal of each of the four BPM buttons is sampled for one clock cycle, and the x and z signals are calculated by analogue circuits in the BPM electronics every four clock cycles (according to eq. 3.1). The internal 10 kHz clock leads to update rates of 2.5 kHz for the x and z outputs, whereas an external 40 kHz clock can achieve 10 kHz update rate.
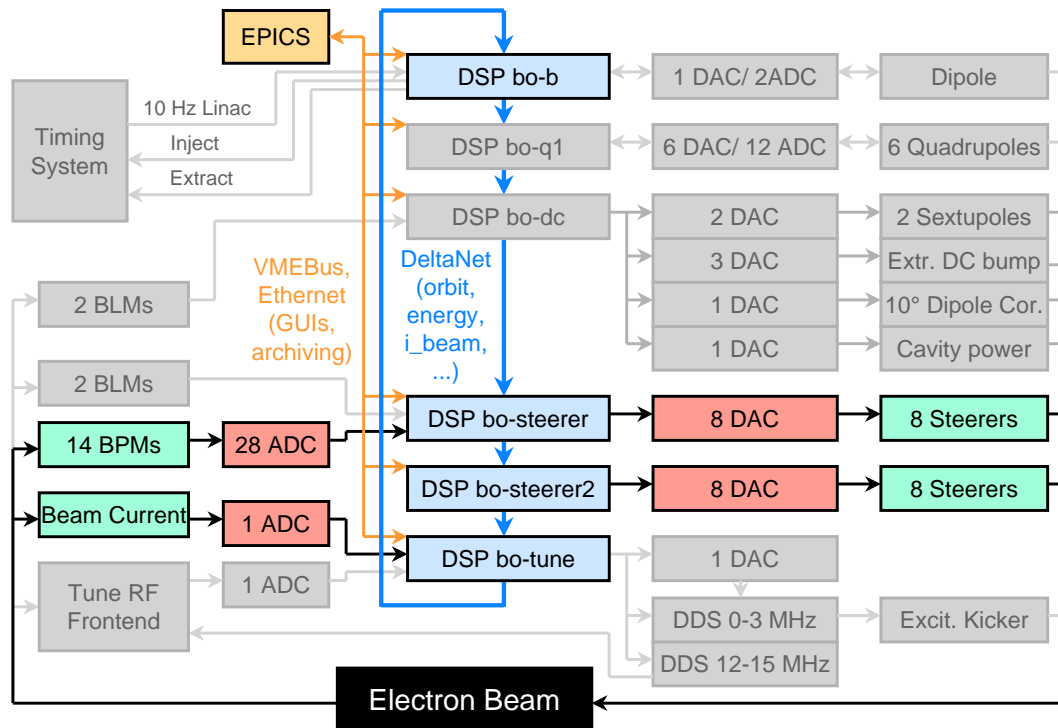
Figure 7.1: *Architecture and data flow of the DSP-based global orbit feedback system of Bodo.*

The x and z voltages are sampled by external ADC boards that were developed in this thesis (see section 5.3.1). Each board contains two 16-bit ADCs that perform conversions at a programmable rate (usually 4 to 20 kHz). The ADC boards are installed in the same BPM racks as the BPM RF front-end electronics, which minimises analogue signal cable lengths (and thus external noise). The ADC board precision corresponds to a position resolution of 300 nm per bit, a monotony error below 600 nm and an RMS noise of 450 nm at 30 kHz analogue ADC board bandwidth and 20 ksamples/s sampling rate.

The ADC data is transferred to DeltaDSP boards by a proprietary synchronous serial bus over distances up to 50 m at a rate of 10 Mbit/s. Up to 32 ADC boards can be connected to one serial bus in a daisy-chain fashion. Each DeltaDSP board can handle four ADC buses, i.e. 128 BPMs. All ADC boards in the DeltaNet ring sample the beam position data synchronously (by synchronised triggers from the DeltaDSP boards). At present, Bodo uses only two BPM ADC buses with 7 BPMs per bus, and the buses are connected to the "bo-steerer" DSP board (see fig. 7.1) that handles 8 of the 16 Bodo steerer magnets. However, hardware and software of the feedback system were designed so that BPM ADCs can be connected to any DSP board in a DeltaNet ring, without affecting the performance of the feedback since DeltaNet distributes the ADC data to all DeltaDSP boards in real-time. The only performance limitation is the number of ADC boards per serial bus. The ADC boards send their data to the DSP board one after another, and the readout requires a base time slightly less than 15 μs (for the ADC conversion etc.) plus 5 μs times the number of ADC boards for the daisy-chained readout. Therefore, 7 ADC boards per bus in Bodo have a maximum readout rate of 20 kHz.

The serial ADC bus transfers the ADC data directly and automatically into the internal memory of the DSPs by their DMA unit, through the synchronous serial interfaces of the DSPs.

At present, the DSP software only uses the two serial ports of the master DSP (the other two ADC buses are connected to the slave DSP), since four buses per DSP board were not yet required. DeltaNet synchronises the sampling triggers and data acquisition for all ADCs in Bodo, with typically less than 300 ns conversion trigger jitter between any two ADCs.

7.2.2 RF Front-End Electronics

In this section, some parameters of the commercially available BPM RF front-end electronics (that are identical for Bodo and Delta) are measured which are important for orbit feedback systems, but that are either not specified in the datasheets or that might be different from the datasheet values due to different environmental conditions (e.g. noise). Parameters like bandwidth or position signal noise level are important for feedback systems since they impose limitations on the maximum closed loop bandwidth and on the maximum orbit stability that can be achieved with an orbit feedback system. The properties of the x and z output signals of the BPM electronics are basically identical, except for a different scaling (in order to account for the shape of the DELTA beam pipe). Therefore the measurement results are usually presented only for the x output.

Bandwidth

Since the datasheet of the BPM RF front-end electronics does not specify its bandwidth [80], it was measured with amplitude-modulated test signals of two RF frequency generators that simulated a periodically moving electron beam. The generators were synchronised in such
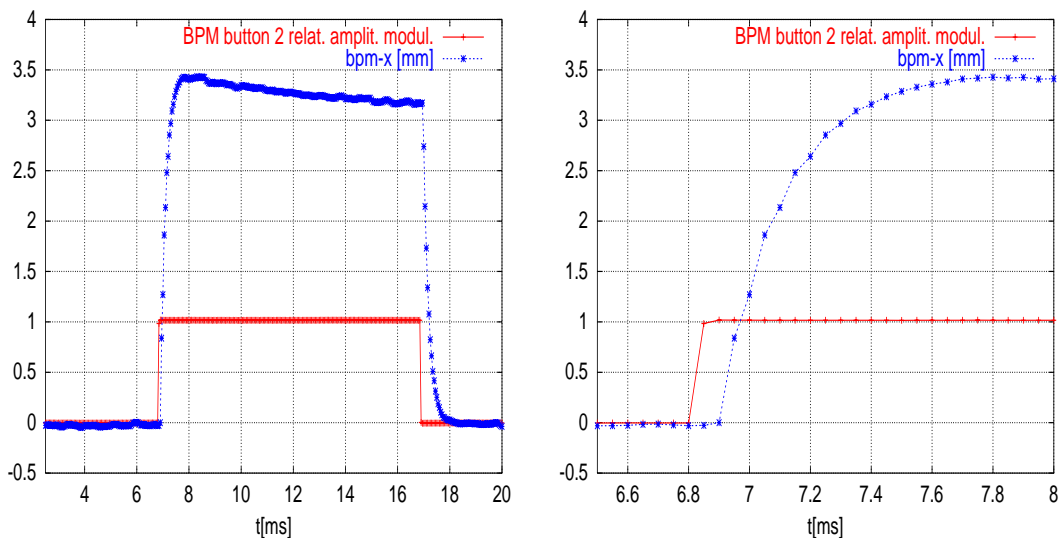
Figure 7.2: *Response of the BPM RF front-end electronics to an instantaneous beam orbit change. The orbit change was simulated by rectangular amplitude modulation of RF generators that were connected to the inputs of the BPM electronics. The red curves show the modulation voltage resp. the simulated beam movement (in arbitrary units), the blue curves are the response signals of the x output of the BPM electronics.*

a way that they had the same frequency and a constant phase relation. The frequencies of the generators were set to 499.820 MHz, which is equal to the usual RF cavity frequency resp. the frequency of the 84th Bodo revolution harmonic that is normally used by the BPM electronics. As mentioned above, the BPM electronics has four RF inputs for the four capacitive pickup buttons of each BPM, and two outputs for the horizontal and vertical beam position. The first RF generator was connected to three of the four inputs by a symmetric power splitter, at a power level of -30 dBm per input (corresponding to about 18 mA average beam current [24]). The second RF generator was connected to the 4th input (that is usually connected to BPM pickup button 2, see fig. 3.6). The input power level of BPM input 2 was modulated, which simulated a moving electron beam and changed the x and z output signals of the BPM electronics according to eq. 3.1. The resulting input power level combination (same levels for BPM inputs 1,3,4, different level for input 2) does actually not occur for real electron beams, but this is not relevant for the measurement of response time and bandwidth.

Figures 7.2 show the step response of the x output of the BPM electronics for an amplitude jump of the second RF generator. Compared to the BPM electronics, the response time of the RF generator to the jump of the modulation voltage was negligible, therefore the modulation voltage is a true measure of the RF amplitude. The amplitude change corresponds to an orbit change of about 3.2 mm. The BPM electronics was operated with an external 40 kHz button sample clock (see above), and the ADC sample rate was 20 kHz. The red curve in fig. 7.2 is proportional to the amplitude modulation, the blue curve is the response of the x output. The figures show that the BPM electronics needs about 0.6 ms to reach the desired x output level, with a slight overshoot that decays in at least 10 ms.
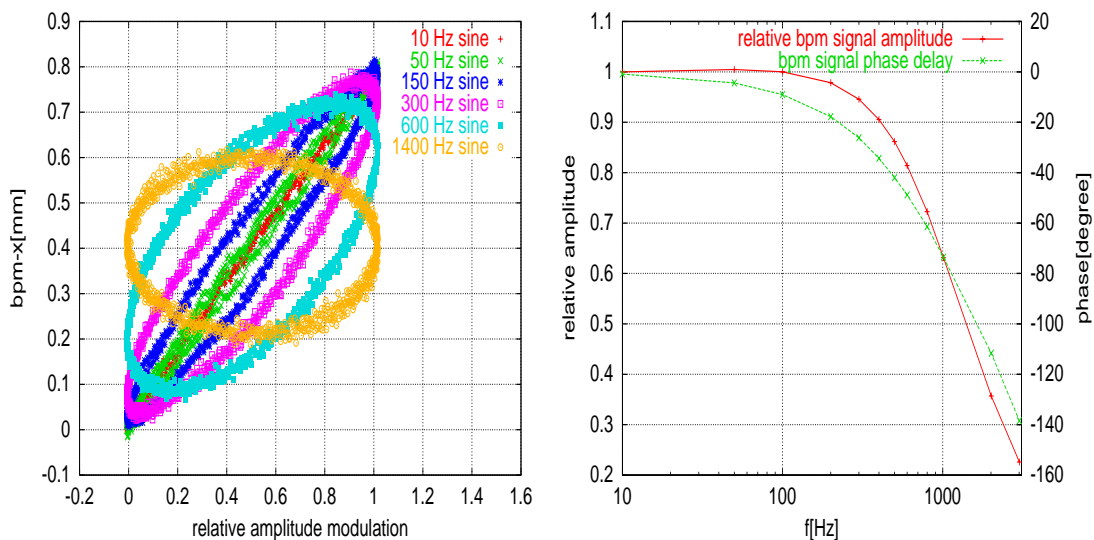

Figure 7.3: *Response of the BPM RF front-end electronics to periodic beam movements. The movements were simulated by sine-wave amplitude modulation of RF generator signals. The left figure shows the x output signal (vertical axis, in [mm] or [V]), plotted against the modulation voltage (in arbitrary units). Both signals were sampled by the BPM ADCs for frequencies between 10 and 1400 Hz. The right figure shows phase delay and relative amplitude attenuation of the x output signal with respect to the simulated beam movement (i.e. with respect to the amplitude modulation voltage).*

Figures 7.3 show the measurement results for a sine-wave amplitude modulation of input 2 of the BPM electronics. The amplitude modulation simulated a periodic orbit perturbation. For low beam movement frequencies, the x output signal of the BPM electronics (left figure, vertical axis) is proportional to the modulation voltage. At higher frequencies, the x output has both a relative amplitude attenuation and a phase delay with respect to the actual beam movement.

Figures 7.4 show the effect of different button sampling frequencies for the BPM RF front-end electronics. The left plot shows a 100 Hz orbit perturbation, the right one a 600 Hz orbit perturbation (both simulated with RF generators as described above). The blue and red curves show the orbit perturbation (i.e. the amplitude modulation voltage for the RF generators) in arbitrary units. The green and magenta curves show the x output of the BPM electronics (i.e. the measured beam position) for 2.5 kHz output update rate, which corresponds to 10 kHz button sampling frequency of the internal sampling clock, and for 10 kHz output update rate, which corresponds to 40 kHz button sampling frequency of an external clock. The signals were sampled by BPM ADCs at 20 kHz, without synchronising ADCs and BPM clocks.

The 100 Hz orbit oscillation looks similar for both button sampling frequencies. The 600 Hz oscillation can only be resolved properly at 10 kHz x output update rate, whereas 2.5 kHz update rate lead to rectangular distortions of the sine curve. The button sampling delay also limits the maximum closed-loop bandwidth of an orbit feedback. In case of Bodo, the BPM electronics usually use the internal button sampling clock (i.e. 2.5 kHz output update rate), which allows a sufficient orbit correction during energy ramps. However, a future orbit feedback in Delta should use an external 40 kHz sampling clock, corresponding to 10 kHz output
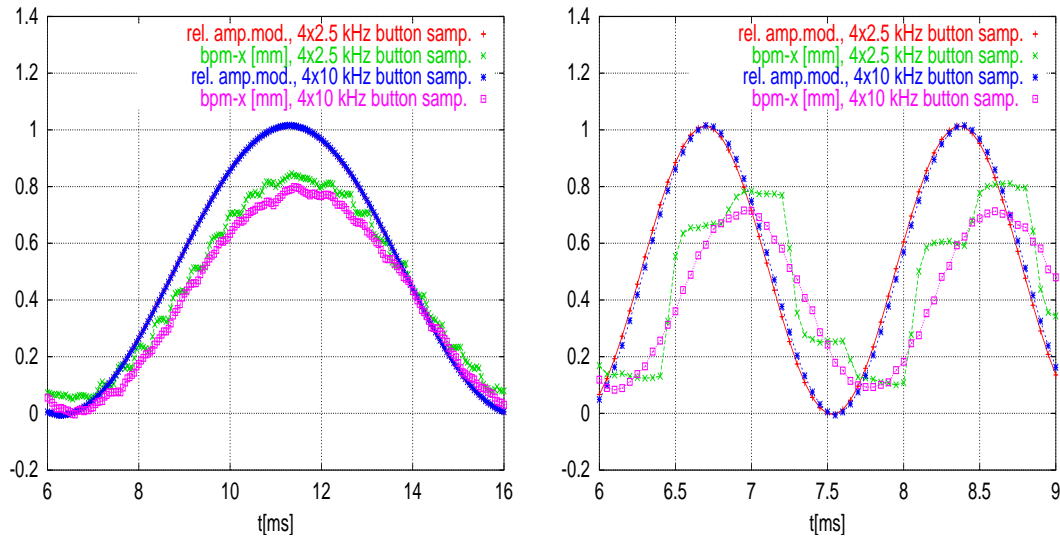

Figure 7.4: *BPM ADC signals for an electron beam that oscillates horizontally with 100 Hz (left figure) and 600 Hz (right figure). Each figure shows the beam position (red and blue curves, in arbitrary units) and the resulting x output signal of the BPM RF front-end electronics for 2.5 kHz (green curve) and 10 kHz (magenta curve) update rate of the x output, corresponding to 10 and 40 kHz BPM button sampling frequency. The oscillating beam was simulated by synchronised amplitude-modulated RF generators. The red curve in the right plot is hidden under the blue curve, since the two amplitude modulations were synchronous.*

update rate for the BPM RF front-end electronics, in order to maximise the feedback band-width.

Power and Phase Dependence

In order to investigate the beam current dependence of the x and z output signals of the BPM RF front-end electronics, the signals were recorded for different input power levels. An RF generator and a symmetrical power splitter simulated a centered beam with different beam currents. Input power levels of -35, -55 and -75 dBm correspond to average beam currents of about 0.1, 1 and 10 mA [24] (neglecting cable damping). Figures 7.5 show the RMS noise and beam position (i.e. x output signal) as a function of the power level. The x output signal was sampled by a BPM ADC board at 20 kHz ADC sample rate and 30 kHz analogue ADC board bandwidth. The noise level of the ADC board was negligible compared to the BPM electronics. Due to component tolerances of the BPM electronics, the x and z output signals are not exactly zero, but have typical offsets in the order of 100 μm.

For input power levels between -40 and 0 dBm, the RMS noise of the x output signal is nearly constant (about 11 μm). Lower levels increase the noise up to 1 mm at -90 dBm. In case of Bodo, these noise levels are sufficient both for normal and test machine operation. However, a future fast orbit feedback in Delta (that uses the same BPM electronics) might require lower BPM noise levels. This might be achievable by replacing the switched-mode power supplies of the Bodo BPM racks with low-noise regulated power supplies, or by improved cabling and better shielding, since the intrinsic noise level of the BPM electronics may be increased by a noisy environment.

Furthermore, a noise reduction to sub-micron level can be achieved by averaging of the ADC samples, at the expense of a smaller bandwidth. The intrinsic noise of BPM systems
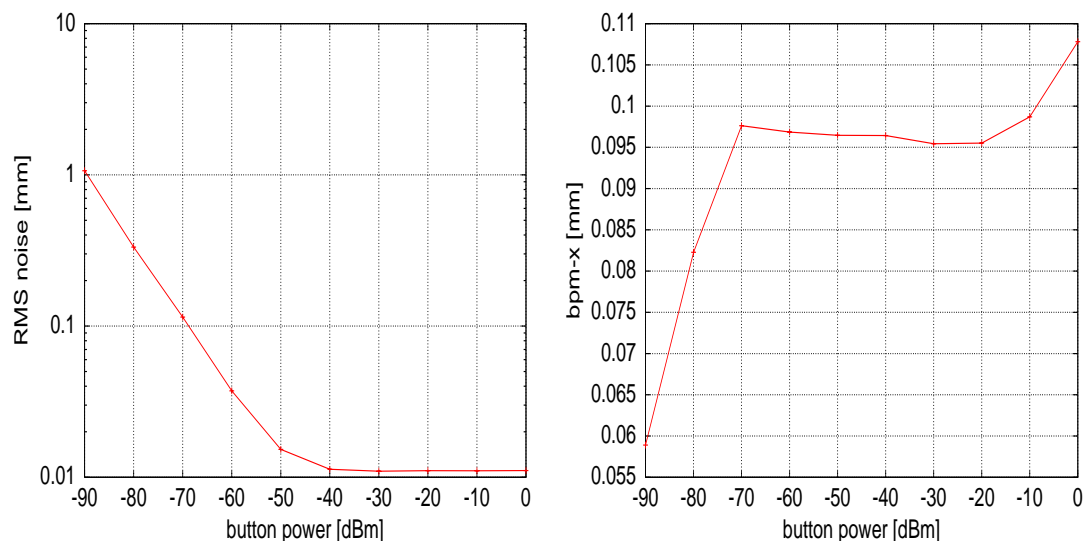

Figure 7.5: *RMS noise (left plot) and averaged output signal (right plot) of the x output of the BPM RF front-end electronics as a function of beam current resp. RF input power. The RF signals of a centered electron beam were simulated by an RF generator and a symmetrical power splitter.*

usually scales with the square root of the system bandwidth [80]. Despite a maximum x and z output update rate of 10 kHz, the BPM RF front-end electronics of Bodo and Delta have a 3 dB bandwidth frequency of less than 1 kHz (see fig. 7.3), possibly due to a low pass filter in the electronics. The variation of the output update rate between 2.5 and 10 kHz had no effect on the noise level of the BPM electronics in case of constant input signals. Therefore, ADC samples must be averaged for more than 1 ms in order to achieve a significant reduction of intrinsic BPM noise. Lower averaging periods will only reduce external noise (e.g. due to long ADC input cables or stray fields of switched-mode power supplies).

Each BPM electronics has an oscillator that locks onto the frequency component of the RF input signal that is closest to the nominal cavity frequency, with a minimum locking power level of about -70 dBm (according to an LED that displays if the oscillator is locked). However, for RF generator frequencies in the order of the usual cavity frequency (499.820 MHz $\pm$ some ten kHz) it was found that the x output showed meaningful position signals both for simulated centered and non-centered beams up to -90 dBm input power level (with the "PLL-locked LED" being off), at the expense of more noise and a systematic position error in the order of 50 μm. The right plot in figures 7.5 shows x output variations of a few micrometers for input power levels between -70 and -20 dBm, while both larger and smaller levels lead to systematic position errors between +10 μm (for 0 dBm) and -50 μm (for -90 dBm).

Figure 7.6 shows the dependence of the x output signal of the BPM electronics on the phase of the RF input signal. RF inputs 1,3 and 4 of the BPM electronics received an RF signal of same phase and amplitude, while input 2 had an different amplitude that resulted in a beam position (i.e. an x output signal) of 5.1 to 5.2 mm. The phase of the signal of input 2 was periodically modulated by 180^o (red curve), which caused x output variations of 80 μm (blue curve). Therefore, length (and thus phase) changes of cables from BPM pickup buttons to the BPM electronics can lead to significant position measurement errors, even if cables with very low attenuation are used and the signal amplitude differences of old and new cables are negligible. This is mainly important for the Delta BPM system, since an orbit change of 80 μm
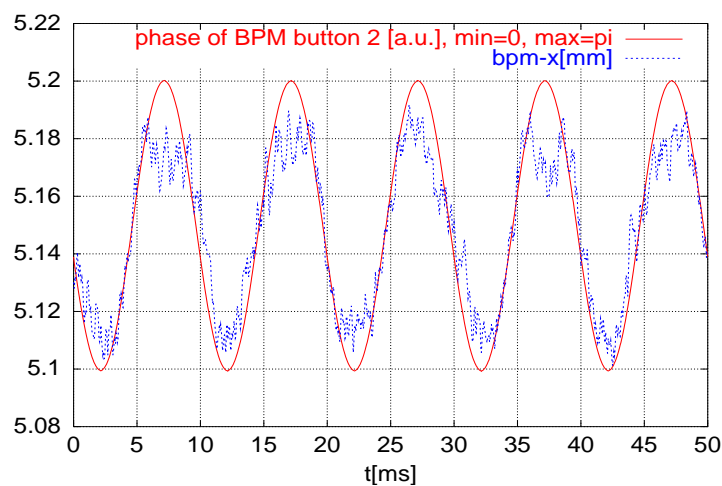

Figure 7.6: *Dependence of the position signal (blue curve) of a BPM RF front-end electronics on the phase of RF input 2 (red curve, with arbitrary offset and units). The 499.820 MHz RF signal of input 2 was phase-modulated periodically by* 180^o *with respect to the other inputs.*

at an insertion device BPM may cause movements of the synchrotron light at the experimental station in the order of millimeters, due to the usually large distance of up to some 10 meters from the experimental stations to the source point.

7.2.3   Steerer Power Supplies and Magnets

The steerer magnets and the beam pipe of Bodo were already described in section 3.2. The magnets are integrated function magnets that consist of additional coil windings on the laminated iron quadrupole magnets of Bodo. The steerer magnet coils are connected to commercially available power supplies which can sink and source ± 12 A current both for negative and positive output voltages up to 36 V. In linear approximation and without saturation correction, a horizontal steerer magnet current I_h or vertical steerer magnet current I_v changes the orbit of an electron beam of energy E in Bodo by an angle [32]

$$\Delta x' = 0.45 mrad \cdot \frac{I_h[A]}{E[GeV]} \tag{7.1}$$

$$\Delta z' = 0.17 mrad \cdot \frac{I_v[A]}{E[GeV]} \tag{7.2}$$

The steerer magnet power supplies have a 3 dB bandwidth of 10 kHz in current regulation mode for purely resistive loads [100, 101], whereas the bandwidth for inductive loads is also limited by the maximum output voltage and may therefore be smaller than 10 kHz, depending on the load inductance and on the required frequency and amplitude of the output current. In case of the Bodo steerers, the inductance limits the current change to about 105 A/s for horizontal and 250 A/s for vertical steerer magnets, assuming an initial current of 0 A. This is illustrated in fig. 7.7, which shows the Bodo beam orbit at BPM 13 (see fig. 3.4, page 23) for current set
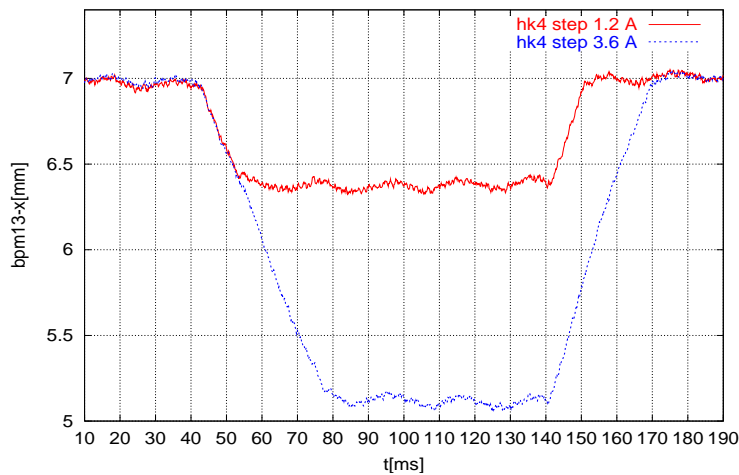

Figure 7.7: *Response of the horizontal beam orbit at BPM 13 in Bodo for current set value jumps of ± 1.2 A (red curve) and ± 3.6 A (blue curve) of the horizontal steerer magnet HK4 (see text). The current set values are changed within 10 μs at $t = 42$ ms and again at $t = 140$ ms. Due to the load inductance, the power supplies cannot change the magnet current immediately, but they apply the maximum output voltage until the set current is reached.*

value jumps of ±1.2 A (red curve) and ±3.6 A (blue curve) of the horizontal steerer magnet HK4. The measurement was made for the "bo-006b" optics at 6 mA average beam current and 1.45 GeV beam energy. The power supplies are not able to change the magnet current instantaneously, they can only apply their maximum output voltage and wait until the current through the magnet inductance has reached the desired value. Therefore, the orbit change for both current changes in fig. 7.7 has the same slope, but the orbit movement for a current change of 3.6 A takes three times longer than for 1.2 A.

The steerer magnet power supply currents are set by external 16-bit DAC boards that are controlled by DeltaDSP boards (see fig. 7.1). The DACs have an output settling time of less than 10 μs, and one DeltaDSP board can write set values to a maximum of eight DAC boards at a rate up to 8 MHz.

The overall closed-loop bandwidth of the orbit feedback system depends on the bandwidth of the analogue BPM RF front-end electronics and on the ADC to DAC delay of the digital system that calculates the DAC settings for the steerers from the BPM readings. Furthermore, the feedback bandwidth depends on the magnet power supply bandwidth and of the geometry, material and orientation of magnets and beam pipe, because eddy currents cause a phase shift and an attenuation of AC magnet fields at higher frequencies. In case of fast orbit feedbacks that are designed before the construction of an accelerator, the closed-loop bandwidth usually has to be obtained from a theoretical model that requires the measurement of the time and phase delays and of the attenuation or gain of the different components of the feedback loop. However, in case of Bodo with its already existing hardware, the frequency response of the closed feedback loop could simply be determined by beam-based measurements at the real machine (see chapter 9). During the work on the Bodo feedback, there was no spare quadrupole magnet available that could have been used for AC magnetic field measurements with and without beam pipe in order to determine bandwidth limitations due to eddy currents. This was another reason for the beam-based approach.

In case of a future fast orbit feedback in Delta that would require new steerer magnet power supplies and possibly new dedicated feedback steerer magnets, a model-based approach could be useful nevertheless, e.g. in order to optimise the costs for individual components of the feedback chain for a certain desired feedback bandwidth. A planned major reconstruction of the Delta vacuum system in the near future (after the work on this thesis had been finished) would allow beam-independent magnetic field measurements, in order to compare the frequency response of the integrated Delta steerer magnets and of the beam pipe to dedicated separated function feedback steerer magnets that might use air coils instead of laminated magnet iron.

7.2.4   Orbit Feedback Algorithm and Graphical User Interfaces

The DeltaDSP boards of the Bodo orbit feedback use the SVD method (that was described in section 2.8) in order to correct the beam orbit. The required beam response matrix is calculated from a theoretical optics model, using the program MAD [26] and a Bodo-specific graphical user interface ("optics GUI") that was developed in this thesis (using the EMW user interface builder that was described in section 5.4.1). The optics GUI is shown in fig. 7.8. The quadrupole and sextupole magnet strengths can be entered in the left upper part of the GUI window (as well as the beam energy and the energy spread). By pressing a button on the GUI, MAD calculates the optical functions and other parameters of the optics. The results are displayed by the GUI,

Figure 7.8: *Graphical user interface for the online calculation of the Bodo beam optics, of the vertical and horizontal beam response matrices and of the SVD-inverted beam response matrices for the orbit feedback system.*

e.g. beta functions and dispersion. Furthermore, MAD calculates the horizontal and vertical response matrices, which are SVD-inverted by a Tcl/Tk program that uses a matrix calculation module developed by M. Grewe [102].

The SVD-inverted matrices are optionally written to the DeltaDSP boards of the feedback system that require the matrices in order to calculate the steerer magnet settings. Each DeltaDSP board receives only the rows of the SVD-inverted horizontal and vertical response matrix that are required to calculate the corrections for the steerer magnets connected to the respective DeltaDSP board. The matrix rows are transferred from a Linux workstation to the DeltaDSP boards via data files and EPICS output records of "string" type (as described in section 5.4.2). The transfer is initiated by writing the predefined filename to the VAL field of the respective record. The transfer via files with predefined names guarantees that the DeltaDSP boards of the feedback system are able to reload the matrix rows automatically after reboot or power failures of their VMEbus rack, so that they do not have to be reloaded manually by the operators in the control room.

The SVD cutoff values (see chapter 2, end of section 2.8) for the horizontal (x) and vertical (z) plane can be adjusted in the left lower part of the optics GUI. A second optional set of cutoff

values allows to compare inverted beam response matrices for different cutoffs. The eigenvalue spectrum of the SVD inversion is plotted against the number of the eigenvalue in descending order in the upper right part of the window. The middle and right part in the lower section of the GUI show 2-dimensional plots of the columns of the beam response matrix and of its SVD inverse. All matrices and data files use the units "mm" and "mrad", and the response matrices are calculated for 1 mrad orbit kicks. The response matrix and its SVD inverse can be optionally displayed as a 3-dimensional plot (using a Tcl/Tk plot widget developed by D. Zimoch). Criteria for the choice of the SVD cutoff value that are based on the interpretation of such plots are described in section 7.3.

The Feedback Magnet PID Loops

The digital part of the feedback (i.e. the DeltaDSP boards) applies global corrections at a rate of usually 4 kHz, while the bandwidth of the analogue parts of the system (e.g. the BPMs, steerer magnets and power supplies) is lower (as mentioned above). Therefore, the DSPs cannot measure the orbit, apply the full steerer current correction that would be required to correct the orbit, and expect the orbit to be perfectly corrected at the next correction cycle. Instead, the DeltaDSP boards use a PID controller [96] for each steerer magnet that determines the regulation characteristics and speed of the orbit correction.

When the feedback is started, the DeltaDSP boards store the DAC settings of the steerer magnets. In each correction cycle, they calculate the required DAC setting for each steerer magnet that is used for the feedback (based on the SVD-inverted beam response matrix, see section 7.3). The difference between the actual DAC set value and the "ideal" DAC set value (that would correct the beam position perfectly in case of BPMs, power supplies, magnets and beam pipe with infinite bandwidth) is called the DAC set value error. The PID loop changes the actual DAC value of each steerer using a sum of three numbers: one is proportional to the DAC set value error ("P component"), one is proportional to the integral of the DAC set value error ("I component"), and one is proportional to its derivative ("D component"). These three components of the DAC correction sum can be scaled individually for each steerer magnet with a GUI ("PID GUI"). The PID GUI also allows to select the steerer magnets that are used for the feedback. The PID sum components of each steerer and its DAC base setting that was stored when the feedback was started are continuously scaled with the beam energy, in order to achieve energy-independent orbit corrections. The beam energy is available to all DeltaDSP boards via DeltaNet. Additional steerer magnet corrections due to saturation at high steerer magnet and quadrupole fields could easily be implemented in the future, but were not yet required since the feedback algorithm can correct saturation effects based on the orbit response (see chapter 9).

The settings of the PID GUI can be saved in files and reloaded by the "Setup" menu of the GUI, which allows fast testing and comparison of different PID settings. The "Tools" menu of the GUI can be used either to scale some or all parameters or to add user-defined offsets to the parameters. These features are common to all EMW GUIs at DELTA, they are not only used to configure the Bodo orbit feedback, but also e.g. to scale magnet currents, to add steerer magnet offsets e.g. in order to create local orbit bumps, and to save and load settings for magnets and other accelerator devices.

Another GUI is used to select the BPMs that are used for the orbit feedback. The GUI allows to adjust the desired beam positions and a weight factor for each BPM. When the feedback

calculates the differences between measured and desired orbit, the elements of the resulting dif-
ference vector are multiplied with the weight factors before the steerer magnet corrections are
calculated. An increase of the weight factor at some BPMs can be used to correct the orbit
at these BPMs faster and with higher priority and accuracy (e.g. if there are more BPMs than
steerer magnets so that the orbit cannot be corrected perfectly at all BPMs).

7.3   SVD-Inversion of Bodo Beam Response Matrices

When the optics GUI (fig. 7.8) calculates the beam optics data via MAD, the beam response
matrix and its SVD inverse can be optionally displayed as 3-dimensional plots. Fig. 7.9 shows
such plots for the "bo-101c" beam optics, with horizontal and vertical beam response matrices
on the left side and the corresponding SVD-inverted matrices on the right side. The SVD cutoff
value is 0.5 for the horizontal and 0 for the vertical plane. These values were also used for the
beam-based tests of the orbit feedback system in chapter 9. The necessity of non-zero cutoff
values in case of a large quotient between largest and smallest eigenvalue was already discussed
in section 2.8.

The four plots in figure 7.10 are SVD-inverted horizontal beam response matrices of the
"bo-101c" optics for different cutoff values, calculated from a theoretical optics model. If a cut-
off value of zero is used (left upper plot), the SVD algorithm tries to correct orbit displacements
at BPM 10, 11 or 12 by undesired zigzag corrections kicks of all steerers, indicated by the two
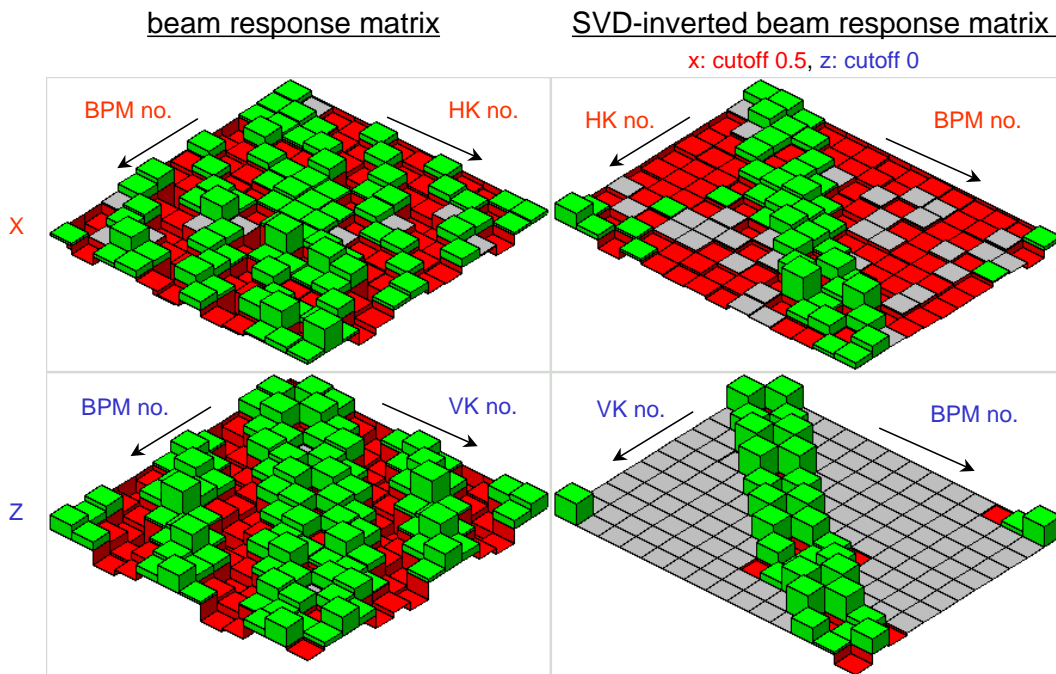
Figure 7.9: *Horizontal ("x") and vertical ("z") beam response matrices (left side) and SVD-
inverted beam response matrices (right side), based on a theoretical model of the "bo-101c"
optics. Green blocks indicate positive values, red blocks negative values, and grey blocks in-
dicate values that are close to zero (with arbitrary but linear scaling). The BPM and steerer
magnet numbers increase from one to their maximum number in the direction of the arrows.*
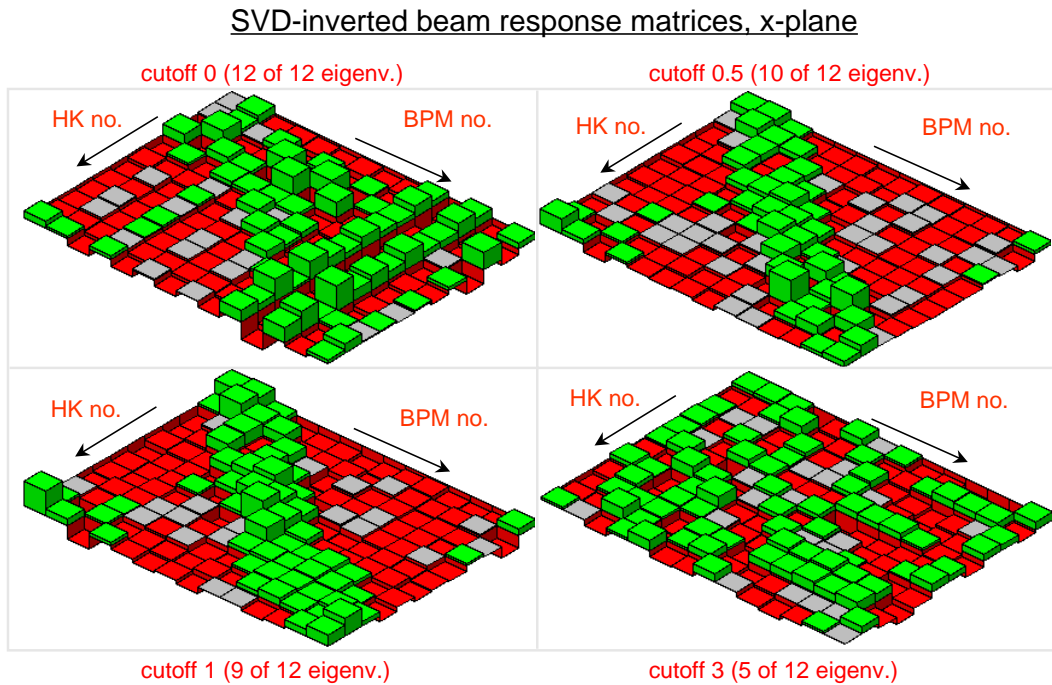
Figure 7.10: *Calculated SVD-inverted horizontal beam response matrix of the "bo-101c" optics for different SVD cutoff values. The number of eigenvalues that are used for the feedback was reduced from a maximum of 12 (left upper plot) to 5 (right lower plot) by increasing the cutoff value from 0 to 3.*

"green" and one "red" columns of the matrix for these BPMs. Such corrections are undesired and result from the non-ideal placement of the Bodo BPMs. All except two BPMs are located in vertically focussing quadrupoles ("QDs"), and all horizontal steerers are located in horizontally focussing quadrupoles ("QFs"). Therefore, the horizontal steerers may kick the beam so that it performs static zigzag oscillations around the ideal orbit. Since the existing BPMs are basically in the middle of two adjacent horizontal steerers, the resulting orbit oscillation amplitudes may be nearly invisible at the BPMs because they are close to the oscillation nodes, but the amplitudes may be very large in the QFs. Furthermore, the horizontal beam envelope is very small in the QDs and very large in the QFs (see figures A.5, A.6, A.7 and A.8 on pages 156, 157, 158 and 159). This holds especially for optics with a small emittance (which deteriorates the effect). The undesired kicks can be removed by increasing the SVD cutoff value from 0 to 0.5, which removes the two lowest eigenvalues (right upper plot). The resulting SVD-inverted matrix has a clear diagonal structure despite the non-ideal BPM positions, and BPM displacements will be corrected mainly with a few steerer magnets around the respective BPM, without the undesired zigzag patters described above. A further increase of the cutoff value (left lower plot, 9 of 12 eigenvalues used) leads to a broadening of the non-zero matrix elements around its diagonal, i.e. the algorithm tends to use more steerers with larger currents in order to correct an orbit displacement at a single BPM. This gets worse when the SVD cutoff value is increased to 3 so that only 5 of 12 eigenvalues are used (right lower plot). In this case, the matrix has no emphasis on its diagonal elements any more, therefore orbit displacements at individual BPMs will not be corrected individually, but only "averaged" over several subsequent BPMs, which is undesired for an efficient feedback system. Therefore, a horizontal cutoff value of 0.5 was
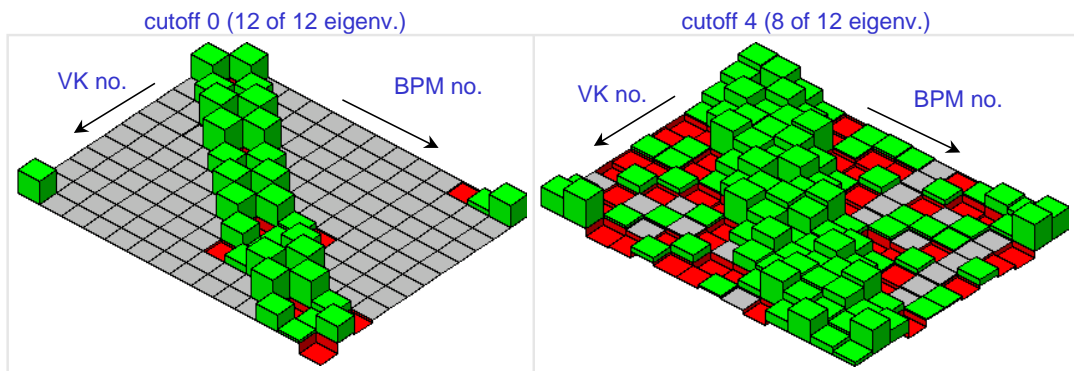
SVD-inverted beam response matrices, z-plane

Figure 7.11: *Calculated SVD-inverted vertical beam response matrices for the "bo-101c" optics. A cutoff value of zero (left plot) results in nearly ideal corrections with usually three (sometimes four) steerers around each BPM. A cutoff value of four is too large (right plot) and results in non-ideal corrections with a larger number of steerers around each BPM.*

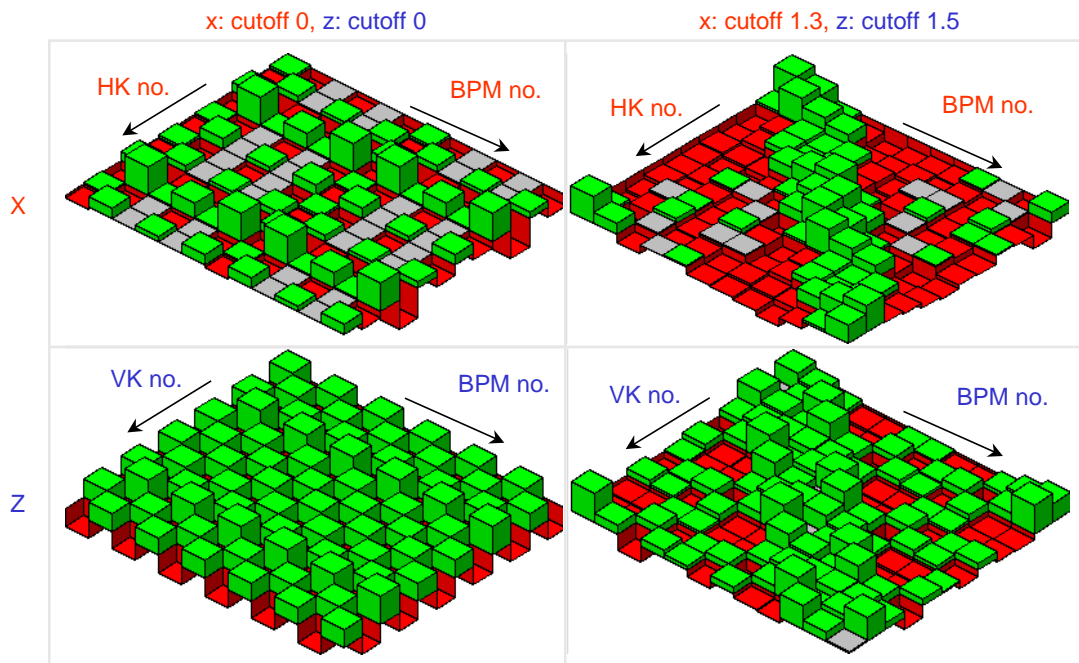SVD-inverted beam response matrices

for <u>badly</u> placed „virtual BPMs": 1 x-BPM per QD, 1 z-BPM per QF

Figure 7.12: *Calculated SVD-inverted horizontal and vertical beam response matrices with different cutoff values (left: 0, right: > 0) for non-existing "virtual" BPMs in Bodo. The BPMs have a worst-case placement: one x-BPM in each vertically focussing quadrupole ("QD"), one z-BPM in each horizontally focussing quadrupole ("QF").*

used for the beam-based test of the feedback system as described in chapter 9.

Figure 7.11 shows calculated SVD-inverted vertical beam response matrices for the "bo-101c" optics. The left plot was made for a cutoff value of zero, leading to ideal corrections via a superposition of local orbit bumps around each BPM. A cutoff value of 3 (right plot) is too large and causes the feedback to use more steerers per BPM orbit displacement than necessary, which leads to suboptimal feedback performance as already described for the horizontal plane. Therefore, a vertical cutoff value of zero was used for the beam-based tests of the feedback in chapter 9.

As described above, a horizontal cutoff value of zero for the "bo-101b"-optics leads to zigzag-shaped orbit "corrections" that might improve the orbit at the BPMs while perturbing it elsewhere. In order to investigate the dependence of this effect on the locations of the BPMs in a circular accelerator, SVD-inverted beam response matrices were calculated not for the existing BPMs, but for "virtual" BPMs that do not exist in Bodo. In order to investigate this effect separately for both transverse planes, the virtual BPMs were divided into horizontal and vertical ones.

Figure 7.12 shows the SVD-inverted beam response matrices for the "worst-case" placement of such virtual BPMs (for the "bo-101c" optics): horizontal BPMs in each QD quadrupole (upper plots) and vertical BPMs in each QF quadrupole (lower plots). A cutoff value of zero (left plots) leads to zigzag orbit corrections in both planes. According to the small eigenvalues, the resulting orbit oscillation is nearly invisible at the BPMs, i.e. a cutoff value of zero could cause the feedback to make the orbit worse at locations where the orbit cannot be measured, namely in QFs for the horizontal and in QDs for the vertical plane. This undesired effect can be eliminated more or less by increasing the cutoff values in both planes (right plots), which leads to the desired diagonal structure of the SVD-inverted matrices. However, there is still a large number of non-zero matrix elements that are far away from the diagonal, which is undesired because it makes the feedback change a large number of steerers significantly in order to correct a single orbit displacement at one BPM.

In contrast to figures 7.12, figures 7.13 show the beam response matrices (left plots) and the SVD-inverted matrices (right plots) for an "ideal" placement of virtual BPMs in Bodo: one x-BPM in each QF quadrupole and one z-BPM in each QD quadrupole. The resulting SVD-inverted matrices show an ideal diagonal structure without the zigzag BPM columns that were visible e.g. in the horizontal matrix for the real BPMs with cutoff zero or in both matrices for the worst-case virtual BPMs.

The only minor drawback even in case of an ideal BPM placement is the large number of "red boxes" in the upper right plot of fig. 7.13. The "green boxes" close to the diagonal of the matrix correspond to local orbit bumps that change the length of the beam orbit (since there is a dipole magnet between most adjacent horizontal steerers in Bodo, see eq. 2.9, chapter 2). The length change forces the beam onto a dispersion orbit outside the bump, and the SVD algorithm corrects this by systematic orbit kicks that compensate the dispersion, but only at the location of the BPMs. If an integer multiple of the RF cavity frequency of a circular accelerator matches its length, the feedback will correct the orbit nevertheless, but only at the expense of a larger number of steerers that are changed in order to perform local corrections. However, if this effect is undesired, fig. 7.14 shows that the undesired global "dispersion correction" (red boxes) in Bodo can be reduced to a small level by a horizontal SVD cutoff value of 2 (for the "bo-101c" optics), which removes the smallest SVD eigenvalue. Another possibility to
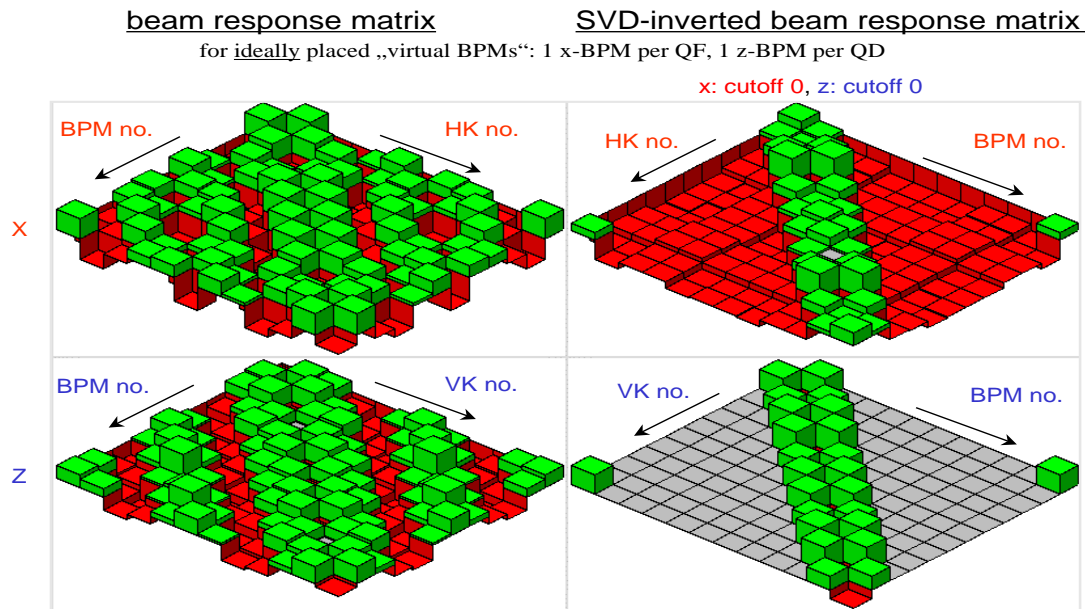
Figure 7.13: *Calculated horizontal and vertical beam response matrices and their SVD inverse for non-existing "virtual" BPMs in Bodo with zero cutoff values. In contrast to the "worst-case" BPM placement that was used for fig. 7.12, the above plots were calculated for an "ideal" placement of the virtual BPMs: one x-BPM in each horizontally focussing quadrupole ("QF"), one z-BPM in each vertically focussing quadrupole ("QD").*
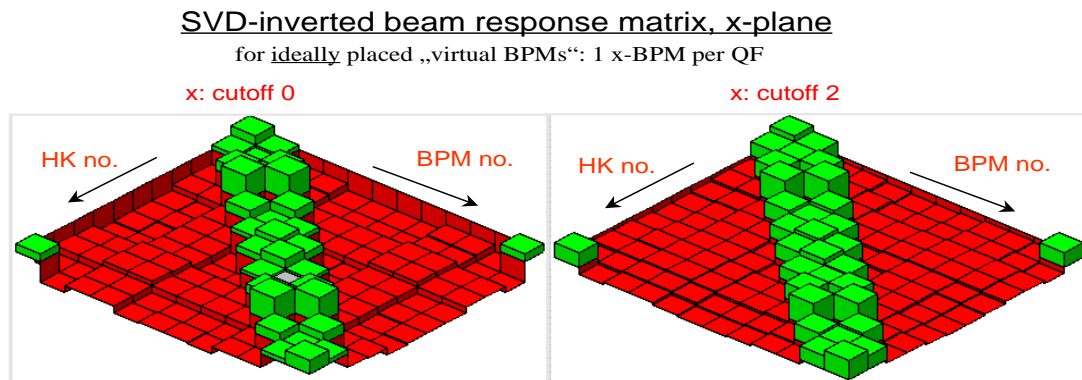

Figure 7.14: *SVD-inverted horizontal beam response matrices for ideally-placed non-existing "virtual" BPMs in Bodo (same placement as fig. 7.13). The matrix rows correspond to local orbit bumps around each BPM (green boxes). Since most orbit bumps change the length of the orbit in the rest of the ring and move the beam onto a dispersion orbit, the SVD algorithm tries to compensate this by systematic kicks of all steerer magnets (red boxes). These kicks generate a dispersion orbit with opposite sign at the location of all BPMs. In case this effect is undesired (because the feedback applies larger corrections to a large number of steerer magnets to achieve local orbit changes), the effect can be decreased by increasing the cutoff value from 0 (left figure) to 2 (right figure), which removes the smallest SVD eigenvalue.*

avoid the undesired "red boxes" in the upper right plot of fig. 7.13 is to integrate the RF cavity frequency into the SVD algorithm, so that length changes of the orbit outside a local bump can be corrected by a suitable frequency adjustment [32]. However, the cavity frequencies of Bodo and Delta are generated by the same RF generator which is already used to correct the Delta orbit, so that the Bodo RF frequency cannot be adjusted independently.

As a conclusion, the existing BPM placement in Bodo is nearly ideal for vertical orbit corrections, but not ideal for horizontal orbit corrections, since 12 BPMs are located in QD quadrupoles (that have vertical orbit correctors) and only two in QF quadrupoles (that have horizontal orbit correctors). This leads to ideal visibility of local vertical steerer magnet orbit bumps on the BPM system, due to BPMs at the beginning, at the top and at the end of the bump, and due to a local maximum of the vertical beam envelope in QDs. In contrast, some linear combinations of horizontal steerer magnet settings can cause zigzag beam oscillations that are nearly invisible at the BPMs, since most BPMs are just between two horizontal correctors and thus close to the nodes of the oscillations where the oscillating orbit crosses the ideal orbit. Furthermore, the small horizontal beam envelope at the QF BPMs amplifies noise and BPM position measurement errors in case of a horizontal feedback. Moreover, magnetic field errors are usually larger in the horizontal plane due to dipole magnet imperfections, especially in ramped storage rings with different kinds of bending magnets, therefore a large number of QF BPMs would have been better than a large number of QD BPMs. In case of a future upgrade of the Bodo vacuum and BPM system, it is recommended to install additional QF BPMs, especially at the location of the injection and extraction septum (see e.g. figures A.5, A.6, A.7 and A.8 on pages 156, 157, 158 and 159), where the septum blades may limit the mechanical beam pipe aperture and where a constant beam position is important in order to achieve high beam charge transfer during beam injection and extraction.

Despite some performance limitations due to the existing BPM distribution, a suitable choice of the SVD cutoff value can avoid the undesired effects of an SVD-based orbit feedback in the horizontal plane that result from the suboptimal BPM placement. The SVD inverse for an optimised cutoff value allows a horizontal orbit feedback in Bodo that corrects the orbit properly at the locations of the existing BPMs (see measurements in chapter 9).

Chapter 8

Betatron Tune Measurement and Feedback at Bodo

Chapter 6 described the architecture and tests without electron beam of the DSP-based betatron tune measurement and feedback system of Bodo. This chapter presents the beam tests of this system with the Bodo synchrotron. The measurements in the following sections investigate the performance of the feedback, e.g. if and how precise and fast betatron tunes can be corrected. Furthermore, section 8.2.1 compares tune measurements that were made with the normal Bodo tune measurement system using BPM pickup buttons (see chapter 6) to measurements of the Bodo betatron tunes by the beam loss monitors ("BLMs") of the DeltaDSP system (see section 5.3.3). The BLM-based tune measurement method allows a beam-based verification of the calibration measurements of chapter 6.

8.1 Measurement of Bodo Betatron Tunes With and Without Feedback

Figure 8.1 shows the nominal Bodo energy ramp for the "bo-101c" optics that is usually used for synchrotron radiation machine shifts. In general, the energy ramp as well as all power supply ramps of Bodo are programmable and arbitrary. The ramp shown in fig. 8.1 was made as short as possible with the available magnet power supplies, in order to maximise the Delta injection rate. It is even slightly shorter than the minimum ramp period for 1.5 GeV ramps that was calculated in section 3.2 (based on the maximum power supply voltages and currents), but
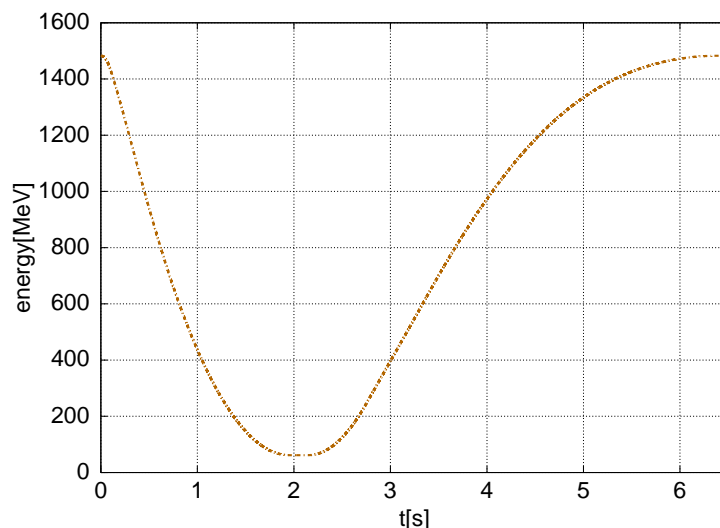

Figure 8.1: *Bodo energy ramp for 1.5 GeV synchrotron light source operation of Delta. The ramp curve starts at 1483 MeV (at $t = 0$ s), ramps down to 62 MeV where the beam is injected (at $t = 2.1$ s) and ramps up again to 1483 MeV where the beam is extracted (at $t = 6.455$ s).*

124

the safety margins of the power supply voltages are still sufficient for reliable operation of the machine. The ramp was chosen for the tests of the tune measurement system since it has nearly the fastest energy change rates that are possible in Bodo and will therefore cause very large energy-related betatron tune drifts (e.g. due to different current regulation systems of dipole and quadrupole magnet power supplies).

The start of the ramp and the injection of the beam are synchronised with a 10 Hz signal that is derived from the 50 Hz power line frequency, therefore the nominal repetition period of the ramp is a multiple of 100 ms. In case of fig. 8.1, the ramp is restarted at a nominal period of 6.5 s. As the 50 Hz power line frequency is not perfectly stable and jitters up to several milliseconds per ramp period, the ramp period changes slightly from cycle to cycle. In order to avoid energy jitter during injection (at $t = 2.1$ s) due to the varying power line frequency, the energy ramp has an injection plateau between $t = 1.95$ s and $t = 2.15$ s. In contrast, the beam extraction is triggered directly by the DeltaDSP board of the bending magnet power supply (without additional power line synchronisation), with negligible timing and energy jitter.

8.1.1 Ramps With Manual Quadrupole Focussing Corrections

Figure 8.2 shows the horizontal and vertical fractional betatron tune frequencies f_x and f_z during the Bodo energy ramp for the "bo-101c" optics (resp. for a very similar optics with slightly different and varying tune frequencies). The red (f_x) and pink (f_z) curves were measured for
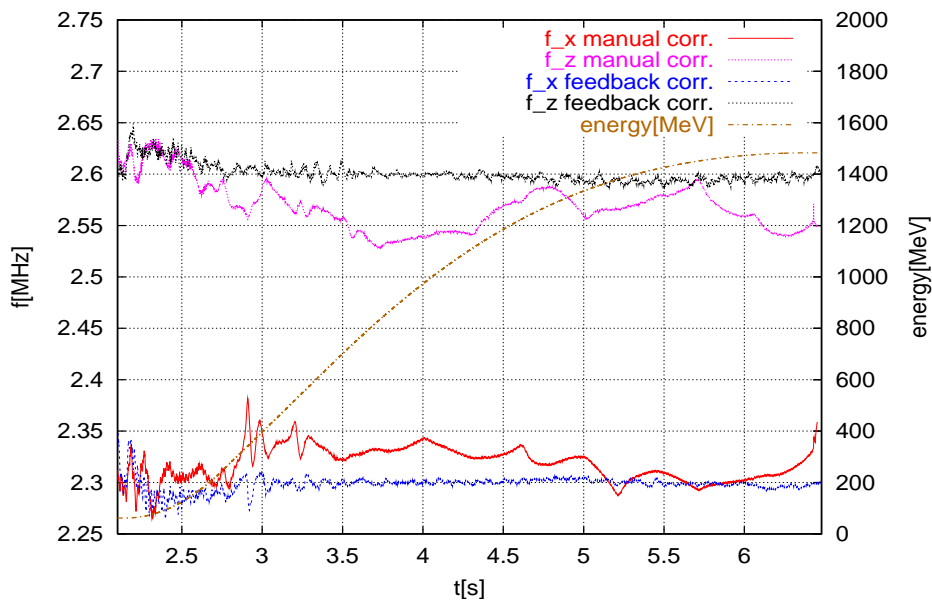

Figure 8.2: *Comparison of a Bodo energy ramp with manually corrected fractional betatron tune frequencies (red curve: f_x, pink curve: f_z) to a ramp where the tunes were corrected in real-time by the DSP-based tune feedback system (blue curve: f_x, black curve: f_z). The feedback used PI control loops to correct f_x and f_z, with nominal values of $f_x = 2.3$ MHz and $f_z = 2.6$ MHz. The beam is injected at $t = 2.1$ s ($E = 62$ MeV) and extracted at $t = 6.455$ s ($E = 1483$ MeV).*

a ramp where the tunes had been corrected manually by changing quadrupole set currents iteratively in the magnet ramp data file (as described in section 6.2.4). The blue and black curves are the measured fractional betatron tune frequencies for the same energy ramp without any manual corrections, but with the corrections being performed in real-time by the betatron tune feedback system. The nominal fractional tune frequencies for the feedback were set to $f_x = 2.3$ MHz and $f_z = 2.6$ MHz. Fig. 8.2 demonstrates that the tune feedback system works, and that the resulting tune corrections are more precise than the manual corrections.

The only requirements for the feedback system are that the real machine tunes at injection are sufficiently close to the desired tunes, so that e.g. no destructive resonance has to be crossed when the feedback moves the tunes to the nominal values. While the difference of the feedback-corrected tune frequencies to the desired values in fig. 8.2 is usually below $\pm$ 10 kHz, the difference is nearly 50 kHz during the first 500 ms after beam injection. This can be explained by the very large relative energy change during the first part of the ramp. Therefore, this part of the ramp is more sensitive e.g. to differences in the current regulation time constants of quadrupole and dipole magnets. Quadrupole current variations as small as 0.1 % lead to tune variations of 34 kHz for horizontally focussing quadrupoles ("QFs") and 23 kHz for horizontally defocussing quadrupoles ("QDs").

8.1.2   Ramps Without Manual Quadrupole Focussing Corrections

Figure 8.3 shows the betatron tune frequencies of a Bodo ramp without any manual quadrupole power supply current corrections, both with and without tune feedback. Without feedback, the
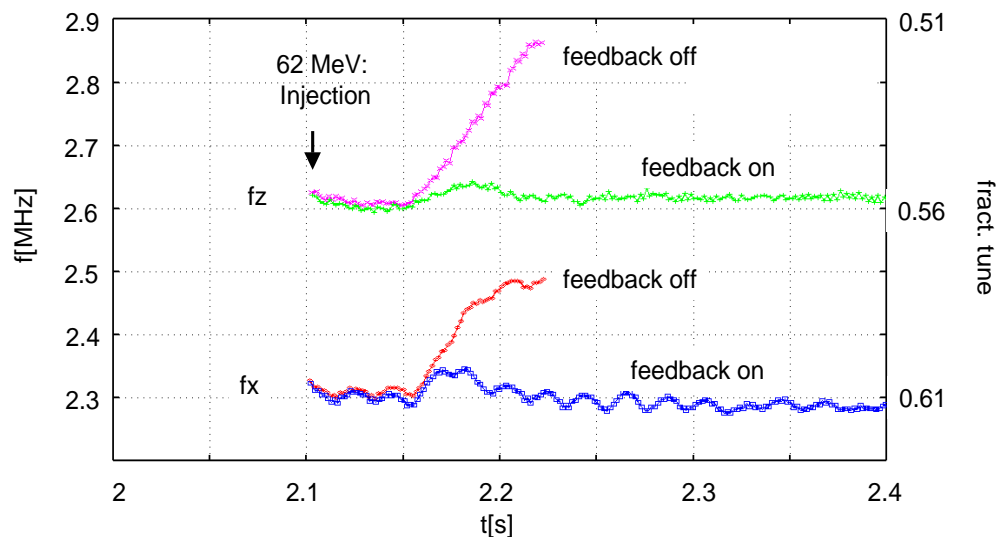

Figure 8.3: *Fractional Bodo betatron tune frequencies for an energy ramp without any manual quadrupole magnet current corrections (as described in section 8.1). Without tune feedback, the betatron tune frequencies drift up to 4 kHz per millisecond. The measurements stops 150 ms after injection (due to low beam current), just before the vertical tune crosses a half integer resonance and the beam gets completely lost. In contrast to fig. 8.2, the feedback used PID control loops, with nominal values of $f_x = 2.3$ MHz and $f_z = 2.6$ MHz. It is able to reduce the tune drift to less than 50 kHz just after injection, and to less than 20 kHz afterwards.*

betatron tunes drift heavily towards resonance lines, and the beam gets lost about 150 ms after injection. The tune feedback manages to keep the beam in the machine, with differences to the desired tune frequencies of less than 50 kHz after injection and less than 20 kHz afterwards. The green f_z curve in fig. 8.3 remains always slightly above the nominal value of $f_z = 2.6$ MHz, which indicates that the feedback has to fight a frequency drift to higher values. Since both betatron tunes of this optics ("bo-101c") are below an integer, increasing fractional tune frequencies are equivalent to decreasing tunes or too small quadrupole focussing strengths. It should be noted that the feedback in fig. 8.3 used a PID control loop, which explains the faster tune correction after injection compared to fig. 8.2, which shows a tune feedback measurement with a PI control loop because the DSP feedback software was not yet complete at the time of measurement.

Fig. 8.3 also shows a 50 Hz modulation of the betatron tunes, which can be caused both by 50 Hz current variations of the quadrupole power supplies and by 50 Hz orbit perturbations due to external stray fields and 50 Hz power supply ripple, since orbit drifts lead to tune drifts in case of sextupole magnetic fields and a non-ideal closed orbit. The resulting 50 Hz tune modulation cannot be compensated by the tune feedback, as the bandwidth of the system is too low. However, the amplitude of the modulation is not critical for synchrotron light source operation of the machine, and may only be a minor limitation for some test machine experiments at Bodo.

Figures 8.4 show tune measurements for quadrupole current ramps without manual corrections, with different tune feedback set values for f_x (2.2, 2.3 and 2.4 MHz) and f_z (2.5, 2.6 and 2.7 MHz). The tunes at the time of injection are identical, and the feedback drives them to their nominal values in some ten milliseconds. The remaining small differences between desired and actual tune frequency are mainly identical for all three f_x resp. f_z curves, e.g. the short undershoot of the f_x curves just after $t = 2.9$ s. These undershoots are caused by a systematic design imperfection in the current regulation systems of all quadrupole power supplies. As mentioned in section 3.2, they are switched-mode power supplies that connect a common input voltage source periodically to its output circuit via a MOSFET. Due to the design of the output circuits, the output voltage is a nonlinear function of the MOSFET duty cycle. When the current increases and the location of largest nonlinearity of this function is reached, the current regulation electronics that was actually designed and optimised for linear functions generates a short voltage overshoot, which results in a tune overshoot and thus a fractional betatron frequency undershoot, since the tunes are below an integer.

While it is possible to inject beam at one betatron tune and to use the feedback to move the beam optics to another betatron tune (as shown in figures 8.4), it is recommended to adjust the quadrupole currents at the time of injection so that the resulting tune is identical to the desired tune (up to $\pm$ 20 kHz or better), at least until the beam has reached some 100 MeV energy. At the time of injection, the beam has a very large emittance and a very large energy spread. Changing tune at low energies usually leads to increased beam loss, e.g. due to a decrease of the energy acceptance (via non-zero chromaticity) or due to a decrease of the transverse acceptance, since changing tunes can change both beam cross section and beam orbit significantly for off-axis beams. Depending on the beam optics, a tune stabilisation of 20 kHz or better is usually sufficient to avoid such effects, and the tune feedback can easily achieve this tune stability. At high energies above 1 GeV, adiabatic and synchrotron radiation ramping in all three beam dimensions make the beam more insensitive to beam optics changes. In case a future optics upgrade of DELTA requires a different Bodo beam optics with different betatron tunes
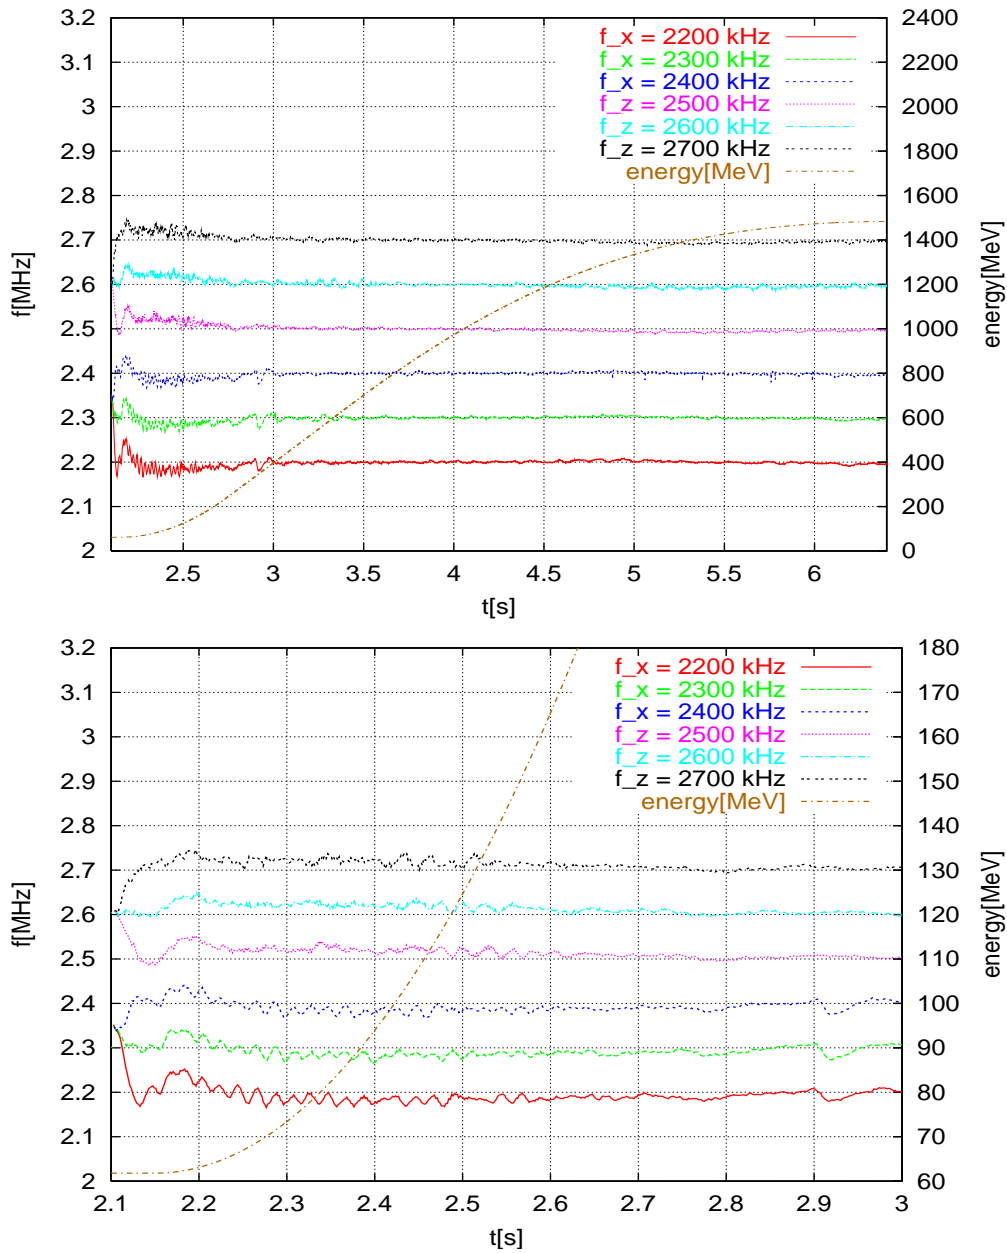
Figure 8.4: *Fractional Bodo betatron tune frequencies for an energy ramp (see fig. 8.1) without any manual quadrupole magnet current corrections. The plots are superpositions of three measurements with activated tune feedback and different set values for f_x (2.2, 2.3 and 2.4 MHz) and f_z (2.5, 2.6 and 2.7 MHz), with the lower plot being a time scale zoom of the upper plot. All measurements were performed for the same uncorrected quadrupole ramp curve files, therefore the tunes at the time of injection ($t = 2.1$ s) are identical, and the feedback drives them to their nominal values within some ten milliseconds after injection.*
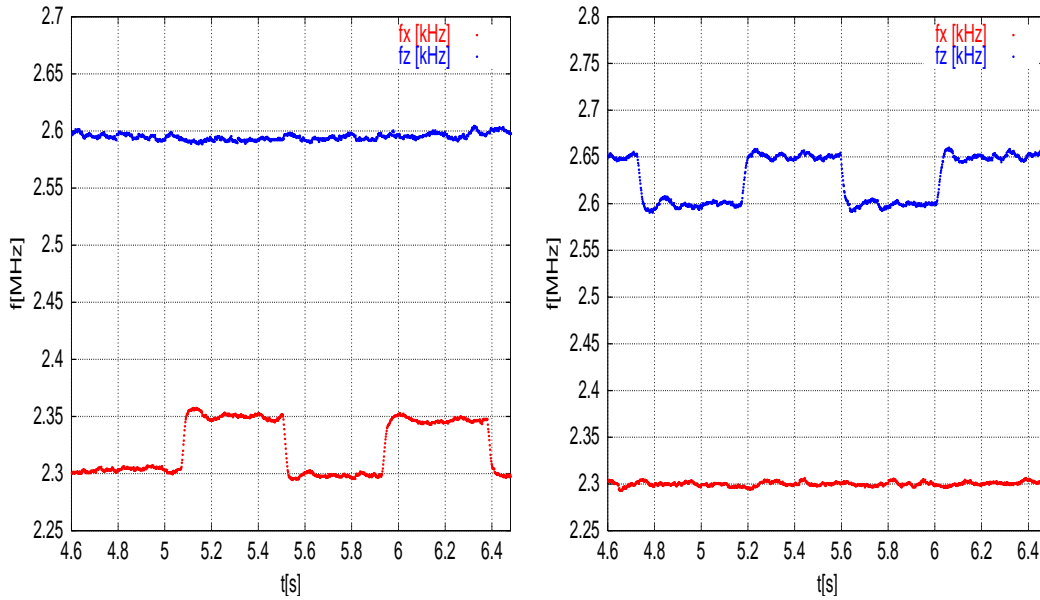
Figure 8.5: *Left plot: Step response of the tune feedback system for f_x set value changes between 2.3 MHz and 2.35 MHz (with a constant f_z set value of 2.6 MHz). Right plot: step response for f_z set value changes between 2.6 MHz and 2.65 MHz (with a constant f_x set value of 2.3 MHz).*

for injection and extraction, it is recommended to keep the injection beam optics and tunes constant up to about 1 GeV and then to apply the extraction optics via suitable interpolation, in contrast to a linear interpolation of the two optics from injection to extraction.

8.1.3   Feedback Response Optimisation and x-z Crosstalk

Figures 8.5 show the step response of the tune feedback system for a sudden change of the fractional tune frequency set values f_x (left plot) and f_z (right plot) of the feedback system. The set values were changed manually between 2.3 MHz and 2.35 MHz for f_x and 2.6 MHz and 2.65 MHz for f_z by a graphical user interface (introduced in fig. 6.14, section 6.3.7) during the usual 1.5 GeV energy ramp for the "bo-101c" optics. The PID parameters of the vertical and horizontal tune feedback loop (see section 6.3.7) were adjusted so that the tune feedback moved the tune frequencies to the desired values at maximum speed (in about 20 ms, see fig 8.6), with minimum overshoots or oscillations of the tune curves after the step. Nearly all measurements in this chapter were made with this optimised PID parameter setting (which is also shown in fig. 6.14, section 6.3.7). The EMW GUI of the tune feedback system can save, restore and scale settings of the GUI, which is useful for fast comparison and optimisation of feedback loop set values.

As discussed in section 6.3.6.2, the tune feedback algorithm uses two separate feedback loops, with f_x being controlled only by the QFs and f_z being controlled only by the QDs (instead of using a 2x2 matrix as described in section 6.3.6.2). According to fig. 6.1 in section 6.2, a change of the focussing strengths of either the QF or the QD quadrupole families changes the betatron tunes in both planes. For the usual Bodo optics, the QFs change mainly f_x and the
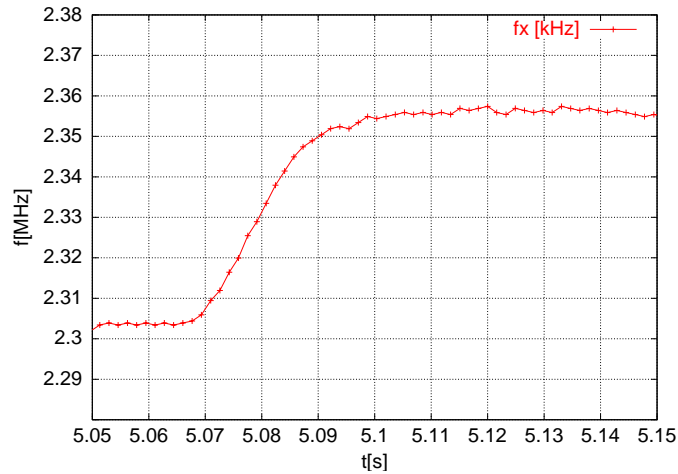
Figure 8.6: *Step response of the tune feedback system for an f_x set value change from 2.3 to 2.35 MHz (with a constant f_z set value of 2.6 MHz). The feedback makes a slight overshoot up to 2.357 MHz and takes about 50 ms to correct this overshoot (see left plot at $t = 5.1$ s in figures 8.5).*

QDs change mainly f_z, therefore the step response of the feedback for sudden tune frequency set value changes in one plane should lead to a small change of the tune frequency in the other plane. Figures 8.5 show that the feedback is able to compensate this, so that a change of the f_x or f_z set value in one plane does not lead to a significant change of the tune in the other plane. Therefore, the optics-independent approach of two separate feedback loops has no disadvantage compared to a model-based approach with coupled feedback loops (that predicts the f_x and f_z changes, see section 6.3.6.2).

Furthermore, the same PID loop setting was used successfully to correct the tunes reliably and efficiently for completely different beam optics, which makes the system a robust and easy-to-use tool both for normal and test machine operation of Bodo. While the previous manual tune correction required many hours, new energy ramps of arbitrary shape can now be tune-corrected automatically in real-time and require only a moderate adjustment of the tunes when the beam is injected. Furthermore, the feedback can be used with manual control of the beam energy (e.g. with a slider on a GUI), which allows test machine operation with stable betatron tunes during arbitrary non-periodic changes of the beam energy.

8.2   Beam Loss Based Betatron Tune Measurement

In order to verify that the tune measurement hardware and software that analyse the coherent Bodo beam oscillations work correctly, tune measurements with the usual "RF-based" DSP-based Bodo betatron tune measurement system (as described in chapter 6) were compared to tune measurements that used the Bodo beam loss monitoring system. When the excitation kicker frequency is identical to one of the fractional betatron tune frequencies, the resulting coherent betatron oscillations cause beam loss if the kicker amplitude is sufficiently large. The beam loss was detected with suitable monitors during periodic frequency sweeps of the kicker. The beam loss monitors ("BLMs") that are installed in Bodo were developed by K.
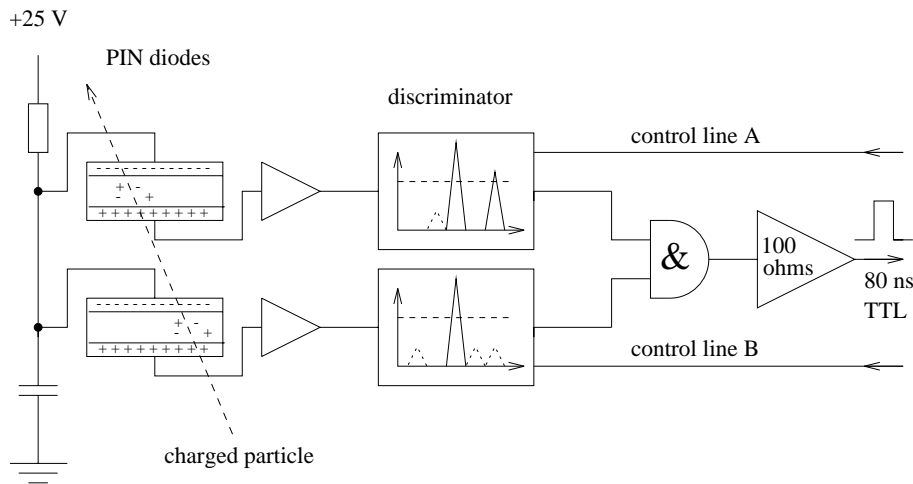
Figure 8.7: *Principle of operation of the beam loss monitors that are used in Bodo. Electrons that hit the Bodo beam pipe generate electromagnetic showers. Electrons and positrons of these showers are detected by coincident generation of electron hole pairs in the depletion zone of two PIN photodiodes. Two discriminators with an "and" gate suppress noise pulses and generate a TTL output pulse of 80 ns length if a charged particle was detected.*

Wittenburg and M. Swars [103]. Their application for the detection of electron beam loss in Bodo was investigated in detail in ref. [60], therefore the following section will only give a brief description of their principle of operation.

8.2.1 The PIN Diode Beam Loss Monitors

If an electron in Bodo hits the stainless steel beam pipe of 3 mm thickness (usually under flat angles of some mrad [60]), it generates an electromagnetic shower of electrons, positrons and photons. As long as the initial hit angle remains small, increasing beam energies in Bodo lead to roughly proportionally increasing numbers of electrons and positions in the electromagnetic shower, while the average energy of the generated ionising particles changes slightly from 13 MeV (at 70 MeV beam energy) to 14 MeV (at 1.5 GeV beam energy) [60].

Figure 8.7 shows the principle of operation of the beam loss monitors that are used in Bodo. They consist of two flat PIN photodiodes (with an active area of 1.5 cm^2) that are connected to a 25 V power supply. The polarity is chosen so that the diodes are not conducting, with a charge depletion zone of about 200 μm thickness. If a charged electron or positron of some MeV energy (also called minimal ionising particle or "MIP") crosses both diodes, the MIP creates electron-hole pairs in the depletion zone that generate a current in the amplifier which is connected to the diode. If the current pulses of both diodes are coincident and have an adjustable minimum height (determined by the discriminators in fig. 8.7), the BLM will generate a TTL pulse of about 80 ns length.

The TTL pulses of the BLMs are detected by counters on FPGAs of two DeltaDSP boards (see fig. 5.5, page 70), and the counter values are sampled (i.e. written to readout registers in the FPGA) synchronously to the sampling of the ADCs (e.g. BPM ADCs) of the DeltaDSP system in Bodo. The resulting beam loss data is stored in the DRAM of the DeltaDSP boards. The data

can be archived in files and visualised as a function of time via EMW GUIs. Since DeltaNet synchronises the global clocks of all DeltaDSP boards, the BLM and ADC measurement data of different boards have a common timebase, which allows the analysis of coincidences in the data at a time resolution below 1 μs.

8.2.2   Tune Measurement with BLMs

In order to compare RF-based tune measurements to BLM-based tune measurements, the amplitude of the excitation kicker of the Bodo tune measurement system was increased until a beam loss monitor on top of the beam pipe in the quadrupole 3QDI (see fig. 3.1, page 18) detected significant beam loss. The tune excitation kicker frequency sweeps were unidirectional (i.e. up sweeps). As shown in fig. A.6 (on page 157) and fig. A.8 (on page 159), the four QDI quadrupoles of Bodo have the largest vertical beam envelope for the "bo-101c" optics that was used for all measurements in this section. According to the BPM system, the vertical orbit perturbation in Bodo had an absolute maximum in the 3QDI quadrupole, therefore electrons of a beam with vertically increasing cross section will preferably hit the beam pipe at the top of this quadrupole. Measurements and beam loss simulations in ref. [60] have demonstrated that electrons which hit the beam pipe in a quadrupole can be localised by a BLM in that quadrupole, although electrons may also be scattered back into the beam pipe and get finally lost several meters downstream [104].

The red curves of both plots in fig. 8.8 show the integrated number of MIPs that were detected by the BLM (as a function of time) at the top of the beam pipe in the 3QDI quadrupole. The pink resp. blue vertical marker lines indicate the time at which the RF-based tune measurement system detected a horizontal (x) resp. vertical (z) tune peak. As expected, the detection of z tune peaks coincides with vertical beam loss, since the excitation kicker generates coherent vertical beam oscillations that increase the vertical electron oscillation amplitudes until the beam scrapes at the 3QDI quadrupole and an increased number of electrons get lost. Consequently, the horizontal tune peaks do not coincide with vertical beam loss.

As explained in section 6.3.4, a test of the RF-based tune measurement system with RF frequency generators showed a tune peak shift that depends on the sweep frequency. The tune measurement DSP board is able to correct this shift, based on calibration measurements with an RF generator (see section 6.3.4). The lower plot in fig. 8.8 shows a time scale zoom of the upper plot. The beam loss per time (i.e. the derivative of the red curve) should have a maximum when the excitation kicker frequency is equal to the fractional betatron tune frequency, since this frequency causes the maximum growth rate of the resulting coherent vertical beam oscillation and thus the largest beam loss. The black (right) vertical marker line indicates the time when the RF-based tune measurement system detected a tune peak without activated tune peak correction, the blue (left) vertical marker line indicates the respective time with activated tune peak correction. As expected, the blue line coincides well with the largest increase of the beam loss rate in the symmetry point of the red curve. This confirms that the RF-based tune measurement works correctly, and it also shows that betatron tunes in Bodo can be measured with BLMs instead of BPMs.
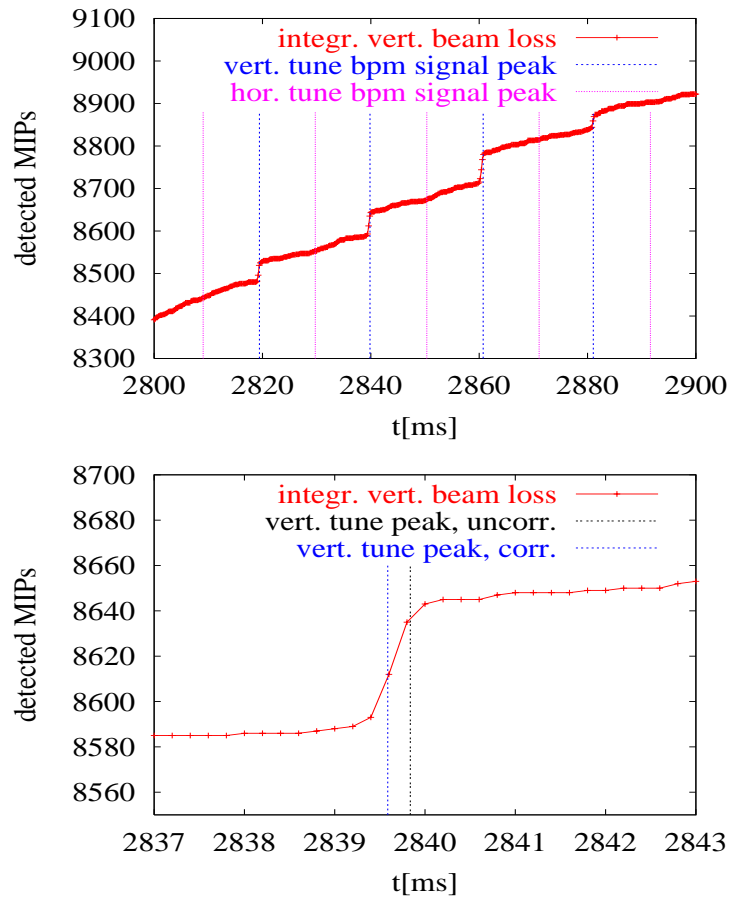
Figure 8.8: *The red curves are the integrated number of MIPs that were detected by a beam loss monitor at the top of the beam pipe in the vertically focussing Bodo quadrupole 3QDI. The lower plot is a time scale zoom of the upper one. The measurement was made while the normal tune measurement system was active, with a very large excitation kicker amplitude in order to generate large beam loss rates when the kicker excited coherent vertical or horizontal beam oscillations. Blue resp. pink vertical marker lines indicate the time when the tune measurement system detected a vertical resp. horizontal tune peak. The beam loss increases only for vertical tune peaks, i.e. when the excitation kicker generates coherent vertical beam oscillations that cause beam loss in the vertical plane. The blue (left) and black (right) vertical marker lines in the lower plot indicate the time when the tune measurement system detected a tune peak with (blue) and without (black) peak frequency correction.*

Chapter 9

Orbit Measurement and Feedback at Bodo

Chapter 7 described the architecture of the global orbit measurement and feedback system of Bodo, as well as some tests of the measurement system with RF generators. This chapter presents orbit measurements and tests of the feedback system with the Bodo electron beam. Section 9.1 contains measurements of the Bodo orbit during 1.5 GeV ramp cycles with and without global SVD-based orbit feedback. When the orbit feedback system changes the orbit of Bodo, sextupole magnet fields cause undesired betatron tune shifts. Section 9.2 demonstrates the correction of such tune shifts by simultaneous operation of the global orbit and the betatron tune feedback system. Section 9.3 investigates if and how well a DSP-based local orbit feedback at 1.5 GeV is able to correct periodic orbit perturbations that originate from a source outside the local orbit bump of the feedback. The frequency response of this local feedback was obtained by perturbing the orbit via sine modulation of a steerer magnet field. This measurement is prototypical for the most simple future application of the DSP system in Delta, namely an orbit feedback that stabilises the beam position at a single beamline.

It should be noted that the beam positions in the figures in this chapter are the scaled output voltages of the BPM electronic front-ends, without corrections of nonlinearities by the DSPs for an extremely off-centered beam. However, the BPM nonlinearities are no problem for the orbit feedback, since it usually corrects the beam positions to zero or values of a few millimeters where the nonlinearities are negligible (see section 3.5). If the feedback tries to correct an extremely off-centered orbit to zero, it is of little relevance if the real beam position is 9 mm or 10 mm, since the feedback will still move the beam position in the right direction, only with slightly different speed (due to the PID algorithm). Once the beam position has been roughly corrected to a few millimeters, the nonlinearities are negligible and the feedback can correct the position precisely to the desired value.

Furthermore, Bodo has no beam-based calibration ("BBC") system for its BPMs like Delta, where such a system can measure the BPM signal offsets with respect to the center of the quadrupole magnets that define the ideal orbit [24]. Such offsets of the measured beam position with respect to the real beam position can be caused by a non-centered beam pipe, mechanical tolerances of BPM pickup buttons and component tolerances of the BPM front-end electronics. In case of Delta, the overall BPM offsets for the horizontal and the vertical plane are typically $dx = 0.4 \pm 0.5$ mm and $dz = 0.5 \pm 0.7$ mm [24], with drifts in the order of 0.1 to 0.2 mm over several months. Since Bodo has the same beam pipe and quadrupole magnet design as Delta, the BPM systems of both rings should have similar offsets. However, the orbit in Bodo is not as critical as in Delta, where a precisely centered beam is important e.g. in order to reduce the emittance coupling and thus improve the brightness of the synchrotron radiation [24]. Therefore, systematic position measurement errors in the order of 0.5 mm in Bodo are tolerable.

9.1 Global Orbit Feedback During Ramp Cycles

Figure 9.1 shows a GUI with the Bodo orbit during a Bodo energy ramp with and without SVD-based global orbit feedback. When the "run" or "single" shot button of the GUI is pressed, the DeltaDSP boards record and store the BPM data during one ramp cycle in one of two internal DRAM memory buffers ("traces") of the DSP board. Trace A (thick lines) in the figure was recorded with operational global orbit feedback, trace B (thin lines) with orbit feedback switched off. The GUI allows to adjust the offset ("T0") and scaling ("DT", time between two BPM ADC samples) of the time axis similar to an oscilloscope. For cost reasons, Bodo is equipped with BPM electronic front-ends that are spare parts of Delta. Since only 13 of them were available, the data of BPM4 in the GUI is missing.

9.1.1 Horizontal Plane

Without orbit feedback, the orbit at BPM9 (upper pink line in fig. 9.1) exceeds $x = 10$ mm. Due to the nonlinearity of the BPM system that the GUI does not account for, the actual x beam position is more than 10 percent larger [24] (see fig. 3.9, page 27). When the feedback is switched on at $t = 2.5$ s, it reduces the horizontal orbit perturbation significantly. Although the set values for the desired BPM positions are zero, the x beam positions with operational feedback are systematically larger than zero. This indicates that the RF frequency of Bodo is too low, so that the feedback cannot achieve a centered orbit, but only force the beam onto a dispersion orbit. While both Bodo and Delta were originally designed for an RF frequency of 499.650 MHz, it turned out that the beam in Delta is centered at about 499.820 MHz. This suggests systematic errors during the alignment of the magnets in Bodo and Delta, with a larger effect on Bodo. A precise measurement of the Bodo RF center frequency would require a beam-based calibration of the BPM system in order to differentiate between systematic measurement errors of the BPM system and real orbit perturbations due to a wrong RF frequency.

At $t = 4.5$ s, the uncorrected horizontal beam positions in fig. 9.1 start to drift systematically, e.g. at BPM11 (lower thin blue line) or BPM13 (upper thin blue line). The beam energy increases by a factor of 19 between $t = 2.1$ s and $t = 4.5$ s ($E = 1.2$ GeV), but only by a factor of 1.24 between $t = 4.5$ s and extraction (see fig. 8.1, page 124). Therefore the drift cannot originate from energy-related effects like steerer magnets that were not ramped proportionally to the beam energy, but that were adjusted empirically in order to maximise the injected beam current. However, $t = 4.5$ s coincides with the beginning saturation of the dipole magnet iron, which indicates that the orbit drift is caused by the two different types of dipole magnets in Bodo (see fig. 3.1, page 18). The 10 degree dipole magnets ("short dipoles") are about half as long as the 20 degree dipoles ("long dipoles"), therefore saturation of the magnet edges may change the absolute effective length of the dipoles in the same way, but the relative length and thus the bending angle varies. A comparison of the measured orbit drift with the theoretical beam response of the short dipoles showed that the dipoles were about 4 ± 0.4 mrad too weak at 1.5 GeV compared to 1 GeV [38]. The value was confirmed by magnetic field measurements [41]. As both short and long dipoles are connected in series, the short dipoles were equipped with additional coils which are connected in series to a dedicated ramped power supply. As the field error is known and systematic, the field correction for the short dipoles is not integrated into the orbit feedback, but the field is usually corrected by a predefined current curve based on the measured orbit drifts.
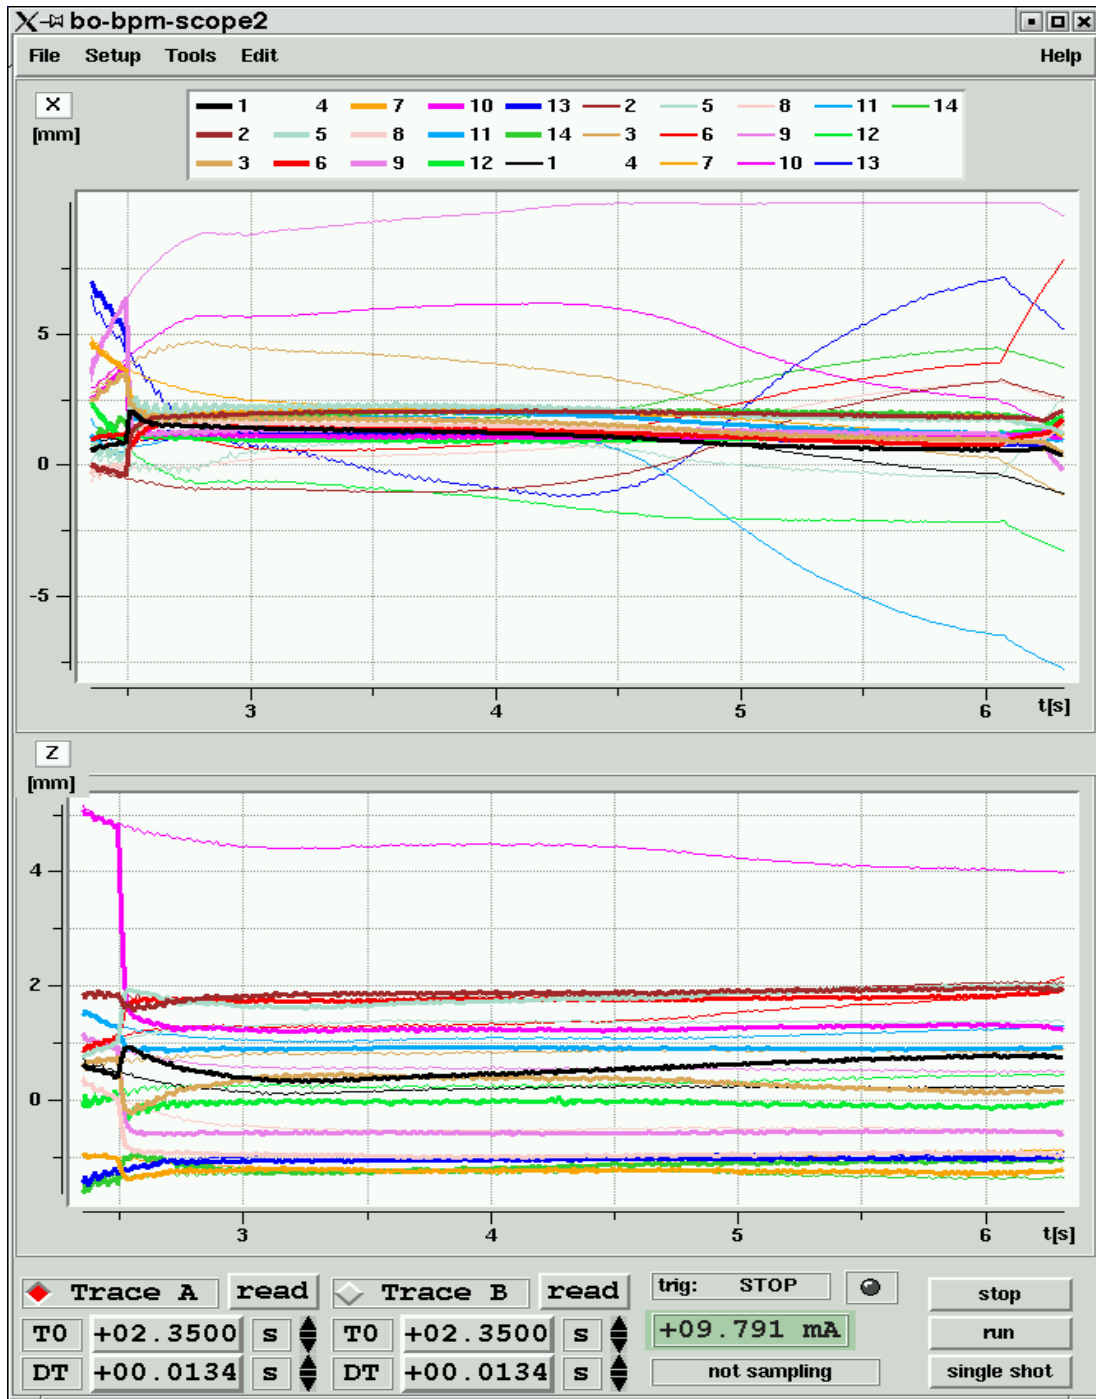
Figure 9.1: *Horizontal (x) and vertical (z) beam positions in Bodo as a function of time, with orbit feedback switched off (thin lines) and with global orbit feedback switched on (thick lines). The orbit was measured during a nominal 1.5 GeV ramp (see fig. 8.1, page 124) with the "bo-101c" optics. The beam is injected at $t = 2.1$ s ($E = 62$ MeV) and extracted at $t = 6.455$ s ($E = 1.483$ GeV). The feedback is switched on at $t = 2.5$ s. All 12 horizontal, but only 3 vertical steerer magnet power supplies (for VK8, VK9, VK10, see fig. 3.4) were available.*

The systematic linear drift of the horizontal orbit without feedback at $t = 6.1$ s in fig. 9.1 is caused by a local orbit bump that is used for beam extraction and that was not switched off during feedback tests. Obviously, the feedback is able to reduce this "orbit perturbation" to a fraction of the initial amplitude.

9.1.2 Vertical Plane

In contrast to the horizontal plane, only 3 of 12 vertical correctors are equipped with power supplies (for cost reasons), which is not sufficient to correct the orbit at all BPMs. Since BPM10 has the largest vertical orbit perturbation (see fig. 9.1, lower plot, upper pink line), the three power supplies were connected to three adjacent steerer magnets (VK8, VK9 and VK10) so that the feedback can correct the orbit at BPM10 (which is close to VK9) with a local orbit bump. It should be noted that the horizontal and vertical plots in fig. 9.1 have different scalings, and that nearly all BPMs are located in QD quadrupoles, so that the measured beam positions show the maxima of the vertical orbit, while the horizontal orbit in QFs without BPMs are typically three times higher than the visible orbit in the QDs (for the "bo-101c" optics, see beam envelopes in fig. A.7 and A.8). The maximum vertical orbit perturbations with feedback are about ± 10 % of the vertical aperture, therefore the remaining vertical orbit perturbations are tolerable, with negligible degradation of the beam current and machine performance with respect to a perfectly centered orbit.

9.1.3 Correction Speed at Feedback Start-Up

Figure 9.2 shows the horizontal and vertical Bodo orbit with operational feedback (thick lines) and with feedback switched off (thin lines). The feedback is switched on at $t = 2.5$ s. The data were recorded for a different steerer magnet setting (with feedback switched off), therefore the orbit is slightly different from the values in fig. 9.1. The horizontal feedback requires about 20 ms in order to correct the main perturbations, the vertical feedback is slightly slower due to a different PID loop adjustment.

9.2 Simultaneous Operation of Orbit and Betatron Tune Feedback

As mentioned in section 6.2.3, orbit drifts may cause betatron tune drifts, e.g. via sextupole components of the Bodo steerer magnets or due to the dedicated sextupole magnets of Bodo. Furthermore, variations of quadrupole focussing strengths cause orbit drifts if the orbit is not centered. These effects may result in undesired crosstalk between the global orbit feedback and the betatron tune feedback of Bodo. In order to investigate the performance and crosstalk of both feedbacks during simultaneous operation, the global orbit feedback was switched off and on while the betatron tune feedback was operational. The measurement was performed during a nominal 1.5 GeV energy ramp with the "bo-101c" optics. When the orbit feedback is switched off, the steerer magnet currents follow a predefined ramp curve, with an immediate change of the current set value. When the orbit feedback is switched on, the feedback PID loops change the currents continuously until they have reached the values obtained from the SVD algorithm. This crosstalk measurement is a worst-case scenario, since the set values for orbit and tunes
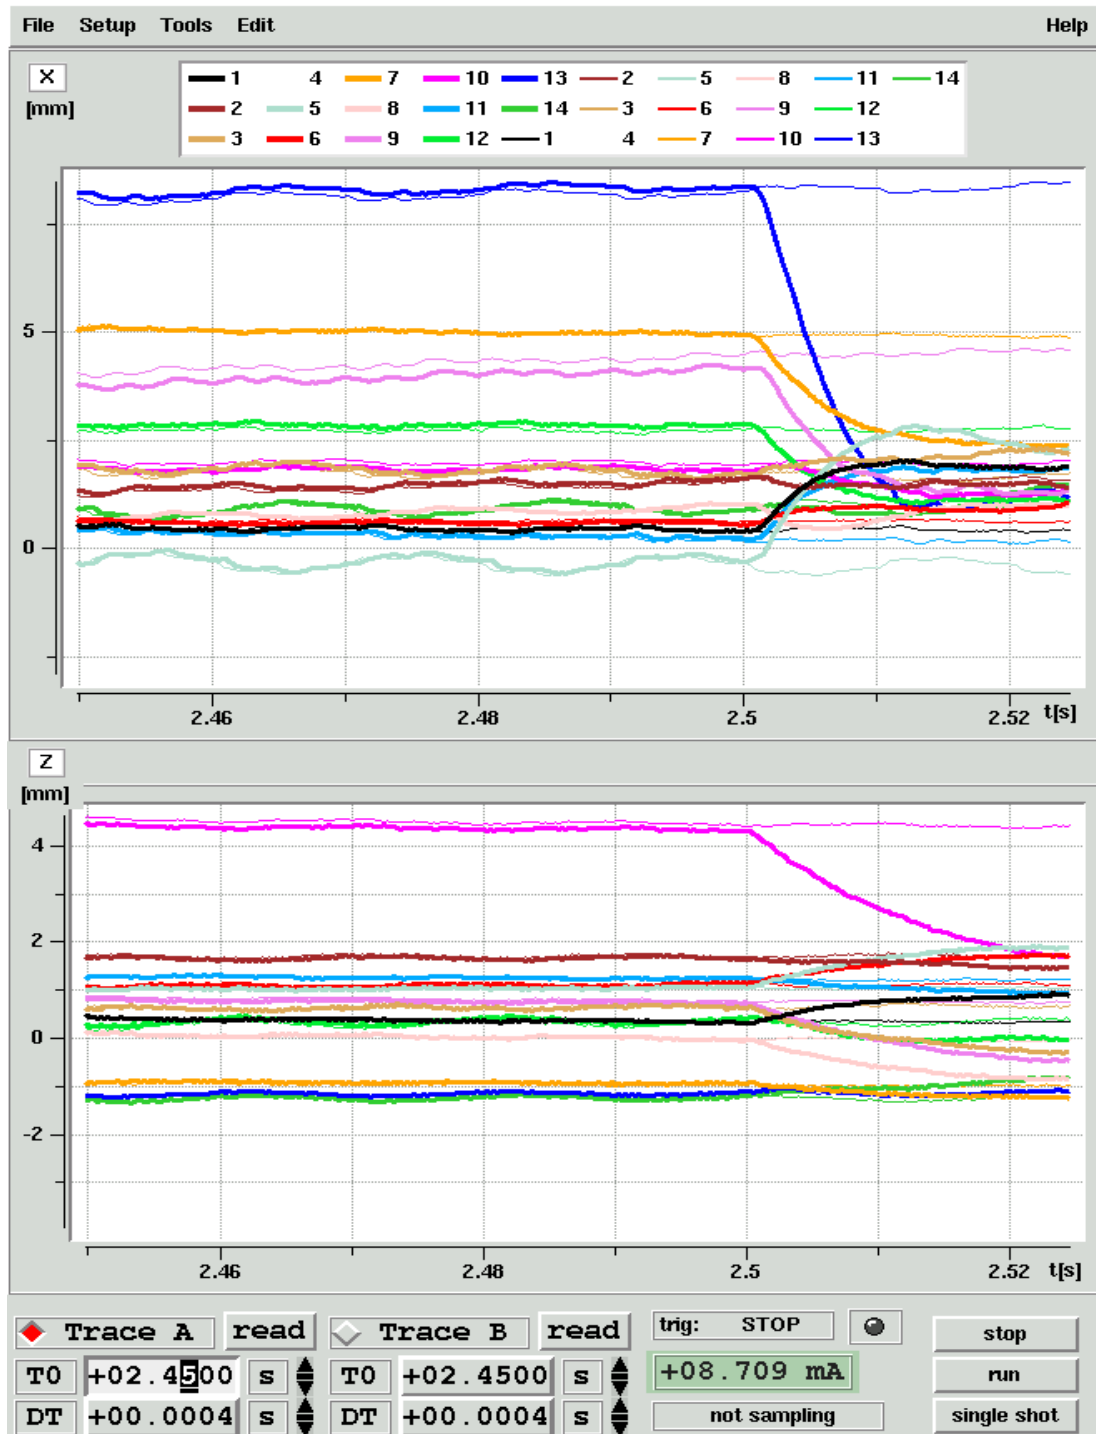
Figure 9.2: *Bodo orbit during a Bodo ramp with global feedback switched off (thin lines) and with feedback switched on at $t = 2.5$ (thick lines), with magnified time scale. The steerer magnet currents at injection (with feedback switched off) are different from fig. 9.1, which explains the slightly different beam orbit.*

are usually either constant or changing more slowly, while the switching of the orbit feedback causes a rapid large change of the beam orbit.

Figure 9.3 shows the fractional tune frequencies and the horizontal orbit at BPM13, which had the largest orbit variation during the measurement. When the global orbit feedback is switched off at $t = 3.87$ s, the orbit movement causes an increase of both fractional tune frequencies. The tune feedback compensates the tune drift in both planes successfully within 30 ms, with a non-critical overshoot of 20 kHz and an undershoot of 37 kHz with respect to the nominal tune frequency. When the orbit feedback is switched on again at $t = 4.53$ s, both feedbacks require about 70 ms until orbit and tunes are stable, again with negligible tune frequency variations of about ± 20 kHz. The stability test of the feedback loops was repeated successfully for various beam energies and beam position set values, as well as for instantaneous changes of the tune frequency set values. Therefore, both feedbacks can operate simultaneously without risk of instability or beam loss due to crosstalk.

9.3 Frequency Dependence

In order to investigate the dependence of the orbit correction on the frequency of the orbit perturbation, the horizontal beam position at BPM13 was corrected by the orbit feedback with a local orbit bump. The feedback used the correctors HK10, HK11 and HK12, with BPM13 being in the center of the bump close to HK11 (see fig. 3.4 on page 23). The current of HK4 was modulated with a sine generator at different frequencies in order to generate a well-defined orbit perturbation that is caused by a source outside the feedback bump. Due to the insufficient number of power supplies for the vertical correctors, the measurement was only performed
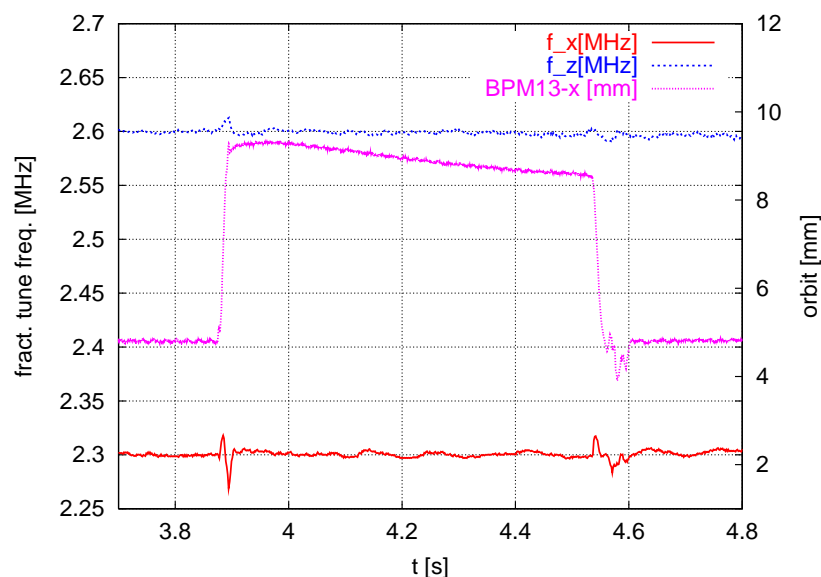

Figure 9.3: *Fractional betatron tune frequencies and horizontal beam orbit at BPM13 during a booster ramp. While the betatron tune feedback is operational (with set values for the fractional tune frequencies of $f_x = 2.3$ MHz, $f_z = 2.6$ MHz), the global orbit feedback is switched off (rising edge of the orbit at BPM13) and on (falling edge).*
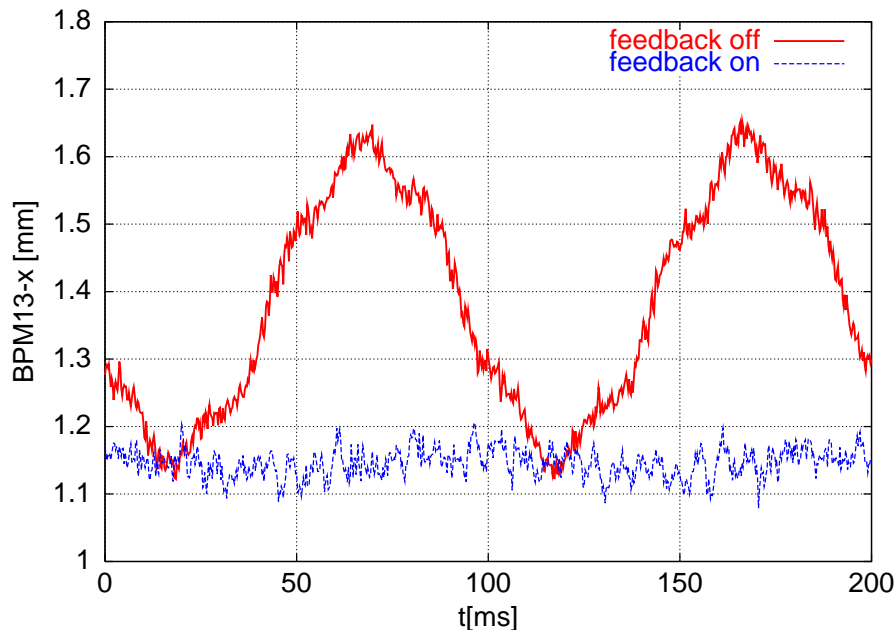
Figure 9.4: *Horizontal beam position at BPM13, with operational orbit feedback (blue curve) and with orbit feedback switched off (red curve). The 10 Hz orbit perturbation was generated with the corrector HK4.*

for the horizontal plane. The use of only three correctors simplified the PID loop adjustment, while the closed-loop bandwidth of the feedback should have little dependence on the number of correctors. The PID parameters of the three steerer magnets were optimised iteratively in order to achieve a stable feedback loop, good suppression of orbit perturbations in the 10 Hz frequency range, and little amplification of orbit perturbations or noise above the maximal frequency that can be damped ("cut-off frequency"). As Bodo and Delta have the same magnets and beam pipe design, the results of the measurement can be used to estimate the performance of a future DSP-based orbit feedback in Delta, e.g. the stabilisation of the orbit at one or several insertion devices. Since Delta usually operates at 1.5 GeV, the measurements in Bodo were also made at this energy.

Figure 9.4 shows the horizontal orbit at BPM13 with operational feedback and with feedback switched off, for a 10 Hz orbit perturbation generated by HK4. The sine generator that modulated the current of HK4 was not synchronised with the measurement, therefore the phases of the two curves are not correlated. The frequency spectrum of BPM13 is plotted in fig. 9.5. The spectrum is an FFT of the BPM data, with 1 Hz frequency resolution. The amplitudes in the plot are the absolute values of the complex Fourier coefficients, i.e. half of the peak-to-peak oscillation amplitude at the respective frequency.

While the 10 Hz component of the perturbation is damped by a factor of 15, the 50 Hz component (that was not generated by HK4, but may originate from power supply ripple and stray fields) is only damped by a factor of 2.5 . The plots in fig. 9.6 are measurements of the orbit oscillation amplitudes for different perturbation frequencies between 1 Hz and 90 Hz for HK4. The plots do not cover the complete Fourier spectrum, but only the Fourier component at the perturbation frequency with operational feedback and with feedback switched off. The upper plot was obtained from FFTs of BPM data with 0.1 Hz FFT resolution, the lower has
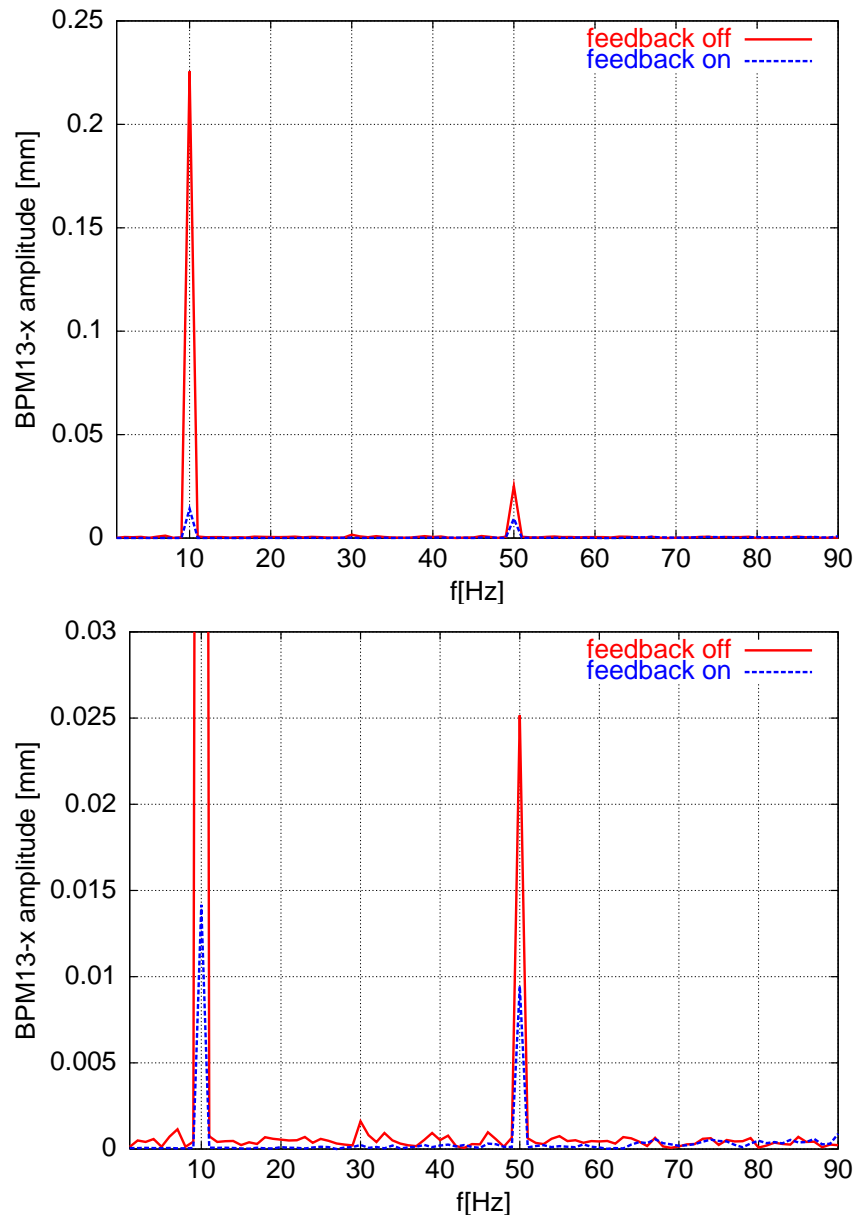
Figure 9.5: *Frequency spectrum of the horizontal beam position at BPM13, with orbit feedback switched on (blue curves) and with orbit feedback switched off (red curves). The lower plot is a magnification of the upper plot.*

1 Hz FFT resolution. The feedback loop can correct perturbations up to 90 Hz. It is able to reduce the 10 Hz frequency component of an orbit perturbation up to a factor of 15 and the 2 Hz component up to a factor of 80.

Therefore, the use of DeltaDSP boards for a fast orbit feedback in Delta may be a useful extension of the existing slow orbit feedback, since the fast feedback could cover the frequency range above 1 Hz that is not corrected by the slow feedback. This would allow the correction of orbit perturbations that are caused by mechanical vibrations, power supply ripple or stray fields from booster or transfer line T2.
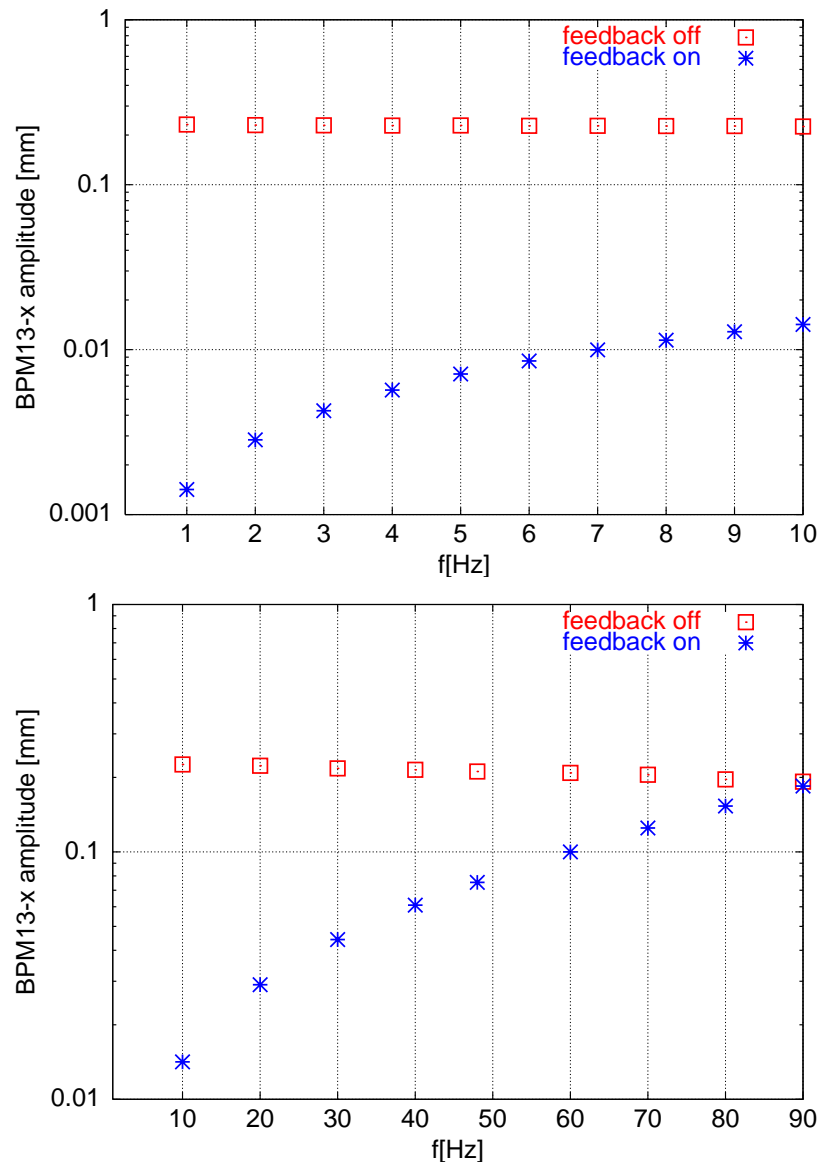
Figure 9.6: *Fourier amplitudes of horizontal orbit perturbations at BPM13, with orbit feedback switched on and with orbit feedback switched off. The orbit perturbations were generated by current modulation of HK4 at frequencies between 1 and 10 Hz (upper plot) and between 10 and 90 Hz (lower plot). In contrast to fig. 9.5, the vertical scales are logarithmic because of the high damping at low modulation frequencies.*

Chapter 10

Future Applications

So far, the DSP-based beam diagnostics and feedback system that was developed in this thesis has only been used at Bodo. This chapter introduces some possible future applications of the system at Delta.

10.1 DSP-Based Orbit Feedback at Delta

Due to the similarity of both rings, a DSP-based fast global orbit feedback for Delta may be built simply by copying the Bodo system and adapting it to the larger number of BPMs and orbit correctors. The flexible generic hardware and software concept of the Bodo system would only require modifications of a small fraction of the DSP software. While the Bodo DSPs perform a variety of different tasks like control of ramped power supplies, beam loss measurement, state machines, archiving of measurement data, injection and extraction control etc., the DSP software of a future orbit feedback at Delta should be optimised only for orbit correction in order to achieve maximum performance. As mentioned in section 5.1, a fast orbit feedback in Delta would require a replacement of the slow CANbus-based BPM readout system by DSP ADC boards, as well as a replacement of the existing slow power supplies of the steerer magnets by faster DSP-controlled ones.

Integration of Slow and Fast Orbit Feedback

Alternatively, one could install dedicated magnets and power supplies for the fast feedback in the vicinity of insertion devices and keep the existing steerer magnets with their slow power supplies, in order to operate a slow and a fast orbit feedback in parallel. This solution has already been implemented successfully at other synchrotron light sources where fast and slow orbit feedback run on different computing systems (e.g. DSP boards and a workstation) [46, 49].

In order to avoid crosstalk between the two feedbacks that may result in feedback instabilities and beam oscillations, fast and slow feedback are usually separated in the frequency domain. Suitable high-pass filters avoid the correction of slow orbit drifts by the fast feedback, and the slow feedback has a low-pass filter to avoid corrections at high frequencies. However, this solution causes a frequency deadband between the frequency ranges of both feedbacks where neither the fast nor the slow feedback achieve optimal orbit correction. In order to avoid this problem at Delta, the DeltaDSP boards have CANbus interfaces. Since steerer magnets and BPMs of the slow Delta orbit feedback are controlled via CANbus, suitable software on the DeltaDSP boards would be able to handle both fast and slow orbit correctors and BPMs simultaneously. Thus, problems with crosstalk could be avoided by integrating slow and fast orbit correction into a single DSP-based feedback system.

Correction Rate and Digital Feedback Loop Delay

As mentioned above, the DeltaDSP boards in Bodo perform a variety of different tasks so that the DSP software is not optimised for a fast orbit feedback. All real-time tasks of a DeltaDSP board in Bodo are performed in an interrupt handler that is executed periodically every 250 μs. If the BPM ADCs are sampled and transmitted via DeltaNet in interrupt cycle n, the corrections are calculated in interrupt cycle $n + 1$ and the DACs of the steerer magnets are set in interrupt cycle $n + 2$. Therefore, the delay from ADC trigger to DAC update is about 510 μs, which is larger than the DAC update period of 250 μs. In contrast, the interrupt handler software of a future global orbit feedback at Delta could be optimised only for this application in order to improve feedback loop delay and closed-loop bandwidth.

Table 10.1 shows that a global SVD-based orbit feedback consisting of 8 DeltaDSP boards, 64 BPMs, 32 horizontal and 32 vertical orbit correctors would be able to achieve a feedback loop delay of 164 μs, corresponding to a maximum rate of about 6 kHz. However, the DeltaDSP boards require some additional time for administration tasks and communication with the EPICS control system, therefore a correction rate of 4 kHz is suitable for this application, compared to 2 Hz for the existing Delta orbit feedback [32]. The calculation assumes that all BPMs and steerers are part of the fast feedback, rather than a combination of a fast and a slow feedback. The calculation was made for a system with 8 DACs and 8 BPM ADCs per DeltaDSP board, with 2 BPMs per synchronous serial ADC bus (i.e. all 4 serial buses have to be used). During the work on this thesis, Delta had only 43 BPMS, 30 horizontal and 26 vertical correctors. Therefore, the assumed correction rate of 4 kHz is pessimistic and leaves a safety margin for the integration of additional orbit correctors and RF or X-ray BPMs into the system.

	Delay[μs]	**Event**
$t_0 =$	0	Sample trigger for BPM ADCs
$+\Delta t_1$	25	BPM ADC data transferred into DSP memory ...
$+\Delta t_2$	8	... and copied to DeltaNet transmission buffer
$+\Delta t_3$	38	DeltaNet transmission begins
$+\Delta t_4$	4	DeltaNet transmission finished
$+\Delta t_5$	13	BPM data received by all DeltaDSP boards
$+\Delta t_6$	70	Power supply set values calculated and written to DACs
$+\Delta t_7$	6	DAC settling time for steerer magnet power supplies
$= t_{sum}$	164	Overall delay from ADC sample trigger to DAC output

Table 10.1: *Estimated worst-case delays in the digital processing chain for a future fast global SVD orbit feedback at Delta that consists of 8 DeltaDSP boards, 64 BPMs, 32 horizontal and 32 vertical steerer magnets. $\Delta t_{...}$ are individual delays of each data processing and transmission stage. t_{sum} is the resulting overall delay. The real-time DeltaNet data transfer requires 55 μs (= $t_3 + t_4 + t_5$) or about one third of the correction cycle in order to transfer all BPM data to all DeltaDSP boards.*

Correction Algorithm

In case a future Delta orbit feedback does not use the SVD algorithm, the correction rate may be different from 4 kHz, depending on the algorithm and its suitability for distributed processing. However, SVD is used successfully by most state-of-the-art synchrotron light sources (e.g. ESRF [46], APS [49], SLS [58] and Bessy II [57]). Despite some initial problems with this algorithm during first tests at Delta [32] it should be possible to use SVD also at Delta by a proper choice of the cutoff values and of the BPMs and steerer magnets that are used for the feedback. One "disadvantage" of SVD is its capability to correct individual BPM positions immediately to a desired value, which makes it sensitive to noise spikes of BPM ADC readings. In contrast, the orbit feedback algorithm that is usually used for Delta [32] may require up to some 10 correction cycles in order perform a desired correction of a single BPM position, which makes it slower and thus less sensitive to single noisy BPMs. However, the typical closed-loop bandwidth of an orbit feedback with 4 kHz correction rate, 230 μs ADC to DAC delay and a 3 mm stainless steel beam pipe is in the order of 200 Hz [46]. Therefore, the resulting ADC oversampling factor of 20 or more will reduce the effect of occasional noise spikes for single BPM readings typically by this factor, while a slow feedback is usually able to apply the full correction immediately. Furthermore, the DeltaDSP boards may use filter algorithms in order to detect and eliminate BPM noise spikes. Finally, the envisaged installation of a larger number of additional BPMs at critical locations in Delta will also improve the performance of an SVD-based orbit feedback by reducing the factor between smallest and largest SVD eigenvalue for a given number of feedback corrector magnets (see sections 2.8 and 7.3). Therefore it is recommended to use the SVD algorithm for a future DSP-based orbit feedback at Delta.

10.2 DSP-Based Betatron Tune Feedback at Delta

The Delta is usually operated with betatron tunes of $Qx = 9.17$ and $Q_z = 3.31$ in order to achieve good beam lifetime and injection rates. Since the vertical betatron tune is close to a destructive vertical resonance, Delta has a PC-based betatron tune feedback that corrects the betatron tunes at a rate of typically 2 Hz [35]. This is sufficient to compensate betatron tune variations that are caused by sextupole magnetic fields during orbit changes by the slow Delta orbit feedback. A fast DSP-based orbit feedback may cause betatron tune drifts on a time scale of milliseconds that cannot be corrected by the existing tune feedback, which might result in beam loss. This problem could be solved by an upgrade of the Delta betatron tune feedback in analogy to the DSP-based Bodo system. Since both machines use the same beam excitation kickers and BPM pickup buttons, the Bodo tune measurement system could be adapted to Delta with some minor hardware and software modifications due to the different beam revolution time. Smaller filter bandwidths, smaller DDS generator frequency steps and a DDS clock signal with improved temperature stability might be useful in order to improve the frequency resolution of the system. The replacement of the quadrupole power supply controllers of Delta by DeltaDSP boards would allow the implementation of a betatron tune feedback in analogy to Bodo. This might require modifications of the regulation units of at least some Delta quadrupole power supplies in order to achieve a response time that is sufficient for fast tune corrections.

10.3   Top-Up Injection

Orbit Perturbations by Booster and Transfer Line

The envisaged implementation of a top-up injection mode for Delta requires sufficiently small orbit perturbations during injection. If the orbit perturbations that are generated by the ramping booster and by pulsed T2 transfer line magnets are too large for certain synchrotron radiation experiments, the beam line users either have to interrupt their measurements for a short time during injection, or the perturbations have to be corrected. This section will briefly discuss preliminary measurements of the orbit perturbations during injection at one BPM in Delta, as well as possibilities to correct the perturbations.

Figures 10.1 show the orbit perturbation at the upstream BPM of the U250 undulator in Delta (BPM11, see fig. 1.1) that is generated by the booster ramp (left figure) and by stray fields of pulsed T2 transfer line magnets (right figure). Since the usual Delta BPM readout system achieves only readout rates of a few Hz, the Delta BPM was connected to a BPM ADC of the Bodo BPM system, and the signal was recorded synchronously to the Bodo ramp.

The orbit perturbation that is caused by the ramping booster (left figure, slow orbit change between $t = 6.5$ s and $t = 13$ s) could either be corrected by a feed-forward method or by a DSP-based orbit feedback similar to the Bodo system. The falling edge of the perturbation
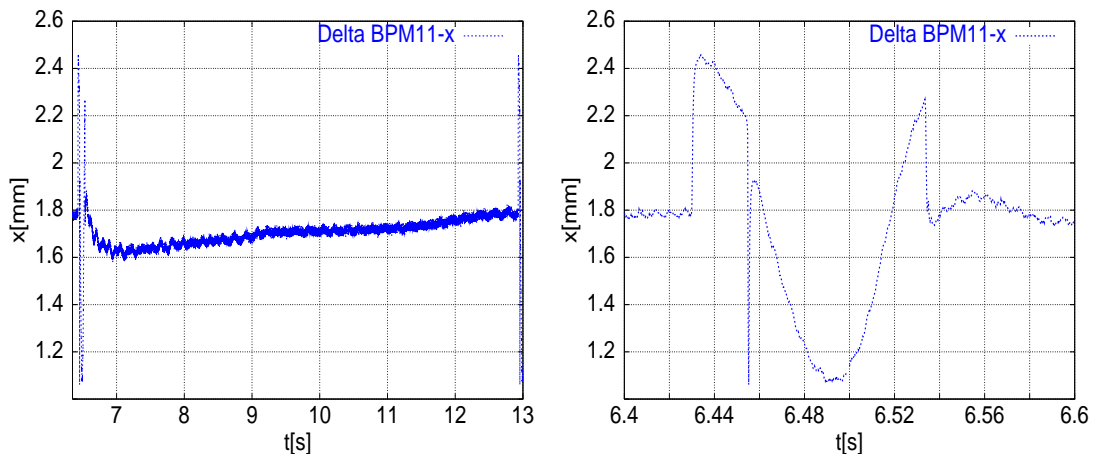

Figure 10.1: *The figures show the the orbit perturbation at the upstream BPM of the U250 undulator in Delta (BPM11, see fig. 1.1) that is generated by the booster ramp of 6.5 s length and by stray fields of pulsed T2 transfer line magnets. The right figure shows the orbit of the left figure with increased time resolution. The Bodo beam is injected into Delta at $t = 6.455$ s. The Bodo ramp energy changes from 1.488 GeV (at $t = 6.5s$) to 62 MeV (at $t = 8.5s$) and back to 1.488 GeV (at $t = 13$ s). The larger one of the T2 dipole magnets has a sine-wave current pulse of 100 ms length. It is triggered at $t = 6.43$ s, 25 ms before injection. The dipole generates an orbit perturbation that is approximately proportional to the derivative of its current. The orbit perturbation by the injection septum stray fields (visible at $t = 6.455$, short BPM signal spike in the right figure) cannot be resolved properly with the available BPM RF front-end electronics, since the septum pulse of 80 µs length is too fast for the time resolution of the BPM system.*

corresponds to a frequency of roughly 1 Hz, and the Bodo feedback managed to damp this frequency component by more than a factor of 100. Due to a correction rate of 2 Hz, the existing slow orbit feedback cannot correct the falling edge of the perturbation. The perturbation by the pulsed T2 dipole magnet (right figure) has the same length than the 100 ms sine wave current pulse of the magnet. However, the perturbation has rise and fall times in the order of 1 ms. Since a feedback-based correction of the complete perturbation with this speed is very difficult, an adaptive feed-forward correction scheme or a combination of feedback and feed-forward with one or several dedicated fast corrector magnets in the vicinity of the T2 transfer line might be a better solution. As long as Delta has no fast global orbit feedback, the correction of the perturbations could be achieved with one or two additional DeltaDSP boards in the booster DeltaDSP network that are connected to some Delta BPMs and some dedicated fast steerers in Delta. Since the Bodo DSP system triggers the extraction and determines the ramp speed, the integration into the Bodo DSP network would simplify the precise synchronisation that is required for a feed-forward correction of the dipole pulse perturbation. However, the first step should be a further reduction of the perturbation source by additional magnetic shielding of the transfer line.

Booster Beam Current Control

In order to keep the Delta beam current during top-up injection mode as stable as possible, the injected beam current should be neither too large nor too small. Due to temperature variations of the Linac RF systems and pulsed Linac and T1 components, the Bodo beam current may vary by a factor of two or more. This evokes the idea of a beam current regulation system that reduces the Bodo beam current to the desired level before extraction, and that inhibits the extraction if the beam current that was injected from the Linac is too low. This could be achieved by an additional regulation loop on the tune measurement DSP board in Bodo. A suitable feedback loop should be able to adjust the excitation kicker amplitude so that the resulting beam loss at low energies reduces the beam current to the desired level. Alternatively, a controlled loss of a certain amount of beam charge could be generated by the orbit feedback system which could reduce the mechanical aperture automatically with a sufficiently large orbit bump, and remove the bump when the beam current has reached the desired level.

Chapter 11

Summary

The subject of this thesis was the development, construction and application of a distributed DSP-based system that is tailored to the measurement and real-time correction of beam parameters like global orbit and betatron tunes at DELTA, with measurement and correction rates of some 100 Hz to some kHz. In order to achieve this, a novel distributed DSP multiprocessing system was developed that may consist of up to 255 VMEbus boards ("DeltaDSP boards") with 2 DSPs each. The DSP boards are connected in a ring by a newly developed 160 Mbaud real-time fibre optics network ("DeltaNet"), with up to 2 km distance from board to board and up to 200 km overall length of the ring. A single fibre optics cable is used for real-time data transfer and for exact synchronisation of the main board clocks of all DSP boards. Furthermore, DeltaNet distributes trigger signals e.g. for the synchronous start of the Bodo ramp.

When being used for a fast global orbit feedback, the newly developed architecture has several advantages over existing DSP-based orbit feedback systems at other light sources. In contrast to centralised architectures with a single DSP [46], the computing power and bandwidth of the distributed DeltaDSP system can be adapted to the number of BPMs and steerer magnets that are handled by the system. The SVD algorithm that was used in this thesis is ideally suited for distributed multiprocessing, while the computing time of a centralised architecture grows proportionally to the number of BPMs and to the number of orbit correctors, which limits the size and bandwidth of the system. A distributed DSP-based feedback system that was developed at the Swiss Light Source ("SLS") in parallel to this thesis consists also of a ring of interconnected DSP boards [45]. However, these boards receive only BPM data from their next neighbours for real-time orbit corrections via direct point-to-point connections of the DSPs. This architecture only allows global feedbacks for certain algorithms like SVD and for a certain structure of the SVD-inverted beam response matrix. In contrast, DeltaNet distributes the measurement data of each DSP board to all other DSP boards in real-time, in short deterministic intervals of typically some 10 μs. This allows global feedbacks for any algorithm and for any beam response matrix. Moreover, the above next-neighbour architecture requires a certain order of the BPMs and DSP boards. The DeltaDSP boards can be arranged in any order in the DeltaNet ring, and BPMs and corrector power supplies can be connected to any DeltaDSP board. This flexibility would simplify the gradual extension of a future DSP-based local orbit feedback in Delta to a global feedback.

The DeltaNet interface was completely implemented in a field programmable gate array ("FPGA"), including the detection of transmission errors via 32-bit cyclic redundancy checksums ("CRCs"). This is essential to avoid beam loss e.g. in global orbit feedbacks due to transmission errors. Existing distributed DSP orbit feedbacks without hardwired checksum calculation [45] would have to use software checksum algorithms to reach the same level of transmission security. This would take about 10 times longer, slow down the network and consume a large amount of DSP processing time.

Furthermore, DeltaNet distributes the main board clock of one DeltaDSP board ("clock master") to all other boards in the ring. The use of fault-tolerant clock recovery phase-locked

loops (PLLs) allows the synchronisation of measurements, ramped power supplies and feed-back algorithms of all DSP boards in the accelerator, while existing feedback systems usually require dedicated cables and hardware to synchronise distributed feedback components.

In addition to circuit and FPGA schematics and printed circuit board layouts of the DeltaDSP board, the work on this thesis also includes the development of the software that was required for the control system integration of the DSP boards, as well as the development of the software for the DSPs. The generic software concept enables future applications of the system e.g. at Delta with minor software modifications.

A system of six DeltaDSP boards was installed at Bodo in October 2001. Since then the boards control all ramped power supplies and RF power. Furthermore, they handle most beam diagnostics systems like beam loss monitors, beam position monitors, beam current measurement and betatron tune measurement, as well as the measurement of power supply currents and voltages, with user-defined sampling rates of 4 to 20 kHz. The beam position signals of the Bodo BPM RF front-end electronics are sampled by a chain of external ADC boards that were developed in this thesis. The proprietary bus system of the ADC chain transfers the BPM data directly into the internal memory of the DSPs, from where they are distributed to all other DSPs via DeltaNet.

In order to demonstrate the operation of a fast orbit feedback, a distributed SVD-based real-time orbit feedback algorithm was implemented on the DeltaDSP boards in Bodo. The DSP boards apply orbit corrections at a rate of 4 kHz. The beam response matrix was calculated from a theoretical optics model. A suitable choice of the SVD cutoff values for both planes allowed the correction and stabilisation of the horizontal orbit during arbitrary booster energy ramps. The remaining systematic horizontal offset of typically 1.5 mm at all BPMs can be explained by a wrong circumference of the ring. Since only three vertical steerer power supplies were available, the vertical Bodo orbit could only be corrected up to ± 2 mm. However, this is sufficient for normal 1.5 GeV operation and does not limit the performance of the booster.

The DSP system was also used to test the response of a local horizontal orbit feedback to orbit perturbations that were generated by modulating the current of an orbit corrector at different frequencies. The feedback achieved damping factors of 150 for the 1 Hz frequency component and 15 for the 10 Hz component of the orbit perturbation. A possible future global orbit feedback in Delta consisting of 64 BPMs and 64 correctors could also achieve a correction rate of 4 kHz by optimised DSP software, even with a 50 percent reduction of the digital delay in the feedback loop with respect to the booster feedback. In combination with optimised BPM RF front-end electronics, a faster BPM button sampling clock and dedicated corrector magnets, this should allow a closed-loop bandwidth in excess of 90 Hz.

In addition to a DSP-based global orbit feedback this thesis also includes the development of a DSP-based measurement and real-time correction system for the betatron tunes of Bodo. The measurement system is based on a DeltaDSP board which excites the electron beam via DSP-controlled direct digital synthesis (DDS) of sine waves for the diagonal excitation kicker magnet. Combining swept-frequency and chirp method, coherent beam oscillations are detected in the frequency domain, with another DSP-controlled DDS generator that sweeps the analysed frequency. While the software and most parts of the hardware were developed in this thesis, the analogue RF front-end for the detection of coherent oscillations was taken from the previous all-analogue system [4, 5]. The DSP-based correction of systematic frequency measurement errors and the combination of swept-frequency and chirp method improved both

the tune measurement rate and the precision by a factor of 10. These improvements and the DeltaNet interface of the tune measurement system allowed the implementation of a real-time tune feedback. This feedback corrects the quadrupole magnet currents so that the betatron tunes remain nearly constant for arbitrary energy ramps of the booster. When using the fastest possible 1.5 GeV energy ramp and the "bo-101c" optics, the deviations from the nominal fractional betatron tune frequencies in both planes are less than ± 50 kHz ($dQ < \pm 0.008$) for the first 500 ms after injection, and less than ± 20 kHz ($dQ < \pm 0.003$) afterwards, with a Bodo revolution frequency of 5.95 MHz. Moreover, the simultaneous operation of betatron tune and orbit feedback was successfully tested, and the measurement of betatron tunes by the Bodo beam loss monitors was demonstrated.

The booster is controlled and its beam parameters are monitored by graphical user interfaces ("GUIs") that were developed in this thesis. The automatic real-time optimisation of beam parameters, the simplified control and fault detection and the improvement of the monitor time resolution up to three orders of magnitude contributed to a major improvement of the performance, handling and reliability of the booster. The DeltaDSP boards proved to be highly reliable, with no failures after more than one year of operation. Despite the complexity of the software with more than 18000 lines of code for DSPs and EPICS integration, the machine downtime due to software problems was only two hours during the first year of operation.

The DeltaDSP system and the improved beam diagnostics allowed the efficient test of new Bodo optics that were developed during the work on this thesis. In contrast to the "bo-006b" optics that was used so far, the new optics required orbit corrections during the booster ramp in order to store the beam and to avoid beam loss. In combination with suitable transfer line optics, one of the newly developed Bodo optics ("bo-101c") achieved a significant increase of the charge transfer rate from Bodo to Delta, with average Bodo beam currents of 8 mA and more. This allowed the reduction of storage ring filling times for 1.5 GeV synchrotron light operation to typically 15-30 minutes.

In addition to the improved performance as 1.5 GeV injector for Delta, the improved diagnostics and the feedback-based control of beam parameters increase the attractiveness and potential of Bodo as a test accelerator for newly developed components, monitors, hardware and software. Furthermore, the DSP-based improvement of beam parameters may serve as a paradigm for the development of cost-effective industrial storage rings that compensate the reduced precision and stability of low-cost components (e.g. power supplies) by DSP-based real-time correction of beam parameter drifts and perturbations.

Finally, the CANbus interfaces of the DeltaDSP boards and the compatibility with the proprietary power supply interfaces of most Delta magnets would allow the seamless future integration of most monitors and devices at DELTA into one or several intercommunicating DeltaDSP networks. This enables the implementation of "interdisciplinary" feedback systems that operate the complete accelerator automatically in analogy to human operators, who optimise the accelerator by evaluating a large number of different monitor systems. An example is a real-time feedback system that corrects the global orbit, but that simultaneously minimises beam loss and radiation levels at insertions devices, measures and corrects betatron tunes and beam optics, avoids high temperatures by misaligned synchrotron radiation beams at outlet-chambers, measures movements and vibrations of the beam pipe at BPM pickup buttons and corrects the measured beam positions accordingly.

Appendix A

Bodo Optics

Section A.1 of this appendix contains simulation results for magnet misalignments and field errors that were referred to but not included in chapter 4. Section A.2 contains plots of theoretical beam cross sections at 60 MeV and 1.5 GeV that were also referred to in previous chapters.

A.1 Simulation of Magnet Misalignment and Magnetic Field Errors

This section contains histograms of magnet misalignment and magnetic field error simulations for different Bodo optics. The simulations were performed with the PC program "Optics" [27]. The program can generate random magnet position misalignments and field errors and calculate the absolute horizontal and vertical peak value of the resulting orbit perturbation. For each optics, this calculation was repeated 10^4 times (with randomly changing distributions for magnet misalignments and field errors) in order to obtain sufficient statistics. However, the resulting distributions of the orbit peak values do not look as smooth as one could expect them to, since the "Optics" program uses the same alignment error or field error distribution several times (more exactly: a random number of times, with the maximum number depending on the number of overall repetitions of the simulation) before generating a new distribution. The number of absolute peak orbit perturbations is integrated over 0.5 mm intervals. If the "Optics" program did not re-use the same alignment error distribution several times, the statistical error of the number n of orbit perturbations per 0.5 mm interval should scale with $\sqrt{n}$. Due to the re-use the statistical error is larger.

Figures A.1 and A.2 show the distribution of horizontal and vertical orbit perturbations (absolute peak values) in Bodo for 10000 simulated transverse random misalignments of all Bodo magnets. Figure A.3 shows the distribution of horizontal orbit perturbations for 10000 simulated random field errors distributions of all Bodo bending magnets. Finally, figure A.4 shows the distribution of absolute vertical peak orbit perturbations in Bodo for 10000 simulated tilt misalignments of all Bodo dipole magnets.

As mentioned in chapter 4, the "bo-101c" optics is least sensitive to magnetic field and alignment errors, while the "bo-103b" optics shows the largest perturbations due to its large focussing strengths and the smaller distance of the betatron tunes to integer resonances.

A.2 Theoretical Beam Cross Sections and BPM Positions

Figures A.5 and A.6 show theoretical horizontal and vertical beam cross sections in Bodo for an emittance of $\varepsilon_x = \varepsilon_z = 1000\ nm\ rad$ and an energy spread of $\Delta p/p = 0.005$. Figures A.7 and A.8 show the respective cross sections for an energy of 1.5 GeV, with 10 % emittance coupling and equilibrium values for emittance and energy spread. The small horizontal cross sections at most BPM locations illustrate the suboptimal placement of the Bodo BPMs with respect to the horizontal plane.
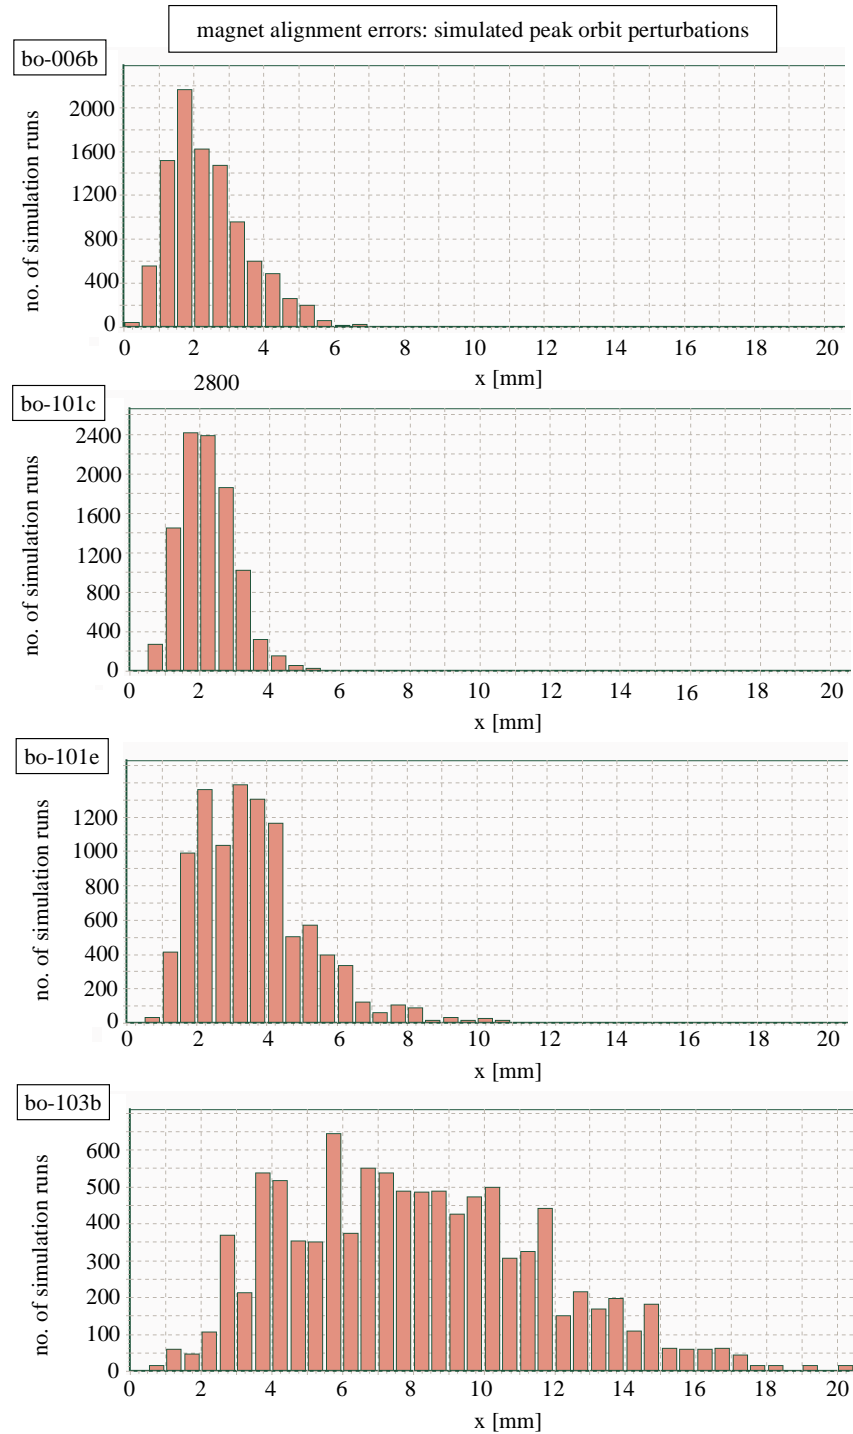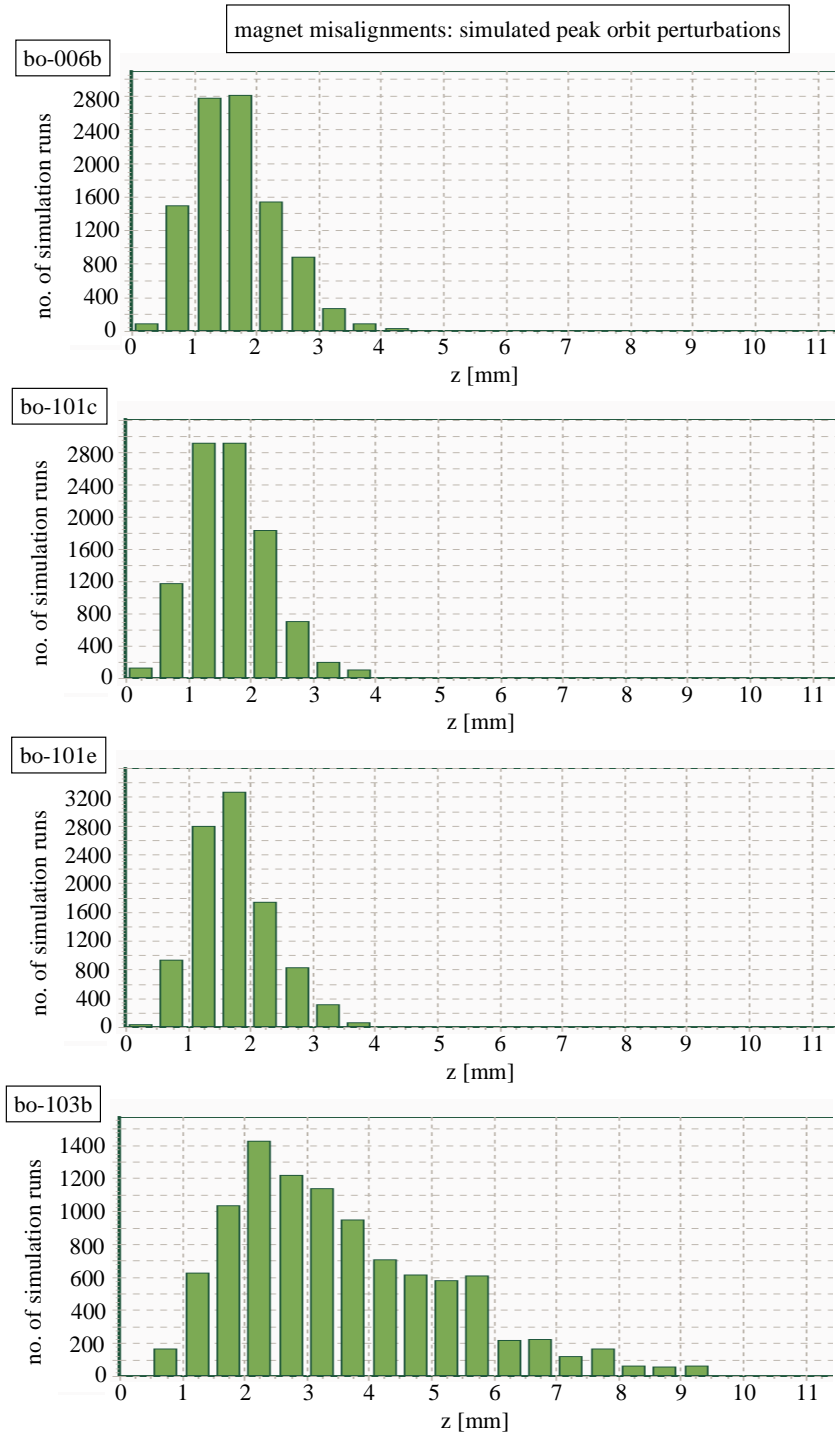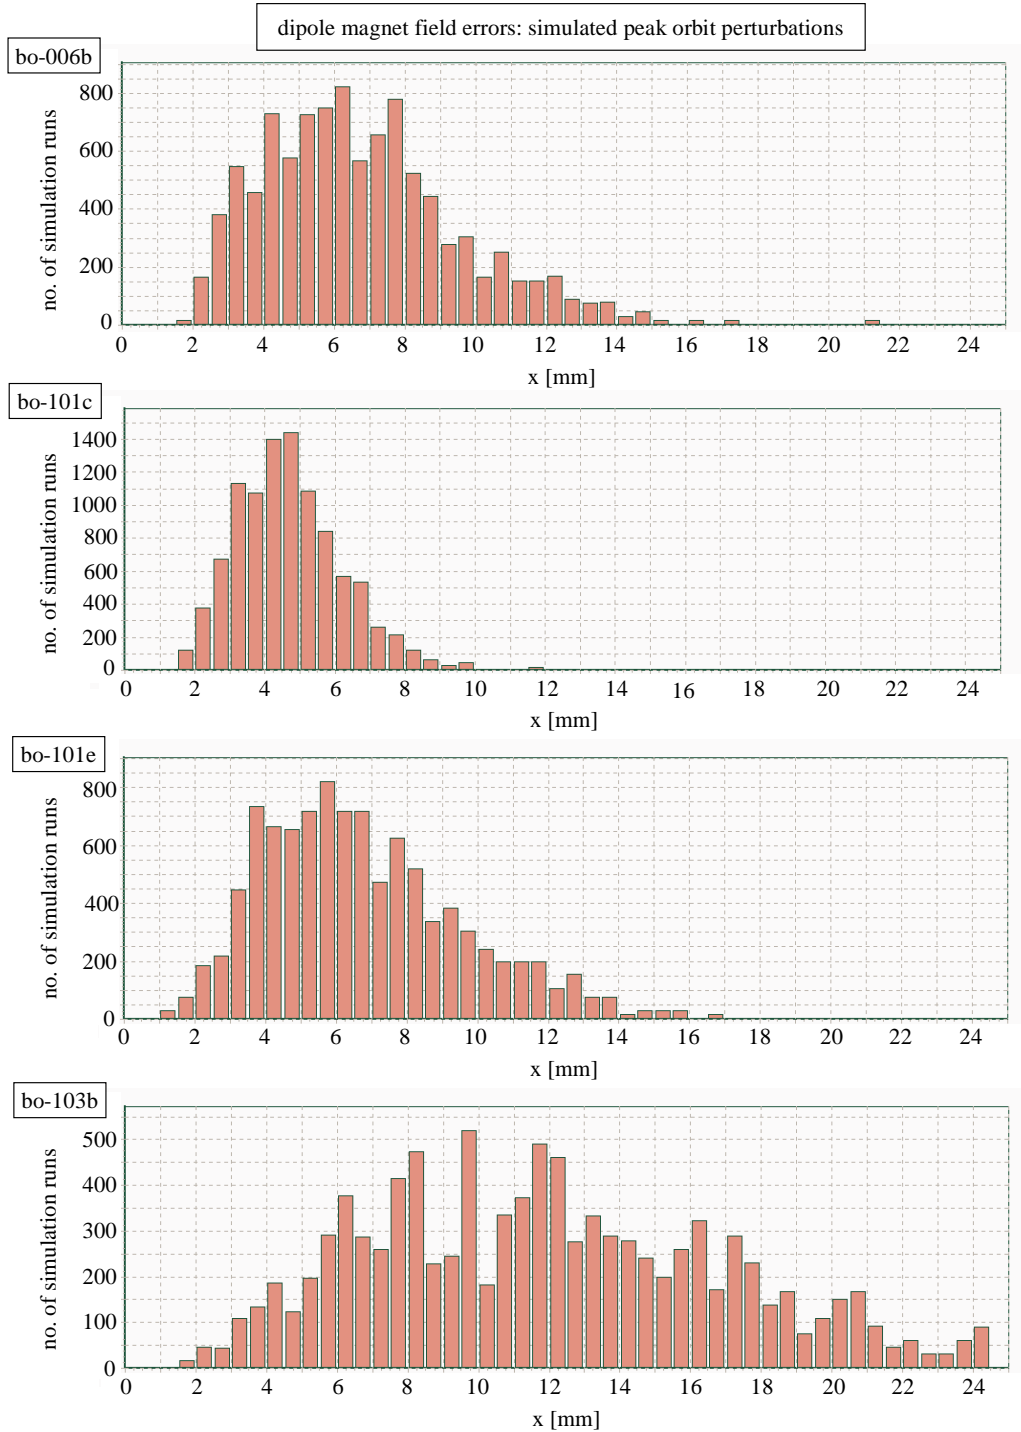
Figure A.1: *Calculated distribution of horizontal orbit perturbations (absolute peak values) in Bodo for* **10000 simulated transverse random misalignments of all Bodo magnets** *(random x and z position offset between −0.3 mm and +0.3 mm for each magnet). From top to bottom: distribution for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics (i.e. ordered by decreasing beam emittance). The horizontal axis is the absolute peak orbit perturbation for the horizontal orbit in* [mm]*, the vertical axis is the number of random misalignments that resulted in the respective peak orbit perturbation (integrated over intervals of* 0.5 mm *length).*
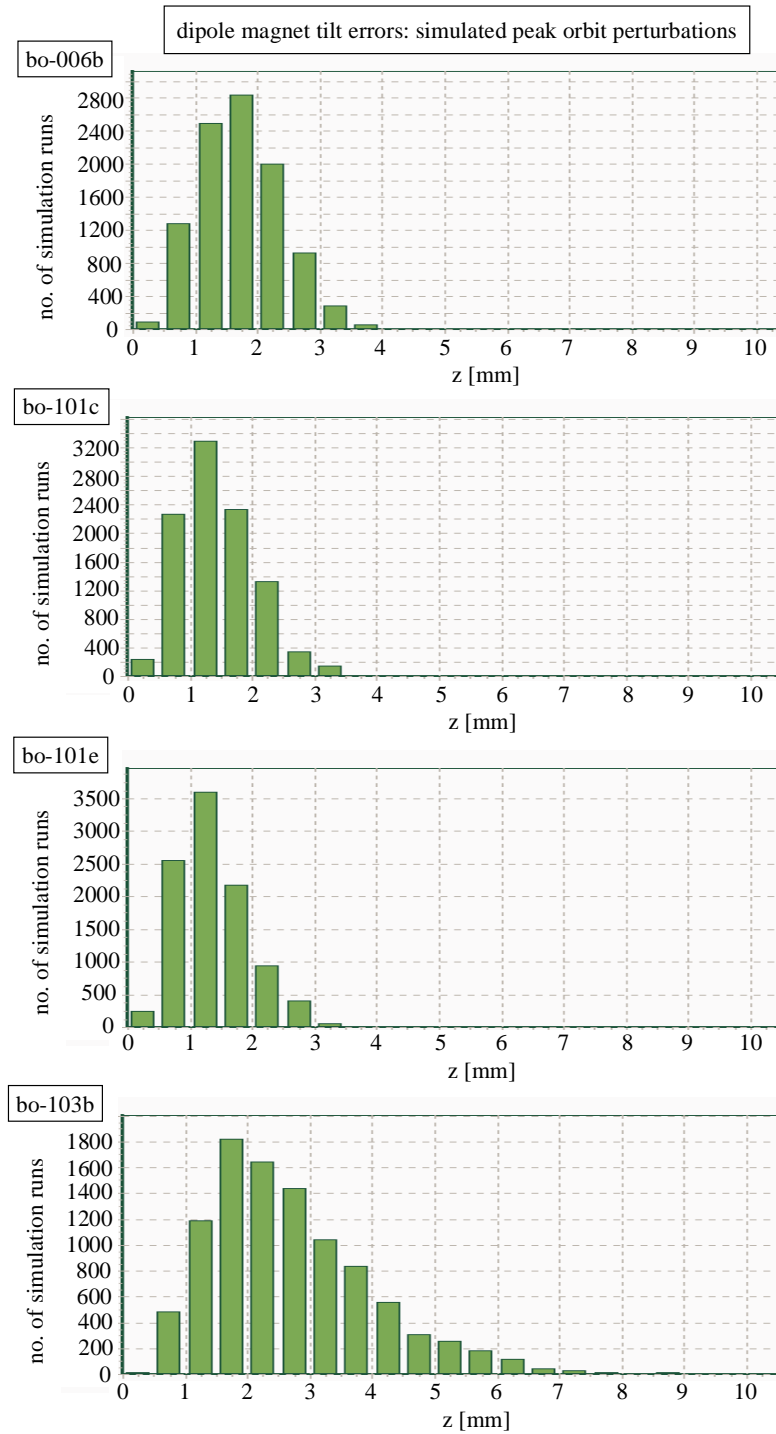
Figure A.2: *Calculated distribution of absolute vertical peak orbit perturbations in Bodo for* **10000 simulated transverse random misalignments of all Bodo magnets** *(random x and z position offset between* −0.3 *mm and* +0.3 *mm for each magnet). From top to bottom: distribution for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics. The horizontal axis is the absolute peak orbit perturbation of the vertical orbit in* [mm]*, the vertical axis is the number of random misalignments that resulted in the respective peak orbit perturbation (integrated over intervals of* 0.5 *mm length).*

Figure A.3: *Calculated distribution of horizontal orbit perturbations (absolute peak values) in Bodo for **10000 simulated random field errors distributions of all Bodo dipole magnets** (random field error between -0.2 % and +0.2 % for each dipole magnet). From top to bottom: distribution for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics. The horizontal axis is the absolute peak orbit perturbation of the horizontal orbit in* [mm], *the vertical axis is the number of random error distributions that resulted in the respective peak orbit perturbation (integrated over intervals of* 0.5 mm *length).*

Figure A.4: *Calculated distribution of absolute vertical peak orbit perturbations in Bodo for* **10000 simulated tilt misalignments of all Bodo dipole magnets***, with a random tilt angle between* $-0.5\ mrad$ *and* $+0.5\ mrad$ *for each magnet. From top to bottom: distribution for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics. The horizontal axis is the absolute peak orbit perturbation of the vertical orbit in* $[mm]$*, the vertical axis is the number of random misalignments that resulted in the respective peak orbit perturbation (integrated over intervals of* $0.5\ mm$ *length).*
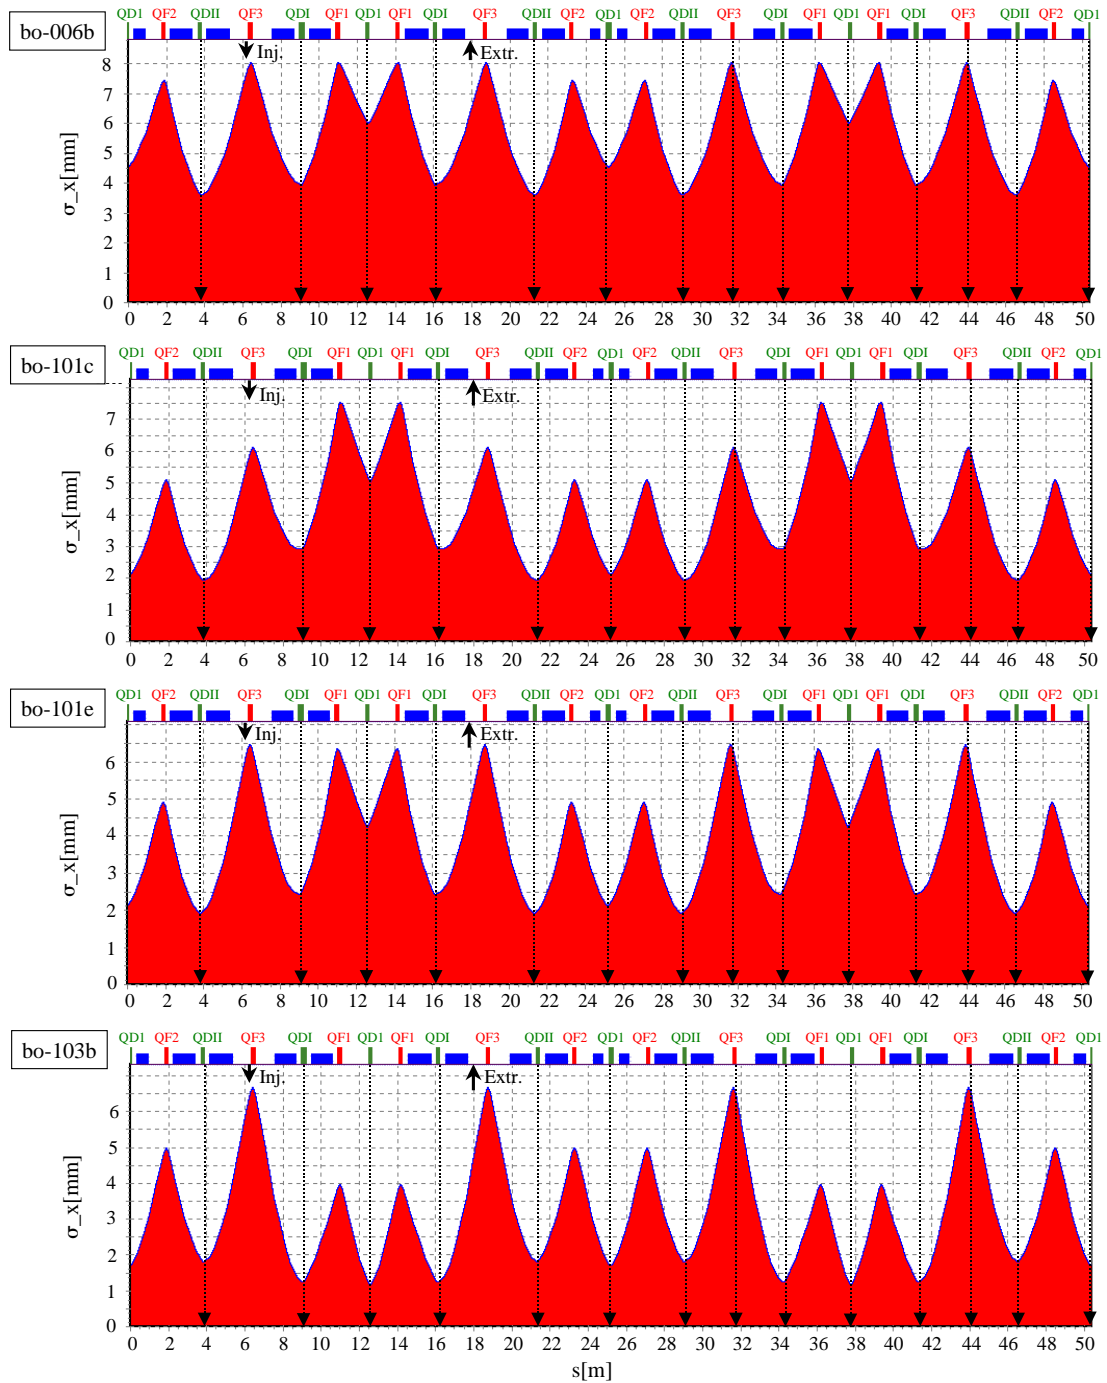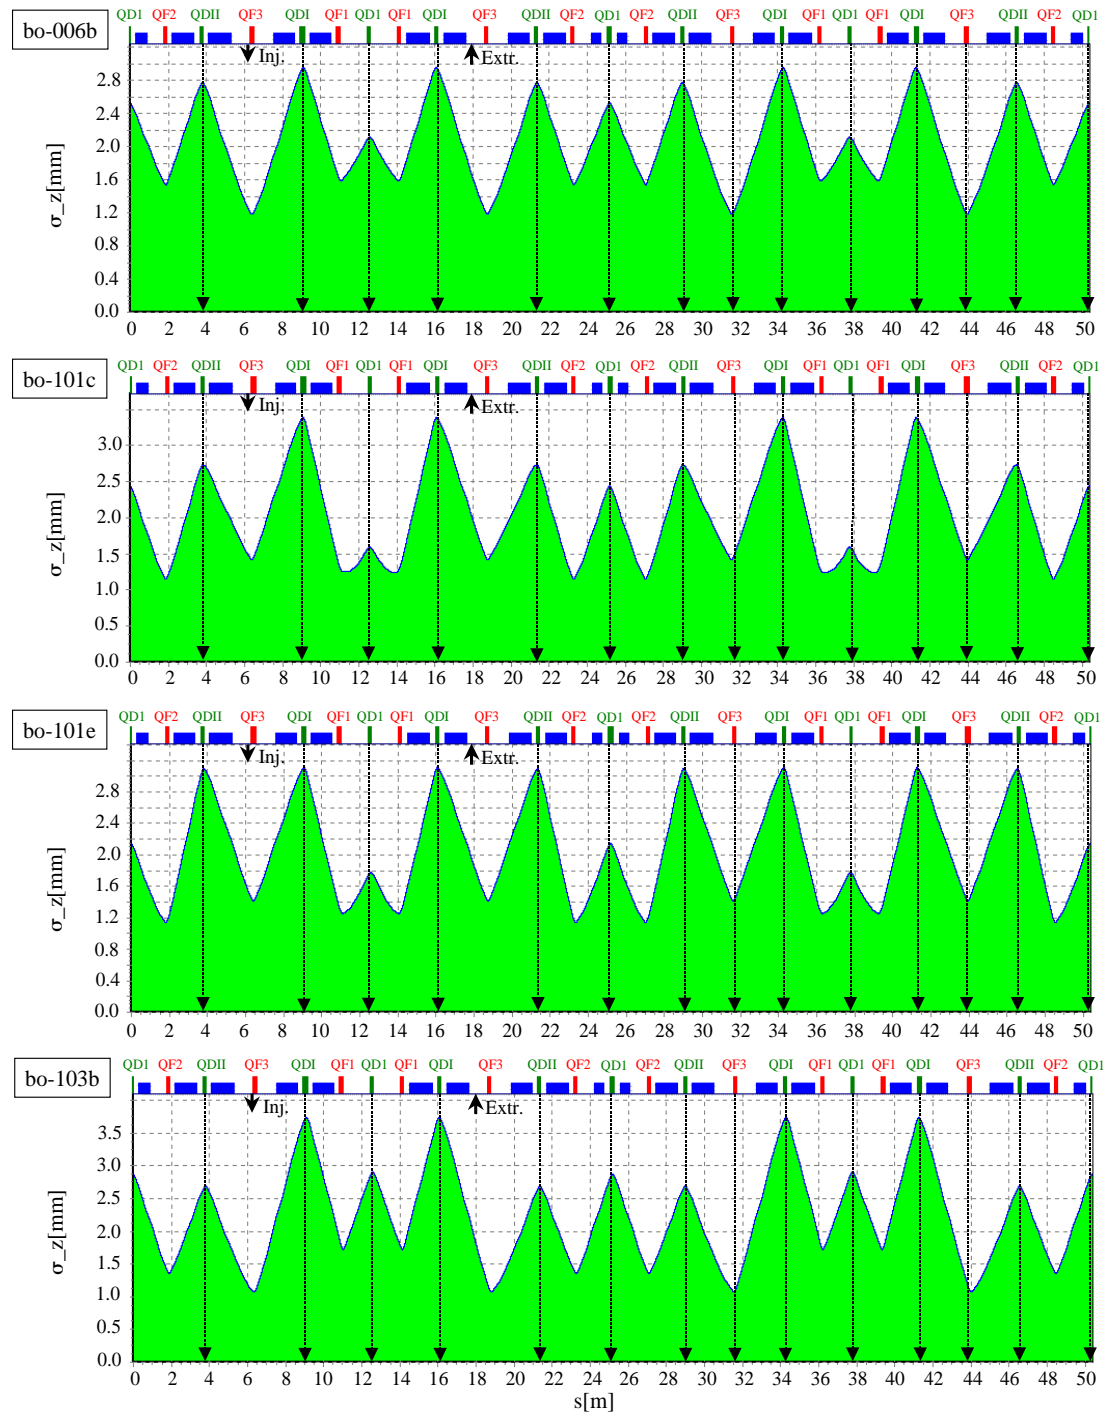
Figure A.5: *Calculated horizontal Bodo beam cross section (1 standard deviation) at 60 MeV. From top to bottom: for "bo-006b", "bo-101c", "bo-101e" and "bo-103" optics, assuming an emittance of $\varepsilon_x = \varepsilon_z = 1000\ nm\ rad$ and an energy spread of $\Delta p/p = 0.005$. Emittance and energy spread are non-equilibrium values that are determined by the Linac, assuming the idealised case of matched optical functions without aperture limitation at the septum. The locations of the 14 BPMs are indicated by black arrows from the respective quadrupole to the s-axis. All except two BPMs are located in QD quadrupoles, at local minima of the beam cross section. As the amplitude of orbit perturbations scales with the beam envelope, the horizontal orbit perturbations that are visible on the BPMs may be much smaller than the maximum orbit perturbations, especially for optics with low emittance.*

Figure A.6: *Calculated vertical Bodo beam cross section (1 standard deviation) at 60 MeV. From top to bottom: for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics, assuming an emittance of $\varepsilon_x = \varepsilon_z = 1000\ nm\ rad$ and an energy spread of $\Delta p/p = 0.005$. The locations of the 14 BPMs are indicated by black arrows from the respective quadrupole to the s-axis. 12 of 14 BPMs are located at local maxima of the vertical beam envelope, therefore the BPMs give a realistic impression of vertical orbit perturbations.*
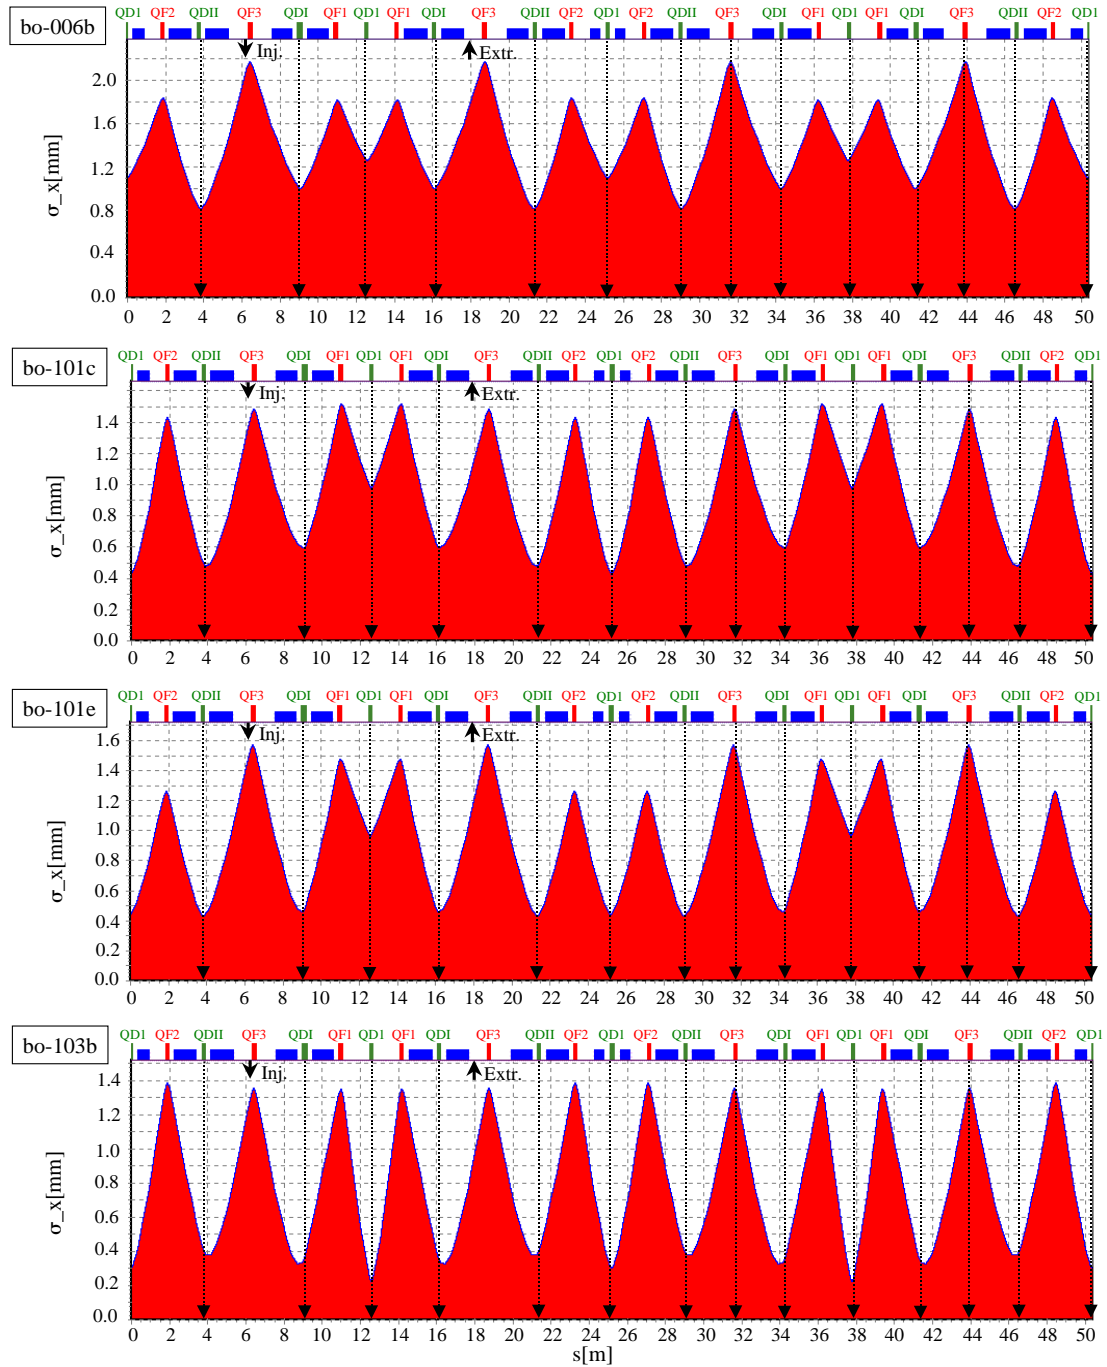
Figure A.7: *Calculated horizontal Bodo beam cross section (1 standard deviation, 10 % emittance coupling) at 1.5 GeV. From top to bottom: for "bo-006b", "bo-101c", "bo-101e" and "bo-103" optics (i.e. ordered by decreasing beam emittance). Emittance and energy spread are equilibrium values that are determined by energy and optics. The locations of the 14 BPMs are indicated by black arrows from the respective quadrupole to the s-axis.*
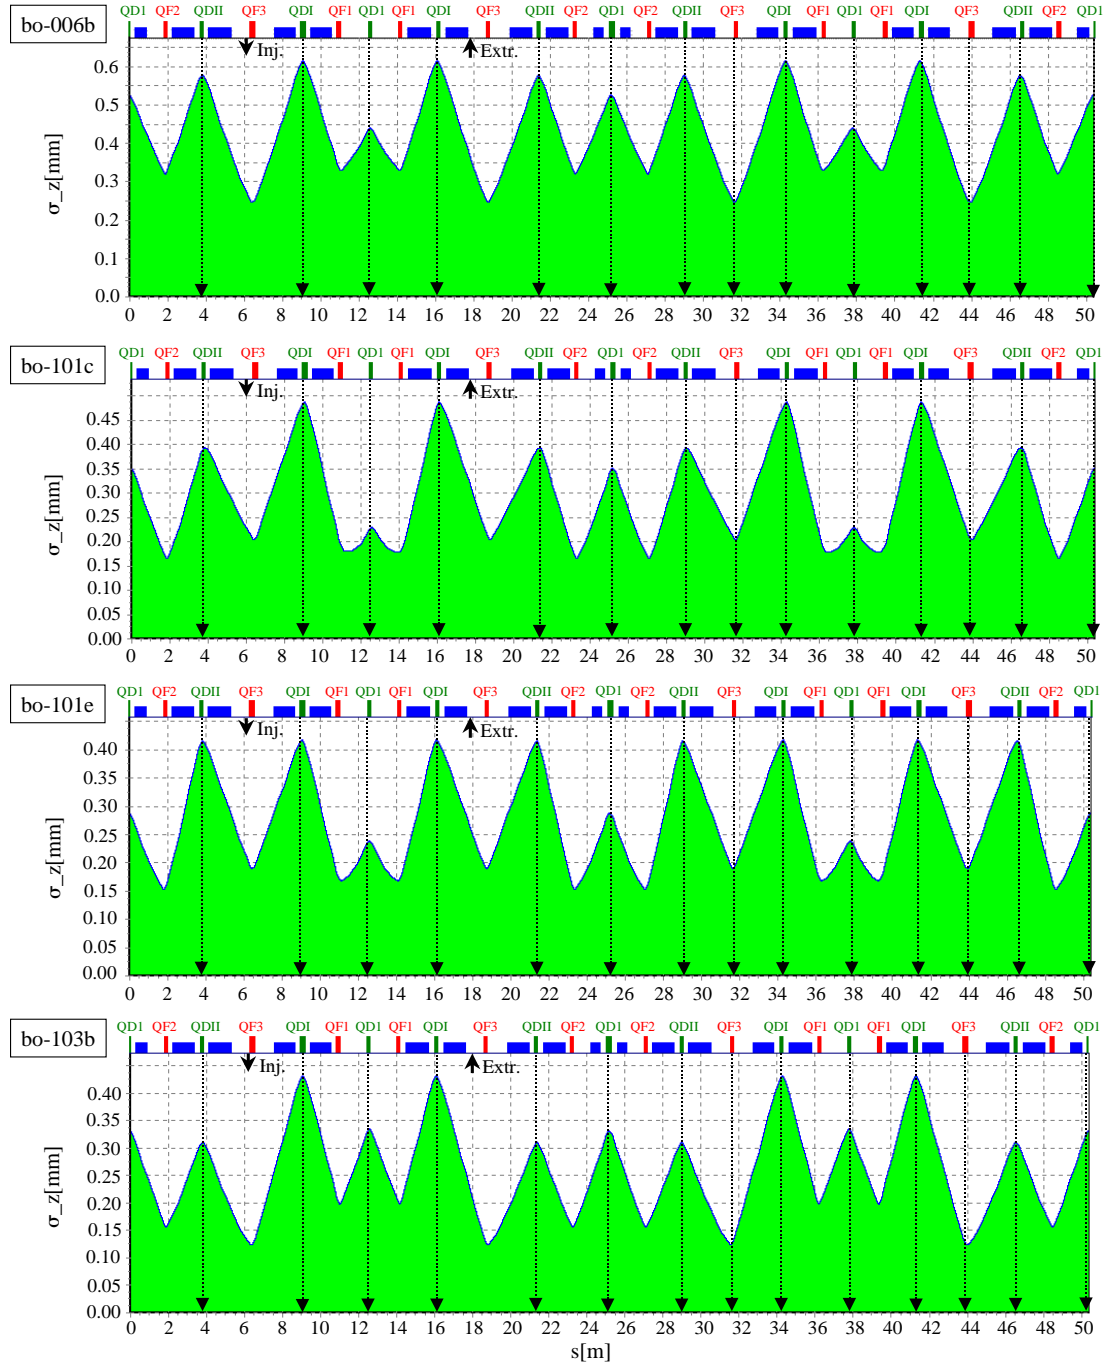
Figure A.8: *Calculated vertical Bodo beam cross section (1 standard deviation, 10 % emittance coupling) at 1.5 GeV. From top to bottom: for "bo-006b", "bo-101c", "bo-101e" and "bo-103b" optics. Emittance and energy spread are equilibrium values that are determined by energy and optics. The locations of the 14 BPMs are indicated by black arrows from the respective quadrupole to the s-axis.*

Appendix B

Aperture and Acceptance Limitations

B.1   Transfer Line T2

Table B.1 contains the theoretical beam cross sections (1 standard deviation) and apertures at the Bodo injection septum, extraction septum (as seen from the stored beam), extraction septum slit (as seen from the extracted beam that moves through the middle of the slit) and at the maxima of optical functions and beam cross section in Bodo at 1.5 GeV. The table also contains the number of standard deviations that fit into the respective aperture.

Parameter / Optics	bo-006b	bo-101c	bo-101e	bo-103b
$\varepsilon_x[nm\,rad]$	429.3	204.8	178.5	131.7
$\varepsilon_z[nm\,rad]$	47.70	22.75	19.83	14.63
$\Delta p/p[10^{-4}]$	6.58	6.73	6.76	6.84
$\beta_{x,\sigma max}[m]$	$=\beta_{x,\beta max}$	$=\beta_{x,Dmax}$	$=\beta_{x,\beta max}$	12.656
$D_{x,\sigma max}[m]$	$=D_{x,\beta max}$	$=D_{x,Dmax}$	$=D_{x,\beta max}$	0.693
$\sigma_{x,\beta max}[mm]$	2.152	1.420	1.561	1.340
$\sigma_{x,Dmax}[mm]$	1.802	1.507	1.561	1.343
$\sigma_{x,max}[mm]$	2.152	1.507	1.561	1.375
$\sigma_{x,isept}[mm]$	1.972	1.327	1.413	1.200
$\sigma_{x,esept}[mm]$	1.703	1.093	1.100	0.880
$apert_{x,max}=37mm/\sigma_{x,max}$	17.2	24.6	23.7	27.6
$apert_{x,isept}=29mm/\sigma_{x,isept}$	**14.7**	**21.9**	**20.5**	**24.2**
$apert_{x,esept}=28mm/\sigma_{x,esept}$	16.4	25.6	25.5	31.8
$apert_{x,esept-slit}=3mm/\sigma_{x,esept}$	**1.76**	**2.74**	**2.73**	**3.41**
$\sigma_{z,\beta max}[mm]$	0.643	0.510	0.434	0.451
$\sigma_{z,esept}[mm]$	0.365	0.309	0.271	0.231
$apert_{z,\beta max}=20mm/\sigma_{z,\beta max}$	**31.1**	**39.2**	**46.1**	**44.3**
$apert_{z,esept-slit}=4mm/\sigma_{z,esept}$	**11.0**	**12.9**	**14.8**	**17.3**

Table B.1: *Theoretical Bodo aperture limitations in Bodo at 1.5 GeV. Emittance and energy spread are theoretical equilibrium values. The $\sigma_{x,...}$ and $\sigma_{z,...}$ values are standard deviations of the resulting beam cross section at different locations in Bodo (i.e. at the injection and extraction septum as seen from the stored beam, at the extraction septum slit, at the maxima of beta functions, dispersion and beam cross section). The lower part of the table contains the number of standard deviations that fit into the respective mechanical aperture. Only 1.76 horizontal standard deviations of the "bo-006b" optics fit through the slit of the extraction septum (5th row from bottom), which limits the booster to storage ring charge transfer efficiency to 92 %. Due to its lower emittance, up to 99.4 % of the electrons (2.74 standard deviations) may pass the septum slit for the "bo-101c" optics.*
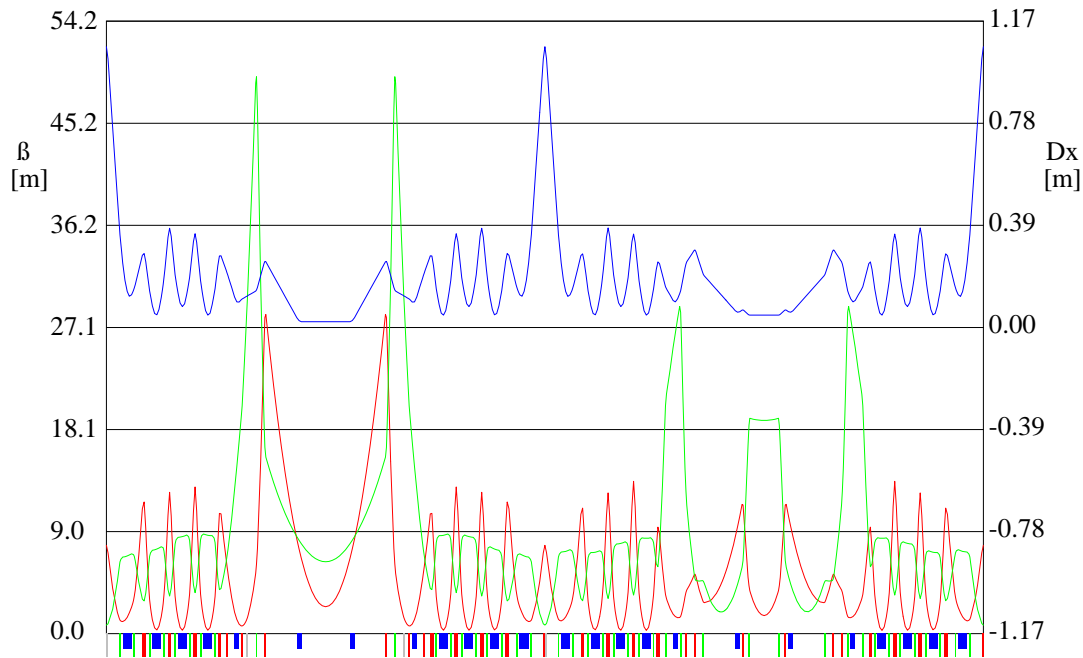
Figure B.1: *Calculated beta functions (red: x, green: z, left axis) and dispersion (blue, right axis) of the Delta optics "del-008a" that is the standard optics for 1.5 GeV operation of the storage ring. The horizontal axis is the s coordinate (left: $s = 0$ m, right: $s = L_{0,Delta} = 115.16$ m).*

B.2 Delta Optics

Figure B.1 shows the calculated optical functions of the Delta optics "del-008" that was operational during the work on this thesis. The optics has a theoretical equilibrium emittance of 21.8 nm rad at 1.5 GeV. Since the Bodo electrons are injected at the fringe of the acceptance ellipse of the stored beam (see fig. 4.8, page 44), they have very large single particle emittances after injection and therefore require a horizontal phase space of some 10 standard deviations. In contrast, their energy spread is very small, and an energy acceptance of 3 standard deviations for the injected beam is sufficient for negligible beam loss during the first turns. Therefore, the acceptance ellipse for the injected electrons is not limited by the largest dispersion (blue curve), but by the two absolute maxima of the horizontal beta function in the left half of the figure (red curve). The resulting theoretical acceptance was used for the plots of the horizontal phase space at the Delta injection septum in section 4.2.4.

Appendix C

Circuit Schematics

The circuit schematics of the DeltaDSP board and its FPGAs consists of 150 pages, therefore it could not be included in this appendix. Instead, the schematics of DeltaDSP board, FPGAs, piggybacks, ADC board and DDS signal generator board (all of which were developed during the work on this thesis) are included in refs. [53, 54, 55].

In order to simplify the development, the design of the DeltaDSP board was organised in a tree that that consists of a number of layers. The top layer of the tree is shown in fig. C.1. Each module (i.e. green block) in the figure defines the interface to a subtree that itself may consists of other subtrees (e.g. FPGAs) etc. By defining interface specifications between the different parts of this hierarchy, the design was split into several sub-designs, which simplified the development and made the design of a subtree independent of the rest of the design as long as the subtree fulfilled its interface specifications.

The DeltaDSP design consists of 25 pages for the mainboard and piggyback printed circuit boards and 125 pages for the 5 FPGAs of the board. Therefore, more than 80 percent of the design can be modified by reprogramming the FPGAs. This will allow the adaption of the board to the requirements of future beam diagnostics and feedback applications without physical modifications of the mainboard. Furthermore, the mainly synchronous design of the FPGAs simplifies the design of future faster versions of the DeltaDSP board.

Figure C.1: *Top level circuit schematics of the DeltaDSP VMEbus board.*

Bibliography

[1] L. Emery, M. Borland, *Top-Up Operation Experience at the Advanced Photon Source*, Proceedings PAC99, New York, 1999

[2] A. Luedeke, M. Munoz, *Top-up Operation Experience at the Swiss Light Source*, Proceedings EPAC2002, Paris, France, 2002

[3] VMEbus International Trade Assiciation (VITA), *VME64 (ANSI/VITA 1-1994)*, PO Box 19658, Fountain Hills AZ 85269-1958, USA

[4] K. Dunkel, *Entwicklung und Test einer Messelektronik zur Arbeitspunktbestimmung am Synchrotron Bodo*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1996, Unpublished (in German)

[5] T. Kamps, *Messung optischer Parameter am Synchrotron Bodo*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1996, Unpublished (in German)

[6] K. Wille, *Physik der Teilchenbeschleuniger und Synchrotronstrahlungsquellen*, Teubner-Verlag, Stuttgart 1992 (in German)

[7] H. Wiedemann, *Particle Accelerator Physics I*, Springer-Verlag, Berlin 1993

[8] H. Wiedemann, *Particle Accelerator Physics II*, Springer-Verlag, Berlin 1995

[9] J.D. Jackson, *Classical Electrodynamics*, 2nd Edition, de Gruyter-Verlag, Berlin 1982

[10] N. Kumagai, H. Ohkuma, K. Soutome, M. Takao, H. Tanaka, *Estimation of Betatron Coupling and Vertical Dispersion for SPring-8 Storage Ring*, Proceedings PAC99, New York, 1999

[11] J. A. Clarke and H. L. Owen, *Measurement of Vertical Dispersion and Coupling in the Daresbury SRS*, Proceedings EPAC96, Sitges, Spain, 1996

[12] D. Robin, E. Forest, C. Pellegrini, A. Amiry, *Quasi-Isochronous Storage Rings*, Physical Review E, 48, p. 2149, 1993

[13] DELTA-Group, *DELTA Status Report*, University of Dortmund, Physics Department, 1990, Unpublished

[14] S.Y. Lee, *Single particle dynamics at synchro-betatron coupling resonances*, Phys. Rev. E 49, p. 5706, 1994

[15] M. Abo-Bakr et al., *The Bessy II Booster Synchrotron*, Proceedings EPAC96, Sitges, Spain, 1996

[16] U. Weinrich et al., *Cycling and Injection Improvements on the ESRF Booster*, Proceedings EPAC96, Sitges, Spain, 1996

[17] G. Guignard, *The General Theory of All Sum and Difference Resonances in a Three-dimensional Magnetic Field in a Synchrotron*, Report No. CERN 76-06, CERN, Geneva, Switzerland, 1976

[18] G. Guignard, *Hamiltonian treatment of synchrobetatron resonances*, CERN Accelerator School: Course on Advanced Accelerator Physics, Vol. 95-06 v 1, CERN, Geneva, Switzerland, 1994

[19] M. Sands, *The Physics of Electron Storage Rings. An Introduction.* SLAC Publication SLAC-121, 1970

[20] T. Nakamura et al., *Low Emittance Operation of the Spring-8 Storage Ring by Damping Partition Control*, Proceedings PAC2001, Chicago, USA, 2001

[21] T. Weis, DELTA Machine Group, Private Correspondence

[22] S. Y. Lee, *Accelerator Physics*, World Scientific, Singapore, 1999

[23] V. Schlott, *Global Position Feedback in Synchrotron Radiation Sources*, Proceedings EPAC2002, Paris, France, 2002

[24] A. Jankowiak, *Strahldiagnose und Closed-Orbit-Charakterisierung mit HF-Strahllagemonitoren am Beispiel der Synchrotronstrahlungsquelle DELTA*, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, 1999, Unpublished (in German)

[25] G. Strang, *Linear Algebra and its Applications*, 2nd Edition, Academic Press, New York, 1980

[26] H. Grote, F.C. Iselin, *The MAD Program (Methodical Accelerator Design), Version 8.19*, CERN, Geneva, Switzerland, 1996

[27] K. Wille, *PC Program "Optics", Version 3.18*, DELTA, University of Dortmund, Germany, 2002

[28] Private Communications, DELTA Power Supply Group

[29] F. Brinker, *Variable integrierte Sextupole für Speicherringe*, Ph.D. thesis, DELTA, University of Dortmund, Physics Department, 1993, Unpublished (in German)

[30] N. Marquardt et al., *The DELTA Vacuum System*, Proceedings EPAC96, Sitges, Spain, 1996

[31] M. Michel, *Wirbelströme in der Vakuumkammer des DELTA-Booster-Synchrotrons*, Internal Report 89-002, DELTA, University of Dortmund, Germany, 1989, Unpublished (in German)

[32] D. Zimoch, *Implementierung eines Orbitkorrektursystems an der Synchrotronstrahlungsquelle DELTA*, Ph.D. thesis, DELTA, University of Dortmund, Physics Department, 2002, Unpublished (in German)

[33] E. Kasel, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, in Preparation

[34] R. Heine, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, in Preparation

[35] A. Kwiatkowski, Diploma Thesis, DELTA, University of Dortmund, Physics Department, in Preparation

[36] G. Schmidt et al., *Position Sensors For Monitoring Accelerator Magnet Motion at DELTA*, Proceedings EPAC2002, Paris, France, 2002

[37] W. Joho et al., *The SLS Booster Synchrotron*, Proceedings EPAC98, Stockholm, Sweden, 1998

[38] B. Keil, K. Wille, *Control of a Ramped 1.5 GeV Electron Storage Ring by a Multiprocessor DSP Network*, Proceedings EPAC2002, Paris, France, 2002

[39] T. Bizen, *Demagnetization of undulator magnets irradiated high energy electrons*, Proceedings of the Accelerator Reliability Workshop, ESRF, 2002

[40] A. Peters, *Neue Elektronenoptiken der Synchrotronstrahlungsquelle DELTA*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1996, Unpublished (in German)

[41] Private Communications, G. Schmidt, DELTA Machine Group

[42] P.J. Bryant, *Beam Transfer Lines*, Proceedings of CERN Accelerator School 94-01, p. 219ff, CERN, Geneva, Switzerland, 1994

[43] M. A. J. Michel, *Simulationen des Verhaltens von Ionen in Elektronenspeicherringen, sowie deren Wechselwirkungen mit dem gespeicherten Strahl*, Ph.D. thesis, DELTA, University of Dortmund, Physics Department, 1994, Unpublished (in German)

[44] C. Steier et al., *Orbit Feedback Development at the ALS*, Proceedings EPAC2002, Paris, France, 2002

[45] M. Böge, M. Dehler, T. Schilcher, V. Schlott, R. Ursic, *Fast Closed Orbit Control in the SLS Storage Ring*, Proceedings PAC99, New York, 1999

[46] E. Plouviez, J.M. Koch, F. Uberto, *A Fast Global Feedback System to Correct the Beam Position Deviation in the ESRF Storage Ring*, Proceedings of the International Conference on Accelerator and Large Experimental Physics Control Systems, Trieste, Italy, 1999

[47] Glenn Decker, Om Singh, *Orbit Feedback Using X-Ray Beam Position Monitoring at the Advanced Photon Source*, Proceedings ICALEPCS2001, San Jose, California, USA, 2001

[48] J.A. Carwardine, F.R. Lenkszus, *Real-Time Orbit Feedback at the APS*, AIP Conference Proceedings No. 451, p. 125, 1998

[49] G. Decker, O. Singh, *Strategy for Achieving True Sub-Micron Orbit Stabilization at the Advanced Photon Source*, Proceedings PAC2001, Chicago, USA, 2001

[50] Analog Devices Inc., *Datasheet ADSP-21062 Rev. C*, Norwood, USA, 2000 (http://www.analog.com)

[51] Analog Devices Inc., *ADSP-2106x Sharc User's Manual, 2nd Edition*, Norwood, USA, 1997

Bibliography 167

[52] B. Keil, *Charge Transfer Optimisation of Bodo and Transfer Line Optics*, Internal Report 2003-001, DELTA, University of Dortmund, Germany, 2003, Unpublished (in German)

[53] B. Keil, *The DeltaDSP Board - Printed Ciruit Board and FPGA Schematics Manual*, Internal Report 2003-002, DELTA, University of Dortmund, Germany, 2003, Unpublished (in German)

[54] B. Keil, *The DeltaDSP DDS Generator Board - Circuit Schematics and Technical Data*, Internal Report 2003-003, DELTA, University of Dortmund, Germany, 2003, Unpublished (in German)

[55] B. Keil, *The DeltaDSP ADC Board - Circuit Schematics and Technical Data*, Internal Report 2003-004, DELTA, University of Dortmund, Germany, 2003, Unpublished (in German)

[56] N. Koch, K. Biermann, *Führungsgrößengenerator für das DELTA-Kontrollsystem*, Internal Report 93-001, DELTA, University of Dortmund, Germany, 1993, Unpublished (in German)

[57] R. Müller et al., *Orbit Control at Bessy II: Present Performance and Plans*, Proceedings EPAC2000, Vienna, Austria, 2000

[58] M. Böge, J. Chrin, P. Pollet, T. Schilcher, V. Schlott, *Orbit Control at the SLS Storage Ring*, Proceedings EPAC2002, Paris, France, 2002

[59] A. S. Tanenbaum, *Computernetzwerke*, Prentice Hall, Germany, 1997 (in German)

[60] B. Keil, *Bau und Test eines Strahlverlustmonitor-Systems für Elektronenspeicherringe*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1996, Unpublished (in German)

[61] O. Walter, *Ein System zur Messung und Stabilisierung des Arbeitspunktes am Booster-Ring von DELTA*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1998, Unpublished (in German)

[62] A. Lüdeke, *Technische Dokumentation des neuen DELTA Timing Systems*, Internal Report 98-12, DELTA, University of Dortmund, 1998, Unpublished (in German)

[63] A. Lüdeke *Entwurf und Aufbau effizienter Zustandsdiagnostik für das Booster-Synchrotron Bodo*, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, 1999, Unpublished (in German)

[64] C. Steier, A. Biocca, E. Domning, S. Jacobson, G. Portmann, Y. Wu, *Design of a Fast Global Orbit Feedback System for the Advanced Light Source*, Proceedings of the 2001 Particle Accelerator Conference, Chicago, USA, 2001

[65] Agilent Technologies Inc., *Datasheet of the HFBR-5205 fibre optics transceiver*, 1999

[66] B. Keil, EMW User's Guide, Internal DELTA Report, 1999, http://www.delta.uni-dortmund.de/controls/pub/doc/emw/

[67] C. Piel, H. Vogel, P. Vom Stein, G. Blokesch , D. Kraemer, *Design and Construction of a Turn-Key 100 MeV Linac for the Swiss Light Source*, Proceedings EPAC2000, Vienna, Austria, 2000

[68] J. K. Osterhout, *Tcl and the Tk Toolkit*, Addison-Wesley, USA, 1995

[69] C. Smith, *incr tcl/tk from the Ground Up*, MCGraw-Hill Professional, 1st Edition, USA, 2000

[70] K. Evans, *An Overview of MEDM*, Proceedings of the International Conference on Accelerator and Large Experimental Physics Control Systems, Trieste, Italy, 1999

[71] L. R. Dalesio et al., *EPICS Architecture*, Los Alamos National Laboratory, Argonne National Laboratory, USA, 1993

[72] S. A. Lewis, *Overview of the Experimental Physics and Industrial Control System: EPICS*, Lawrence Berkeley National Laboratory, USA, 2000

[73] B. W. Kerninghan, D. M. Ritchie, *The C Programming Language, Second Edition, ANSI C*, Prentice Hall Inc., London, 1988

[74] B. Stroustrup, *The C++ Programming Language, Second Edition*, AT & T Bell Laboratories Inc., New Jersey, USA, 1993

[75] *VxWorks Version 5.3*, Wind River Systems Inc., Alameda, CA, USA, 1997, http://www.windriver.com

[76] G. Schmidt, E. Kasel, K. Wille, M. Körfer, H. Henschel, J. Kuhnhenn, *Optical Fibre Beam Loss Monitors for Storage Rings at DELTA*, Proceedings EPAC2002, Paris, France, 2002

[77] A. Hofmann, *Tune Shifts from Self-Fields and Images*, Proceedings of CERN Accelerator School 94-01, p. 329ff, CERN, Geneva, Switzerland, 1994

[78] U. Scholz, *Design und Ansteuerung eines Kickers für das transversale Feedbacksystem von DELTA*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1994, Unpublished (in German)

[79] L. Tosi, V. Smaluk, D. Bulfone, E. Karantzoulis, M. Lonza, *Diagnostics and Analysis of Instabilities with the Digital Transverse Multibunch Feedback at Elettra*, Proceedings PAC2001, Chicago, USA, 2001

[80] *Beam Position Monitor User's Manual*, Rev. 1.5.2, BERGOZ Precision Beam Instrumentation, Crozet, France, 1999

[81] H. Schmickler, *Diagnostics and Control of the Time Evolution of Beam Parameters*, Proceedings DIPAC97, Frascati, Italy, 1997

[82] K.H. Mess, *Time and Frequency Domain Measurements*, Proceedings DIPAC97, Frascati, Italy, 1997

[83] A. Stillman, *Wavelet Analysis and Accelerator Signals*, Proc. of the Beam Instrumentation Workshop BWI94, Vancouver, Canada, 1994

[84] D. Heins et al, *The Transverse Damping System with PLL Tune Measurement for HERA-p*, Proceedings EPAC96, Sitges, Spain, 1996

[85] O.Berrig, M.Jonker, K.D.Lohmann, G.Morpurgo, *The Q-Loop : A Function Driven Feedback System for the Betatron Tunes During the LEP Energy Ramp*, Proceedings EPAC98, Stockholm, Sweden, 1998

[86] P. Cameron, P. Cerniglia, R. Connolly, J. Cupolo, W.C. Dawson, C. Degen, A. DellaPenna, J. DeLong, A. Drees, A. Huhn, M. Kesselman, A. Marusic, B. Oerter, J. Mead, C. Schultheiss, R. Sikora, J. van Zeijts, *Tune Feedback at RHIC*, Proceedings PAC2001, Chicago, USA, 2001

[87] G. Wu, B. C. Brown, D. P. Capista, R. H. Flora, D. E. Johnson, K. S. Martin, *Tune Control in the Fermilab Main Injector*, Proceedings PAC99, New York, 1999

[88] J. Dietrich, I. Mohos, *Broadband FFT Method for Betatron Tune Measurements in the Acceleration Ramp at COSY-Jülich*, Proceedings of the Beam Instrumentation Workshop BIW98, SLAC, USA, 1998

[89] D. Heins, K.H. Matthiessen, J. Pätzold, D. Renken, K. Wille, *The Transverse Feedback System at PETRA*, DESY M-79/06 Internal Report, Hamburg, Germany, 1979

[90] H.Koziol, *Q Measurement for the PS Booster*, Internal Report CERN-MPS/Int. BR/74-14, CERN, Geneva, 1974.

[91] Analog Devices Inc., *AD9850 datasheet, Rev. E*, Norwood, USA, 1999

[92] Bergoz Instrumentation Inc., *Modular Parametric Current Transformer User's Manual, Rev. 1.4*, Crozet, France, 1999

[93] D. Zimoch, *Entwicklung von Strahldiagnosetools unter Berücksichtigung objektorientierter Ansätze*, Diploma Thesis, DELTA, University of Dortmund, Physics Department, 1996, Unpublished (in German)

[94] M. Körfer *Entwurf und Aufbau einer Bestrahlungsanlage für die Röntgentiefenlithographie mit periodischer Bewegung der Synchrotronstrahlung am Speicherring Delta*, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, 1997, Unpublished (in German)

[95] J. Doyle, B. Francis, A. Tannenbaum, *Feedback Control Theory*, Macmillan Publishing Co., 1990

[96] K. Ogata, *Modern Control Engineering*, Prentice Hall, 4th edition, 2001

[97] D. Bulfone, V. Forchi, M. Lonza, L. Zambon, M. Dehler, *The Elettra Digital Multi-Bunch Feedback Systems*, Proceedings EPAC2002, Paris, France, 2002

[98] S. Khan, T. Knuth, W. Barry, J. Byrd, G. Stover, *Commissioning Results of the Transverse Feedback System at Bessy II*, Proceedings EPAC2000, Vienna, Austria, 2000

[99] M. Dehler, V. Schlott, D. Bulfone, M. Lonza, R. Ursic, *Current Status of the ELETTRA/SLS Transverse Multibunch Feedback*, Proceedings EPAC2000, Vienna, Austria, 2000

[100] *Instruction Manual for BOS/S Power supply, Model BOS/S 36.12.2.0841*, Electronic Measurements Inc., Neptune, NJ, USA, 1996

[101] *Instruction Manual BOP 36-12M Bipolar Operational Power Supply*, Kepco Inc., Flushing, NY., USA, 2000

[102] M. Grewe, *Tcl/Tk Matrix Calculation Package Version 3.0*, Internal DELTA Control System Documentation, 2002, Unpublished

[103] S. Schlögl, *Einsatz von PIN-Photodioden als Protonen-Strahlverlustmonitore bei HERA*, Diploma Thesis, University of Hamburg, Germany, 1992, Unpublished (in German)

[104] U. Weinrich, *Mastering Beam Losses on Small Gap Vacuum Chambers in Synchrotron Light Sources*, Ph.D. Thesis, DELTA, University of Dortmund, Physics Department, 1999, Unpublished (in German)

Danksagung