Development and commissioning of a digital rf control system for the S-DALINAC and migration of the accelerator control system to an EPICS-based system

Vom Fachbereich Physik
der Technischen Universität Darmstadt

zur Erlangung des Grades
eines Doktors der Ingenieurwissenschaften
(Dr.-Ing.)

genehmigte

D i s s e r t a t i o n

angefertigt von

Dipl.-Phys. Martin Konrad
aus Hanau

Darmstadt 2013
D 17

Referent:	Professor Dr. rer. nat. Dr. h. c. N. Pietralla
Korreferent:	Professor Dr.-Ing. H. Klingbeil
Tag der Einreichung:	12. Februar 2013
Tag der Prüfung:	15. April 2013

Zusammenfassung

Die am supraleitenden Darmstädter Elektronenlinearbeschleuniger S-DALINAC durchgeführten hochauflösenden Streuexperimente erfordern eine hohe Energieschärfe des Strahls von $(\Delta E/E) \approx 1 \times 10^{-4}$. Dies erfordert die Stabilisierung von Amplitude und Phase des elektrischen Felds im Inneren der Beschleunigungsstrukturen auf $(\Delta A/A)_{\mathrm{rms}} = 8 \times 10^{-5}$ und $(\Delta\varphi)_{\mathrm{rms}} = 0.7°$. Gegenstand dieser Arbeit ist die Entwicklung und Inbetriebnahme der dafür notwendigen digitalen Hochfrequenzregelung.

Am S-DALINAC werden zwei Arten von Kavitäten verwendet. Die normalleitenden Chopper- und Buncher-Kavitäten benötigen lediglich Korrekturen für langsame Temperaturschwankungen und können durch einen Algorithmus geregelt werden, bei dem die Kavität von einem Generator getrieben wird. Die supraleitenden Beschleunigungsstrukturen weisen hingegen eine sehr hohe Güte auf, wodurch sie sehr empfindlich für Vibrationen sind. Daher werden sie in einer selbsterregten Schleife betrieben.

Die Hochfrequenzregelung basiert auf hausintern entwickelter Hardware, welche das Hochfrequenzsignal ins Basisband heruntermischt, es digitalisiert und in einen FPGA einspeist. In diesem FPGA arbeitet ein einprogrammierter digitaler Signalprozessor den Regelalgorithmus ab. Die resultierende Korrektur wird wieder auf das Hochfrequenzsignal moduliert und zur Kavität geschickt.

Alle Beschleunigerkomponenten werden von einem zentralen Kontrollraum ferngesteuert. Da komplexe und reprogrammierbare Geräte vom bestehenden Kontrollsystem nicht gut unterstützt werden, war der Entwurf und der Aufbau eines neuen Kontrollsystems ebenfalls Teil dieser Arbeit. Entscheidende Aspekte sind dabei die Erweiterbarkeit, Bedienbarkeit und Wartbarkeit des Systems. Das neue Beschleunigerkontrollsystem verwendet das EPICS-Framework als Basis, da dieses bereits einen Großteil der Basisfunktionalität wie graphische Benutzeroberflächen und flexible und schnell anpassbare Kontrollserver bereitstellt. Dies ermöglicht die Umsetzung komplexer Funktionalität wie z. B. eine umfassende Auslese und Diagnose für das HF-Regelungssystem.

Die erfassten Daten können mit einem Software-Oszilloskop und einer Spektrumanalysesoftware dargestellt werden. Weiterhin werden kontinuierlich RMS-Fehler

berechnet, die genutzt werden können, um die Regelparameter sehr präzise zu optimieren und die Leistung der Regler zu überwachen.

Messungen ergaben, dass die Stabilität der Hochfrequenzregelung gegenüber dem analogen System um eine Größenordnung verbessert werden konnte, sodass nun eine Phasenstabilität von $(\Delta\varphi)_{\mathrm{rms}} = 0.8^\circ$ und eine Amplitudenstabilität von $(\Delta A/A)_{\mathrm{rms}} = 7 \times 10^{-5}$ erreicht wird. Damit wird die Spezifikation im Rahmen der Messgenauigkeit erfüllt.

Die beschriebene Hochfrequenzregelung wurde vor zwei Jahren in Betrieb genommen und wird seitdem erfolgreich für den Strahlbetrieb verwendet. Während dieser Zeit erwies sich die Hochfrequenzregelung als wesentlich stabiler und zuverlässiger als das alte analoge System.

Abstract

The high resolution scattering experiments conducted at the superconducting Darmstadt electron linear accelerator S-DALINAC call for a small energy spread of $(\Delta E/E) \approx 1 \times 10^{-4}$ of the beam. This requires stabilization of amplitude and phase of the electric field inside the accelerating cavities to $(\Delta A/A)_{\mathrm{rms}} = 8 \times 10^{-5}$ and $(\Delta\varphi)_{\mathrm{rms}} = 0.7°$. The design and the commissioning of a new digital rf control system is the subject of this thesis.

At the S-DALINAC two types of cavities are in use. The normal-conducting chopper and buncher cavities only need corrections for slow temperature drifts and can be controlled by a generator-driven resonator control algorithm. The superconducting accelerating cavities have a very high quality factor and thus are very susceptible to vibrations. Therefore they are operated in a self-excited loop.

The rf control system is based on in-house developed hardware that converts the rf signal down to the baseband, digitizes it and feeds it into an FPGA. Inside this FPGA, a soft digital signal processor executes the control algorithm. The resulting correction is modulated onto the rf signal again and sent back to the cavity.

All accelerator components are remote-controlled from a central room via an accelerator control system. Since complex and re-programmable devices are not supported well by the existing in-house developed control system, the design and implementation of a new accelerator control system is also subject of this thesis. Further important aspects are expandability, usability and maintainability of the system. Therefore the new accelerator control system uses the EPICS framework as a basis since it already provides much of the basic functionality like graphical user interfaces and flexible control servers that can be customized rapidly. This allowed the implementation of more advanced functionality like extensive read-out and diagnostics for the rf control system.

The read out data can be visualized with a software oscilloscope and a spectrum analyzer software. Additionally the system provides on-line rms errors that can be used to optimize the control parameters very precisely and to monitor the performance of the controllers.

Measurements show that the performance of the rf control system has been improved by one order of magnitude compared to the analog system, yielding a phase stability of $(\Delta\varphi)_{\mathrm{rms}} = 0.8^\circ$ and an amplitude stability of $(\Delta A/A)_{\mathrm{rms}} = 7 \times 10^{-5}$ and thus meeting the specification.

The described rf control system has been commissioned and successfully used for beam operation for two years. During this time the system has proven to be significantly more stable and reliable than the old analog system.

Contents

1 Introduction

1.1 S-DALINAC

The superconducting Darmstadt linear accelerator S-DALINAC is used for experiments in nuclear physics and astrophysics since 1987 [1]. It provides electron beams within an energy range of 2 MeV to 130 MeV. The beam current can be adjusted from fractions of a nanoampere up to 60 μA. Since the accelerator is operated in continuous wave (cw) mode it is particularly suited for coincidence experiments. Superconducting accelerating cavities made from high-purity niobium are used to allow cw operation at the required beam energies. For stable operation these cavities have to be cooled to a temperature of 2 K.

Figure 1.1 shows an overview of the S-DALINAC. The electron beam is produced by one of two electron guns. On the one hand there is a thermionic gun that accelerates the beam electrostatically to an energy of 250 keV. On the other hand a new photoelectric gun providing a spin-polarized electron beam of up to 100 keV has been commissioned recently [3, 4].

The continuous electron beam from the electron gun is prepared for radio frequency (rf) acceleration by chopper and buncher cavities. Subsequently it is injected into the superconducting accelerating cavities of the injector linac. In the injector linac the electrons are accelerated to an energy of up to 10 MeV. The electron beam can be used at an experimental area for nuclear resonance fluorescence experiments [5, 6].

Alternatively the beam can be directed into the main linac where it gains up to 40 MeV. The beam can be fed back into the main linac twice by transporting it back with recirculation beam lines. This allows to use the energy gain of the main linac up to three times resulting in a maximum energy of 130 MeV. Electrons of these three different energies are separated by a special dipole magnet at the end of the main linac and recombined at the end of the recirculation beam lines.

When the desired beam energy is reached, the beam is extracted and transported to the experimental hall (see fig. 1.2) where it can be used for scattering experiments at the 169° high resolution scattering facility [7], the 180° QCLAM spectrometer [8, 9, 10] or the photon tagger NEPTUN [11, 12].

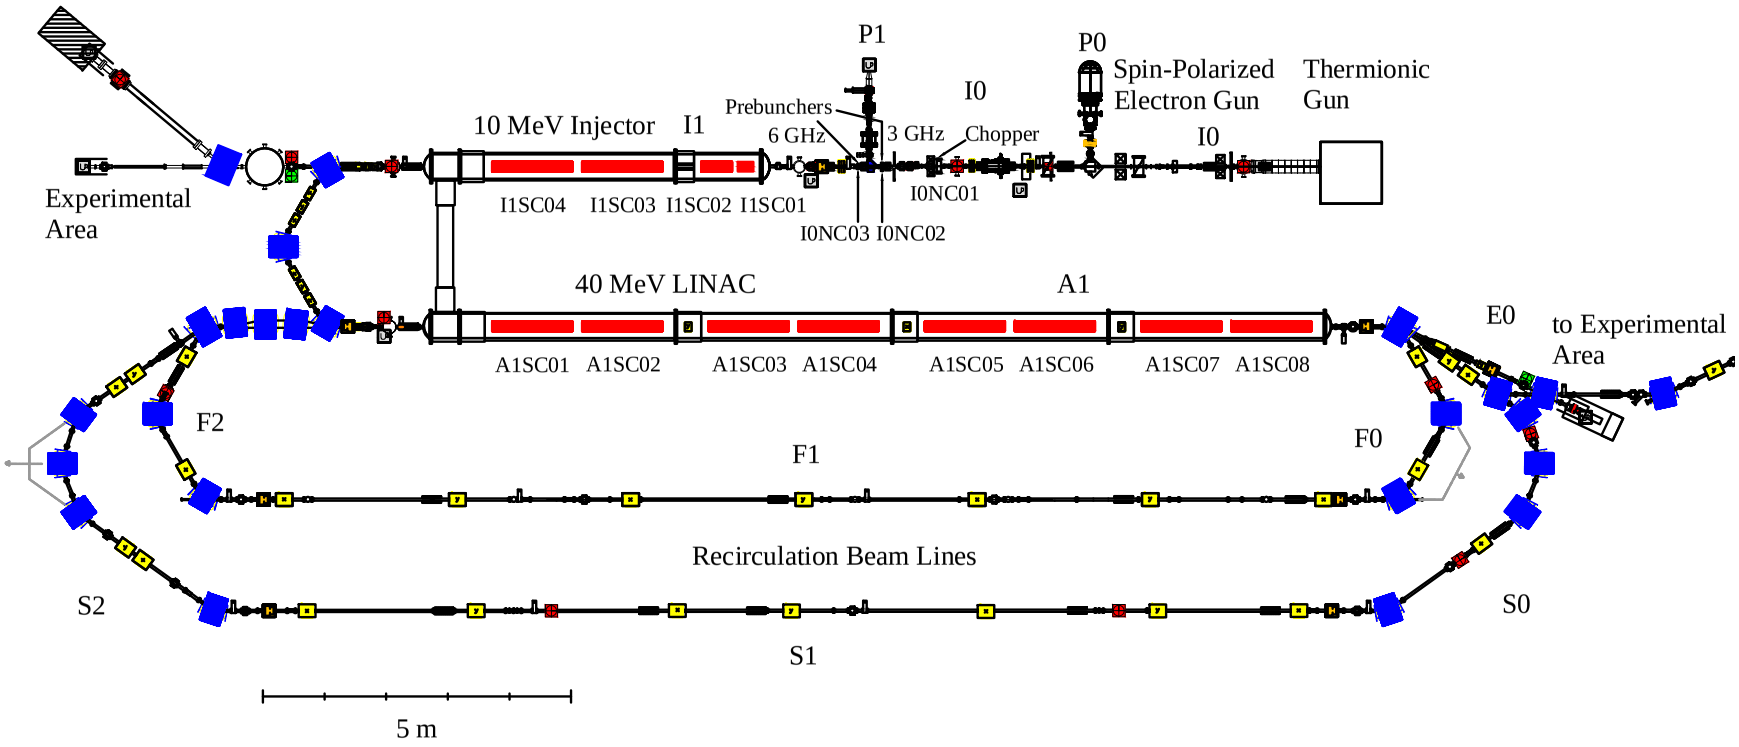

Figure 1.1: Floor plan of the superconducting Darmstadt electron linear accelerator S-DALINAC (based on [1, 2]).

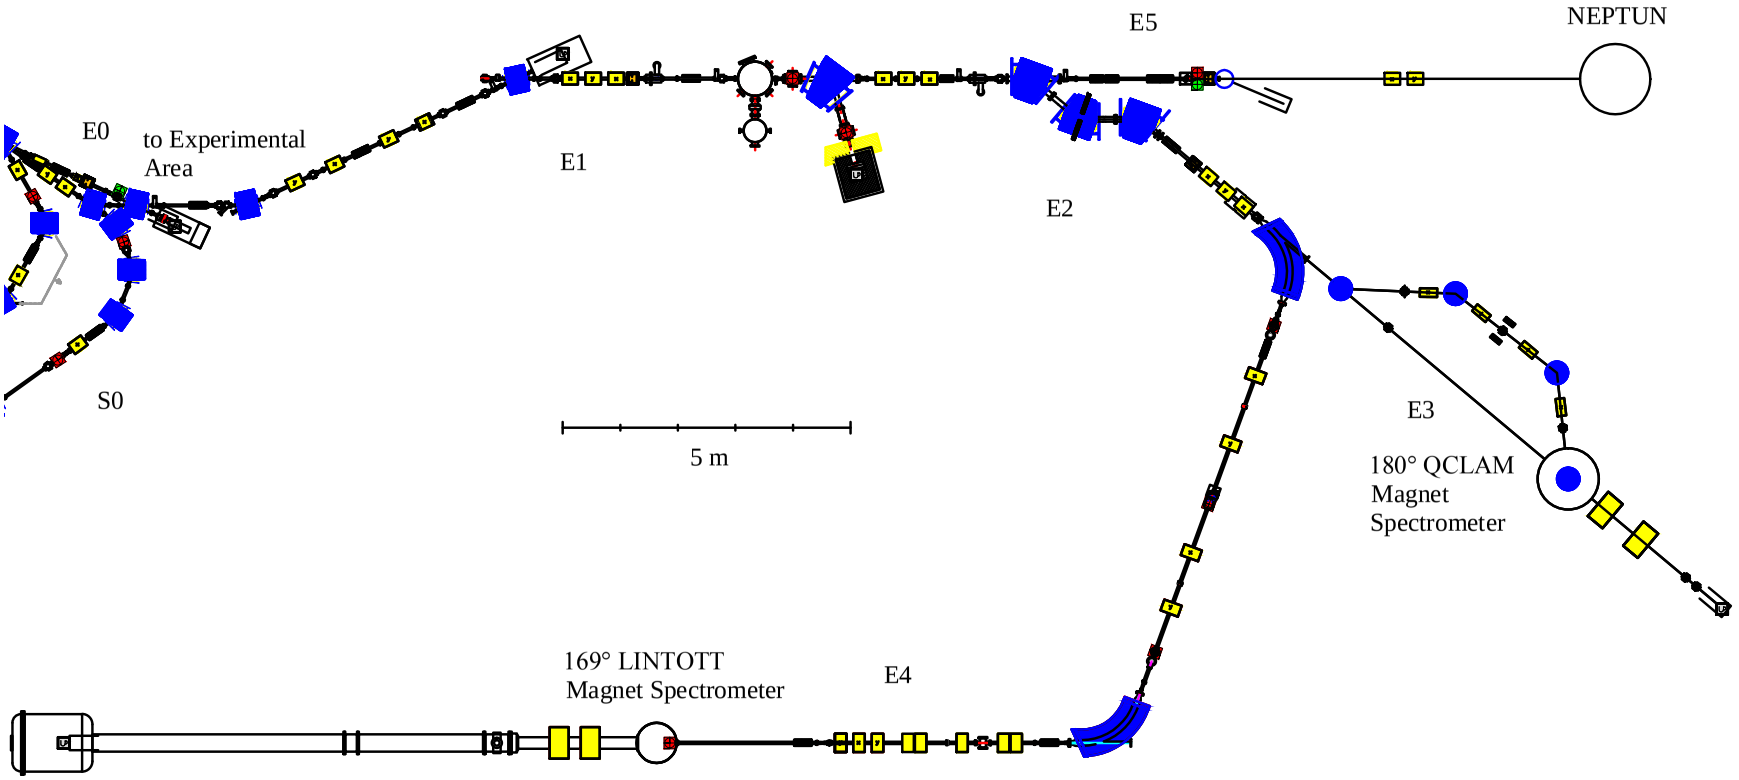

Figure 1.2: Floor plan of the extraction beam-line and the experimental areas of the S-DALINAC (based on [2]).

The analysis of nuclei with a high level density in scattering experiments requires a high energy resolution. One of the most important factors affecting the resolution is the energy spread of the initial electron beam.

1.2 RF acceleration

The electron energy of the S-DALINAC is too high to be reached by acceleration in a static electric field. Instead, the electrons are accelerated by an rf field that oscillates inside the accelerating cavities. These cavities consist of multiple elliptical cells that make up a set of coupled electric resonators. Each of these resonators has a set of eigenmodes from which only the transverse magnetic TM_{0np} modes have a longitudinal electric field on the beam axis that can be used for acceleration. Like most other accelerators the S-DALINAC uses the TM_{010} mode for acceleration because it has the lowest eigenfrequency. The cavities of the S-DALINAC are made of 20 cells and thus the TM_{010} mode is split into 20 eigenmodes. From these modes the π mode is used at the S-DALINAC because it provides the highest acceleration. In this mode the field in consecutive cells is always directed in opposite directions.

When an electron enters the first cell of the cavity it is accelerated by the field (see fig. 1.3). While the electron passes the iris to the next cell the sign of the field changes, leading to an acceleration in the next cell as well. Thereby the electron is accelerated in each cell of the cavity. For energies above 1 MeV the electrons are highly relativistic and essentially do not further gain speed due to acceleration but only kinematic mass. Thus the same type of 20-cell cavity can be used for every beam energy in this relativistic energy range. For capturing the low energy beam from the electron gun a 2-cell cavity optimized for lower electron speed and a 5-cell cavity are used [13].

RF cavities can only accelerate short bunches of electrons. Hence the continuous electron beam from the electron gun has to be chopped before it reaches the superconducting cavities. This is achieved by deflecting the beam to a circular path with a revolution frequency of 3 GHz with an rf chopper cavity (further referred to as `I0NC01`). A mask is used to cut out a segment of this circle resulting in pulses of approximately 36° with respect to the rf phase [4]. The length of these pulses is reduced to approximately 2° by a 3 GHz bunching cavity (`I0NC02`). An

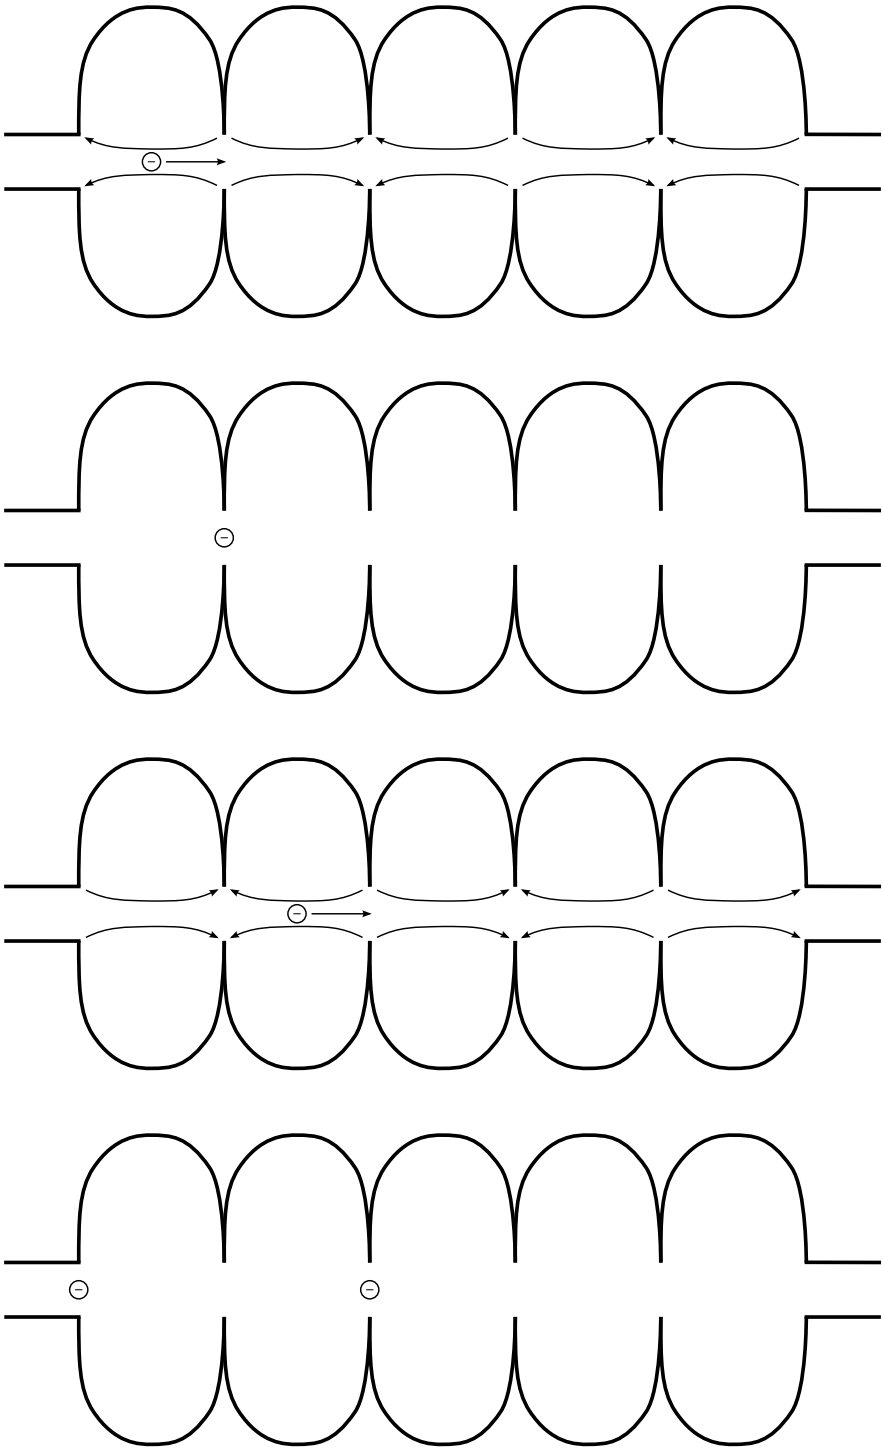

Figure 1.3: RF acceleration in an elliptical cavity.

Cavity name	I0NC01	I0NC02	I0NC03
Material		High purity copper	
Operating temperature (°C)	20.0 ± 0.1	20.0 ± 0.1	20 ± 5
Operating frequency (GHz)	2.997	2.997	5.994
Loaded quality factor Q_L	5278 ± 3.6 [14]	3182 ± 3.4	1018 ± 4 [14]
Power dissipation 100 keV (W) [15]	70	4	0.65
Power dissipation 200 keV (W) [15]	95	42	7
Water cooled	yes	yes	no

Table 1.1: Parameters of the normal-conducting chopper and buncher cavities.

additional 6 GHz buncher (I0NC03) is needed to achieve this for the lower beam energy of the source for spin-polarized electrons. Table 1.1 lists the parameters of the normal-conducting cavities.

The highest electron energy can be achieved if the electron bunches pass the cavity cells when the electric field reaches its maximum (on-crest acceleration). The energy gained by an electron in a cavity is

$$E = E_{\max} \cos \varphi \tag{1.1}$$

where $E_{\max}$ is the maximum energy gain and φ is the phase between the electron and the electric field. Fluctuations in amplitude and phase lead to an error in the beam energy. Since the errors in amplitude $\Delta E_{\max}$ and phase $\Delta\varphi$ are partly correlated, the error in beam energy caused by a single cavity is smaller than the maximum error

$$\left(\frac{\Delta E}{E_{\max}}\right)_{\text{cav,max}} = \left|\frac{\Delta E_{\max}}{E_{\max}}\right| + |1 - \cos \Delta\varphi| \tag{1.2}$$

but larger than the error obtained for uncorrelated errors:

$$\left(\frac{\Delta E}{E_{\max}}\right)_{\text{cav,rms}} = \sqrt{\left(\frac{\Delta E_{\max}}{E_{\max}}\right)^2 + (1 - \cos \Delta\varphi)^2} \tag{1.3}$$

Material	Niobium
Operating temperature	2 K
Operating pressure	35 mbar
Acceleration mode	TM_{010}, π mode
Operating frequency	2.997 GHz
Number of cells	2, 5, 20
Loaded quality factor Q_L	3×10^7

Table 1.2: Design parameters of the superconducting accelerating cavities.

An additional error is introduced by the length of the bunches in the accelerating cavities of 2°. The mean normalized energy gain of an electron passing the cavity

$$m = \frac{1}{2^\circ} \int_{-1^\circ}^{+1^\circ} \cos\varphi \, \mathrm{d}\varphi \tag{1.4}$$

is slightly lower than the energy gain for an electron arriving at $\varphi = 0$. The error in beam energy introduced by the bunch length is

$$\left(\frac{\Delta E}{E_{\max}}\right)_{\text{bunch,rms}} = \sqrt{\frac{1}{2^\circ} \int_{-1^\circ}^{+1^\circ} (m - \cos\varphi)^2 \, \mathrm{d}\varphi} = 4.5 \times 10^{-5} \quad . \tag{1.5}$$

This error cannot be compensated by the rf control system and thus determines the best energy spread the S-DALINAC can provide in conventional mode of operation. With $(\Delta E/E)_{\text{rms}} \approx 1 \times 10^{-4}$ the target specification for the rf control system has been defined in the same order of magnitude. In addition to the normal on-crest acceleration mode a non-isochronous recirculation mode with phase focusing can be used to reduce the energy spread of the beam at the price of a slightly reduced energy [2, 16].

Continuous wave operation at reasonably high accelerating gradients is only possible with accelerating cavities that have a very high quality factor (and thereby low losses). The S-DALINAC uses superconducting cavities whose design parameters are listed in table 1.2. While the cavities themselves have a design quality factor of $Q_0 = 3 \times 10^9$, their loaded quality factor that takes damping from power couplers into account is $Q_L = 3 \times 10^7$ [17]. This loaded quality factor

corresponds to a very small 3 dB bandwidth of the resonators of 100 Hz at the operating frequency of 2.997 GHz.

1.3 RF control

When an electric resonator is operated slightly off-resonance, amplitude and phase of the field inside the cavity change with respect to the drive signal (see fig. 1.4). Unfortunately the high operating frequency makes the S-DALINAC's superconducting cavities very susceptible to eigenfrequency changes caused by mechanical deformation. The 20-cell design as well as the thin niobium cavity walls lead to cavities that are not very stiff. Therefore even small vibrations can result in massive disturbances. The amplitude and phase fluctuations caused by these vibrations have to be compensated by an rf control system.

Pressure changes in the liquid helium bath used to cool the cavities also result in deformations of the cavity causing a shift of its eigenfrequency according to [18]:

$$\frac{\Delta f}{\Delta p} = -15\,{}^{\mathrm{Hz}}/_{\mathrm{mbar}} \tag{1.6}$$

Although the pressure in the helium tank is stabilized at 35 mbar the pumps that are used to maintain this pressure can only react on a very slow timescale (in the order of minutes). All remaining fluctuations have to be compensated by the rf control system.

In addition to these effects a cavity can be detuned by the force of its own field (Lorentz force detuning, [18, 17]) according to

$$\frac{\Delta f}{E_{\mathrm{acc}}^2} = -4\,{}^{\mathrm{Hz}}/_{(\mathrm{MV/m})^2} \quad . \tag{1.7}$$

Lorentz force detuning is constant during operation since the S-DALINAC is operated in cw mode. But the rf control system has to deal with it whenever the amplitude of a cavity is changed by the operator or a cavity is switched off and on again.

Since all described disturbances occur in a frequency range of up to a few kHz, they are often referred to as microphonics. In the presence of microphonics an rf control system is required to stabilize the fields inside the cavity.

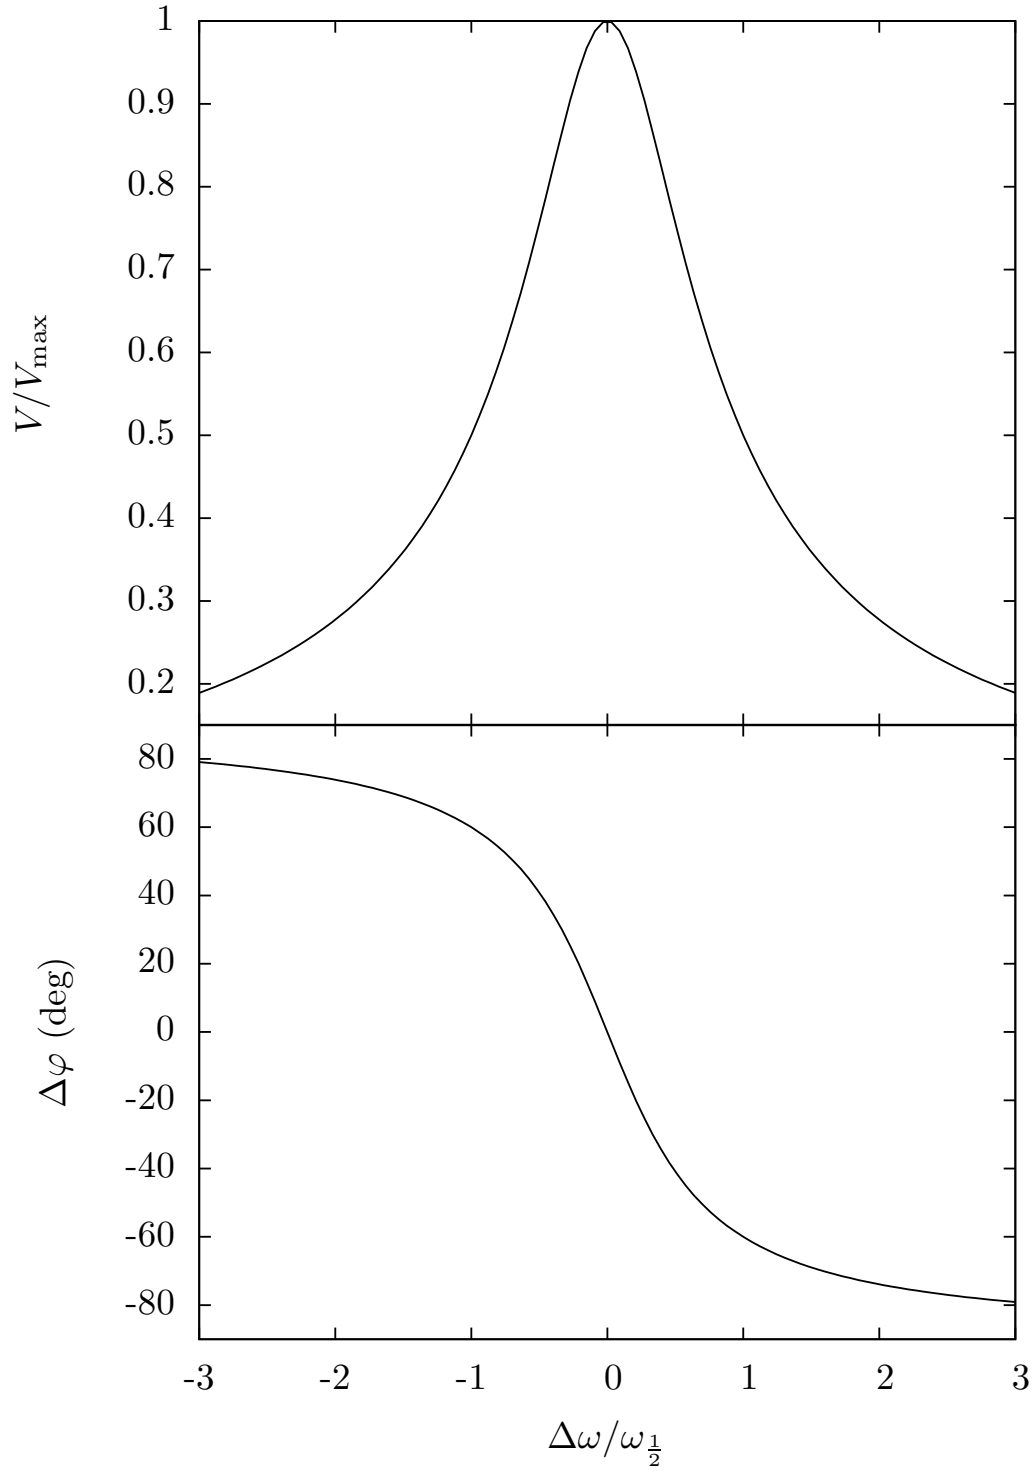

Figure 1.4: Transfer behavior of the cavity in proximity to the resonance frequency. The upper plot shows the normalized amplitude transfer function, the lower plot the phase transfer function. The frequency axis has been scaled to show multiples of the resonator bandwidth.

Normal-conducting cavities are not affected by microphonics because of their much larger bandwidth. Instead, temperature changes of the copper resonators can lead to phase drifts. To keep temperature effects low, the cooling water for the chopper and 3 GHz bunching cavity is stabilized to (20 ± 0.1)°C. Due to temperature effects in the amplifiers, the normal-conducting cavities still need an rf control system to keep amplitude and phase stable.

The target specification for the stability of the accelerating field is a relative error in beam energy of 1×10^{-4} rms. This can be achieved if the error in phase of the oscillating electromagnetic field inside the cavities is kept below 0.7° rms and the relative error in amplitude below 8×10^{-5} rms for on-crest acceleration.

Since 1989 an analog rf control system has been used which never achieved these values [19] and became more and more unreliable due to aging of the components. Additionally, the old analog system provided only very limited means for diagnostics and was quite hard to use by the operators. To meet the requirements with respect to accuracy, availability, and diagnostics, this former rf control system has been replaced by a state of the art digital rf control system. The first part of this thesis covers development and commissioning of this system. It can be used with the existing 3 GHz cavities as well as the 6 GHz buncher cavity needed for the spin-polarized electron gun [20]. Moreover, the new digital system provides extensive diagnostics and a user-friendly operator interface. Parts of the work described in part I of the thesis have already been published in [21, 22, 23, 24, 25].

1.4 Accelerator control system

The rf control system is tightly connected to the accelerator control system which allows to control the components of the accelerator from a central room. The existing accelerator control system developed in-house [26] turned out not to support re-programmable devices like the rf controller boards very well. Every change of the parameter set of the control algorithm in the FPGA would have resulted in the need of changing the source code of the control server as well as the clients. Moreover, the existing control system was lacking modern tools to build graphical user interfaces, to archive data, and to implement finite-state machines. Since modern accelerator control system frameworks provide all these

tools, prototypes of the rf control system have been implemented using the Experimental Physics and Industrial Control System (EPICS).

Due to the high reliability of the many ready-to-use EPICS modules, the implementation and maintenance of the prototype control system turned out to be much easier and faster than maintaining and extending the existing control system. This finally led to the decision to stick with EPICS for the final control system as well. After the rf control system had proven to be very flexible and reliable, it was decided to abandon the in-house development and migrate existing components to EPICS as well. The second part of this thesis describes the architecture of the new accelerator control system. Parts of this work have also been published in [24, 27, 28, 29].

Part I
Digital rf control system

2 RF control principles

The basic functionality of a control system using feedback is illustrated by Figure 2.1. The controlled system (plant) is disturbed by one or several disturbance signals. The controlled variable is compared to the set-point by calculating the difference between the plant output and the set-point. The resulting difference is called control error. The controller derives the control signal from the control error and sends it to the actuator to control the system.

Considering the rf control system of the S-DALINAC, the controlled system consists of the rf amplifiers (pre-amplifiers and klystrons or solid-state amplifiers), the cavity, and the cables to and from the cavity. In addition to an input power coupler all cavities used at the S-DALINAC have a pick-up antenna that couples out a small amount of the field. The obtained signal is proportional to the field inside the cavity and is used as an input signal for the rf control system.

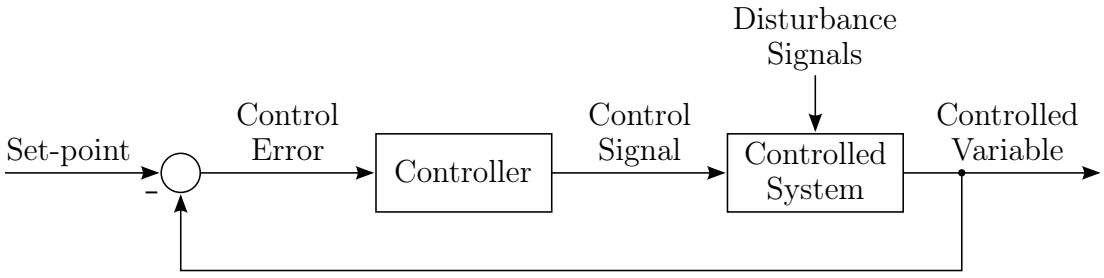

Figure 2.1: Block diagram of a control loop. The controlled signal is fed back into the controller.

2.1 Generator-driven resonator

There are different approaches to rf control. A simple way is to use the generator-driven resonator (GDR) approach which uses the amplified signal of an rf generator and sends it to the cavity (see fig. 2.2).

The signal from the pick-up antenna is used to compare the phase of the cavity output signal with the phase set-point provided by the operator. In a simple analog system this can be achieved by using the pick-up signal and the phase-shifted signal of the generator as input signals for an rf mixer. In combination with a

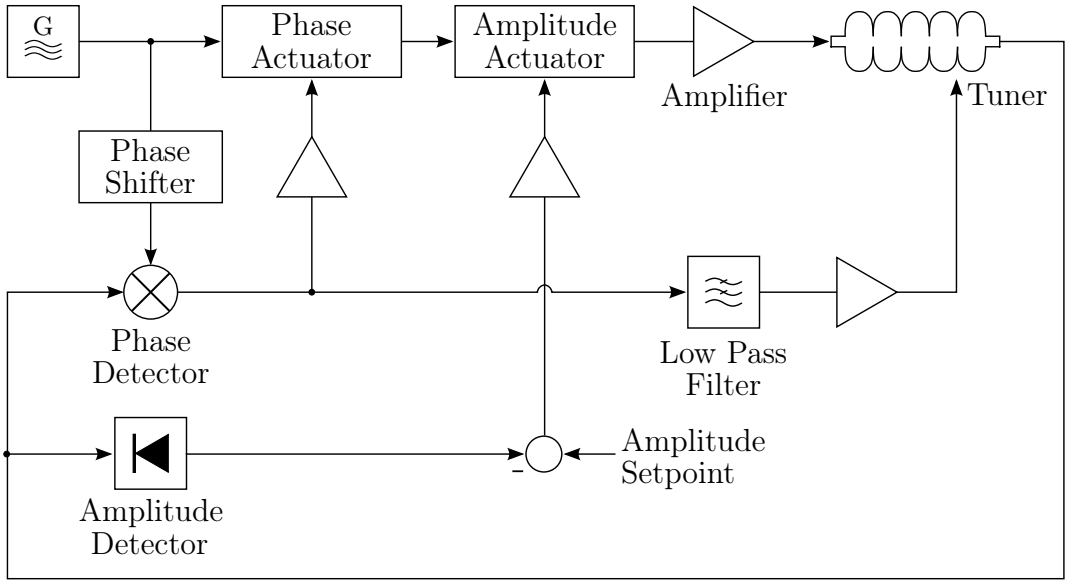

Figure 2.2: Simplified block diagram of an rf control system with a generator-driven resonator.

low-pass filter, this device acts as a phase detector. Its output voltage corresponds to the phase difference between the inputs and thus can be used as the control error. After multiplying with an appropriate control gain, the signal is used to drive the tuner. To make sure the tuner does not excite mechanical eigenmodes of the cavity, the signal has to be low-pass filtered. Hence only slow disturbances can be compensated by the tuner. Fast disturbances have to be compensated by a fast electric phase actuator.

The magnitude of the cavity output signal is measured by an amplitude detector. Its output signal is subtracted from the amplitude set-point provided by the operator. The resulting amplitude error signal is amplified by the amplitude control gain to obtain the amplitude control signal. This signal is used to drive the amplitude actuator that changes the magnitude of the signal that is sent to the cavity.

The tuner controller is essential to keep the cavity on resonance. Without eigenfrequency control the cavity might cause more attenuation than the amplitude controller can compensate.

2.2 Self-excited loop

Another approach is to operate the cavity in a self-excited loop (SEL, see fig. 2.3, [30]). Instead of driving the cavity by the generator, the cavity is operated in a closed rf loop. This loop contains an amplifier, the cavity, and a phase shifter. In the presence of enough gain the loop starts oscillating from noise. The cavity acts as a narrow band-pass filter and thereby filters out everything but its resonance frequency. Thus the loop starts oscillating freely on a frequency determined by the eigenfrequency of the cavity and the loop phase. The frequency the loop can oscillate on is one at which the total loop phase shift is a multiple of 2π.

When the loop is oscillating, its frequency and phase can be locked to the master oscillator by a technique called phase-locked loop (PLL, [31]). Therefore the phase of the pick-up signal is being compared to the phase of the master oscillator shifted by the phase set-point. The resulting phase error signal is used to drive the tuner and an electric phase actuator in the same way as described for the generator-driven resonator approach. But in case of the self-excited loop the phase controller compensates for the loop phase change caused by the cavity. This way it keeps the loop phase as well as the frequency constant.

To prevent the amplitude from increasing indefinitely, amplitude control is crucial for the self-excited loop. It is implemented in the same way as for the generator-driven resonator approach.

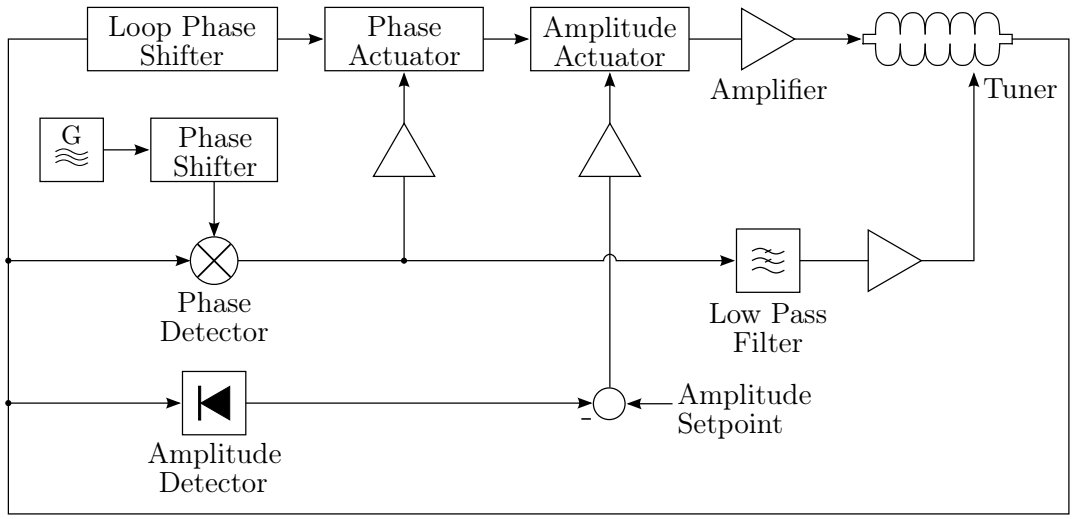

Figure 2.3: Simplified block diagram of a cavity operated in a self-excited loop.

3 Cavity model

The first step in designing a control system is to analyze the properties of the system that is to be controlled. Since the behavior of superconducting cavities is complex, a simplified model will be used to describe them. This model can be used to understand and to simulate the transfer behavior of the cavities.

3.1 Differential equation of the cavity

A single-cell cavity is an electric resonator with a behavior equivalent to a parallel RLC circuit. In this circuit the capacitor C, the inductor L, and the resistor R_L are all in parallel with an external current source (see fig. 3.1).

The currents through each of the three elements can be obtained from the corresponding constitutive equations:

$$Q(t) = C\,V(t) \quad \Longleftrightarrow \quad \dot{I}_C(t) = C\,\ddot{V}(t) \tag{3.1}$$

$$V(t) = L\,\dot{I}_L(t) \quad \Longleftrightarrow \quad \dot{I}_L(t) = \frac{V(t)}{L} \tag{3.2}$$

$$I_R(t) = \frac{V(t)}{R_L} \quad \Longleftrightarrow \quad \dot{I}_R(t) = \frac{\dot{V}(t)}{R_L} \tag{3.3}$$

According to Kirchhoff's current law, the sum of the currents through the three components matches the external current:

$$I_C(t) + I_L(t) + I_R(t) = I(t) \tag{3.4}$$

Differentiating this equation and substituting eqs. (3.1) to (3.3) yields the governing differential equation for the RLC circuit:

$$\ddot{V}(t) + \frac{1}{R_L C}\,\dot{V}(t) + \frac{1}{LC}V(t) = \frac{1}{C}\dot{I}(t) \tag{3.5}$$

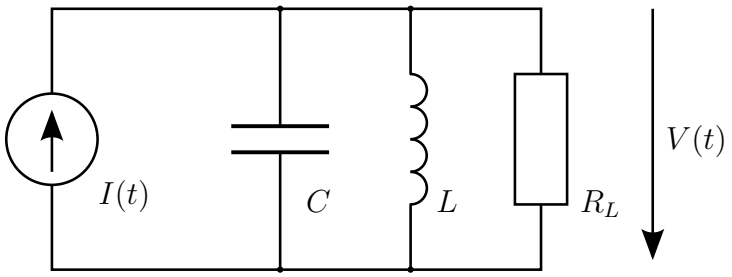

Figure 3.1: Equivalent circuit of a single-cell cavity resonator.

For cavity resonators the parameters R_L, L, and C are difficult to measure. In the following these parameters will be matched with cavity quantities that are easier to determine.

3.1.1 Quality factor

An important quantity used to characterize the losses in a cavity is its quality factor. It is defined as 2π times the ratio of stored energy U to the energy dissipated in one period of the oscillation $T \times P_{\text{diss}}$ [32]:

$$Q := \frac{\omega U}{P_{\text{diss}}} \quad . \tag{3.6}$$

The intrinsic quality factor Q_0 describes the losses caused by dissipation in the cavity walls (ohmic losses). In normal operation of a cavity, power couplers are needed to drive the cavity and to monitor the field inside. These couplers lead to a power leakage out of the cavity. The loaded quality factor Q_L includes these external losses together with the intrinsic losses and, hence, is smaller than Q_0. Since an rf control system cannot be used without couplers, only Q_L will be used in the following.

At maximum voltage V_{max} the stored energy in the resonator is stored completely in the electric field:

$$U = \frac{1}{2} C V_{\text{max}}^2 \tag{3.7}$$

In the RLC circuit the dissipated power is converted into heat at the resistor

$$P_{\text{diss}}(t) = \frac{V^2(t)}{R_L} \quad . \tag{3.8}$$

Note that R_L in the RLC circuit models the intrinsic as well as the external losses caused by the couplers.

Assuming that $V(t)$ is a harmonic signal the average dissipated power over one period is

$$P_{\text{diss}} = \frac{V_{\text{max}}^2}{2R_L} \quad . \tag{3.9}$$

3.1.2 Resonance frequency

For an RLC circuit that oscillates freely without damping, the resonance frequency is exactly

$$\omega_r = \frac{1}{\sqrt{LC}} \quad . \tag{3.10}$$

This frequency changes slightly for a driven and damped resonator. As all resonators used at the S-DALINAC (including the normal-conducting cavities) have a high operating frequency and a high quality factor, this correction can be neglected and eq. (3.10) can be used for them as well.

Substituting eqs. (3.6), (3.7), (3.9) and (3.10) into eq. (3.5) yields a version of the differential equation in parameters that are well known for cavity resonators:

$$\ddot{V}(t) + \frac{\omega_r}{Q_L}\dot{V}(t) + \omega_r^2 V(t) = \frac{\omega_r R_L}{Q_L}\dot{I}(t) \tag{3.11}$$

3.2 Transfer function

If we assume the parameters Q_L, R_L, and ω_r to be constant over time, a cavity is a linear time-invariant system. Hence a transfer function describing the relation between the input and the output of the cavity can be deduced. This can be accomplished by Laplace-transforming eq. (3.11) and setting all initial conditions to zero:

$$G(s) = \frac{V(s)}{I(s)} = \frac{\omega_r\, R_L\, s}{Q_L\, s^2 + \omega_r\, s + Q_L\, \omega_r^2} \tag{3.12}$$

The amplitude transfer function can be obtained by setting $s = j\omega$ and taking the absolute value of the transfer function:

$$|G(j\omega)| = \frac{R_L}{\sqrt{1 + Q_L^2 \left(\frac{\omega}{\omega_r} - \frac{\omega_r}{\omega}\right)^2}} \tag{3.13}$$

Defining

$$\omega = \omega_r + \Delta\omega \tag{3.14}$$

yields

$$|G(j\omega)| = \frac{R_L}{\sqrt{1 + Q_L^2 \left(\frac{2\omega_r\Delta\omega + \Delta\omega^2}{\omega_r^2 + \omega_r\Delta\omega}\right)^2}} \quad . \tag{3.15}$$

At proximity to the resonance frequency one has $\Delta\omega \ll \omega_r$ which leads to

$$|G(j\omega)| \approx \frac{R_L}{\sqrt{1 + 4Q_L^2 \left(\frac{\omega-\omega_r}{\omega_r}\right)^2}} \tag{3.16}$$

The phase transfer function reads

$$\begin{aligned} \angle G(j\omega) &= \arctan\left(\frac{\mathrm{Im}(G(j\omega))}{\mathrm{Re}(G(j\omega))}\right) && (3.17)\\ &= \arctan\left(Q_L \frac{-2\omega_r\Delta\omega - (\Delta\omega)^2}{\omega_r(\omega_r + \Delta\omega)}\right) && (3.18)\\ &\approx \arctan\left(-2Q_L\frac{\omega-\omega_r}{\omega_r}\right) \quad . && (3.19) \end{aligned}$$

The amplitude and phase transfer functions are shown in fig. 1.4.

3.3 State space representation

The operation frequency of the cavities is several orders of magnitude higher than the frequency of the mechanical disturbances. Therefore the complete information of the signals can be obtained from the envelope and the phase of the rf signal. This is equivalent to signals with complex amplitudes:

$$V(t) = \left(V_{\mathrm{re}}(t) + jV_{\mathrm{im}}(t)\right) e^{j\omega t} \tag{3.20}$$

$$I(t) = \left(I_{\mathrm{re}}(t) + jI_{\mathrm{im}}(t)\right) e^{j\omega t} \tag{3.21}$$

Inserting this ansatz into the differential eq. (3.11), dividing by $e^{j\omega t}$ and neglecting all second order derivatives (since the dynamics of the envelope is slow) yields:

$$\begin{aligned} 2\left(\dot{V}_{\mathrm{re}}(t) + j\dot{V}_{\mathrm{im}}(t)\right) j\omega - \left(V_{\mathrm{re}}(t) + jV_{\mathrm{im}}(t)\right)\omega^2 + \frac{\omega_r}{Q_L}\left(\dot{V}_{\mathrm{re}}(t) + j\dot{V}_{\mathrm{im}}(t)\right) \\ + \frac{\omega_r}{Q_L}\left(V_{\mathrm{re}}(t) + jV_{\mathrm{im}}(t)\right) j\omega + \omega_r^2\left(V_{\mathrm{re}}(t) + jV_{\mathrm{im}}(t)\right) \\ = \frac{\omega_r R_L}{Q_L}\left(\dot{I}_{\mathrm{re}}(t) + j\dot{I}_{\mathrm{im}}(t)\right) + \frac{\omega_r R_L}{Q_L}\left(I_{\mathrm{re}}(t) + jI_{\mathrm{im}}(t)\right) j\omega \end{aligned} \tag{3.22}$$

This equation can be separated into one equation for the real and the imaginary part each:

$$\dot{V}_{\mathrm{im}}(t) + \frac{\omega}{2} V_{\mathrm{re}}(t) - \frac{\omega_r}{2Q_L}\frac{1}{\omega}\dot{V}_{\mathrm{re}}(t) + \frac{\omega_r}{2Q_L} V_{\mathrm{im}}(t) - \frac{\omega_r^2}{2\omega} V_{\mathrm{re}}(t) = -\frac{\omega_r}{2Q_L}\frac{R_L}{\omega}\dot{I}_{\mathrm{re}}(t) + \frac{\omega_r}{2Q_L} R_L I_{\mathrm{im}}(t) \tag{3.23}$$

$$\dot{V}_{\mathrm{re}}(t) - \frac{\omega}{2} V_{\mathrm{im}}(t) + \frac{\omega_r}{2Q_L}\frac{1}{\omega}\dot{V}_{\mathrm{im}}(t) + \frac{\omega_r}{2Q_L} V_{\mathrm{re}}(t) + \frac{\omega_r^2}{2\omega} V_{\mathrm{im}}(t) = \frac{\omega_r}{2Q_L}\frac{R_L}{\omega}\dot{I}_{\mathrm{im}}(t) + \frac{\omega_r}{2Q_L} R_L I_{\mathrm{re}}(t) \tag{3.24}$$

The term $\frac{\omega_r}{2Q_L}$ corresponds to the half bandwidth of the resonator $\omega_{\frac{1}{2}}$ (HWHM). Since terms with ω in the denominator are small compared to the others, they can be neglected:

$$\dot{V}_{\mathrm{im}}(t) + \omega_{\frac{1}{2}} V_{\mathrm{im}}(t) + \left(\frac{\omega}{2} - \frac{\omega_r^2}{2\omega}\right) V_{\mathrm{re}}(t) \approx \omega_{\frac{1}{2}} R_L I_{\mathrm{im}}(t) \tag{3.25}$$

$$\dot{V}_{\mathrm{re}}(t) + \omega_{\frac{1}{2}} V_{\mathrm{re}}(t) - \left(\frac{\omega}{2} - \frac{\omega_r^2}{2\omega}\right) V_{\mathrm{im}}(t) \approx \omega_{\frac{1}{2}} R_L I_{\mathrm{re}}(t) \tag{3.26}$$

In proximity to the resonance frequency the term in parenthesis is

$$\frac{\omega}{2} - \frac{\omega_r^2}{2\omega} = \frac{(\omega + \omega_r)(\omega - \omega_r)}{2\omega} \approx -\Delta\omega \quad . \tag{3.27}$$

This corresponds to a linearization around the resonance frequency and leads to two coupled first order differential equations

$$\dot{V}_{\mathrm{im}}(t) + \omega_{\frac{1}{2}} V_{\mathrm{im}}(t) - \Delta\omega V_{\mathrm{re}}(t) \approx \omega_{\frac{1}{2}} R_L I_{\mathrm{im}}(t) \tag{3.28}$$

$$\dot{V}_{\mathrm{re}}(t) + \omega_{\frac{1}{2}} V_{\mathrm{re}}(t) + \Delta\omega V_{\mathrm{im}}(t) \approx \omega_{\frac{1}{2}} R_L I_{\mathrm{re}}(t) \quad . \tag{3.29}$$

These equations can also be written in state space notation [33]

$$\dot{\vec{x}} = A\vec{x} + B\vec{u} \tag{3.30}$$

$$\vec{y} = C\vec{x} + D\vec{u} \tag{3.31}$$

with

$$\vec{x} = \begin{pmatrix} V_{\mathrm{re}} \\ V_{\mathrm{im}} \end{pmatrix} \qquad \vec{u} = \begin{pmatrix} I_{\mathrm{re}} \\ I_{\mathrm{im}} \end{pmatrix} \tag{3.32}$$

$$A = \begin{pmatrix} -\omega_{\frac{1}{2}} & -\Delta\omega \\ \Delta\omega & -\omega_{\frac{1}{2}} \end{pmatrix} \qquad B = \begin{pmatrix} \omega_{\frac{1}{2}} R_L & 0 \\ 0 & \omega_{\frac{1}{2}} R_L \end{pmatrix} \tag{3.33}$$

$$C = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad D = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \tag{3.34}$$

The state equation (3.30) describes the dynamics of the state vector $\vec{x}$. The system matrix A contains the dynamics of the system whereas the input matrix B describes the coupling of the input vector $\vec{u}$ to the system. The output equation (3.31) describes how the output vector $\vec{y}$ is made up from the current state of the system and its input. C is called the output matrix and D is the feed-through matrix. The state vector of a cavity can be observed directly resulting in C being the identity matrix. There is no direct feed through from the input coupler to the output coupler which results in D being a zero matrix.

The discrete version of the state-space equations with step size τ reads

$$\vec{x}(n+1) = A'\vec{x}(n) + B'\vec{u}(n) \tag{3.35}$$

$$\vec{y}(n) = C\vec{x}(n) + D\vec{u}(n) \tag{3.36}$$

with

$$A' = \mathbb{1} + \tau A = \begin{pmatrix} 1 - \omega_{\frac{1}{2}}\tau & -\Delta\omega\tau \\ \Delta\omega\tau & 1 - \omega_{\frac{1}{2}}\tau \end{pmatrix} \tag{3.37}$$

$$B' = \tau B \qquad = \begin{pmatrix} \omega_{\frac{1}{2}} R_L \tau & 0 \\ 0 & \omega_{\frac{1}{2}} R_L \tau \end{pmatrix} \tag{3.38}$$

This form is especially suited for a numerical simulation of the resonator behavior. In contrast to the transfer function (3.12) the state space representation allows to describe the system in the low frequency domain which speeds up computation by several orders of magnitude.

3.4 Multi-cell cavities

The model derived above only describes single-cell cavities with a single resonance frequency. In contrast the multi-cell superconducting cavities used at the S-DALINAC have multiple resonances. Since the number of resonances corresponds to the number of cells, this results in 2, 5 or even 20 resonance frequencies for the accelerating cavities. Due to the high quality factor of these resonators the resonance frequencies are clearly separated. Thus in proximity to a resonance the S-DALINAC's accelerating cavities show the same behavior as single-cell cavities. Since the cavities are operated in a very small frequency band around the π mode, a model describing the behavior of the cavity in this frequency range is sufficient. Therefore the much simpler single-cell model can be used to describe all cavities at the S-DALINAC.

4 Hardware

The state-space model of the cavity exploits the fact that the relevant dynamics carried by an rf signal from a cavity is much slower than the rf signal itself. This section will illustrate how this slow dynamics can be extracted with technical means and how this has been used to implement the controller hardware. Systematical errors originating from the chosen technique will also be discussed in detail.

4.1 Up and down-conversion with a mixer

The operating frequencies of 3 and 6 GHz are too high to be controlled directly. The perturbations on the other hand are much slower than the frequency ω of the rf signal and effectively only slowly change amplitude and phase of the envelope:

$$s(t) = A(t)\cos\big(\omega t + \varphi(t)\big) \tag{4.1}$$

Therefore the rf signal can be down-converted to a lower frequency without losing information. This step can be accomplished by an electronic mixer (see fig. 4.1). An ideal mixer is a nonlinear device that multiplies two input signals:

$$s_1(t) \times s_2(t) = A_1(t)\cos\big(\omega_1 t + \varphi_1(t)\big) \times A_2(t)\cos\big(\omega_2 t + \varphi_2(t)\big) \tag{4.2}$$

Applying the trigonometric product-to-sum identity shows that the mixer generates a resulting signal that consists of a superposition of two signals with the frequencies $|\omega_1 - \omega_2|$ and $\omega_1 + \omega_2$:

$$\begin{aligned} s_1(t) \times s_2(t) = \frac{A_1(t)A_2(t)}{2} \quad & \Big[\cos\big(\,(\omega_1 - \omega_2)\,t + \varphi_1(t) - \varphi_2(t)\big) \\ & + \cos\big(\,(\omega_1 + \omega_2)\,t + \varphi_1(t) + \varphi_2(t)\big)\Big] \end{aligned} \tag{4.3}$$

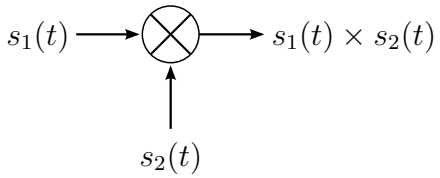

Figure 4.1: An ideal frequency mixer multiplies its input signals.

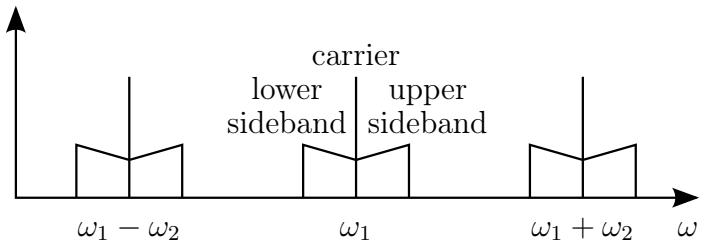

Figure 4.2: Input signal with center frequency ω_1 mixed with a frequency ω_2. The resulting signal is a superposition of the signals around $|\omega_1 - \omega_2|$ and $\omega_1 + \omega_2$.

In most cases these frequencies are clearly separated and the undesired frequency can be removed by a filter. This is the standard way of mixing signals up and down in telecommunication applications.

If the phase of the input signal varies over time this also changes the frequency of the input signal producing an upper sideband for increasing $\varphi_1(t)$ and a lower sideband for decreasing $\varphi_1(t)$. Thus not only the center frequency has to be mixed down correctly but also the sidebands have to be reproduced properly as sketched in fig. 4.2.

Unfortunately this does not work in the right way if $\omega_1 - \omega_2$ becomes negative for any frequency of the sideband. This happens in particular if the center frequency of the signal is mixed down to a frequency of 0 Hz (the signal is mixed down to the baseband). In this case a modulation frequency $-\omega_{\mathrm{mod}}$ is mapped onto its mirror frequency of $+\omega_{\mathrm{mod}}$ and both modulation frequencies cannot be distinguished anymore. Thus the rf control system cannot determine if the cavity is oscillating on a frequency higher or lower than the master oscillator, which renders rf control impossible.

4.2 Quadrature (de)modulation

There are two common ways, an analog and a digital one, to convert an rf signal down to the baseband without losing the information to which sideband the signal belongs. Both transform the input signal into a Cartesian coordinate system that rotates in the complex plane with the frequency of the local oscillator (see fig. 4.3). The coordinates in the rotating coordinate system are called I (in-phase component) and Q (quadrature component). This technique is known as

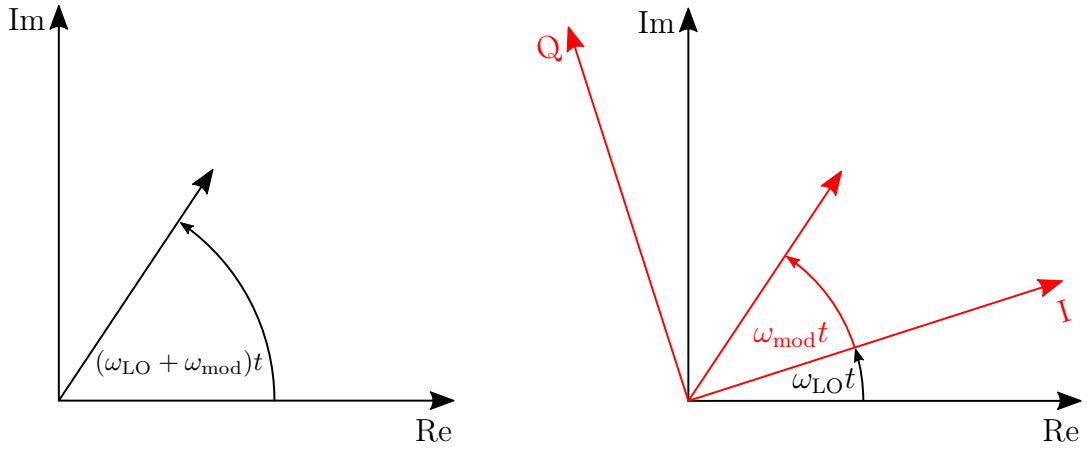

Figure 4.3: RF signal with a frequency of $\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}}$ (a) in the complex plane and (b) in the rotating I/Q coordinate system.

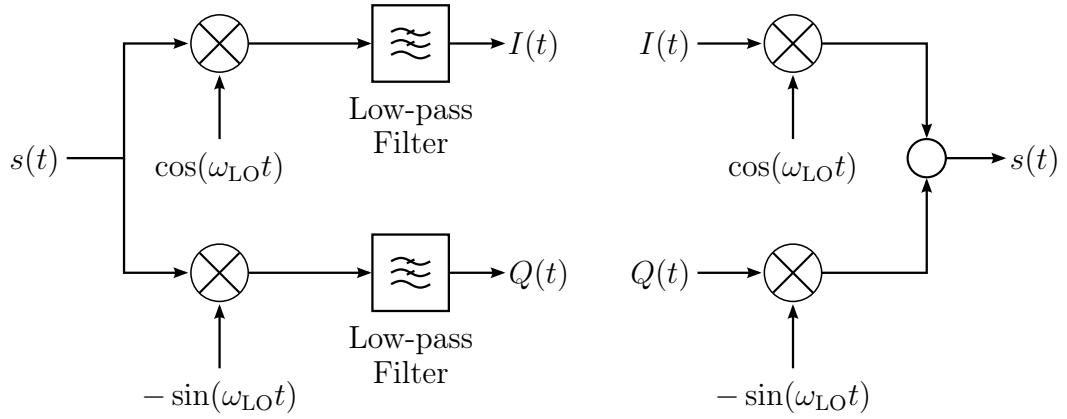

Figure 4.4: Block diagram of (a) a quadrature demodulator and (b) a quadrature modulator.

quadrature demodulation whereas the process of converting a signal given in the rotating coordinate system back into an rf signal is called quadrature modulation.

4.2.1 Analog quadrature (de)modulation

The traditional way of performing quadrature demodulation is by using a hardware quadrature demodulator that mixes the input signal down with the signal of the local oscillator and a 90° phase-shifted signal in parallel (see fig. 4.4a). Filtering out the sum of the frequencies directly leads to the I and Q components of the

input signal:

$$
\begin{aligned}
I(t) &= s_{\mathrm{rf}}(t) \times \cos(\omega_{\mathrm{LO}} t) \\
&\overset{(4.3)}{=} \frac{A(t)}{2} \left(\cos\left[(\omega_{\mathrm{rf}} - \omega_{\mathrm{LO}})\, t + \varphi(t) \right] + \underbrace{\cos\left[(\omega_{\mathrm{rf}} + \omega_{\mathrm{LO}})\, t + \varphi(t) \right]}_{\text{filtered out}} \right) \qquad (4.4) \\
Q(t) &= s_{\mathrm{rf}}(t) \times \big(-\sin(\omega_{\mathrm{LO}} t) \big) \\
&\overset{(4.3)}{=} \frac{A(t)}{2} \left(\sin\left[(\omega_{\mathrm{rf}} - \omega_{\mathrm{LO}})\, t + \varphi(t) \right] - \underbrace{\sin\left[(\omega_{\mathrm{rf}} + \omega_{\mathrm{LO}})\, t + \varphi(t) \right]}_{\text{filtered out}} \right) \qquad (4.5)
\end{aligned}
$$

The quadrature modulator (also called direct modulator, fig. 4.4b) is the same in reverse: The I and Q signals are mixed with the local oscillator frequency and the 90° phase-shifted local oscillator frequency, respectively, before they are added to form the rf output signal

$$
s_{\mathrm{rf,ideal}}(t) = I(t) \times \cos(\omega_{\mathrm{LO}} t) - Q(t) \times \sin(\omega_{\mathrm{LO}} t) \quad . \qquad (4.6)
$$

For the special case of a rotating vector in I/Q coordinates the output signal yields:

$$
\begin{aligned}
I(t) &= A \cos(\omega_{\mathrm{mod}} t) && (4.7) \\
Q(t) &= A \sin(\omega_{\mathrm{mod}} t) && (4.8) \\
s_{\mathrm{rf,ideal}}(t) &= \frac{A}{2} \Big[\cos\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{mod}})t\big) + \cos\big((\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}})t\big) \\
&\quad - \cos\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{mod}})t\big) + \cos\big((\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}})t\big) \Big] && (4.9) \\
&= A \cos\big((\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}})t\big) && (4.10)
\end{aligned}
$$

As can be seen from eq. (4.9) the image frequency terms of the two mixers cancel out and the ideal quadrature modulator produces a pure single-sideband signal.

For a real quadrature (de)modulator the amplitudes of the local oscillator signal are not exactly the same for both mixers, which can be described by the gain mismatch $\alpha = A_Q/A_I$. In addition, the phase shift of the local oscillator signals for the two mixers might differ from the ideal 90° by the phase mismatch Φ. For a demodulator this results in a slight coupling of both components of the (I, Q) input vector. For a real modulator the rf signal is given by:

$$
s_{\mathrm{rf,md}}(t) = I(t) \times \cos(\omega_{\mathrm{LO}} t) - \alpha Q(t) \times \sin(\omega_{\mathrm{LO}} t + \Phi) \qquad (4.11)
$$

In this case the mirror frequency is not suppressed completely (this effect is called modulation distortion). For the special case of the rotating vector one obtains:

$$\begin{aligned} s_{\mathrm{rf,md}}(t) =& \frac{A}{2}\Big[\cos\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{mod}})t\big) + \cos\big((\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}})t\big) \\ &+ \alpha \cos\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{mod}})t + \Phi\big) + \cos\big((\omega_{\mathrm{LO}} + \omega_{\mathrm{mod}})t - \Phi\big)\Big] \end{aligned} \tag{4.12}$$

According to eq. (4.11) dc offsets on the I and Q signals caused by technical imperfections lead to a leakage of the carrier frequency to the rf output signal. Offsets, gain mismatch as well as phase mismatch can be compensated (see section 4.4). But as mixers and thereby analog quadrature (de)modulators are non-ideal nonlinear components they are also affected by higher order errors that are difficult to compensate.

4.2.2 Digital quadrature (de)modulation

Another more modern quadrature demodulation technique is digital demodulation [34]. A digital quadrature demodulator (see fig. 4.5) demodulates an intermediate frequency signal that is obtained from the rf signal by down-conversion with a single mixer. The intermediate frequency signal is bandpass-filtered to remove the high-frequency component and to avoid aliasing effects. The filtered intermediate frequency signal is sampled directly with an analog-to-digital converter (ADC) operating at a frequency that is four times the intermediate frequency. Thus the time period between samples is 90°. This means if the first sample is defined as I, the next sample is Q, the following sample is $-I$, and the fourth sample is $-Q$. This sequence of values is separated into a data stream of I and Q values, respectively, by a digital demultiplexer. The alternating signs are removed

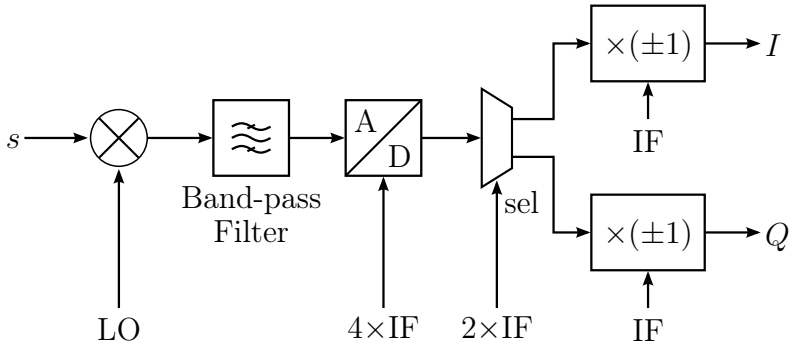

Figure 4.5: Block diagram of a digital quadrature demodulator.

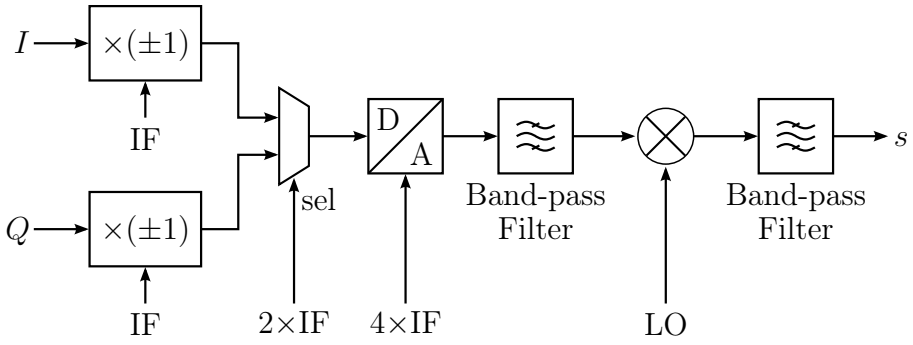

Figure 4.6: Block diagram of a digital quadrature modulator.

by multiplying each stream by $+1$ and -1 alternately. The output signals of the digital quadrature demodulator correspond to the I and Q values that are obtained by a conventional I/Q demodulator.

The digital quadrature modulator (see fig. 4.6) uses the inverse concept: The I and Q signals are multiplied by $+1$ and -1 alternately before they are multiplexed and converted to an analog signal at the intermediate frequency. A bandpass filter removes the dc offset of the digital-to-analog converter (DAC) as well as harmonics before the intermediate frequency signal is mixed to the operating frequency. A final bandpass filter removes the image frequency of this mixing step.

In contrast to an analog quadrature (de)modulator a digital I/Q (de)modulator has no parallel branches with analog signal processing components. Therefore it is not affected by problems like gain mismatch and quadrature phase errors. Offsets of the ADCs and DACs are removed by the bandpass filters.

4.2.3 (De)modulation at the S-DALINAC

In the last years the intermediate frequency approach has been chosen for more and more rf control systems because it does not suffer from errors caused by offsets and gain or phase mismatch of the two analog mixer branches of a conventional quadrature (de)modulator (see e. g. [35, 36, 37]). Additionally, the controller boards do not need to be calibrated which is an advantage if hundreds of channels are required. For the small number of 16 cavities at the S-DALINAC the last aspect is not vital.

The intermediate frequency technique also has disadvantages: To suppress the image frequency produced by the mixer of a digital modulator to a tolerable mag-

nitude, the bandwidth of the rf bandpass filter has to be considerably smaller than the intermediate frequency. For high operating frequencies as for the S-DALINAC this can only be accomplished by filters with high quality factor (and thereby high price) and high intermediate frequencies. The intermediate frequency in turn is limited by the sampling frequency of the ADC and the digital signal processing speed. That is why the limited performance of the modulator output bandpass filter is in many cases compensated by the (in most cases superconducting) cavity that also acts as a bandpass filter. The downside of this approach is that the image frequencies have to be amplified by the klystron, too, which reduces the effective output power at the operating frequency. In addition the occurrence of image frequencies with their sidebands complicates the separation of the different eigenmodes of the cavity. This is particularly an issue with the S-DALINAC's cavities with their little spacing between eigenmodes.

Another challenge is the synchronization of the digital clock to the rf local oscillator with very low jitter which requires a sophisticated design of the printed circuit board as well as of the configuration of the field programmable gate array (FPGA, see section 6). Generating a local oscillator frequency for cavities that are operated at a harmonic of the base frequency also becomes more complicated. For the baseband approach the 6 GHz local oscillator signal can be generated from the 3 GHz local oscillator signal by frequency doubling. In case of the intermediate frequency approach the needed frequencies are $f_{\text{base}} - f_{\text{IF}}$ and $2 \times f_{\text{base}} - f_{\text{IF}}$ which makes the use of straight frequency doubling impossible. Using two different local oscillators leads to higher overall phase noise even if they are synchronized.

Last but not least the much faster ADC has a lower precision (usually 14 bits at this time) than ADCs used for baseband applications (up to 18 bits) which reduces the accuracy of the rf control system. Part of the lost accuracy can be recovered by averaging over several I/Q samples but this comes at the cost of higher latency which means amplitude accuracy is still limited to $\approx 2 \times 10^{-4}$ [38].

There are ways to address some of these problems of the digital (de)modulation technique but things get a lot more complicated with higher operating frequencies and higher demands on accuracy. This is why the baseband approach has been chosen for the S-DALINAC's rf control system with its high operating frequency and its high requirements with respect to stability.

4.3 Controller-board hardware

In the following an overview of the controller-board hardware in use at the S-DALINAC is given. The hardware of the rf control system is split into two parts: A frequency-dependent rf board containing the quadrature (de)modulator and a frequency-independent FPGA board processing the baseband signals (see fig. 4.7). All hardware components have been developed in-house. Prototypes of these boards have been developed by [39, 40, 41] and have also been described in [42]. Improved revisions of both the rf board (see fig. 4.8) and the FPGA board (fig. 4.9) have been developed and commissioned in the context of this thesis.

In addition to the quadrature (de)modulator the rf board contains a separate power-detector chip (Analog Devices `ADL5500`) which improves the accuracy of the amplitude measurement. It provides a high accuracy of $\approx 1\,\mathrm{dB}$ over a range of 30 dB as well as an excellent temperature stability. The output signal of this detector is proportional to the envelope of the input signal and thereby proportional to the field in the cavity. The amplitude detector chip has a much better linearity than the I/Q demodulator. Furthermore the measured power does not depend on the power of the local oscillator signal.

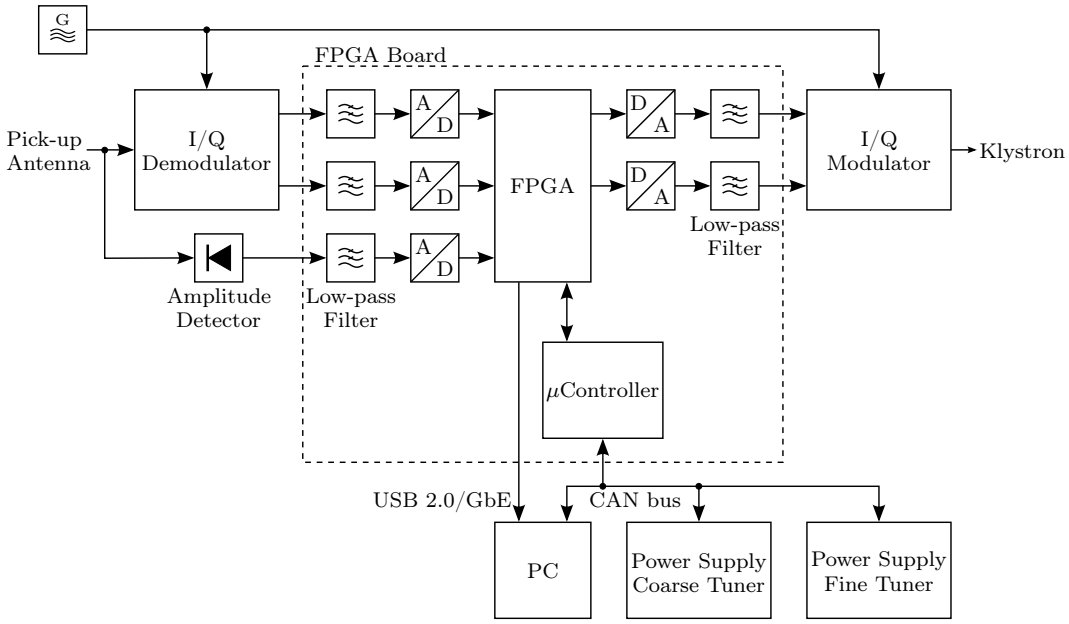

Figure 4.7: Overview of the hardware components of one channel of the rf control system.

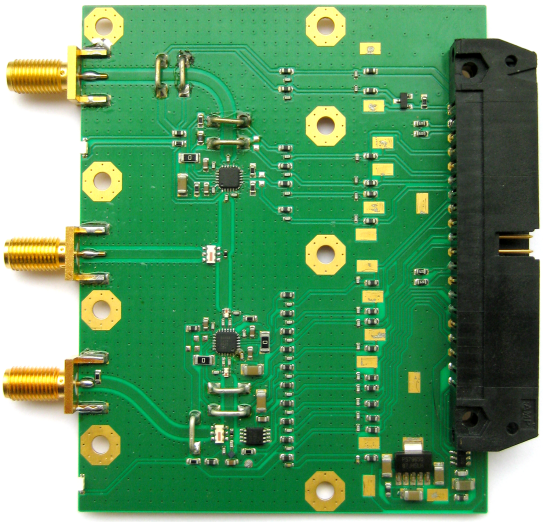

Figure 4.8: Picture of the rf board (`HF-ADL8575C Rev. 1`) containing all frequency dependent components. The rf jackets on the left side allow to connect the input signal, the local oscillator frequency, and the output signal. The baseband signals are transferred from the rf board to the FPGA board and back via the connector on the right side.

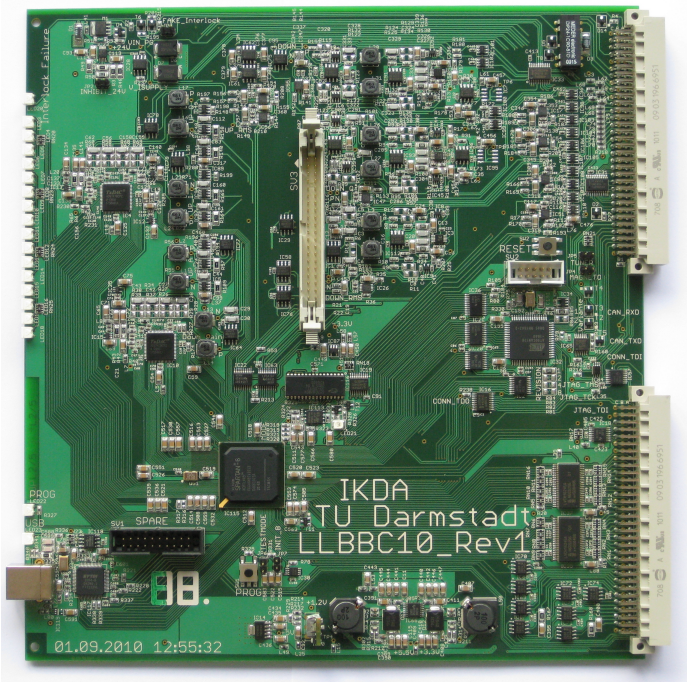

Figure 4.9: Picture of the FPGA board (`LLBBC10 Rev. 1`) containing all baseband signal processing components.

To avoid aliasing effects the analog signals provided by the demodulator and the amplitude detector are low-pass filtered on the FPGA board before they are digitized. The filters used for this purpose are third-order π filters with an edge frequency of $\approx 100\,\mathrm{kHz}$. Signals at the Nyquist frequency of $f_\mathrm{N} = 500\,\mathrm{kHz}$ are attenuated by more than 40 dB relative to the wanted low-frequency signals. This filtering step is also very important because it suppresses the $19/20\,\pi$ mode of the 20-cell cavities which is only separated by $f_{\frac{19}{20}} = 700\,\mathrm{kHz}$ from the π mode [43] used for acceleration. As can be seen from the Bode diagram (fig. 4.10) the chosen filter design ensures that the adjacent mode is attenuated by 50 dB while the phase shift introduced by the filter is small for microphonic frequencies (up to 10 kHz). A derivation of the transfer function can be found in appendix A.

High-linearity 18 bit ADCs (Analog Devices `AD7982`) are used to meet the specification with respect to accuracy. They can be operated with sampling rates of up to $1\,{}^{\mathrm{MS}}/_{\mathrm{s}}$. This keeps latency low and allows the system to detect the signal of a detuned cavity operated in self-excited loop mode.

There are basically two different types of integrated chips available for digital signal processing: digital signal processors (DSPs) and field programmable gate arrays (FPGAs). DSPs are microprocessors with an architecture optimized for digital signal processing which means that they can only process one command at a time[1]. In an FPGA on the other hand user-defined digital circuits can be implemented that can process many tasks in parallel. This can speed up signal-processing performance significantly [44]. Digital rf control systems have been implemented using both techniques [36, 45, 46, 47]. An FPGA has been used for the S-DALINAC's rf control system because DSPs are usually only available for 16 bit or 32 bit arithmetic. 16 bit arithmetic is not sufficient to process the data in full 18 bit resolution whereas 32 bit arithmetic is very slow compared to 18 bit arithmetic. By implementing a custom DSP that executes the control algorithm in the FPGA the advantages from both approaches can be combined (see section 6.5). Discrete logic can be used for infrastructure tasks like digital read-out and ADC/DAC operation whereas the development of the control algorithm benefits from the easy re-programmability of a DSP.

[1] If pipelining or a multi-core CPU is used, a small number of commands can be processed in parallel.

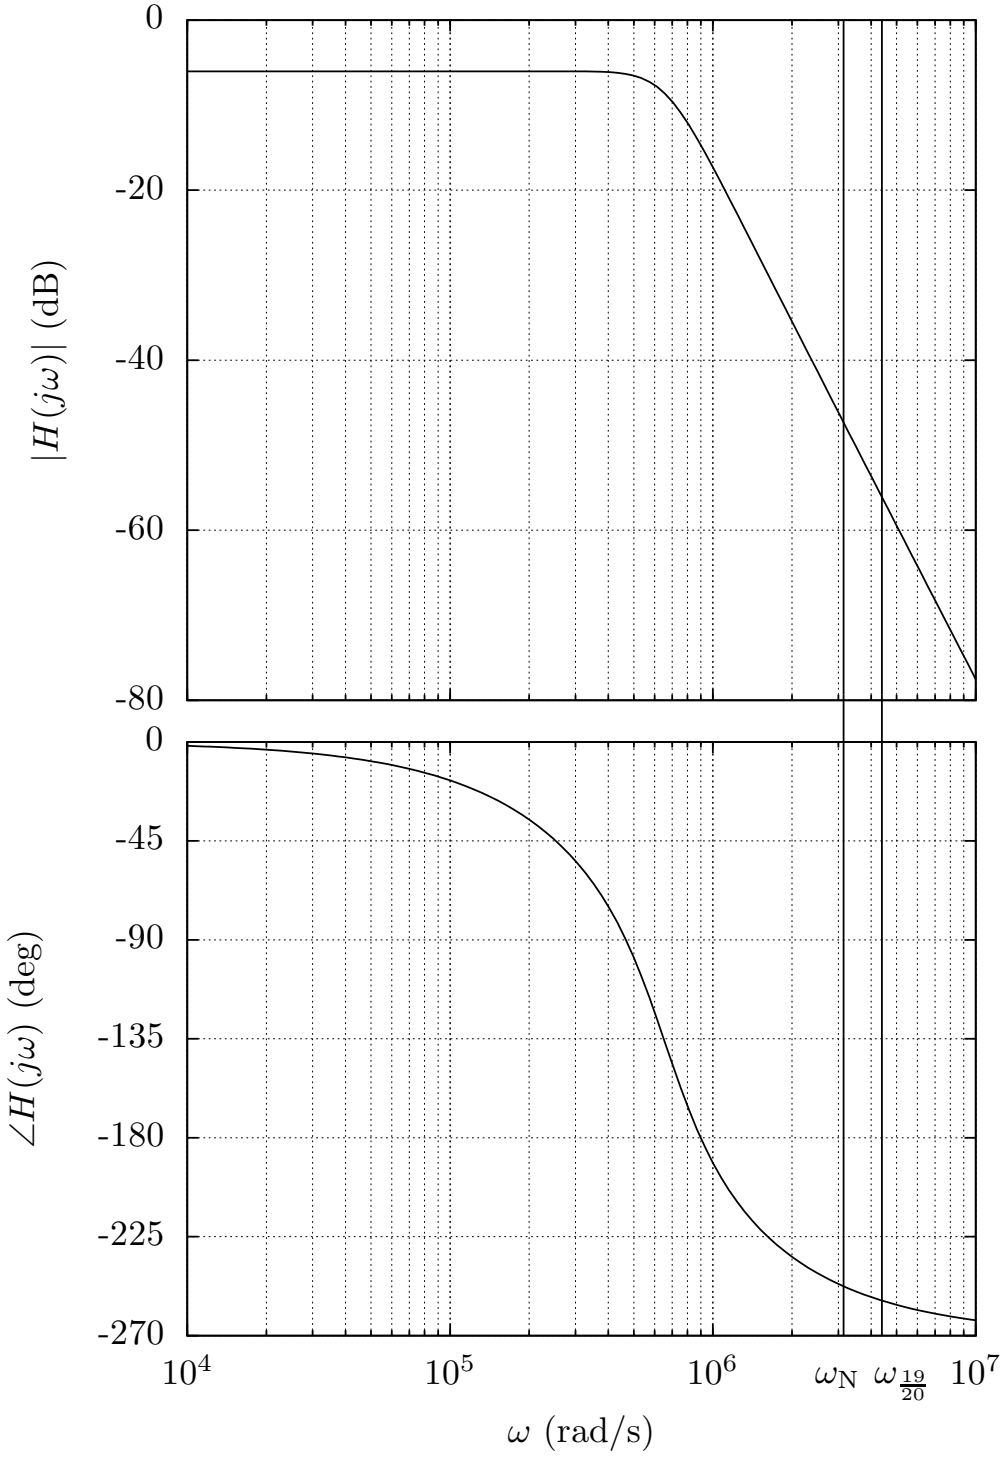

Figure 4.10: Bode diagram of the baseband filters used as anti-aliasing filters and reconstruction filters. The diagram has been plotted from the transfer function of the filters derived in appendix A. The $19/20\,\pi$ mode of the superconducting cavities is attenuated by 50 dB while the phase shift is negligible for frequencies in the low kHz range.

The FPGA used on the controller boards is a Xilinx Spartan 6 (`XC6SLX45`). This FPGA provides 58 18×18 bit hardware multipliers that significantly speed up the signal processing used in control algorithms. The Spartan 6 FPGA family is intended for high-volume applications. Thus these FPGAs can be expected to be commercially available for a much longer time than the high-performance Virtex series that has been used for other rf control systems [36, 47]. The Spartan 6 FPGAs have been chosen since the availability of components is an important factor for a production system that has to be maintained over many years and the additional resources of the Virtex FPGAs are not needed for the control algorithms used at the S-DALINAC. Up to now roughly 25% of the FPGA's resources are used leaving enough space for future extensions of the algorithm.

After the digital signal processing steps inside the FPGA the digital I/Q output signals are converted back to analog using DACs (Analog Devices `AD9747`). Low-pass filters identical to the anti-aliasing filters are used to construct smooth analog signals that are modulated on the rf signal and sent to the rf amplifier.

A micro-controller located on the FPGA board allows setting the parameters of the control algorithm via CAN bus (Controller Area Network) from a PC. Furthermore, it can send commands to the tuner power supplies via CAN bus. A USB 2.0 connection allows to continuously stream out signals from the FPGA to the PC for monitoring.

4.4 Error analysis and calibration

4.4.1 Demodulator

As described in section 4.2.1 the I and Q output signals of the demodulator suffer from dc offsets that may result in wrong values for the input phase. To minimize this error, the demodulator has to be calibrated. As offsets have the biggest impact on small input signals, offset calibration is performed without an input signal. Instead the rf input is terminated by a $50\,\Omega$ load. Since the dc offsets are small (typically 0.5%), they can be eliminated digitally inside the FPGA.

In addition to dc offsets, the gains of the I and Q paths of the demodulator might differ slightly. An appropriate way to measure the gain mismatch is by applying an rf signal to the input of the rf board that differs by approximately 1 kHz from

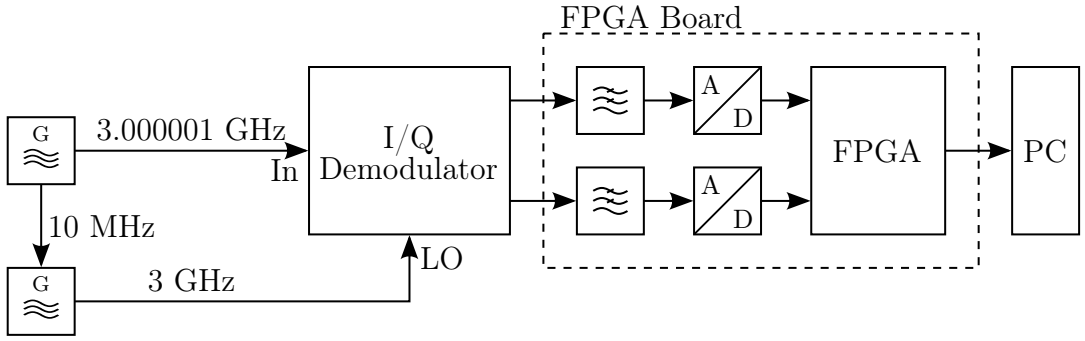

Figure 4.11: Setup for measuring the gain mismatch, the quadrature phase error, and the rms phase error of the demodulator. The two rf generators provide two frequencies that differ by 1 kHz. This results in a rotating vector in the I/Q plane that can be analyzed.

the local oscillator (see fig. 4.11). Frequency drifts between the two oscillators can be prevented if the devices are coupled by a 10 MHz synchronization signal during calibration. The input signal results in an ellipse in the I/Q plane that can be visualized on a PC using a software oscilloscope. Its eccentricity is compensated for by adjusting the gain factor of the signals inside the control algorithm.

The described errors are very small and therefore hard to determine exactly with the software oscilloscope. A MATLAB® [48] script has been developed to increase the accuracy and the reproducibility of the calibration process. It uses a recorded time series of the digital I and Q values as input and automatically fits functions of the form

$$s_I(t) = A_I \cos(\omega t + \varphi_I) + C_I \quad \text{and} \tag{4.13}$$

$$s_Q(t) = A_Q \sin(\omega t + \varphi_Q) + C_Q \tag{4.14}$$

to this data. In addition to the offsets C_I and C_Q that can directly be used for calibration, the script also calculates the gain mismatch $\alpha = A_Q/A_I$ as well as the quadrature phase error $\Phi = \varphi_Q - \varphi_I$. These values can be used to compensate the input signal by applying the following linear mapping to the input data:

$$\begin{pmatrix} I(t) \\ Q(t) \end{pmatrix} = \lambda \begin{pmatrix} 1 & -\alpha \sin\Phi \\ 0 & \alpha \cos\Phi \end{pmatrix}^{-1} \begin{pmatrix} I'(t) \\ Q'(t) \end{pmatrix} - \begin{pmatrix} C_I \\ C_Q \end{pmatrix} \tag{4.15}$$

The coefficient

$$\lambda = \frac{1}{\max\left\{1, 1 + |\tan\Phi|, \left|\frac{1}{\alpha\cos\Phi}\right|\right\}} \tag{4.16}$$

is needed for technical reasons described in section 6.3. It makes sure none of the matrix elements exceeds the allowed range of $[-1, 1]$. Since λ only changes during calibration, this scaling does not change the behavior of the pre-distortion stage during normal operation.

In addition to the calibration values, the MATLAB® script also computes the rms phase error of the digitized signal with respect to the fitted signal. This can be considered as a measure for the quality of the phase detection that includes static deviations as well as phase noise. With typically $\Delta\varphi \approx 0.5°$ rms the phase error of the new rf boards is only half the error of the previous boards (typically $\Delta\varphi \approx 1.0°$ rms).

For a more thorough analysis the script also provides a plot that shows the dependency of the phase error vs. the fitted phase (which corresponds to the input phase). This allows to distinguish between static deviations and phase noise. Figure 4.12a shows this data plotted for one of the prototype rf boards developed in the context of [41, 42]. The phase error shows a clear dependency on the input phase. Thus the rf control system has to deal with nonlinear phase input data. Additionally, if the operator shifts the phase of two consecutive cavities by the same value, this can result in a shift relative to each other. If this effect is not compensated for by readjusting the phases of both cavities, the performance of the accelerator will decrease.

Figure 4.12b shows the same data for one of the new rf boards that do not suffer from this effect anymore. The phase error does not show any systematic dependency on the input phase, instead it is dominated by phase noise. Most of this phase noise is caused by the two rf generators that have been used for the measurement. Although they are coupled by a 10 MHz synchronization signal, they contain independent rf oscillators that are phase-locked to the 10 MHz synchronization signal (phase-locked loop operation).

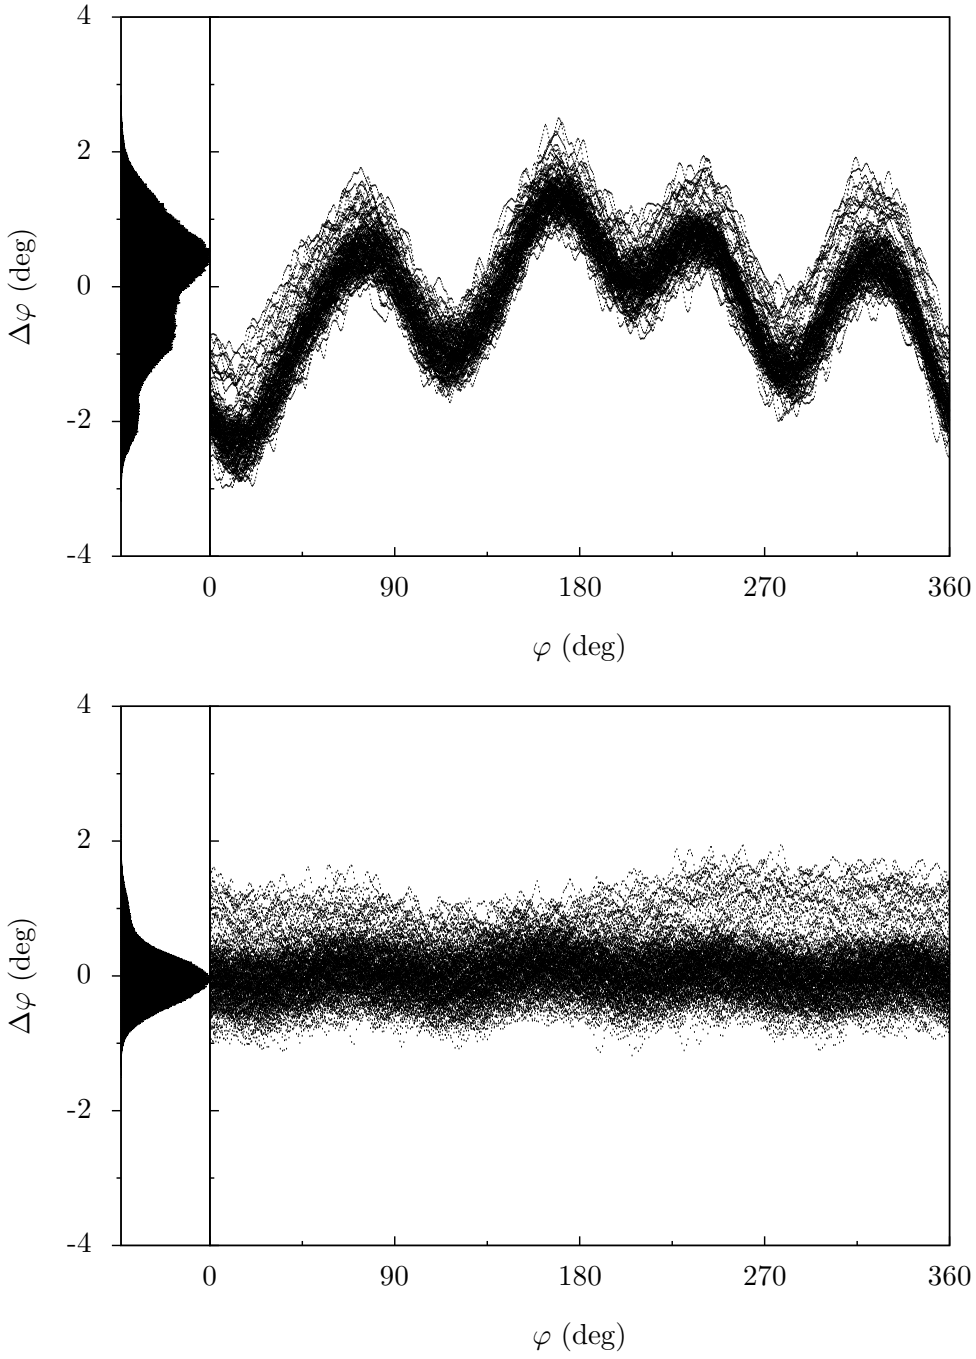

Figure 4.12: Phase error of the digitized signal measured for (a) one of the prototype rf boards and (b) one of the new rf boards. In case of the prototype board the deviation of the measured phase has a clear dependency on the phase of the input signal whereas the second plot is dominated by phase noise. The plots have been created from one million data points that have been acquired during approximately one second. The small plots on the left show the projection of the data to the $\Delta\varphi$ axis.

4.4.2 Modulator

As described in section 4.2.1, dc offsets on the I and Q inputs of a vector modulator lead to carrier leakage (see fig. 4.13a). If the carrier signal is strong enough, this can make self-excited loop operation impossible because the loop locks to the phase of the carrier rather than to oscillate freely (self-locking). This can be avoided by modulator calibration.

The procedure used for modulator calibration is partly based on [49] and [50]. The results obtained with iterative optimization of the I and Q offsets and the gain mismatch and phase mismatch respectively are considerably better than the ones obtained with the more sophisticated procedure described by [51]. This is probably caused by compression effects of the modulator that are not taken into account by the model that is the basis for the last procedure.

Comparison of the magnitude of the carrier and the signal is achieved by a single sideband sine modulation. To this end, a rotating vector in the I/Q plane is modulated with $f \approx 1\,\mathrm{kHz}$ on the carrier while observing the rf output with a spectrum analyzer. The rotating vector is generated in the FPGA using a counter providing a sawtooth signal as phase input for a CORDIC block (see section 6.4). Offset calibration is done by iteratively fine tuning the I and Q offsets for optimal carrier suppression.

In a second step the gain mismatch and phase mismatch of the modulator are compensated. This is done by alternately adjusting the phase mismatch Φ and the gains for I and Q until the magnitude of the image frequency is minimized. A typical result of these calibration steps is shown in fig. 4.13b.

The calibration values obtained by the described procedure slightly depend on the magnitude of the output signal. This is mainly caused by compression effects for high output amplitudes. To minimize the error introduced by these effects, the modulator is calibrated at 50% of its output power which is close to the expected working point of the rf control system.

All correction steps are done completely digitally, and the results can be saved by the accelerator's control system. Each combination of rf module and FPGA module has to be calibrated individually after fabrication. So far, recalibration has not been necessary.

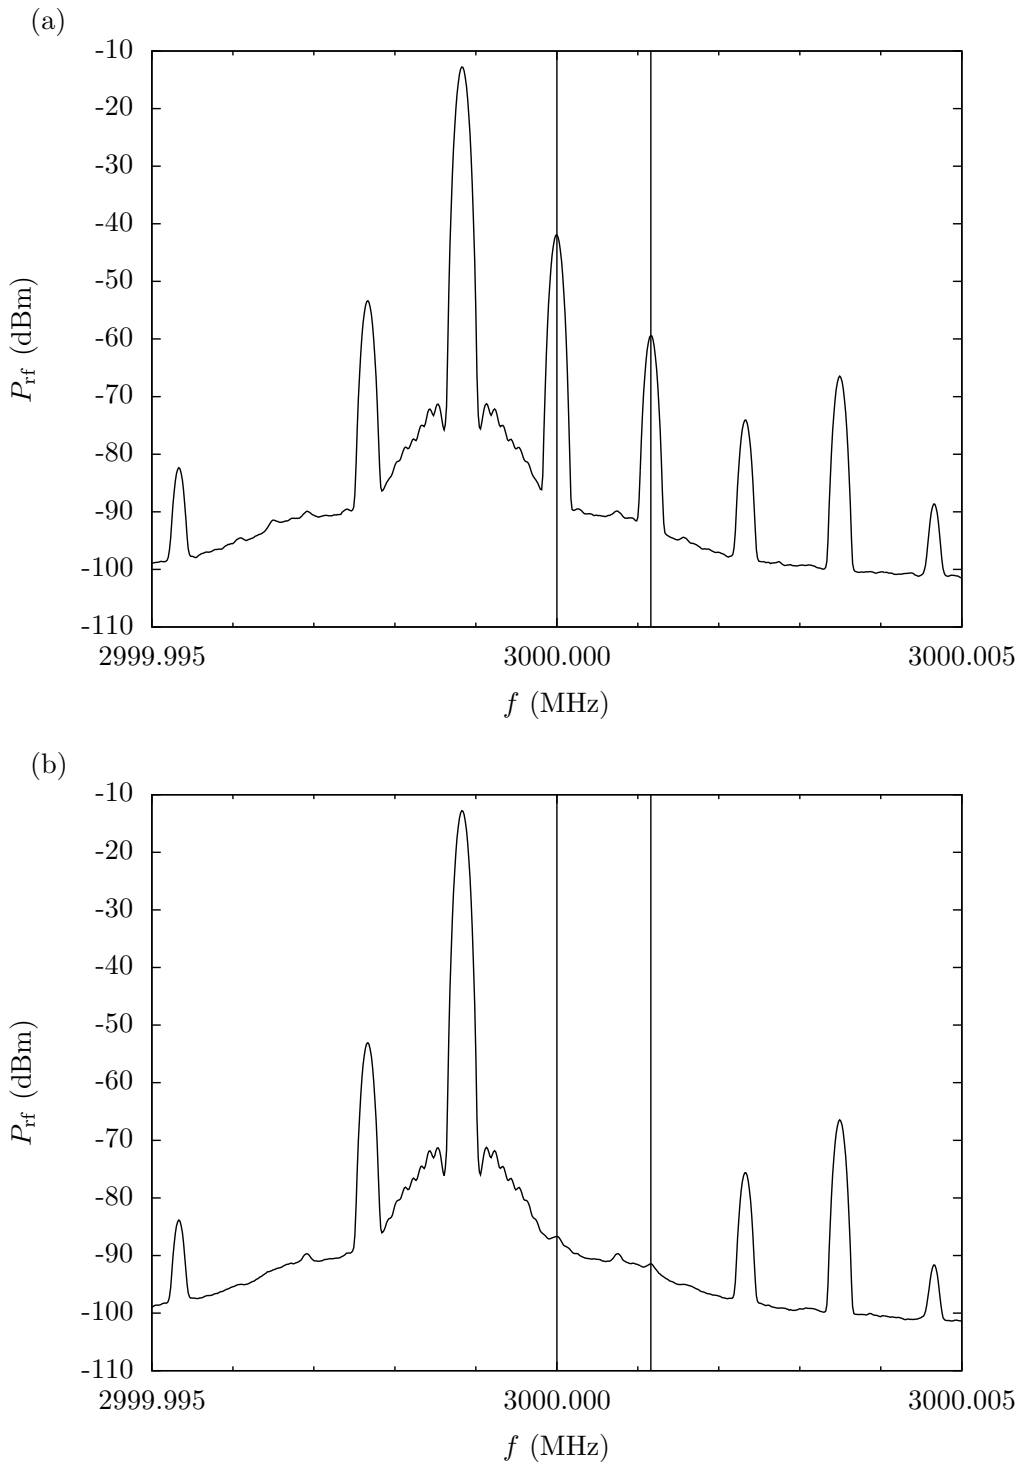

Figure 4.13: Measured spectrum of the modulator output signal (a) without calibration and (b) after calibration. The calibration procedure significantly improves carrier rejection (3000.000 MHz) and decreases modulation distortion ($\approx$ 3000.001 MHz). The remaining peaks are caused by higher order errors of the modulator.

4.5 Reference signal

The rf reference signal for all cavities is generated by the master oscillator. The new rf control system uses the same master oscillator as the old analog system since its accuracy is comparable to modern commercial devices. The master oscillator generates a harmonic rf signal with a frequency of 2.997 GHz. After amplification this signal is split and distributed to all controller boards.

The 5.994 GHz signal for the $2f$ harmonic buncher is not created by a second rf generator since the phase noise of both oscillators would cause this cavity to jitter relative to all other cavities. Instead the signal is created by doubling the frequency of the 2.997 GHz signal [20]. Since frequency doubling is a passive process, it has negligible influence on phase jitter. All cavities are affected by the phase noise of the master generator in the same way.

4.6 Tuners

During operation the eigenfrequency of the superconducting cavities has to be kept close to the operating frequency. Therefore the eigenfrequency can be shifted by means of a tuning system. A magnetostrictive actuator allows continuous fine tuning whereas a motor tuner provides a larger tuning range.

4.6.1 Fine tuner

The fine tuners used with the superconducting accelerating cavities consist of a nickel rod that is surrounded by a superconducting coil. When a magnetic field is applied to the rod, it varies in length changing the length of the cavity (magnetostrictive effect).

The superconducting coil is powered by a current-controlled power supply. To keep the number of different components low, the fine tuners are driven by the same type of `CPS05-7` power supply that is used to power small magnets at the S-DALINAC.

The fine tuners are designed for a tuning range of 1 kHz [19]. In practice their tuning range varies from approximately 500 Hz to 5 kHz. Figure 4.14a shows the characteristics of one of them. As can be seen from the figure the magnetostrictive

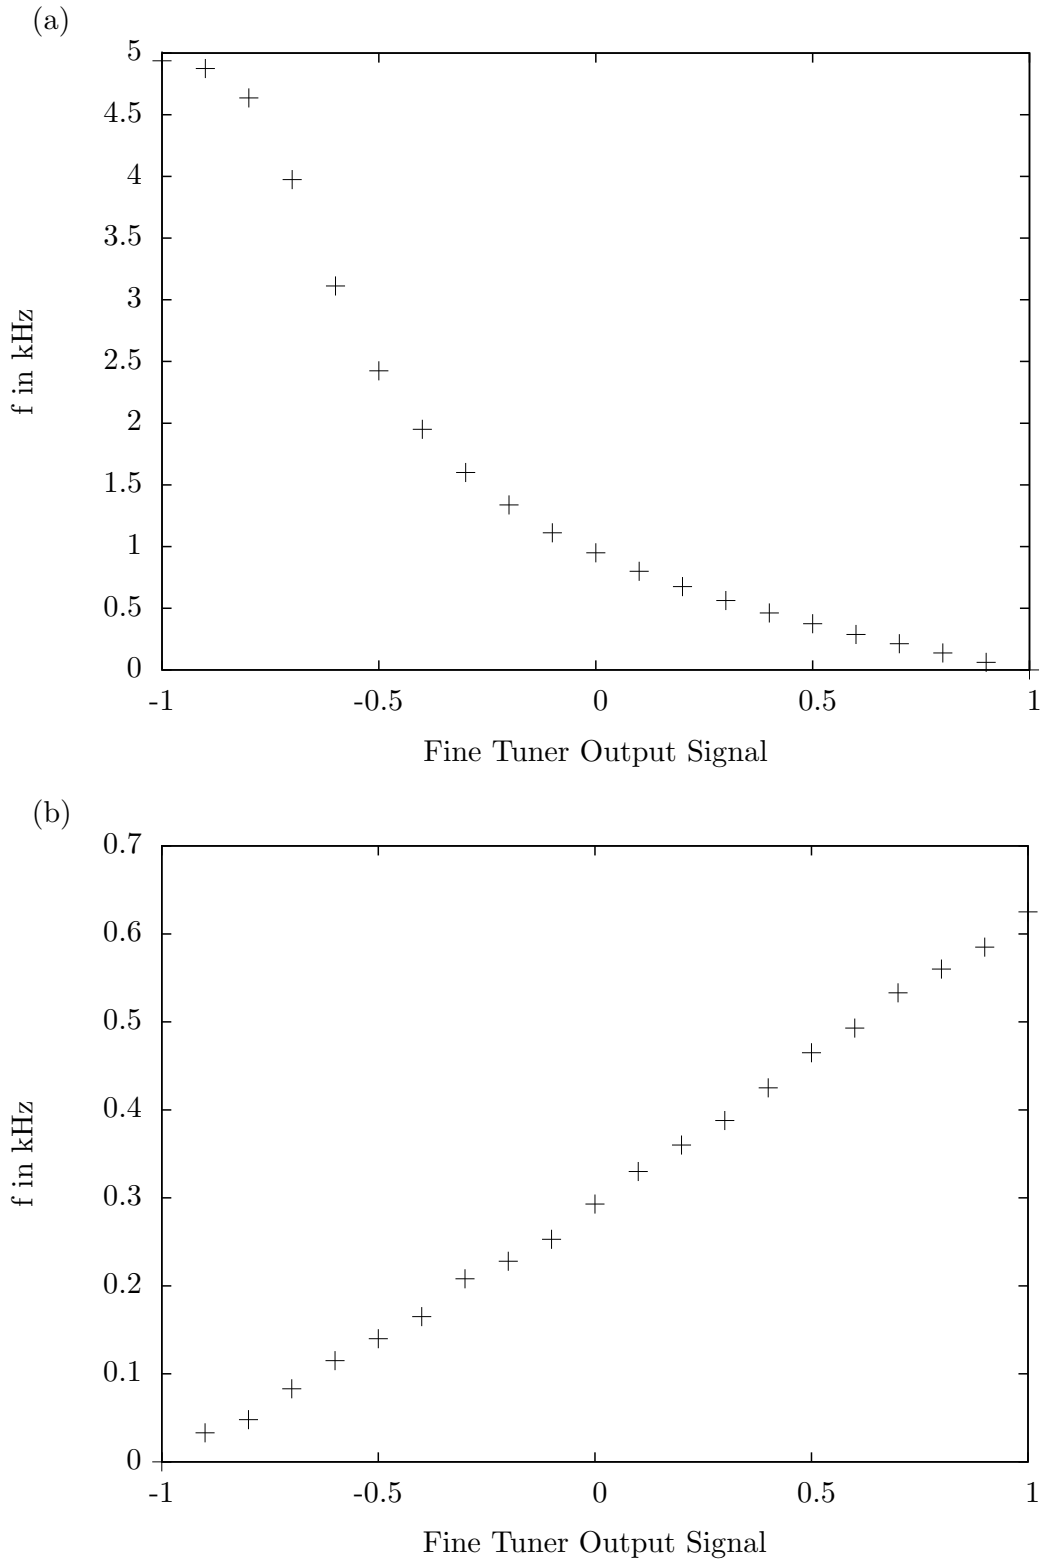

Figure 4.14: Characteristic of (a) one of the magnetostrictive fine tuners and (b) a prototype of a piezoelectric fine tuner. Note that for technical reasons only 83% of the range of the piezoelectric tuner could be used for the measurement.

fine tuners have a nonlinear transfer behavior. This is mainly caused by the remanence of the ferromagnetic rod [19].

The magnetostrictive fine tuners of the S-DALINAC are mounted very closely to the superconducting cavities. Thus their magnetic field can reduce the performance of the cavities [32]. Therefore it is planned to replace them with piezoelectric fine tuners that do not suffer from this problem [52]. Figure 4.14b shows the characteristics of a prototype of a piezoelectric tuner measured in a vertical bath cryostat. In contrast to the magnetostrictive fine tuners a `CPS08GUN` voltage power supply has been used to generate the control voltage for the piezo driver.

The magnetostrictive fine tuners need to be degaussed before warming up the cavities to the normal-conducting state. Otherwise magnetic flux can be frozen inside the superconductor during the next cool down which significantly increases the residual resistance of the cavity. This results in an increased heating load of the cavity and thereby can decrease the overall performance of the accelerator. The degaussing of the fine tuners has to happen while the tuner is cold since the tuners use superconducting coils. While the old rf control system needed a special device for degaussing the fine tuners, the new system provides a software degaussing feature (see section 13.1).

4.6.2 Coarse tuner

The tuning range of the fine tuner is too small to tune the cavities to the operating frequency after cool down. Therefore a motor tuner extends the tuning range to 1 MHz [19]. Via a gear mechanism the motor operates a lever system that stretches or squeezes the cavity. Like the fine tuner the dc motor is powered by a `CPS05-7` power supply but here it is operated in voltage controlled mode. This allows to tune the cavity with constant speed independent of the torque.

Operating the motor tuner introduces vibrations that increase microphonics. Therefore its usage should be kept to a minimum.

4.7 Backplane and crate controller

In addition to stand-alone operation the controller boards can also be installed into crates. Two crates are needed to house the 16 controller boards used at the

S-DALINAC. The first crate contains the controllers for the laser system of the spin-polarized electron gun, the three normal-conducting chopper and buncher cavities, and the four superconducting cavities of the injector linac. The second crate houses the controllers for the main linac. A custom backplane connects the boards with a crate controller card. The crate controller powers the controller boards and connects them to the PC. When used with a crate controller no direct CAN and USB connections between the PC and each individual controller board is needed.

The two crates are coupled via cables connecting their crate controllers. In dual-crate operation one of the crate controllers acts as master and the other one as slave. All controller boards in both crates are running synchronously. The clock signal for all FPGA boards is generated by the master crate controller and is distributed via the backplanes. Synchronous operation of all boards ensures that all boards sample data at the same time and at the same rate. This e. g. allows to study the correlation of microphonics in neighboring cavities. Synchronous sampling also simplifies fast data read-out.

4.8 Interlock

There are different reasons why it can be necessary to shut off the cavities. Therefore the FPGA boards are equipped with an integrated interlock function that switches off the rf output if no "interlock ok" signal is applied to the interlock input. This function is implemented without programmable components to ensure fail-safe operation. For each of the controller boards a separate interlock input connector is available at the backplane. This allows to implement different interlock conditions for the cavities. For example the interlock input of the chopper cavity's controller is connected to a vacuum gauge to protect the cavity from damage caused by bad vacuum (machine protection interlock). The superconducting cavities are automatically shut off if one of the doors to the accelerator hall is open (radiation safety interlock).

An additional interlock output can be used to connect e. g. amplifiers. This output is switched off if the interlock input is off (hard wired). In addition to this it can also be switched off by the FPGA or micro-controller if e. g. a hardware malfunction is detected (soft interlock).

4.9 Installation of the hardware

All hardware components of the rf control system have been installed at the S-DALINAC in the beginning of 2011. In addition to the controller boards this includes the master oscillator, the local oscillator distribution, the frequency doubler for the 6 GHz buncher, the power supplies for fine and coarse tuners as well as a PC server that connects the hardware to the accelerator control system. Due to the compact design all components of the rf control system fit into a single rack (see fig. 4.15). To ensure long-term stability of the rf control system the rack is cooled by temperature-stabilized water of (20 ± 1)°C. Temperature sensors on the rf boards, the FPGA boards, and the tuner power supplies allow continuous monitoring of the temperature.

During the two years of operation the hardware has proven to be very reliable. Although one controller board failed during this time, downtime due to hardware defects of the rf control system has been reduced dramatically compared to the previous analog system.

Figure 4.15: All components of the rf control system have been installed into a single rack.

5 Control algorithms

The different loaded quality factors of the cavities used at the S-DALINAC make different control algorithms necessary. The low-Q_L room-temperature copper resonators have a broad resonance and are relatively rigid: Only slow drifts of the field caused by fluctuations of the ambient air and cooling water temperature are relevant. Therefore the relatively simple generator-driven resonator control algorithm is sufficient.

In contrast the superconducting cavities have a high loaded quality factor. Due to their 20-cell design they have low flexural stiffness and thus are very susceptible to microphonics. Hence the controller has to operate much faster. In addition it is important that a low-bandwidth superconducting cavity can be excited even if it is detuned by several bandwidths. This is accomplished with the self-excited loop algorithm. The following sections describe how these control algorithms have been implemented in the context of this thesis. Both algorithms are designed to run on the soft DSP.

5.1 Generator-driven resonator

Figure 5.1 shows a simplified flow-chart of the generator-driven resonator (GDR) algorithm described in section 2.1 as it is used for the copper cavities. The I and Q input signals are transformed into polar coordinates and back to Cartesian

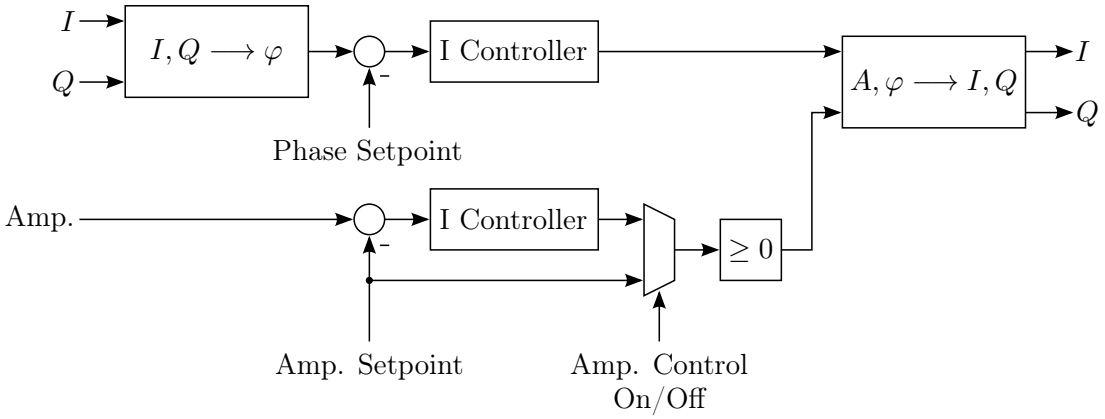

Figure 5.1: Simplified flow-chart of the generator-driven resonator control algorithm used with the room-temperature cavities. Amplitude and phase of the output signal are generated by the integral controllers.

coordinates in the FPGA by the CORDIC algorithm which will be described in section 6.4. In theory it would also be possible to implement the controller directly in Cartesian coordinates using two integral controllers for I and Q. But this requires the phase shift of the controlled system to be a multiple of 2π for an on-resonance cavity which could be achieved by adding a phase shifter to the plant. However, this does not work properly with the normal-conducting cavities at the S-DALINAC because the phase shift introduced by the system drifts[2]. An implementation in polar coordinates has also been preferred since this makes behavior and operation similar to the self-excited loop algorithm.

After conversion to polar coordinates phase control is achieved by an integral controller only. The integrator is necessary to eliminate steady-state offsets. Since there are no fast perturbations, but only slow drifts, no proportional controller is needed for these cavities.

The amplitude controller follows a similar design, but only positive amplitude values must be fed into the output CORDIC, as negative values are unphysical and would cause ambiguities after the CORDIC transformation.

A multiplexer allows to switch between the following modes of operation:

1. Constant output amplitude: In this mode only the phase controller is active while amplitude feedback is turned off. Instead the output magnitude can be adjusted directly by the operator.

2. Amplitude control switched on: Phase and amplitude feedback are both active. This mode is used for accelerator operation.

Unlike the self-excited loop algorithm there is no need for complicated commissioning: A cavity on standby can be switched on straight into mode 2 without any further user interaction.

5.2 Self-excited loop

In contrast to the GDR algorithm a self-excited loop (SEL, [30, 37]) oscillates freely at a frequency that is determined by the eigenfrequency of the resonator and the loop phase. As described in section 2.2 this frequency can be locked to the

[2]This is mainly caused by a temperature-dependent phase shift introduced by the amplifiers.

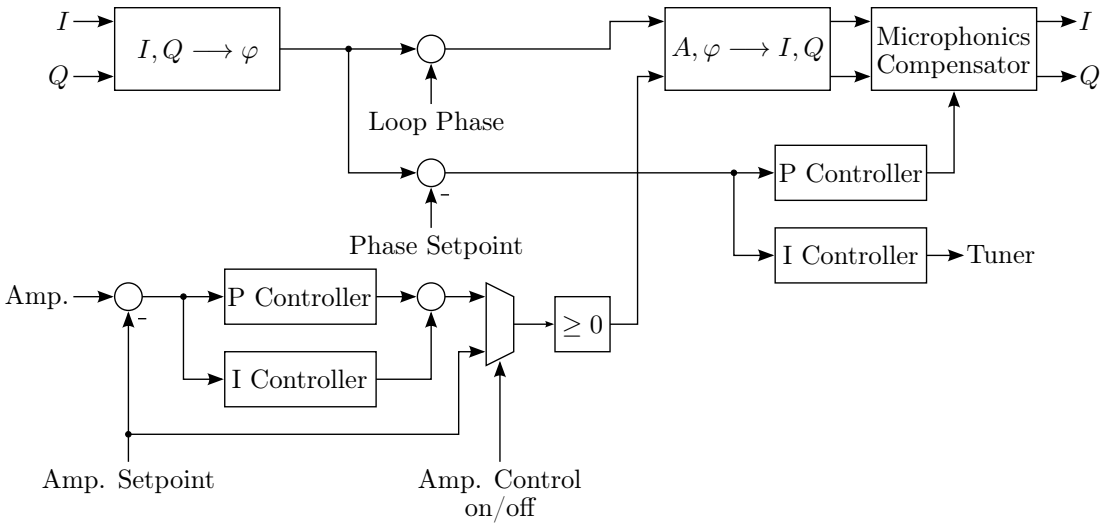

Figure 5.2: Simplified flow-chart of the self-excited loop control algorithm used with the superconducting cavities. In contrast to the GDR algorithm the SEL algorithm passes through the phase-shifted input signal. A correction to this signal is applied by the microphonics compensator.

frequency of the master oscillator by a controller that tunes the cavity's resonance frequency and/or applies an additional phase shift (phase-locked loop).

The advantage of the SEL is that it immediately excites the cavity even if the cavity's eigenfrequency is detuned by many bandwidths. Furthermore the controller can recover from a breakdown even in the presence of Lorentz force detuning that might prevent a GDR from restarting oscillation.

Figure 5.2 shows a block diagram of the chosen implementation of the SEL algorithm. Again the I and Q signals are transformed into polar coordinates and back to Cartesian coordinates by the CORDIC algorithm. Polar coordinates allow distinct controllers and parameters for amplitude and phase. Note that while the controller is implemented in polar coordinates the phase actuator is implemented in Cartesian coordinates.

To avoid excitation of mechanical eigenmodes the eigenfrequency of the resonator can only be tuned slowly. An integral controller is used for eigenfrequency control because it eliminates steady-state frequency offsets by tuning the SEL to exactly match the frequency of the master oscillator. The integral controller is also able to deal with the remanence of the magnetostrictive fine tuners. Fast perturbations are compensated for electrically. Instead of an actuator that shifts the phase [41], a

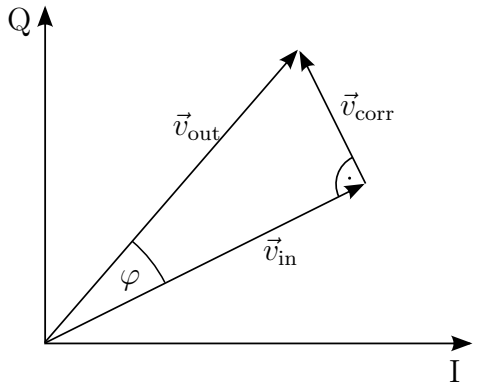

Figure 5.3: The microphonics compensator adds a correction vector in quadrature to its input vector and thereby corrects for phase and magnitude error caused by a detuned cavity.

microphonics compensator [30] is used in this work. This block adds an orthogonal correction vector to the input vector (see fig. 5.3):

$$\vec{v}_{\text{out}} = \begin{pmatrix} 1 & -k\,\varphi_e \\ k\,\varphi_e & 1 \end{pmatrix} \vec{v}_{\text{in}} \tag{5.1}$$

If the length of this correction vector corresponds to the phase error φ_e, the microphonics compensator has the inverse transfer function of the resonator. This can be shown by combining the amplitude transfer function of the cavity (eq. (3.16)) and the phase transfer function (eq. (3.19)). This yields the amplitude transfer function as a function of the phase shift

$$|G(\varphi)| = \frac{R_L}{\sqrt{1+\tan^2\varphi}} = R_L \cos\varphi \quad . \tag{5.2}$$

According to fig. 5.3 the length of the output vector of the microphonics compensator is

$$|G_{\text{mc}}(\varphi)| = \frac{|\vec{v}_{\text{out}}(\varphi)|}{|\vec{v}_{\text{in}}|} = \frac{1}{\cos\varphi} \tag{5.3}$$

Thus phase and amplitude errors, which always occur correlated if they are caused by detuning of the cavity, are corrected in a single step. The microphonics compensator is driven by a proportional controller.

In addition to the microphonics compensator an amplitude controller is needed to compensate for perturbations of the field amplitude. Again this is done by a proportional controller for fast perturbations that cannot be removed completely

by the microphonics compensator, whereas an integral controller removes the steady-state offset. The amplitude controller also compensates for changes of the accelerating field strength caused by beam-current fluctuations.

If eigenfrequency control is activated, the actuator signal for the fine tuner is sent to the power supply by the micro-controller via the CAN bus on a regular basis. The motor tuner is equipped with a three-step switching controller that automatically activates the motor as soon as the fine tuner approaches its limits. The motor tuner causes vibrations and increases the amount of microphonics the rf control system has to compensate. Therefore a hysteresis ensures that the cavity is always tuned to the middle of the fine tuner's controlling range once the motor has been activated. Due to this behavior the motor is rarely activated. In fact, it usually does not happen during beam-time.

The activation of the motor tuner is inhibited while the rf signal is switched off by the operator or the rf interlock to avoid that small offsets or noise detune the cavity. In addition the values of the integrators are held to reduce the recovery time after short pauses of beam operation. This is particularly important for the output value of the fine tuner.

The following modes of operation are possible:

1. Constant output magnitude: The rf control system acts as a limiter. The phase is passed through without any feedback. This mode can be used to test if the SEL oscillates freely (no self-locking).

2. Microphonics compensation: The microphonics compensator is switched on while the output amplitude is otherwise kept constant. This mode is only useful in situations when the cavity is tuned to the frequency of the master oscillator.

3. Microphonics compensation and resonance control: Like mode 2, but the cavity is tuned to the reference frequency automatically.

4. Microphonics compensation, resonance control, and integral amplitude feedback: Like mode 3, but the integral amplitude controller regulates the electric field strength to meet the set-point on average.

5. Full phase and amplitude feedback: In addition to mode 4 the proportional amplitude controller is activated. This mode is used for accelerator operation.

When operated for the first time, the cavity needs to be tuned to resonance manually. This is done by using a network analyzer synchronized to the master oscillator. Alternatively, the cavity's eigenfrequency can be tuned slowly in mode 1 while keeping the transmitted power at its maximum by adjusting the loop phase. In this case, no further equipment is needed. The progress of the tuning process can be monitored by the frequency of the I or Q signal. After tuning the cavity to the reference frequency the loop phase is adjusted for maximum transmitted power. This corresponds to a total phase shift along the open loop that is a multiple of 2π. As long as the operating frequency of the cavities remains the same there is no need to re-tune the loop phase later on.

The next step is to activate the microphonics compensation (mode 2). This compensates for a possible small deviation of the resonator from the reference frequency by increasing the output power as well as shifting the phase of the output signal to stabilize the phase of the field. After switching to mode 3, the integral phase controller ensures the cavity's eigenfrequency matches the reference frequency. This removes the steady-state offset on the input of the microphonics compensator and thereby reduces the amount of rf power that is needed to drive the cavity.

Finally amplitude feedback is activated (modes 4 and 5). The proportional controller stabilizes the field on a fast timescale whereas the integral controller removes the steady-state offset.

The described procedure has to be followed only at first-time operation. After a normal rf shutdown (leaving the loop phase and tuning unchanged) the controllers recover automatically.

The performance achieved with the implementation of both control algorithms will be discussed in section 7.

6 FPGA implementation

The soft DSP that executes the control algorithms as well as the technical infrastructure for reading out the ADCs, controlling DACs etc. have been implemented in the FPGA in the VERILOG hardware description language [53]. The details of this implementation will be discussed in the following.

6.1 Synthesis of the HDL code

The VERILOG code has been synthesized using XILINX ISE WEBPACK [54]. This software translates the VERILOG description of the signal flow from one register to the next (register transfer level description) to a hardware representation. During this process abstract VERILOG elements are translated into a gate level representation of the corresponding circuit along with net-lists describing the interfaces between the components of the system (e. g. between FPGA and the AVR micro-controller). This is done by mapping the VERILOG description to resources in the FPGA like logic cells and I/O blocks (map process). In the next step they are fitted into a suitable location in the gate array and wired with the rest of the logic (place and route process). At the end a bit-level description of the FPGA configuration is generated (bit-stream generation). During the whole process the design software checks that all timing constraints are met.

To configure the FPGA, the bit-stream is written into its memory. In case of the `LLBBC10` FPGA board, the bit-stream can also be written into a separate flash memory to make it persistent. When power is turned on, the bit-stream is loaded into the FPGA automatically. The bit file can be programmed into the virgin FPGA or the flash memory using XC3SPROG [55]. This software allows programming, reading back and verifying the FPGA bit-stream via its JTAG interface. This interface allows in-circuit programming and in-circuit testing of all re-programmable components. XC3SPROG has also been extended to allow access to the JTAG chain via the USB connector on the front of the boards or via the crate-controller card. After the controller board has been programmed for the first time, it can also be updated via CAN bus using the boot-loader of the AVR micro-controller [56].

6.2 Overview of the HDL code

The VERILOG code is organized in several functional modules (see table 6.1). This improves maintainability of the code and eases replacement of components like ADCs during further development of the board.

The correctness of all important VERILOG modules including the soft DSP has been verified by unit tests. These tests can be run on the PC in a VERILOG simulator. With IVERILOG [57] and CVER [58] two open source simulators have been used during the development process. Both allow very fast unit testing on the command line.

Module name	Description
`ad7982_ctl`	Controller for the fast ADCs
`ltc2484_ctl`	Controller for the high-precision ADCs
`hf_alu`	Arithmetic logic unit of the soft DSP
`hf_ctl`	Controller for the arithmetic logic unit of the DSP
`cordics`	CORDIC cores for coordinate transform
`llbbc_reg`	Provides access to parameters and variables from the AVR micro-controller
`usbsource`	Controller for the USB interface
`ft2232h_fifo`	Data FIFO for the USB interface
`sumofsquares`	FPGA part of the rms error calculation
`llbbc_backplane`	Connects and synchronizes the board to the back-plane
`clock_gen`	Clock generation for all components
`frequency_supervisor`	Supervision of the clock frequencies
`sin_nco`	Generates signals for hardware test
`spi_flash`	Provides access to the flash memory from the AVR micro-controller

Table 6.1: Overview of the modules the VERILOG code is split into.

6.3 Arithmetic used for signal processing

The control algorithms described in section 5 are executed by the custom digital signal processor implemented in the FPGA. Thus this DSP has to provide all arithmetic operations needed in these control algorithms. This comprises addition, subtraction and multiplication.

Fixed-point arithmetic has been preferred over floating-point arithmetic because it is faster and needs less resources in the FPGA. In addition to that the input signals from the ADCs and the output signals to the DACs are specified as fixed-point numbers by the vendor of these components.

Since the control algorithms use positive as well as negative values for signals (e. g. the I and Q input signals), the DSP expects all numbers representing signals to be in two's complement representation [59]. This has the advantage that the required basic arithmetic operations can be implemented in the same way they are for unsigned numbers. This makes the implementation very simple and efficient. The value of a number a given in N bit two's complement representation with the bits a_i can be obtained from the following formula:

$$a = -a_{N-1}2^{N-1} + \sum_{i=0}^{N-2} a_i 2^i \tag{6.1}$$

Since the ADCs have a resolution of 18 bit and the FPGA provides 18×18 bit hardware multipliers, 18 bit arithmetic is used in the DSP as well. The values of all signals are normalized to the maximum magnitude in the following to make them more intuitive:

$$a' = \frac{-a_{N-1}2^{N-1} + \sum_{i=0}^{N-2} a_i 2^i}{2^{N-1}} \tag{6.2}$$

This results in values in the interval $[-1.0, +1.0[$.

When control algorithms are implemented digitally, special care has to be taken of arithmetic overflows. Consider e. g. adding the numbers 14_{10} and 21_{10} using 6 bit two's complement arithmetic:

$$\begin{array}{rrr} 001110 & 14 & 0.43750 \\ +010101 & +21 & +0.65625 \\ \hline 100011 & -29 \; \lightning & -0.90625 \; \lightning \end{array} \tag{6.3}$$

Note that in addition to producing the wrong absolute value the most significant bit indicating the sign has flipped. Wrap-arounds like this result in discontinuities which can seriously disrupt operation of a control loop and therefore must be prevented.

The easiest way to eliminate overflows is to use wider result registers. For many consecutive operations this leads to a high resource consumption in the FPGA. In addition to that there are many operations like the correction of a small offset that lead to a poor utilization of the added bits. This can reduce the effective resolution of the output DACs. Moreover it is not possible to use result registers of arbitrary size with a DSP.

Another way to prevent overflows is using saturation arithmetics. This means that the result signal of e. g. a sum $s_1 + s_2$ is calculated according to

$$s_{\text{sum}} = \begin{cases} s_{\max} & \text{if } s_{\max} \leq s_1 + s_2 \\ s_1 + s_2 & \text{if } s_{\min} < s_1 + s_2 < s_{\max} \\ s_{\min} & \text{if } s_1 + s_2 \leq s_{\min} \end{cases} \quad . \tag{6.4}$$

The result of a multiplication is 36 bit wide. This result is also reduced to 18 bits using the same technique. The control algorithms described in section 5 make use of this saturation arithmetic for signals like I, Q and amplitude.

In contrast to amplitude values, phase signals have a range of $[-\pi, +\pi[$ resulting in a slightly different normalization of

$$p = \pi\, a' \quad . \tag{6.5}$$

While saturation arithmetic makes sense for amplitude signals, overflows are needed to represent the modulo 2π behavior of angle variables. Hence the DSP needs different operations to add amplitudes (using saturation arithmetic) and phases (using overflowing arithmetic).

6.4 Coordinate transform

The control algorithms for normal-conducting as well as superconducting cavities make use of coordinate transforms from Cartesian coordinates to polar coordinates and back. These transformations require the calculation of trigonometric functions

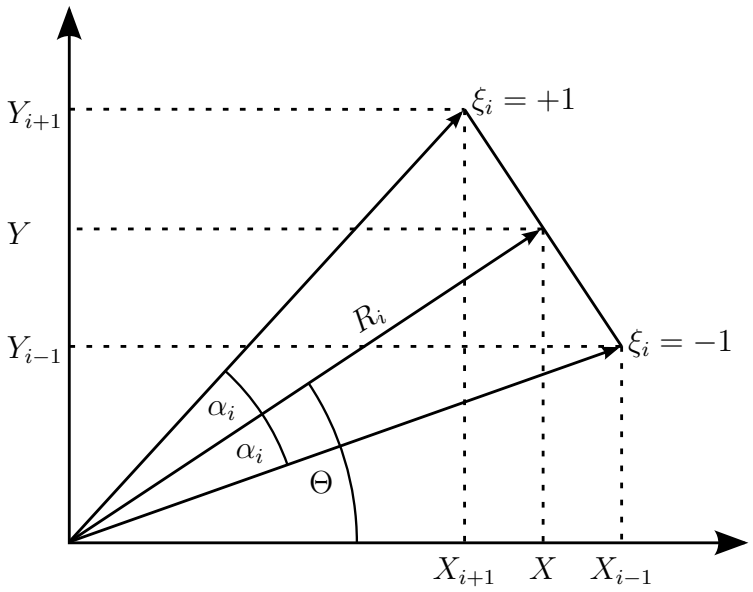

Figure 6.1: Pseudo-rotation step performed by the CORDIC algorithm. Depending on the sign of ξ_i the vector R_i is "rotated" either clockwise or counterclockwise by angle α_i.

which are difficult to implement in an FPGA. The COordinate Rotation DIgital Computer (CORDIC) algorithm [60] solves this problem by calculating vector rotations with an iterative algorithm based on basic arithmetic operations. The algorithm applies a series of pseudo rotations to the input vector. For each of these pseudo rotations an orthogonal vector is added to the input vector (see fig. 6.1). The pseudo-rotated vector is given by

$$X_{i+1} = X_i - \xi\, 2^{-(i-2)} Y_i \tag{6.6}$$

$$Y_{i+1} = Y_i + \xi\, 2^{-(i-2)} X_i \quad . \tag{6.7}$$

The vector is rotated by an angle α_i with each pseudo rotation. Their angles α_i are chosen such that they can be calculated very efficiently:

$$\alpha_i = \tan^{-1} 2^{-(i-2)} \tag{6.8}$$

The length of the orthogonal vector is divided by 2 in every step which can be implemented very efficiently in an FPGA by a bit-shift. It can be shown that any rotation can be approximated by such a series of pseudo rotations with decreasing angle. Each of them rotates clockwise ($\xi_i = -1$) or counter-clockwise ($\xi_i = +1$), respectively. Since no step is skipped the number of pseudo rotations

is independent of the angle. Figure 6.1 also shows that each pseudo rotation increases the magnitude of the vector slightly. Since the CORDIC algorithm always applies the same number of pseudo rotations the increase of the vector length can be precomputed. The precomputed value can be used to compensate for the change in length at the end.

The CORDIC algorithm provides different modes of operation. In the context of this work the ROTATION and the VECTORING mode are used. The ROTATION mode rotates a vector in Cartesian coordinates by a given angle of rotation. Using a vector $(m, 0)$ as input and rotating it by an angle φ results in the Cartesian representation of a vector with magnitude m and angle φ in polar coordinates. This technique is used to convert from polar to Cartesian coordinates.

The conversion from Cartesian to polar coordinates is performed by successively applying pseudo rotations in a way that brings the vector closer to the I axis with each step. At the end the vector is parallel to this axis and the length of the vector can be obtained directly from the I component whereas the angle of the input vector can be calculated by the sum of the pseudo rotations α_i that have been applied.

6.5 Soft digital signal processor

A highly optimized custom digital signal processor has been designed specifically for the task of executing rf control algorithms. The control algorithm is implemented in form of a program that is executed linearly from the beginning to the end. If the DSP reaches the last command, it will proceed with the first one (infinite loop). The absence of an operating system as well as the nonexistence of branching commands renders all DSP programs real-time capable. To exploit the full accuracy of the ADCs the DSP provides 18 bit saturation as well as overflowing arithmetic. The DSP has been stripped down to the functionality needed to implement rf control algorithms (see appendix B for a detailed description of the instruction set) to achieve a high clock rate. Throughput is improved by using a two-staged buffered, synchronous instruction pipeline [59]. With these measures the DSP runs at a clock frequency of 80 MHz. This is 84 times the sampling rate of the ADCs which means that the DSP can carry out up to 84 instructions in its instruction

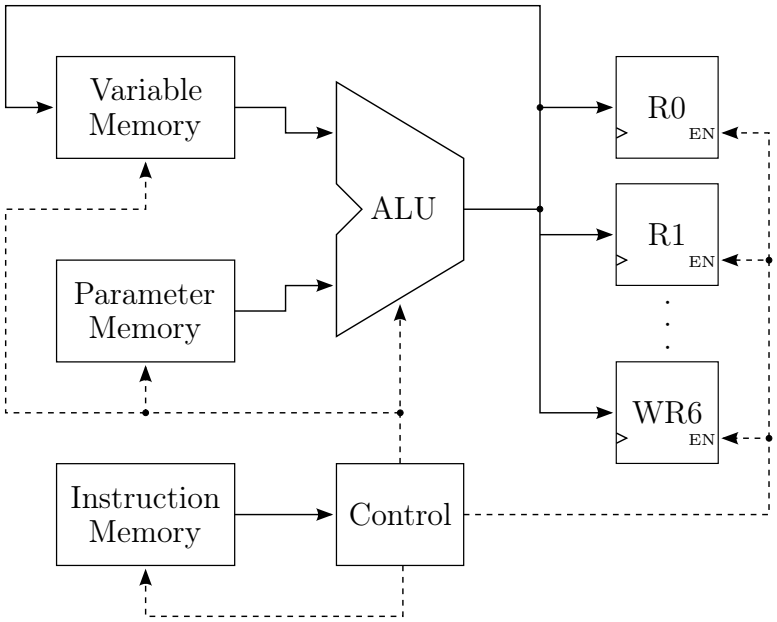

Figure 6.2: Simplified diagram of the digital signal processor implemented in the FPGA. Straight lines show the data flow between the components whereas control lines are dashed.

loop without slowing down the ADCs. This is sufficient to implement the control algorithms described in section 5.

The DSP does not use complex memory access but only variable registers and parameter registers (see fig. 6.2). Parameter registers are used for control parameters that are adjusted by the operator. They are read-only for the soft DSP but read/write for the AVR micro-controller which connects the controller board to the PC. In contrast, variable registers used for intermediate results are read-only from the outside but read/write for the soft DSP.

Most instructions operate on a variable and a parameter. This can e. g. be used to subtract an offset from a signal or to multiply with a control loop amplification that can be adjusted by the operator. Using two variables as operands is possible by copying the value of one of them to one of three 18 bit registers (`Rx`). Values from these registers can be used in the same way parameters are used.

In addition to the 18 bit registers the DSP provides four 36-bit accumulating registers (named `WRx` in fig. 6.2) that can be used for implementing integral controllers or low-pass filters. Only the 18 most significant bits are used for further

processing, but all 36 bits are considered for the accumulated sum. This allows integral controllers with time constants of up to $\approx 1\,\mathrm{s}$ to be realized.

Using a DSP to execute the control algorithm has advantages in diagnostics because all variables (and thereby all intermediary results) are stored in fixed registers (block RAMs) where they can be read out easily to allow for diagnosis on bit level. In particular there is no need for 18 bit data paths and huge multiplexers throughout the whole FPGA that would consume a great deal of the resources. In addition, the soft-DSP approach speeds up the development of the control algorithms. No time consuming synthesis of the VERILOG code is necessary after the control algorithm has been changed. On the other hand the serial processing of the data reduces the performance. For most practical purposes this penalty is small since the latency between ADC and DAC is determined mainly by the CORDIC blocks. The latency caused by the complete control algorithm is approximately $1\,\mu\mathrm{s}$.

7 Measurements

The performance of the rf control system was studied in several measurements to make sure it meets the specifications. Since the new revision of the rf boards is still in production, all measurements presented in this section have been carried out using prototype boards of the type `HF-ADL-5374 Rev. 1`.

7.1 Cavity measurements

It is essential to take the full frequency range of perturbations into account to get the total error. Pressure fluctuations of the helium bath in the sub-Hz range have to be included as well as microphonic perturbations and noise in the kHz range from digital components and switching power supplies. Commercially available rf measurement devices that can quantify the performance independently of the controller hardware cannot provide the needed relative amplitude accuracy of 10^{-5} rms. For that reason the performance of the controller has been determined by analyzing its own digital signals.

The data acquisition software writes the data from the FPGA to hard-disk for an off-line analysis with the full sampling rate of the ADCs over a period of one minute. An appropriate indicator for the performance is the residual error of the controllers. Since residual errors are represented by very small numbers, data with the full resolution of 18 bits has been used for further analysis.

Figure 7.1 shows one second of a typical recorded error signal. The histogram of the signal has been computed from the full data set (60 seconds). It shows that the signal is nearly Gaussian distributed with a deviation of 2.4° rms.

A good visualization of the different perturbation frequencies can be obtained by calculating the "integrated" spectrum [61]

$$(\Delta\varphi(f_n))_{\mathrm{rms}} = \frac{\sqrt{\sum_{i=1}^{n} \left|\mathcal{F}[\Delta\varphi(t)]_i\right|^2}}{\sqrt{2}}, \tag{7.1}$$

where $(\Delta\varphi(f_n))_{\mathrm{rms}}$ is the spectrum of the time series "integrated" up to the perturbation frequency f_n. As the signal is discrete, the "integration" is carried out by cumulating the coefficients of the single-sided discrete Fourier transform $\mathcal{F}[\Delta\varphi(t)]_i$ up to the perturbation frequency of index n. The integrated spectrum is

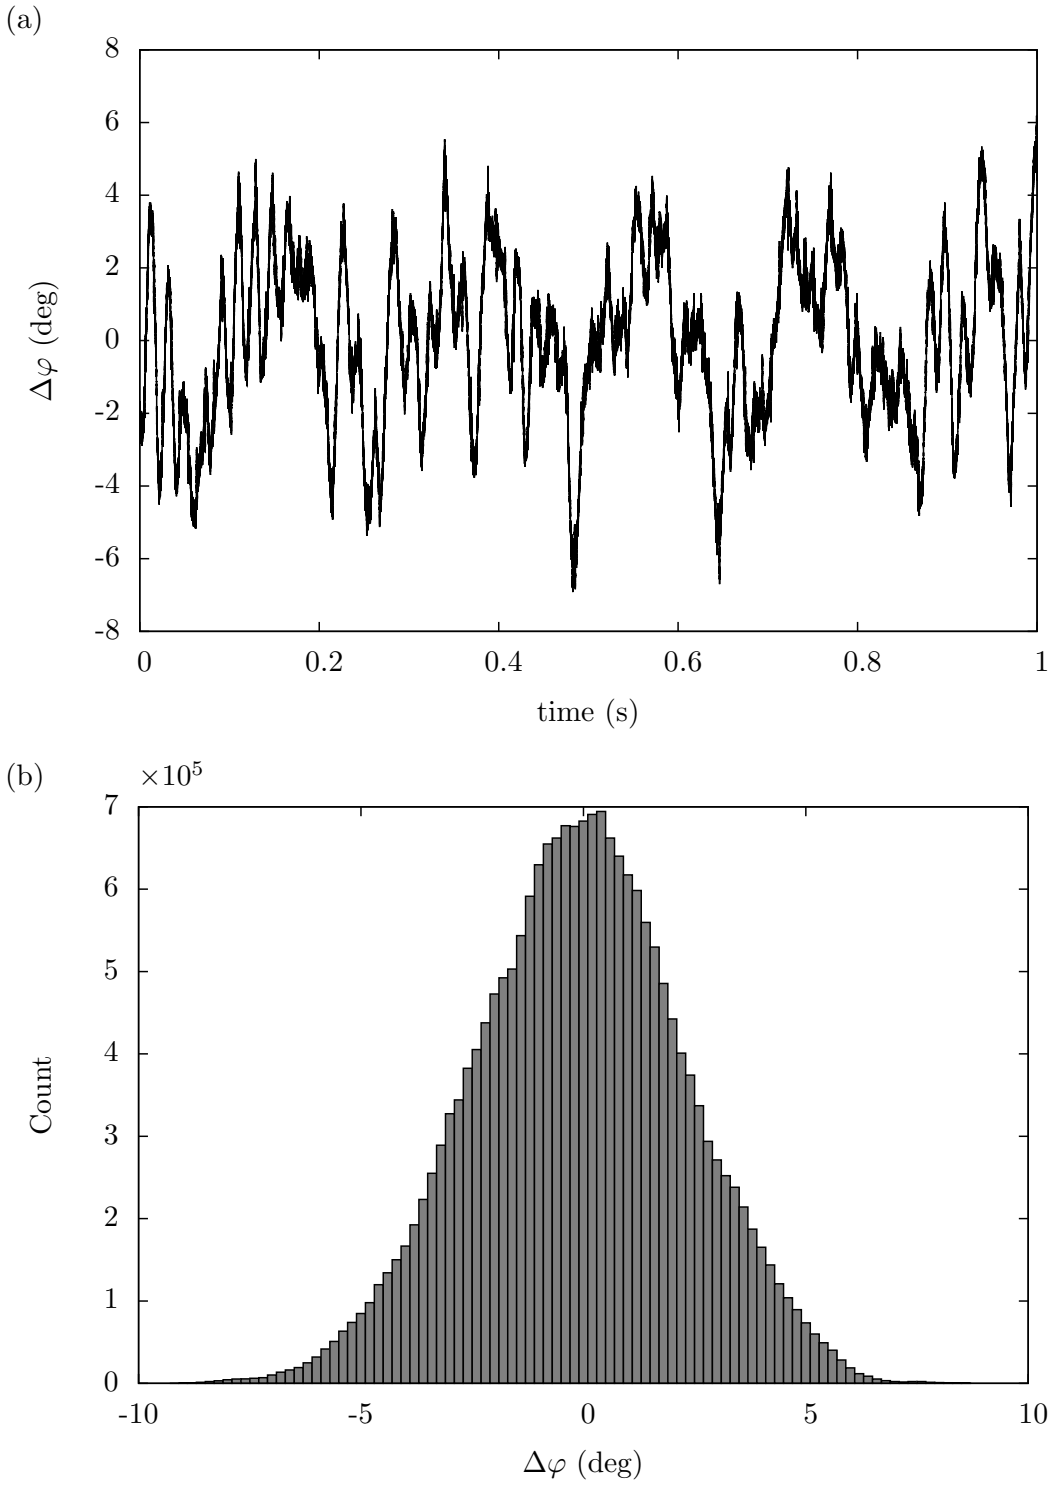

Figure 7.1: Phase error of a freely oscillating superconducting cavity in SEL mode. Sinusoidal structures of several frequencies are superposed resulting in a nearly Gaussian distribution of the phase error.

given in fig. 7.3a. In contrast to the normal spectrum this representation makes it much easier to see how the total rms error is distributed over different frequencies, especially if the spectrum contains many very sharp peaks.

7.1.1 Generator-driven resonator

The performance of the GDR algorithm used for the normal-conducting cavities has been measured with the S-DALINAC's copper resonator that is used for beam chopping (`I0NC01`). The measurements have been performed with the same parameters that are used for beam operation.

Integrated spectra of the residual errors are given in fig. 7.2. Phase and amplitude feedback have both been activated while recording the data. The total phase error of 0.016° rms is very close to the phase noise of the controller hardware itself and is much better than our requirements. The increase of the integrated spectrum at frequencies above 10 kHz is caused by noise. The plateau above 100 kHz results from the anti-aliasing filter that damps higher frequencies.

The relative amplitude error of 1.6×10^{-4} rms is higher than the results achieved with superconducting cavities because of the much bigger bandwidth of the normal-conducting cavity that results in more transmitted noise. Although the amplitude error is larger than with superconducting cavities it still meets the specification.

Results achieved with the 6 GHz control loop are similar to the performance at 3 GHz [20]. Table 7.1 gives an overview of the residual errors obtained at both operating frequencies.

Cavity	Control algorithm	Operating frequency	$(\Delta\varphi)_{\mathrm{rms}}$	$(\Delta A/A)_{\mathrm{rms}}$
`I0NC01`	GDR	3 GHz	0.016°	1.6×10^{-4}
`I0NC03` [20]	GDR	6 GHz	0.012°	3.7×10^{-4}
`I1SC03`	SEL	3 GHz	0.78°	7.2×10^{-5}

Table 7.1: Residual phase and amplitude errors accomplished with both phase and amplitude feedback loops closed.

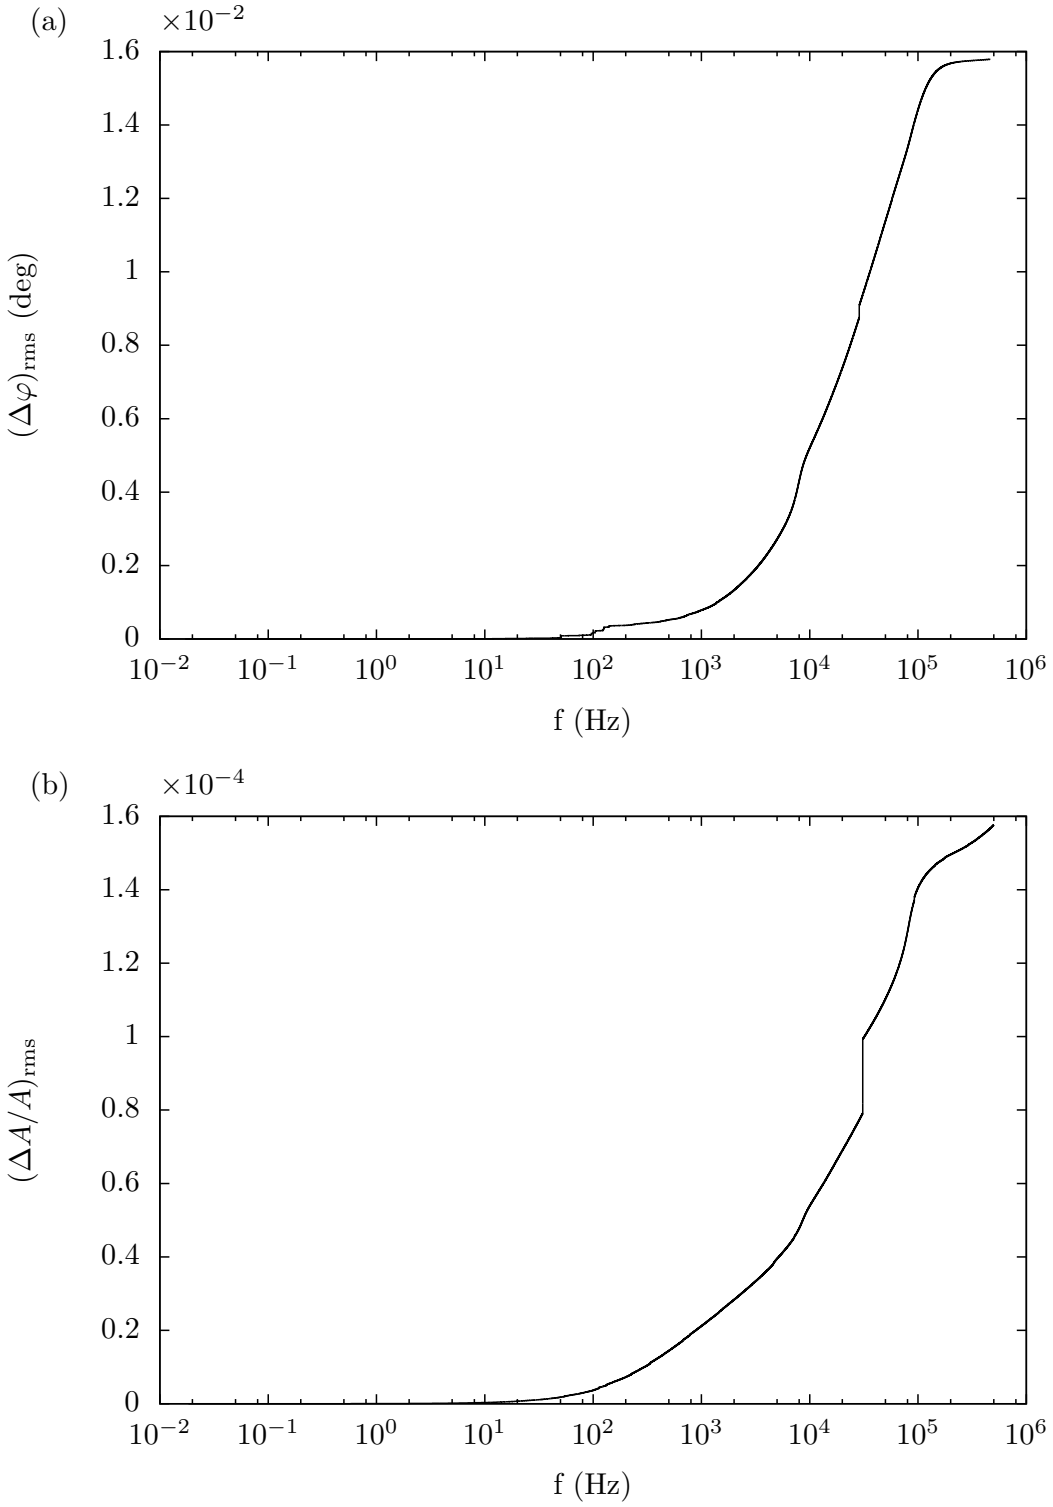

Figure 7.2: Integrated spectra of (a) phase error and (b) amplitude error of the GDR at the chopper cavity `I0NC01`. Amplitude and phase feedback are both active.

7.1.2 Self-excited loop

The performance of the SEL algorithm used for the superconducting cavities has been measured operating the 20-cell cavity `I1SC03` in the cryo-module of the S-DALINAC at a typical field strength of 4 to 5 MV/m.

Figure 7.3a shows the integrated phase error during SEL operation while all controllers are switched off. The total magnitude of the phase fluctuations is 2.4 ° rms. By turning on the phase controller these fluctuations can be reduced to a value below 0.3 ° rms [21].

If both phase and amplitude controllers are active it becomes more difficult to reduce phase as well as amplitude fluctuations. A trade-off between phase and amplitude errors has to be made. In fig. 7.3b the parameters of the controllers have been chosen to damp the phase fluctuations to a value of roughly 0.7 ° rms which is the target specification. Figure 7.4b shows the corresponding amplitude error that has been recorded at the same time. Compared to fig. 7.4a showing the amplitude errors without amplitude control, the amplitude controller reduces the amplitude fluctuations significantly to $\Delta A/A = 7.2 \times 10^{-5}$ rms.

The source of the noticeable high amplitude error in fig. 7.4a at approximately 10 kHz is still unknown. The perturbations only appear if the phase controller is active (with the same parameters that were optimized for small residual errors with both controllers enabled), while the amplitude controller is switched off.

The presented measurements show that the controllers accomplish to stabilize the amplitude as well as the phase of the rf input signal of the controller. This applies to the normal-conducting cavities used in generator-driven mode as well as to the superconducting cavities operated in a self-excited loop.

Assuming 100% correlated amplitude and phase errors results in an error of $(\Delta E/E)_{\mathrm{max}} = 1.6 \times 10^{-4}$ (cf. eq. (1.2)) whereas the error is $(\Delta E/E)_{\mathrm{rms}} = 1.2 \times 10^{-4}$ in the case of uncorrelated errors. The correlation between amplitude and phase errors depends on the control parameters. In closed-loop operation the errors are dominated by noise and can be expected to be only weakly correlated. This leads to a Gaussian distribution of the amplitude error (see fig. 7.5) and makes the rms error a more realistic estimate for the actual error. In terms of the measurement accuracy the rms error is consistent with the specification of $(\Delta E/E)_{\mathrm{rms}} = 1.0 \times 10^{-4}$. With

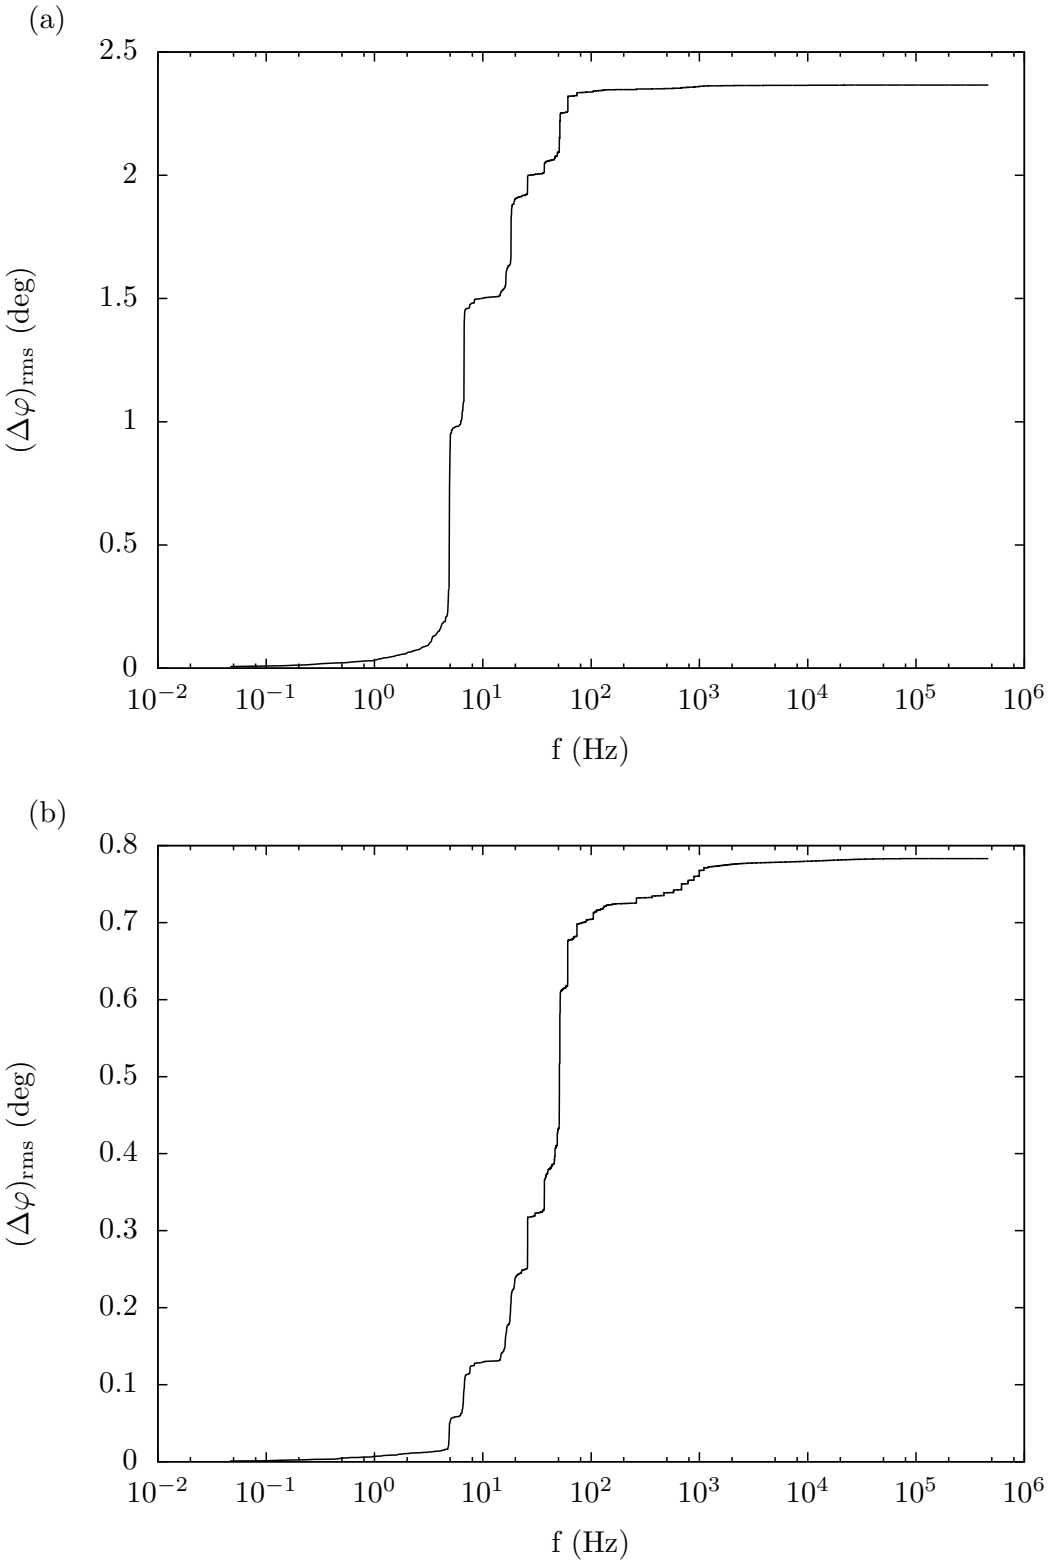

Figure 7.3: Integrated magnitude spectra of phase error of the SEL with (a) amplitude and phase controllers deactivated and (b) both feedback loops closed.

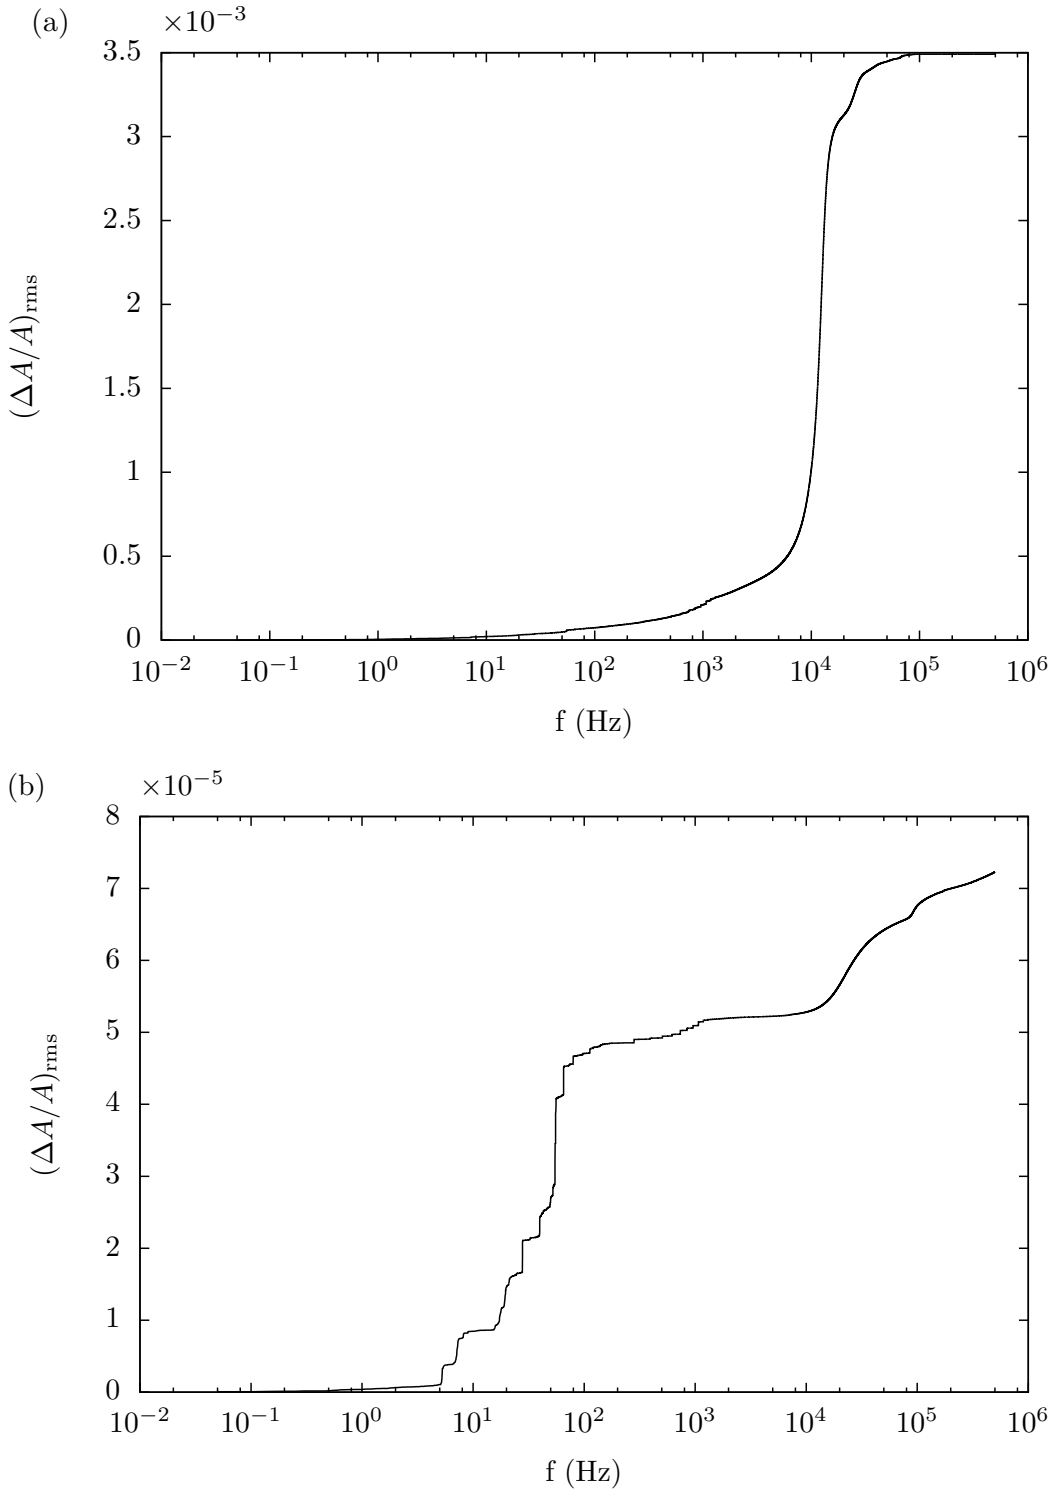

Figure 7.4: Integrated magnitude spectra of relative amplitude error of the SEL with (a) phase controller activated but without amplitude feedback and (b) with both feedback loops closed. Parameters have been optimized for small residual errors with both controllers enabled.

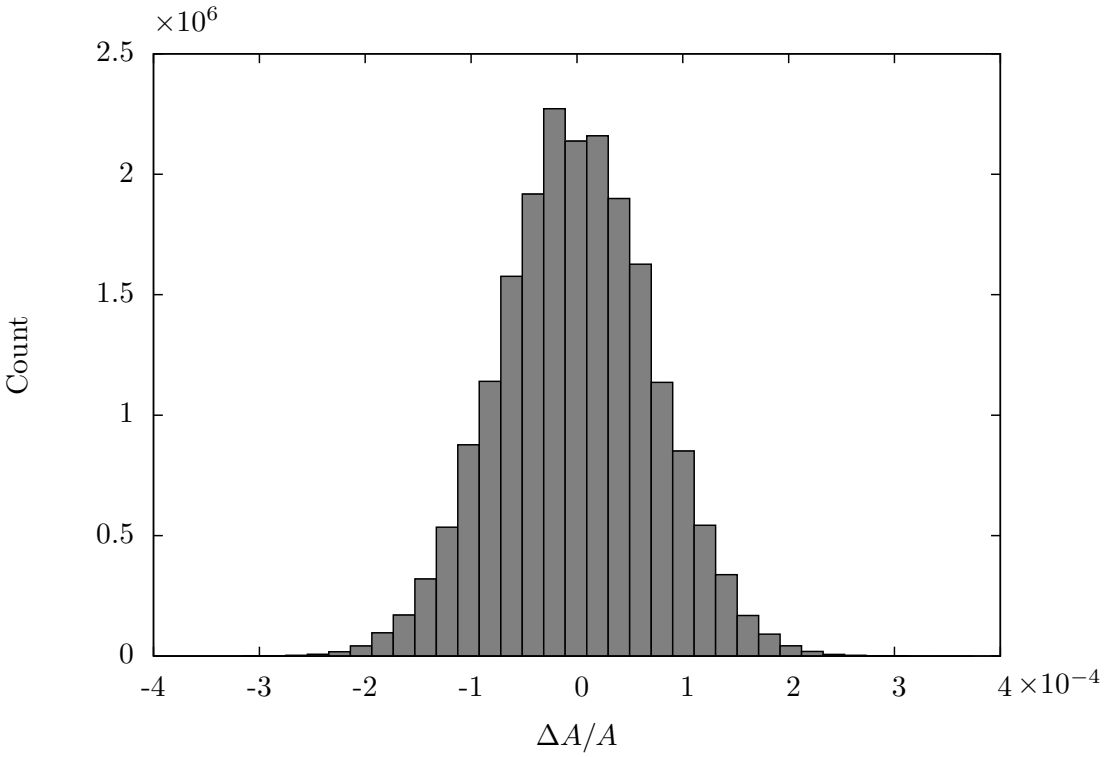

Figure 7.5: Histogram of the amplitude error signal with all controllers active. The signal is kept very close to zero by the amplitude controller. Each bin corresponds to a single step of the ADC.

the new revision of the rf boards the errors can be expected to be slightly better because of their higher linearity.

The residual errors achieved in this work are an improvement of more than a factor of 2 compared to the prototype presented in [41] (see table 7.2). The results are directly comparable since the measurements have been performed using the same cavity. The errors obtained by [41] are only based on 16 bit data while the full 18 bit resolution has been used to calculate the values determined here. By omitting the 2 least significant bits the actual error of the prototype system is slightly underestimated[3]. Compared to the analog rf control system the performance has been improved by more than one order of magnitude [19].

Compared to the results obtained in the context of the CEBAF [63] 12 GeV upgrade project the residual errors achieved in this work are inferior by a factor

[3]The fluctuations that are measured in closed-loop operation are very small and use only a few steps of the ADC (cf. fig. 7.5). With a resolution of 16 bits the number of bins is decreased by a factor of 4 compared to full 18 bit resolution.

Controller	$(\Delta A/A)_{\mathrm{rms}}$	$(\Delta\varphi)_{\mathrm{rms}}$ (deg)	$(\Delta E/E)_{\mathrm{rms}}$
Analog rf control system [19]	2×10^{-3}	0.38	2.5×10^{-3}
Prototype of digital system [41]	2.5×10^{-4}	0.28	2.5×10^{-4}
This work	7.2×10^{-5}	0.78	1.2×10^{-4}
CEBAF 12 GeV upgrade [62]	4.5×10^{-5}	0.11	4.5×10^{-5}

Table 7.2: Performance figures of different rf control systems at the S-DALINAC and other accelerators. All listed control systems are using self-excited loops. The last column contains the rms error in energy calculated from amplitude and phase error under the assumption that these errors are uncorrelated.

of 2.7. Because of the higher number of cells per cavity (20 instead of 7), which makes the cavities more flexible, and the higher operating frequency (3 GHz instead of 1.497 GHz), the S-DALINAC's rf control system has to deal with much more microphonics than at CEBAF (2.4° rms compared to 0.75° rms). Taking this into account the performance of the S-DALINAC's rf control system is roughly comparable to the new CEBAF rf control system.

7.2 Measurements on the electron beam

The measurements presented above are solely based on the error signals calculated by the control algorithm in the FPGA. The described techniques for off-line and on-line diagnosis provide a very direct way to monitor the performance of the control loops. But before using the errors calculated by the rf control system for optimization of the control parameters it is important to ensure that minimizing these values actually leads to a minimal energy spread of the beam. Therefore the energy distribution of the beam has been measured at the S-DALINAC's 169° high-resolution electron scattering facility [64, 65].

It is not possible to measure the energy distribution of the electron beam directly using this spectrometer. Instead electrons that have been elastically scattered at a ^{197}Au target have been measured. The spectrometer detected electrons at a scattering angle of $\Theta = 93°$ in transmission geometry (see fig. 7.6).

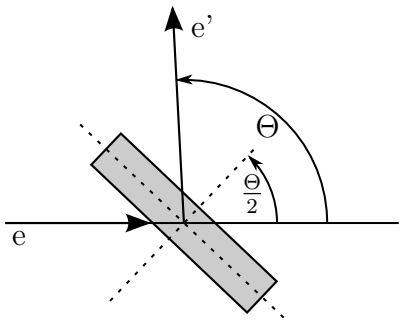

Figure 7.6: Elastically scattered electrons are detected by the spectrometer at a scattering angle of $\Theta = 93°$. The target is mounted in transmission geometry at an angle of $\Theta/2$.

^{197}Au has been chosen as target material because it facilitates the fabrication of thin targets. This keeps broadening effects of the elastic peak caused by multiple scattering inside the target small. The area density of the target used for the measurements is $t = 1\,\mathrm{mg/cm^2}$. The effective thickness of the target is defined by the geometry (cf. fig. 7.6) as

$$t_{\mathrm{eff}} = \frac{t}{\cos\frac{\Theta}{2}} = 1.45\,\mathrm{mg/cm^2} \quad . \tag{7.2}$$

A GEANT4 [66] simulation has been performed to quantify the broadening of the elastic line caused by the target. In this simulation a mono-energetic electron beam with the same energy as used during the measurements ($E_b = 29.5\,\mathrm{MeV}$) has been sent through a ^{197}Au target with the same effective thickness as the target used for the measurements. Figure 7.7 shows the energy distribution of the electrons transmitted through the target. The full width at half maximum of the electron distribution is $(\Delta E)_{\mathrm{target}} = 1.0\,\mathrm{keV}$. Thus the relative error introduced by the target is

$$\left(\frac{\Delta E}{E}\right)_{\mathrm{target}} = 3.4 \times 10^{-5} \quad . \tag{7.3}$$

This error has to be taken into account during the analysis of the results. Note that this error only describes the energy loss in the target. It does not comprise the divergence of the beam caused by the target.

All measurements presented in this context have been performed in on-crest acceleration mode. The alternative non-isochronous recirculation mode using phase focusing [2, 16] cannot be used for the measurements since it reduces the

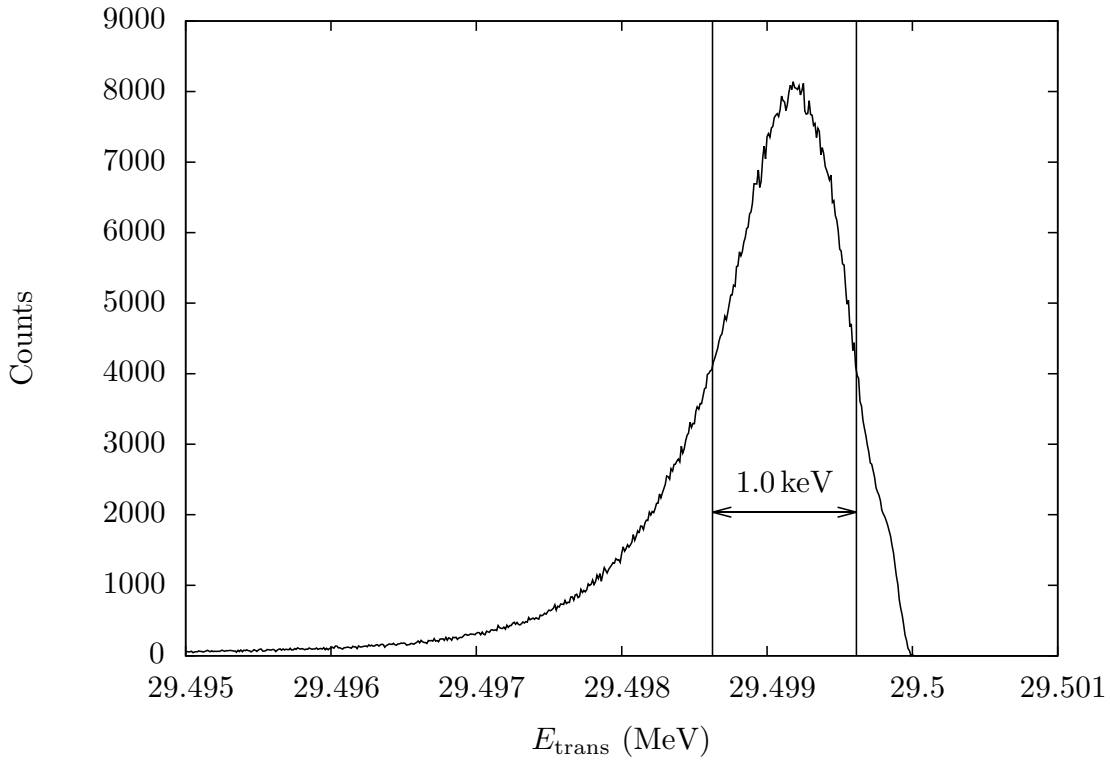

Figure 7.7: Simulated energy distribution of a mono-energetic electron beam passing through a ^{197}Au foil with the same effective thickness as the target used for the analysis of the electron beam. Electron–electron scattering and bremsstrahlung broaden the line to a width of 1.0 keV.

energy spread that is to be determined. In on-crest acceleration mode on the other hand the full energy spread caused by the cavities can be measured at the experimental area. Measurements in the context of preparation for non-isochronous beam recirculation revealed that the adjustment range for the path length of the first recirculation beam-line of the S-DALINAC is too small to tune it to its optimal working point [16]. The resulting phase mismatch between the bunches and the rf field can lead to an increased energy spread. To avoid these effects all measurements have been performed in single-pass operation of the accelerator.

The impact of the injector cavities on the energy spread of the beam is complex because at lower electron energies fluctuations of a cavity's voltage gradient can result in relevant speed changes of the bunches. This leads to a phase shift relative to the next cavity and thereby to a reduced energy gain in consecutive cavities. In the main linac on the other hand the speed of the electrons is nearly constant and

the behavior of each cavity can thus be studied independently. That is why only cavities in the main linac have been taken into account for detailed measurements.

All measurements have been conducted at a beam energy of $E_b \approx 29.5\,\mathrm{MeV}$. To make sure that the energy spread is not reduced unintentionally by cutting off part of the beam in the beam line the main linac as well as the extraction beam-line have been carefully optimized for maximum transmission. Within the precision of the beam current measurement and according to the beam-loss monitoring system no beam losses have been detected in this area during the measurements.

To simplify the measurements, the rf control loops have been optimized for a small phase error. Table 7.3 lists the average amplitude and phase errors of all cavities during a typical run. The data implies that energy fluctuations due to the phase error are negligible compared to energy fluctuations caused by the amplitude error. Since the amplitude error follows a Gaussian distribution (cf. fig. 7.5), the beam energy can also be expected to be normally distributed.

Figure 7.8 shows a typical spectrum of the electron line resulting from elastic scattering off the ^{197}Au target and into the spectrometer. The peak shows a distinct tail at the lower energy side. This tail is caused by energy losses inside the target due to bremsstrahlung and electron–electron scattering. To extract the width of the peak it has been fitted with a function that contains a phenomenological description of these effects [67]:

$$y(x) = y_0 \times \begin{cases} \frac{A}{(x-B)^\gamma} & \text{if } x \leq x_0 \\ \exp\left(-\frac{(x-x_1)^2}{\sigma_l^2}\right) & \text{if } x_0 < x \leq x_1 \\ \exp\left(-\frac{(x-x_1)^2}{\sigma_r^2}\right) & \text{if } x_1 < x \end{cases} \tag{7.4}$$

The peak itself is described by a split Gaussian function whereas a hyperbolic function takes into account the radiative tail. Connecting the three function parts to a continuously differentiable function eliminates two of the three parameters A, B, and γ. The remaining parameter has been fitted to the data together with the amplitude of the peak y_0, the energy of the peak x_1, the starting point of the radiative tail x_0 and the standard deviations σ_l and σ_r of the two Gaussian functions. The arithmetic mean of both standard deviations is used as an estimate

Cavity name	$\left(\frac{\Delta A}{A}\right)_{\mathrm{rms}}\,(10^{-4})$	$\Delta\varphi_{\mathrm{rms}}$ (deg)
I0NC01	3.21	0.03
I0NC02	1.83	0.04
I1SC01	7.88	0.17
I1SC02	3.47	0.19
I1SC03	2.27	0.14
I1SC04	2.43	0.36
A1SC01	2.48	0.12
A1SC02	4.88	0.13
A1SC03	3.21	0.12
A1SC04	3.44	0.15
A1SC05	1.00	0.16
A1SC06	1.63	0.12
A1SC07	N/A	N/A
A1SC08	3.90	0.15

Table 7.3: Average errors of the cavities during a typical run. Cavity A1SC07 has been switched off while data has been taken. The shown data corresponds to the leftmost data point in fig. 7.9b.

for the total standard deviation of the peak:

$$\sigma := \frac{\sigma_l + \sigma_r}{2} \tag{7.5}$$

Unfortunately most of the fits result in the radiative tail including the most part of the left side of the peak (cf. fig. 7.8). This results in high uncertainties in the width of the left Gaussian part of the fit function.

The energy calibration has been done using the momentum acceptance of the spectrometer $\delta_E = 4.04\%$ [65], the beam energy E_b, and the total number of

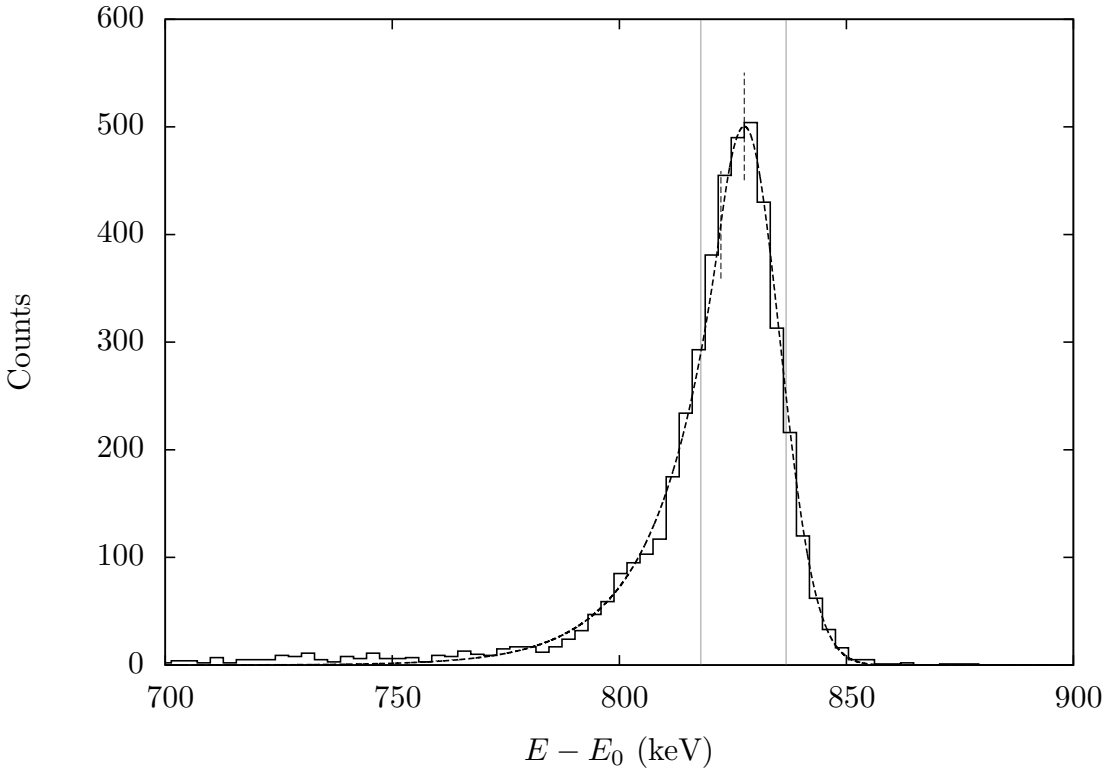

Figure 7.8: Spectrum measured at the 169° high-resolution electron scattering facility. Note that in contrast to the convention used in nuclear physics the energy of the detected electrons increases to the right. The three parts of the function fitted to the elastic peak are joined at the positions of the dashed markers. The straight gray lines indicate the width of the peak obtained by the fit (half width at half maximum of the fitted Gaussian functions).

channels[4] $N_{\mathrm{ch}} = 415.5$:

$$E(c) = \frac{\delta_E \, E_b}{N_{\mathrm{ch}}} \, c + E_0 \tag{7.6}$$

This equation maps the channel number c onto the energy $E(c)$. Note that the energy offset E_0 cannot be determined exactly because the beam energy can vary slightly between runs[5]. Thus the energy axis of fig. 7.8 has been calibrated relative to the edge of the detector plane.

[4]Each of the spectrometer's 384 Si microstrip detectors covers 1/415.5 of the focal length of the spectrometer. The detectors are clustered to four groups that are separated by gaps. The magnet current has been adjusted to move the elastic peak into the center of one detector group.

[5]During the measurement of the last data point of fig. 7.9a a reduced transmission of only 95% has been achieved. This means that the outer 5% of the beam are lost in the beam-line. During the other runs no beam losses have been detected. This means that variations in beam energy have always been significantly smaller than 22 keV.

Error caused by	Value	Type of error
Bunch length (cf. eq. (1.5))	4.5×10^{-5}	rms
Multiple scattering in target (cf. eq. (7.3))	3.4×10^{-5}	fwhm
Beam spot size [64]	1.2×10^{-4}	fwhm
Intrinsic resolution of spectrometer [64]	1.2×10^{-4}	fwhm
Resolution of detector plane [65]	1.0×10^{-4}	fwhm

Table 7.4: Quantitatively known errors that lead to a broadening of the elastic line measured with the spectrometer.

Figure 7.9 shows the dependence of the total standard deviation of the elastic peak on the rms error of a single cavity. To measure this behavior, the gain of the amplitude controller for the single cavity has been reduced successively resulting in an increased rms error. The plot clearly shows that the energy spread of the beam increases as the relative amplitude error of the cavity raises. For relative amplitude errors larger than 1×10^{-3} the energy spread of the beam increases nearly proportional to the relative amplitude error. For small amplitude errors, however, the energy spread of the beam is nearly constant. Hence the measured energy spread of the beam is dominated by the errors of other cavities, by beam dynamic effects like capture effects in injector cavities or by the energy resolution of the spectrometer. Instabilities of the electron gun can also cause a higher energy spread of the beam. The quantitatively known errors are listed in table 7.4.

Since the errors introduced by different cavities σ_i are nearly independent, they contribute to the total error according to

$$\sigma_{\mathrm{tot}} = \sqrt{\sum_{i}^{N} \sigma_i^2} \quad . \tag{7.7}$$

The errors listed in table 7.4 can be taken into account the same way. Equation (7.7) has been fitted to the data for the special case of $N = 2$. This means to distinguish the error introduced by the cavity being varied from all other errors induced by the remaining cavities as well as the other sources mentioned above. Figure 7.9 shows that the data is consistent with eq. (7.7). The resulting total values from

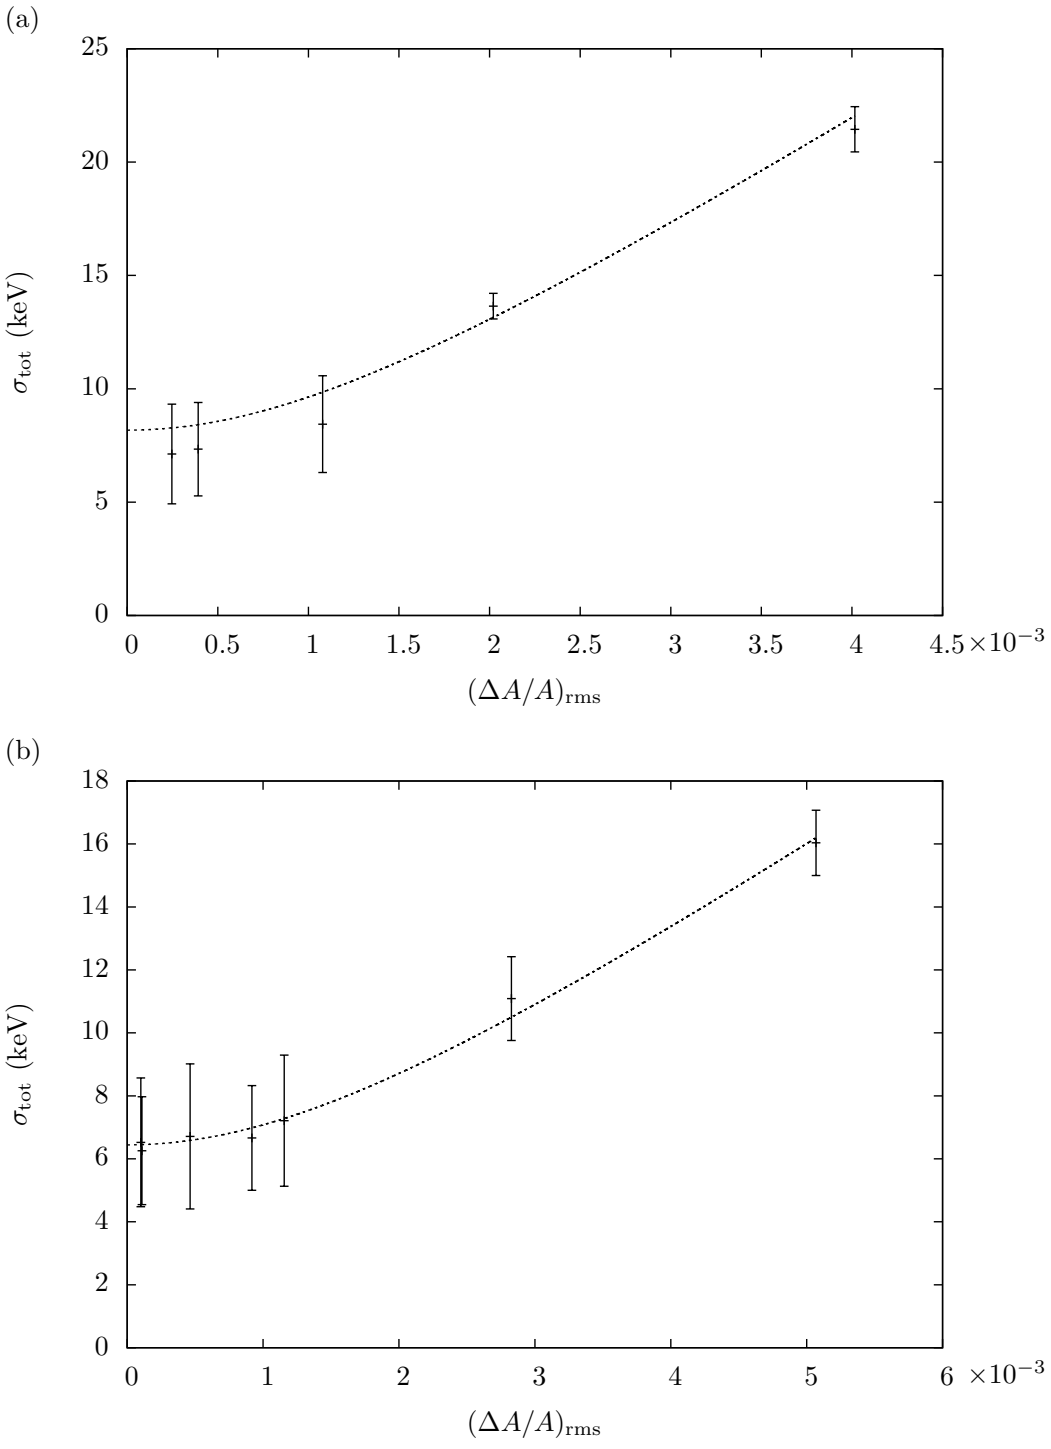

Figure 7.9: Total standard deviation of the elastic peak measured at the experimental area vs. the amplitude error calculated by the rf control system for cavity (a) `A1SC01` and (b) `A1SC05`, respectively. During the measurement of the last data point of (a) about 5% of the beam were lost in the beam-line resulting in a slightly too low value for the line width σ_{tot}.

the fits for $(\Delta A/A) = 0$ are:

$$\sigma_{\mathrm{tot}} = 8.18\,\mathrm{keV} \qquad \text{for } \texttt{A1SC01} \tag{7.8}$$

$$\sigma_{\mathrm{tot}} = 6.44\,\mathrm{keV} \qquad \text{for } \texttt{A1SC05} \tag{7.9}$$

Subtracting the known errors listed in table 7.4 according to eq. (7.7) yields:

$$\sigma_{\mathrm{res}} = 7.67\,\mathrm{keV} \qquad \left(\frac{\Delta E}{E}\right)_{\mathrm{res,rms}} = 2.6 \times 10^{-4} \qquad \text{for } \texttt{A1SC01} \tag{7.10}$$

$$\sigma_{\mathrm{res}} = 5.79\,\mathrm{keV} \qquad \left(\frac{\Delta E}{E}\right)_{\mathrm{res,rms}} = 2.0 \times 10^{-4} \qquad \text{for } \texttt{A1SC05} \tag{7.11}$$

Since these residual errors also comprise other errors like instabilities of the electron gun and beam dynamic effects during the capture process in the injector they can be considered as an upper limit for the error introduced by the rf control system.

The measurements show that the rms errors of the cavities provided by the rf control system are an appropriate tool not only to monitor the performance of the controllers but also as a rough estimate for the energy spread of the beam.

8 Summary and outlook

Based on the work of [20, 39, 40, 41] a new digital rf control system has been developed and commissioned at the S-DALINAC. The new system supports operation of normal-conducting and superconducting cavities at an operating frequency of 3 GHz as well as at the first harmonic frequency of 6 GHz. A generator-driven resonator control algorithm is used for the normal-conducting cavities whereas the high-Q superconducting cavities are operated in a self-excited loop. The self-excited loop algorithm uses a microphonics compensator to compensate for amplitude and phase fluctuations caused by detuning of the cavity at the same time. A soft DSP has been implemented in the FPGA which allows fast and flexible modification of the control algorithm as well as fast read-out. Extensive on-line diagnostics have been implemented based on this read-out. This includes the possibility of analyzing signals from the FPGA with a software oscilloscope as well as a with a spectrum analysis software. Additionally, the system continuously calculates on-line rms errors that can be used by the operator to optimize the control parameters and to monitor the performance of the cavities. Measurements of the electron beam have proven that these rms errors can be used to rate the performance of the controller and the cavity and thus to estimate the energy spread of the beam.

Also, measurements have shown that the rf control system yields a phase stabilization of $(\Delta\varphi)_{\mathrm{rms}} = 0.8°$ and an amplitude stabilization of $(\Delta A/A)_{\mathrm{rms}} = 7 \times 10^{-5}$. This is an improvement of one order of magnitude compared to the previous analog rf control system. Compared to the prototype described in [41] not only the performance but also reliability and usability of the system have been improved significantly. The presented results demonstrate that the baseband approach represents a viable option for rf control systems even in the case of high demands on accuracy.

The complete system has been commissioned and successfully used for two years. Stability and reproducibility have been drastically improved compared to the old analog system. After being the primary source of accelerator breakdowns, the rf control system has become one of the most reliable subsystems of the S-DALINAC. Currently the S-DALINAC's rf control system is used as a basis for the development of a pulsed rf control system for the p-Linac test facility at

the Facility for Antiproton and Ion Research (FAIR, [68]). The system is also considered for an upgrade of the rf control system of the superconducting ion accelerator ALPI at Laboratori Nazionali di Legnaro (INFN LNL).

Part II
Accelerator control system

9 Previous accelerator control system

During beam operation hundreds of devices like magnet power supplies, diagnostic elements etc. need to be controlled at the S-DALINAC. Due to radiation safety regulations and the size of the facility most of these distributed devices have to be remote-controlled. The accelerator control system allows the operators to control and monitor all of these components from a central control room (see fig. 9.1).

Since its commissioning in 1987 the S-DALINAC has been controlled by an in-house developed accelerator control system. Over the years this system has been enhanced and updated several times. The latest innovations include implementation of a client/server infrastructure based on the TCP/IP protocol, XML configuration for the control servers, and support for asynchronous communication to devices over serial interfaces [26]. In addition new devices have been developed in-house to replace old and unreliable ones. These include controller hardware for beam diagnostic targets [69], a high precision data acquisition system [41], hundreds of power supplies, and finally the new digital rf control system described in part I. All these devices are controlled by micro-controllers that are connected to control system PCs via a field bus (Controller Area Network, CAN, [40]).

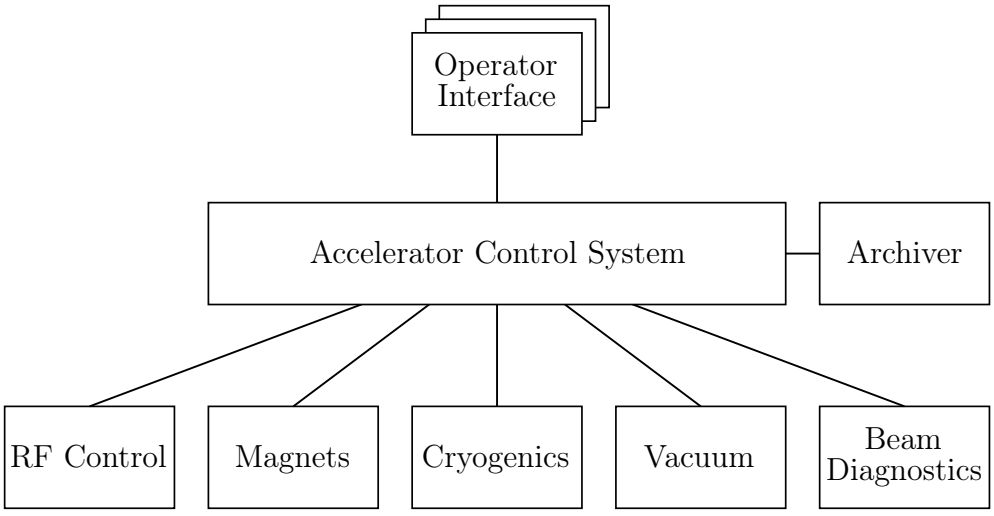

Figure 9.1: Basic architecture of the accelerator control system of the S-DALINAC. The devices of all major accelerator subsystems are connected to the control system and can be operated remotely from an operator console. The archiver stores data from the control system for later analysis.

With the previous accelerator control system, implementation of support for new hardware needed high development effort by experienced C and C++ programmers. Due to a lack of programmers only the most important functions could be realized. Many diagnostic features were still missing. In many cases the lack of diagnostic functions made operation and maintenance of the machine much more difficult and time consuming. Moreover, the asynchronous communication via serial ports did not reach the desired stability.

Although the control system featured a relational database for central configuration management [70], this feature was only used for small parts of the control system. In particular it was not used to configure the control system clients and servers that had to deal with hundreds of magnet power supplies. The relational database was also used to archive process data of the helium liquefier [70]. But since this data was collected directly from the hardware by the archiving software, it could not be monitored on-line via the usual accelerator control system interface.

The data acquisition system, used for measuring the beam current, and the thermionic gun were controlled by their own software that was incompatible with the rest of the accelerator control system [41, 71].

Overall the previous accelerator control system was too inflexible for the digital rf control system with its re-programmable FPGAs. Changes to the control algorithm would always require programming work on the control system side. Moreover development and maintenance of the control system for a hardware providing several thousand channels would have exceeded the available resources. Thus a control system allowing a faster implementation of new devices was needed.

The described shortcomings of the previous control system led to the evaluation of EPICS as a basis for the digital rf control system. At first it was intended to be used only for the prototype of the rf control system. However, due to the much faster implementation and the very high stability of this prototype control system it was decided to migrate the rest of the accelerator control system to EPICS as well. This migration process will be described in detail in the following. It has been started in the context of this thesis and is continued by [72] and [73].

10 EPICS

The Experimental Physics and Industrial Control System (EPICS, [74, 75, 76]) is a framework used to build control systems. It is designed for distributed control systems that comprise a large number of computers connected via a network. Like the previous control system used at the S-DALINAC the EPICS architecture is client/server-based. Control server software (also called input output controller, IOC) performs input and output operations to hardware and provides access to device parameters via a network protocol called Channel Access (CA, see fig. 10.1). Clients connect to the IOCs via the Channel Access protocol to access device data in form of process variables (PVs) published by the IOCs. The only information necessary to connect to a PV is its name. In

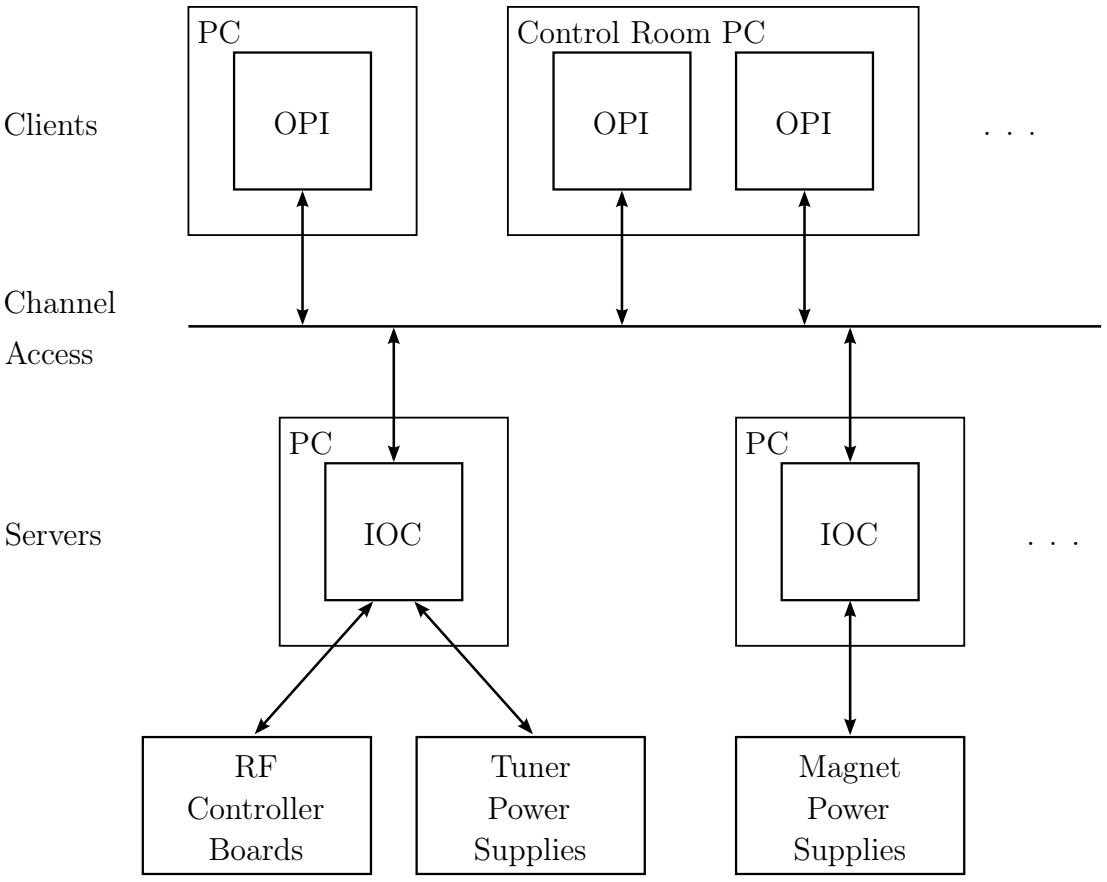

Figure 10.1: The EPICS architecture is client/server-based. Input/output controllers (IOCs) act as servers that publish data from the hardware. Operator interface (OPI) software acts as client that connects via the network to the IOCs using the Channel Access protocol.

Record name	Data type	I/O direction
`ai`	float	input
`ao`	float	output
`bi`	binary	input
`bo`	binary	output
`calc`	depending on input records	input
`calcout`	depending on input records	output
`longin`	32 bit integer	input
`longout`	32 bit integer	output
`mbbi`	enumeration	input
`mbbo`	enumeration	output
`waveform`	array of configurable type	input

Table 10.1: The most important EPICS record types used at the S-DALINAC.

contrast to the previous control system the clients do not need to know which server provides the data. Thus it is possible to move PVs from one IOC to another without changing the clients which makes the system more flexible.

In contrast to the previous control system, which relied mostly on the intelligence of its clients, EPICS features much more server side functionality. EPICS IOCs e. g. convert bit patterns and numbers read from the hardware into meaningful units of measurement (e. g. ampere, tesla etc.). Thus replacing a broken device with a new one providing the same functionality but with a different interface (e. g. from a different vendor) often requires only changes to the control server but does not affect the clients at all.

IOCs store the data that is connected with each PV in a structure called a record. All records on the network are held in a distributed real-time database[6]. The EPICS database supports different types of records. The most important record types in the context of the S-DALINAC control system are listed in table 10.1. Input records typically read data from the hardware while output records are used

[6]The EPICS database is very different from most other databases and should not be confused with the relational database used to store configuration and archive data.

to send data to the hardware. Reading or writing occurs during record processing. Each record can be configured to process on a regular basis (e. g. once a second) or when it is triggered by another record. If the input value is outside predefined limits, some records like `ai` can switch to an alarm state. This information can be used by clients to mark the corresponding widget on the graphical user interface to draw the operator's attention to the problem. Similar functionality can also be used to signal that a reading is invalid and cannot be trusted.

The value held by a record can be passed on to another record. Similarly processing of a record can trigger processing of other records. In addition to records that communicate with the hardware there are also records that use data from records to calculate new values (e. g. `calc` and `calcout` record).

The described features allow to implement complex behavior like feedback loops on the IOC side with a few simple record types. No C/C++ programming skills are necessary to configure the EPICS database. Template techniques and graphical editors help to create and maintain databases with a huge number of records (see fig. 10.2, [77]).

In contrast to the previous control system which has been maintained only for Microsoft Windows®, EPICS supports many different architectures including embedded devices. The so called device support acts as a connector between the EPICS records and the hardware. For each device family a corresponding device support is necessary. EPICS already provides very stable and flexible support for asynchronous communication with serial devices. For the CAN bus and USB devices used at the S-DALINAC no adequate device support has been available in the framework.

On the client side a wide range of software and tools is available. With Control System Studio a very powerful application is available which allows an easy implementation of graphical user interfaces. Beyond that, client libraries provide access to EPICS servers for a wide variety of programming languages.

Since EPICS is open source software, special features that might be needed for the S-DALINAC in the future can be implemented. A further important advantage for university use with many student developers is the big and active community.

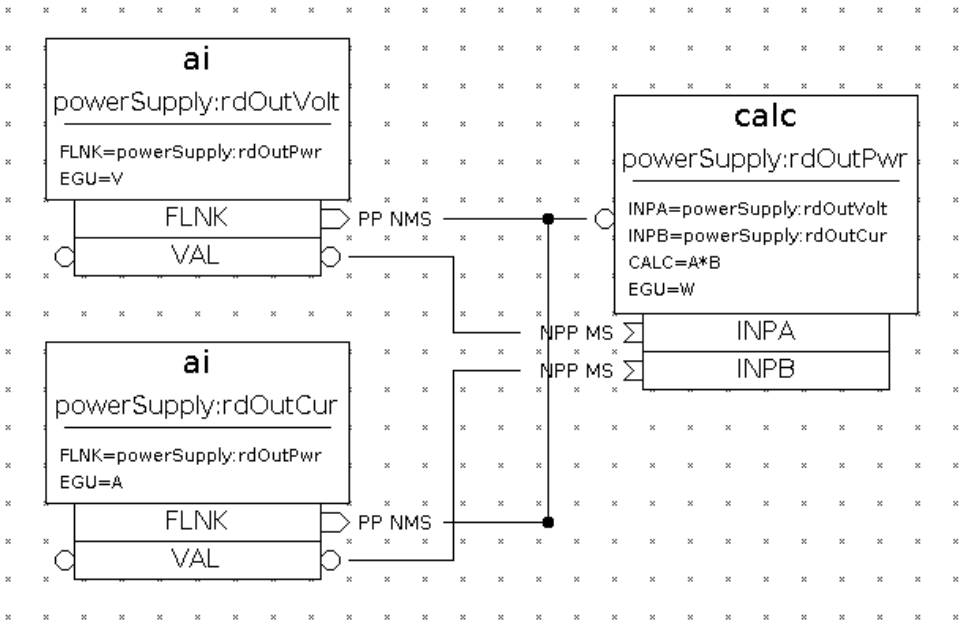

Figure 10.2: Screenshot of three EPICS records shown in the graphical editor VisualDCT. The `calc` record on the right reads the values from the records connected to its inputs and calculates the power as the product of current and voltage. Both `ai` records trigger processing of the `calc` record via a forward link.

11 Slow control

During machine optimization the operator manually adjusts parameters of the magnets or the rf system while monitoring the beam spot on a target screen. These adjustments take place on a relatively slow time scale. To provide the operator with an instantaneous feedback, the time between issuing the command in the control room and the reaction of the device should not exceed 50 ms [40]. Reply times on the accelerator network (Gigabit Ethernet) typically are in the sub-ms range and thus can be neglected. Hence an adequate response time can be obtained using CAN bus or serial ports like RS-232. To keep latency as low as possible all devices are operated at the maximum transfer rate (1 Mbit/s for the CAN devices).

11.1 CAN bus communication

The in-house developed CAN devices use their own address scheme and protocol. To use them with EPICS a custom device support has been developed in the context of this thesis. It uses the SocketCAN network stack [78] included in recent Linux kernels (since version 2.6.25) which acts as an abstraction layer making the device support independent of a specific CAN card or hardware vendor: All CAN cards having a SocketCAN driver are supported. In contrast to normal CAN drivers that only allow a single application to access the CAN bus, SocketCAN provides a network device (BSD socket) which can be accessed by multiple applications at the same time. This allows to attach debugging software like a CAN sniffer to the CAN bus while the IOC is running.

Under heavy load clients can send data faster than it can be transmitted to the devices via the CAN bus. With the previous control system this led to an excessively growing message queue and an overshoot behavior of the PV. To prevent this, the EPICS device support uses an intelligent queuing concept. It exploits the fact that if multiple CAN frames are sent to the same CAN address, only the last one is relevant because the latest data overwrites older data in the device anyway[7]. Thus the SocketCAN device support overwrites queued items in the transmit queue if more recent data is added to the queue. The updated

[7]Note that this is only the case with the high-level CAN protocol used at the S-DALINAC.

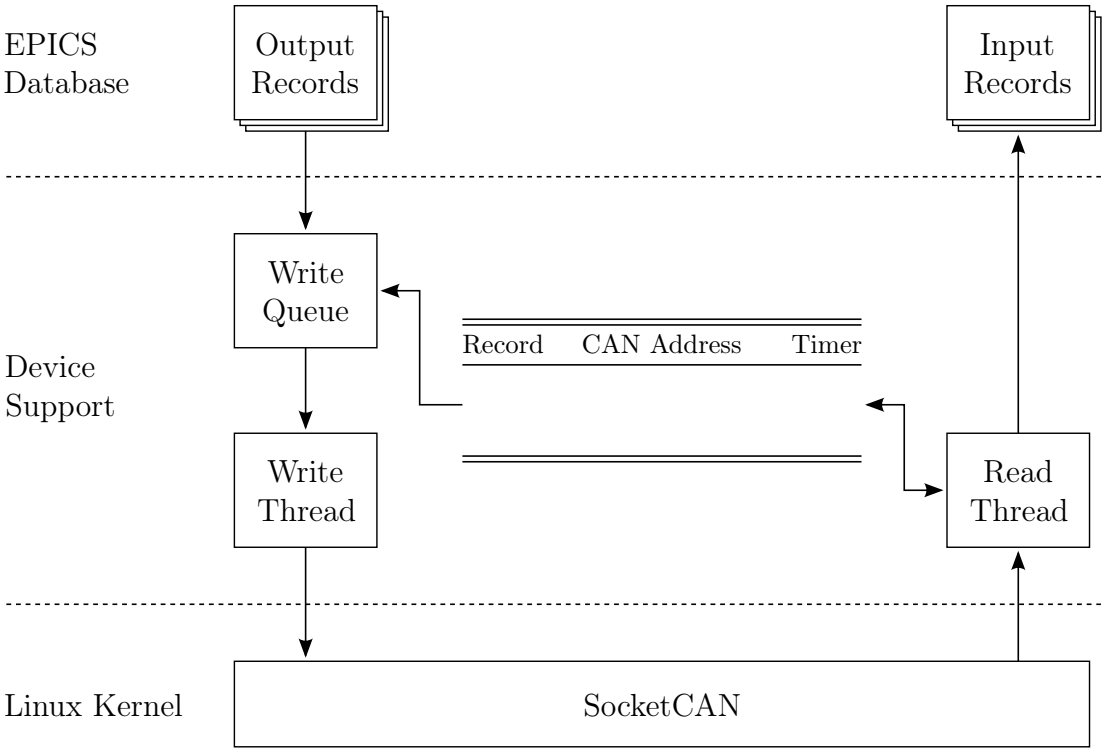

Figure 11.1: Data flow through the SocketCAN device support. Separate threads are used to read data from the CAN bus and to process the write queue. The data from the record table is used to send the output data to the correct CAN address and to map incoming data to the right records. If a predefined time elapses without an incoming CAN frame resetting the timer the corresponding records are switched to a timeout state.

item still remains at the same position in the queue resulting in a fair distribution of bus bandwidth between different CAN commands. This queuing algorithm throttles the rate of very fast sending records where required while still using the highest possible rate the bus provides to keep the latency as low as possible. This ensures an efficient utilization of CAN bus resources under all load conditions. As a consequence of this design the length of the transmit queue is limited by the number of records. Buffer overruns or growing buffer sizes are eliminated by design.

Dedicated threads process the send queue and read data from the SocketCAN device, respectively (see fig. 11.1). For each received CAN frame the read thread determines the relevant records from the list of records. It then copies the received

data to the records and initiates record processing. If a CAN frame contains data items for multiple records this can result in multiple records being processed.

The IOC reads data from the hardware by sending a request command for the data and waiting for the corresponding response from the device. To make sure a broken device is detected the power status of each device is requested regularly (e. g. once per second). An independent record processes the response from the device. The separation of sending requests and processing the response allows to process unsolicited status reports sent by the device. An example for such an unsolicited status report is a status message that is sent when the rf controller board is switched off because of an interlock failure. If a receive record does not receive data for a predefined time, the corresponding record is switched to an error state (timeout).

The CAN address scheme also provides broadcasts which allow to address all devices in a crate or all devices on a CAN segment (usually a rack). This feature is used to save bandwidth when sending status requests to all devices. Furthermore some devices also provide an auto-transmit feature that periodically sends the data without receiving a request. This feature is also used where available to save bandwidth.

11.2 Commercial devices

Most commercial devices used at the S-DALINAC are equipped with serial ports according to the RS-232, RS-422 or RS-485 standard. These devices are connected to the accelerator network via device servers that make the serial port accessible via network. Thus a single IOC can control all commercial devices with serial ports via the network. This IOC uses the EPICS StreamDevice [79, 80] device support which allows communication with hardware that is controlled by sending and receiving character strings.

In contrast to the device driver implementation of the previous control system, StreamDevice supports fully asynchronous communication with timeouts and error handling. Thanks to these features the IOC can deal with unexpected input data and cannot get stuck waiting endlessly for an answer from a non-responsive device. Thus it is much more robust against device failures and events the programmer did not expect than the old control server.

Vendor	Device	Type of device
Cryoelectra	CRE3161	RF amplifier
Oerlikon Leybold	Ionivac IM540	Vacuum gauge
SMC	HRS018-AF-20	Thermo chiller
LakeShore	Model 241	Liquid Helium level meter
Heinzinger	PNC 10000-6ump	High voltage power supply for Wien filter
Heinzinger	PNChp 125000-5neg	High voltage power supply for cathode of polarized gun

Table 11.1: Commercial devices for which EPICS support has been implemented in the context of this work.

The EPICS IOC configures all device servers that support the COM Port Control Protocol [81] remotely. This centralizes configuration and simplifies replacement of broken device servers. It also simplifies the backup of the configuration of the device servers.

In the context of this thesis control system support for the devices listed in table 11.1 has been implemented.

11.3 Graphical user interface

For many years most graphical user interfaces for EPICS control systems have been created using traditional display manager software like the Extensible Display Manager (EDM, [82]) or the Motif Editor and Display Manager (MEDM, [83]). These programs allow to create operator interface windows containing buttons, dials, meters etc. with a graphical editor that can be used as soon as the display is switched to run mode. Separate programs were used for plotting PV data, for alarm handling etc. Most of these programs only run on Unix machines and are based on out-of-date widget toolkits like Motif making them more and more difficult to maintain on state-of-the-art operating systems.

A more recent project providing graphical user interfaces for control systems is Control System Studio (CSS, [84, 85, 86]). It provides an abstraction

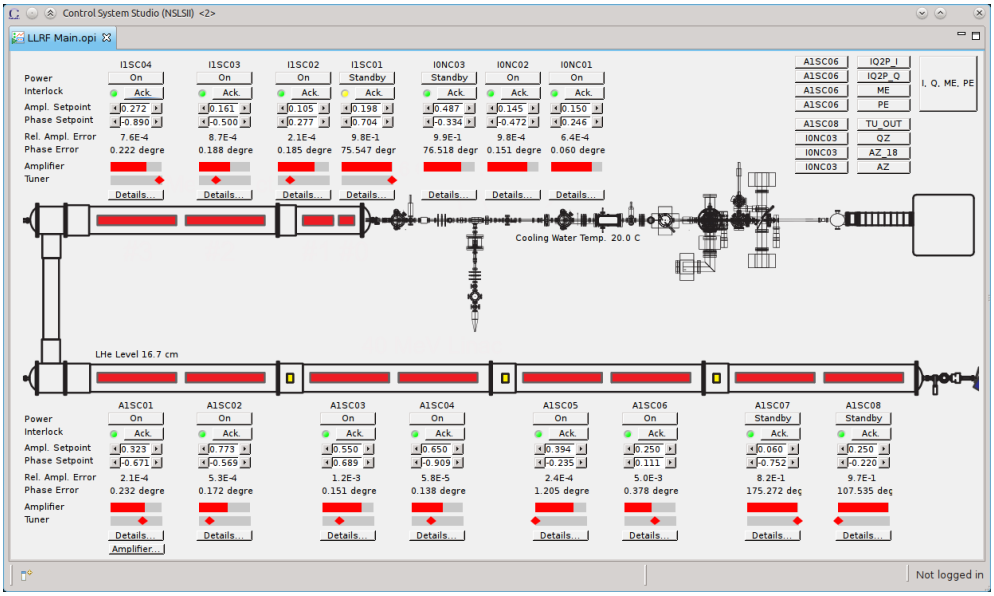

Figure 11.2: Screenshot of the main OPI screen of the rf control system implemented with the BOY plug-in of Control System Studio. It provides an overview of all cavities and allows to switch them on and off and to change their amplitude and phase set-points.

layer which supports different accelerator control systems (EPICS, TANGO [87], TINE [88] etc.). Control System Studio is implemented using the Eclipse Rich Client Platform framework [89, 90] which allows to build platform-independent Java applications with loosely-coupled plug-ins. Available are plug-ins providing operator interfaces (Best OPI Yet, BOY, [91]), functionality for plotting PV and archiving data (Databrowser, see section 14), alarm handling etc. The Eclipse Rich Client Platform technology allows the plug-ins to exchange data with each other which greatly improves usability compared to using separate applications. For example the context menu of an OPI widget allows to directly open a plot of the corresponding PV data.

Control System Studio has been chosen as the primary graphical user interface at the S-DALINAC. All operator interfaces for slow control tasks have been created with BOY. In the context of this work OPIs for the rf control system, the thermionic gun, the in-house developed magnet power supplies (`CPS05-7` and `CPS05-10`), and for the devices listed in table 11.1 have been created. Figure 11.2 shows the overview screen for the rf control system. This screen allows the operator to switch cavities on and off and to adjust amplitude and phase of the

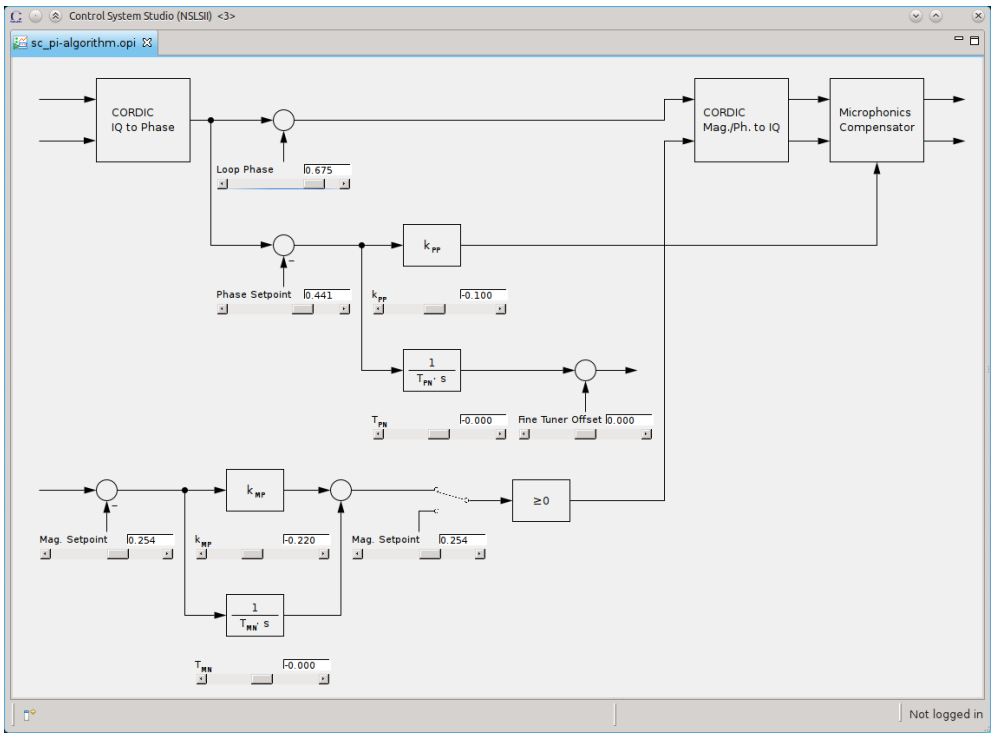

Figure 11.3: Screenshot of the OPI window providing access to the most important rf control parameters of one of the superconducting cavities.

cavities. The buttons on the upper right are used to select the diagnostic channels. Clicking on the "Details" button next to a cavity brings up an OPI window with the most important parameters of the control algorithm (see fig. 11.3). For normal-conducting cavities this opens the OPI for the generator-driven resonator algorithm whereas for superconducting cavities an OPI for the self-excited loop algorithm is opened.

In contrast to most of the clients of the previous control system the new OPIs use graphical elements to visualize the structure of the controllers and the arrangement of the components. This significantly improves usability which is especially important for the student operators at the S-DALINAC.

12 RF diagnostics

The previous rf control system used analog oscilloscopes for diagnostics. In contrast a fully digital read-out of all signals has been implemented for the digital rf control system in the context of this thesis. Eight of the signals can be read out via the USB 2.0 interface of the controller card or via the crate controller, respectively. Data is available with the full sampling rate of the ADCs of up to 1 $^{\mathrm{MS}}/_{\mathrm{s}}$ and an accuracy of 16 bits. Signals can also be read out with the full accuracy of 18 bits if two channels are combined to a 32 bit channel. For most purposes the default accuracy of 16 bits is sufficient.

In addition to this "fast" read-out the crate controller provides a "slow" read-out with a sampling rate of $\approx 2\,^{\mathrm{kS}}/_{\mathrm{s}}$ via a second USB port. It provides all 64 signals of all controller boards (resulting in 1024 channels in total) with the full resolution of 18 bits.

12.1 USB device support

The two USB connections provide a data-rate of 128 Mb/s for the fast read-out and 64 Mb/s for the slow read-out. These data rates are too high to feed the data stream into scalar EPICS records. Instead a USB device support has been implemented that streams the data very efficiently to connected clients (see fig. 12.1).

Dedicated threads read data from each of the USB ports and copy it to ring buffers. In the context of the development of the USB device support the Linux USB library `libusb1` has been extended to support efficient streaming of data by a member of the in-house electronic workshop. A connection manager waiting for incoming connections spawns an own thread for each client connection. This thread reads data from one of the ring buffers and delivers it to the network.

Multiple threads are reading data from the same ring buffer while data is written to it by another thread at the same time. Therefore a thread-safe ring buffer class with support for multiple readers has been implemented in the context of this work. Dedicated USB reader threads, ring buffers, and network connection threads are used for the fast read-out and the slow read-out. The multi-threaded concept of the device support allows very high data rates while distributing load

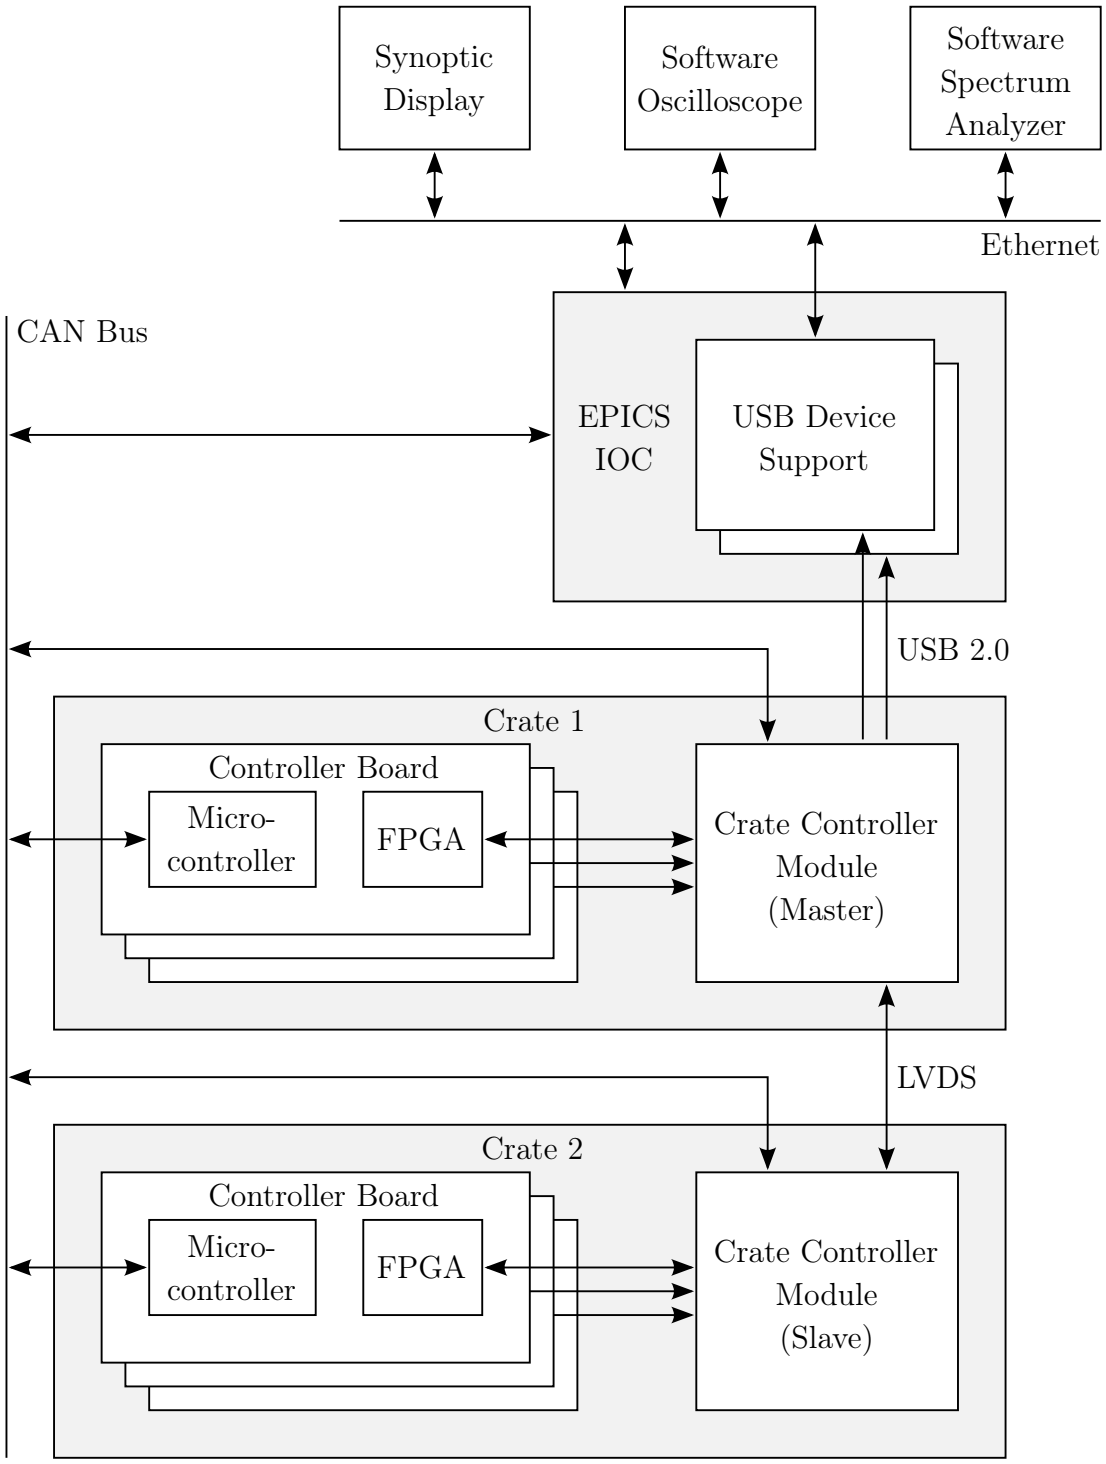

Figure 12.1: Block diagram showing the read-out of the diagnostic data from the rf control system. The crate controllers combine the data from the controller boards while at the same time reducing the data rate.

to multiple CPU cores of the PC. Since no control data is transmitted with the data it can easily be recorded for analysis (cf. section 7) using standard Unix tools like NETCAT.

The ring buffers store about 1 s of data. If one of the clients is reading slower than data is arriving from the hardware, the ring buffer runs over as soon as the client is lagging behind a full buffer size. When this happens, the buffer for this client is emptied completely to make sure that buffer overruns occur as seldom as possible. This results in long sequential data sections leading to oscilloscope images that are distorted only seldom even in the case of an overloaded network while at the same time assuring that the data is not older than one second. Under normal load conditions no data is lost and only a small fraction of the buffer size is needed leading to minimal latency.

In addition to streaming the data directly to the clients, the slow data stream is also read from the ring buffer, separated into the 1024 channels and copied into 1024 EPICS `waveform` records. The array data from these `waveform` records is then scaled and reduced to a lower data rate of approximately 10 S/s by averaging using `aSub` records. The resulting data can be used for monitoring and archiving. For example the output amplitude data is shown on the OPI to inform the operator of output saturation.

The described approach combines the performance of direct data streaming to the clients with the availability of EPICS database records which allow for flexible monitoring. In total the rf control IOC running at the S-DALINAC provides more than 17.000 records for the 16 channels, most of them related to diagnostics.

12.2 Graphical user interface

The eight signals of the fast data stream can be displayed on-line using the open source software oscilloscope OSQOOP (see fig. 12.2, [92]). This oscilloscope software has been extended in the context of this thesis with a data-source plug-in that reads data from the diagnostic server process via the network. The oscilloscope features adjustable vertical and time axes, a configurable trigger, and plug-ins for math operations and signal filtering. It also provides an X/Y mode that can be used to view a pair of signals in the I/Q plane (see fig. 12.3). The stability and

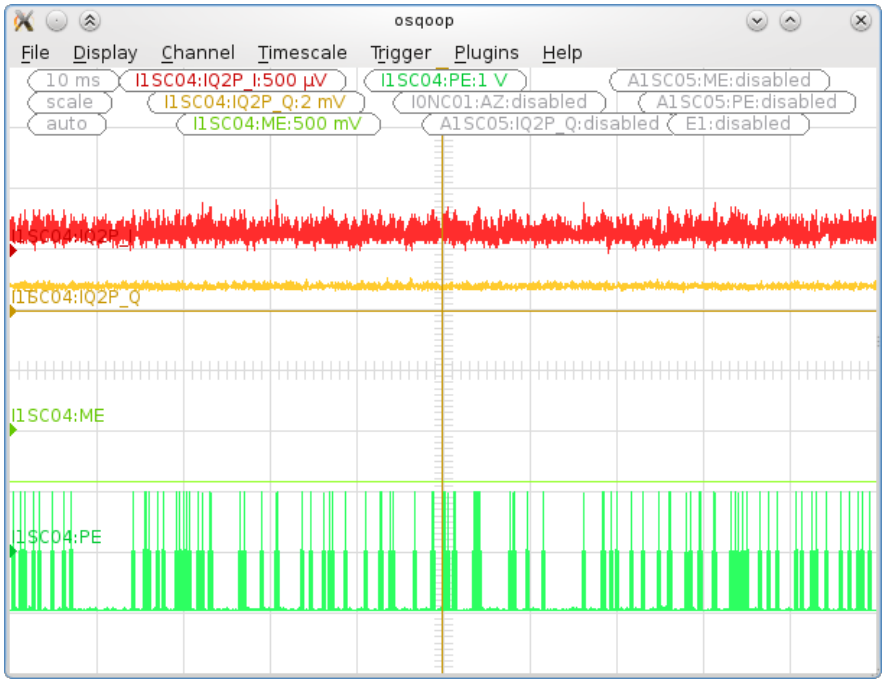

Figure 12.2: Screenshot of the software oscilloscope. The scale of the signals can be adjusted with the widgets on the top. The names of the signals are automatically read from the IOC.

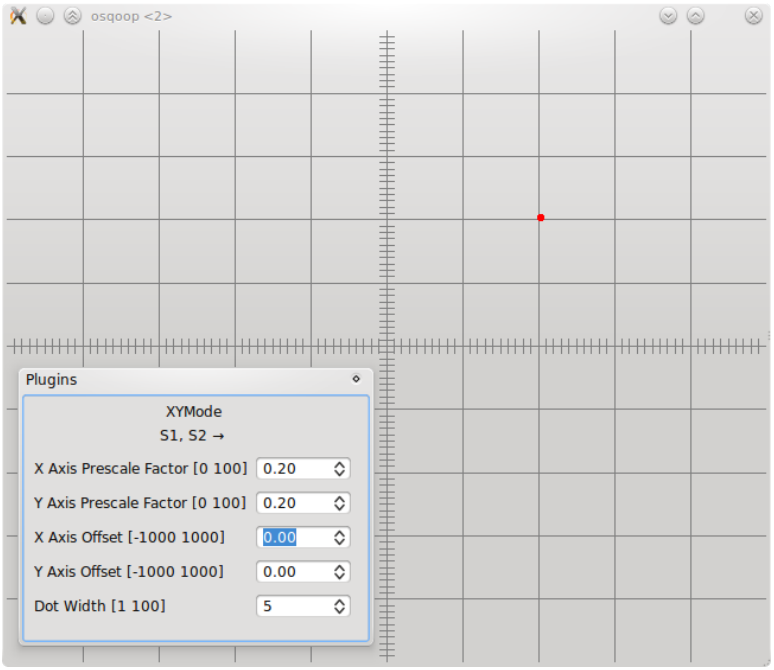

Figure 12.3: Screenshot of the X/Y mode of the software oscilloscope showing the I/Q representation of the FPGA input signal with active controllers and a phase set-point of 45°.

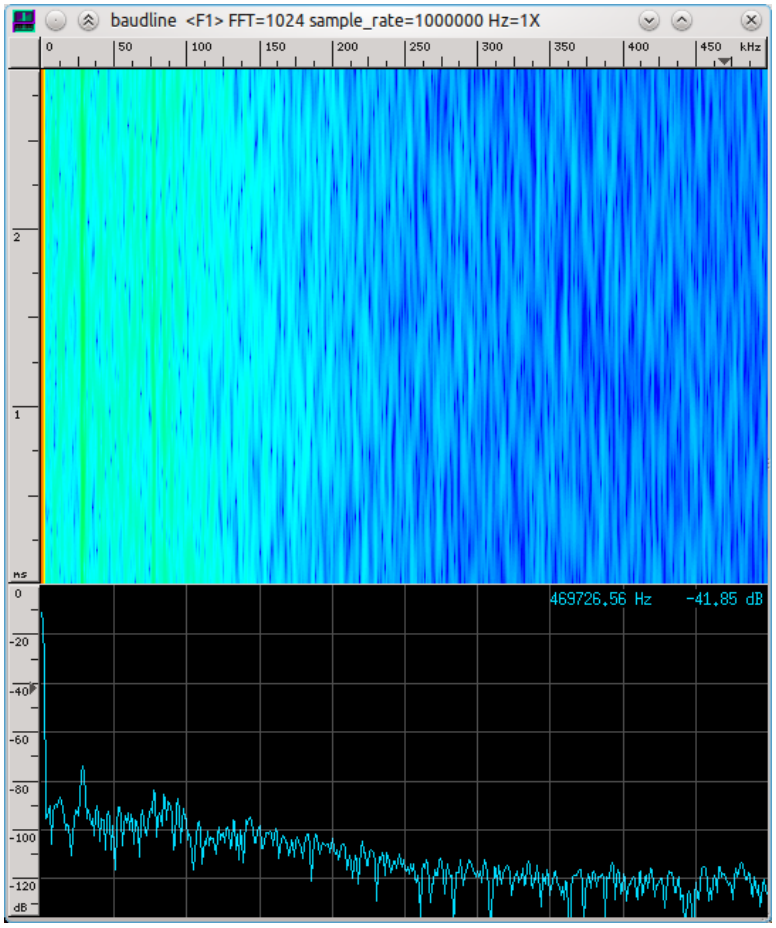

Figure 12.4: Screenshot of the spectrum analyzer software.

usability of the X/Y mode has been improved considerably in the context of this thesis.

The signals shown on the oscilloscope can be selected on the main rf control OPI. The software oscilloscope has been extended by [72] to automatically read the channel names from the EPICS IOC via Channel Access and display them next to the signal.

In addition to the oscilloscope application the software spectrum analyzer Baudline [93] can be used for frequency-domain analysis of the rf control system. It calculates the discrete Fourier transform of a signal and displays it in form of a spectrum and a waterfall diagram (see fig. 12.4).

The digital read-out improves the accuracy of the oscilloscope images significantly. Compared to digital storage oscilloscopes that usually offer a resolution of 8 bit the software oscilloscope of the digital rf control system provides a much higher

accuracy of 16 bits. This gain in resolution is important to diagnose small errors on the signals while feedback is active. With the software spectrum analysis tool the control system provides the possibility to analyze signals in the frequency-domain. This feature can e. g. be used to identify the source of microphonics disturbances. The waterfall diagram has proven to be very handy to identify intermittent disturbances caused by digital components on previous revisions of the hardware.

12.3 On-line rms errors

The data from the slow readout available in the EPICS database can be used to detect if a controller is out of lock, but it does not provide precise information about the performance of the cavity and its controller. In order to quantify the variations of the field in the cavity the rms errors of the amplitude and phase error signals are calculated. Instead of determining the errors from the fast streaming data at the PC, thereby occupying all eight diagnostic channels, the rms calculation has been moved partly into the FPGAs of the controller boards. They have the whole data stream available and as a side-effect relieve the PC from processing work by carrying out the most time-consuming part of the calculation. The FPGAs calculate the squares of the signals for each ADC sample and sum 2^{17} of them up which corresponds to a time of roughly 0.1 s. The result is read out by the PC via CAN bus. The EPICS IOC calculates the square root of the sums and scales the values to finally get an absolute error in degrees for the phase or a relative error in case of the amplitude, respectively.

The resulting values are displayed on the main rf control OPI. The rms errors are very sensitive to changes of the control parameters and thus can be used to find the optimal control parameters. This allows to adjust the control parameters much more accurate than with the analog rf control system leading to a better energy spread of the beam. During beam operation the rms errors are used to monitor the performance of the cavity. This e. g. allows to detect increasing control errors caused by beam losses in the superconducting cavities.

13 Finite-state machines

Simple actions in the context of an accelerator control system can be automated using plain EPICS records. An example for such an action is shutting off rf power if the liquid helium level falls below the level needed for stable operation. Other high level processes on the other hand are too complex to be implemented with EPICS records in a convenient and maintainable way. An example is the sequence of commands needed for degaussing the magnetostrictive fine-tuners of the superconducting cavities. But nevertheless these processes would benefit from being reproducibly executable. For most processes the sequence of commands can be visualized in a state diagram (see fig. 13.1). Each of the circles represents one state. As the number of states in this model is finite and the machine is in only one state at a time, the machine is called a finite-state machine or finite-state automaton. It can switch from one state to another when a triggering event occurs. These transitions are depicted in the state diagram by arrows with the corresponding triggering condition next to them. Commands are always executed in conjunction with transitions from one state to another. In the state diagram these actions are put in parentheses.

In an EPICS environment finite-state machines can be implemented as a stand-alone client program that connects to the IOC via the network or as part of the IOC. The client approach has the drawback that each Channel Access connection the state machine uses might fail. This demands for rather sophisticated error handling. An implementation of a state machine on the IOC-side on the other hand in many cases requires no access to PVs via the network because all needed PVs reside on the same IOC. Error handling for network errors is only needed if the finite-state machine accesses PVs provided by other IOCs.

The most widely used tool for implementing finite-state machines in an EPICS control system is the RUN-TIME SEQUENCER [94, 95, 96]. This extension uses the State Notation Compiler to translate a description of the finite-state machine in the State Notation Language (SNL) to C code, which in turn is compiled into a program that can be run on an IOC. The SEQUENCER has been used to implement all PC-side state machines in the S-DALINAC accelerator control system because it allows to run them either as independent applications on the client side (which simplifies testing) or as a part of the IOC application. Other reasons for using the

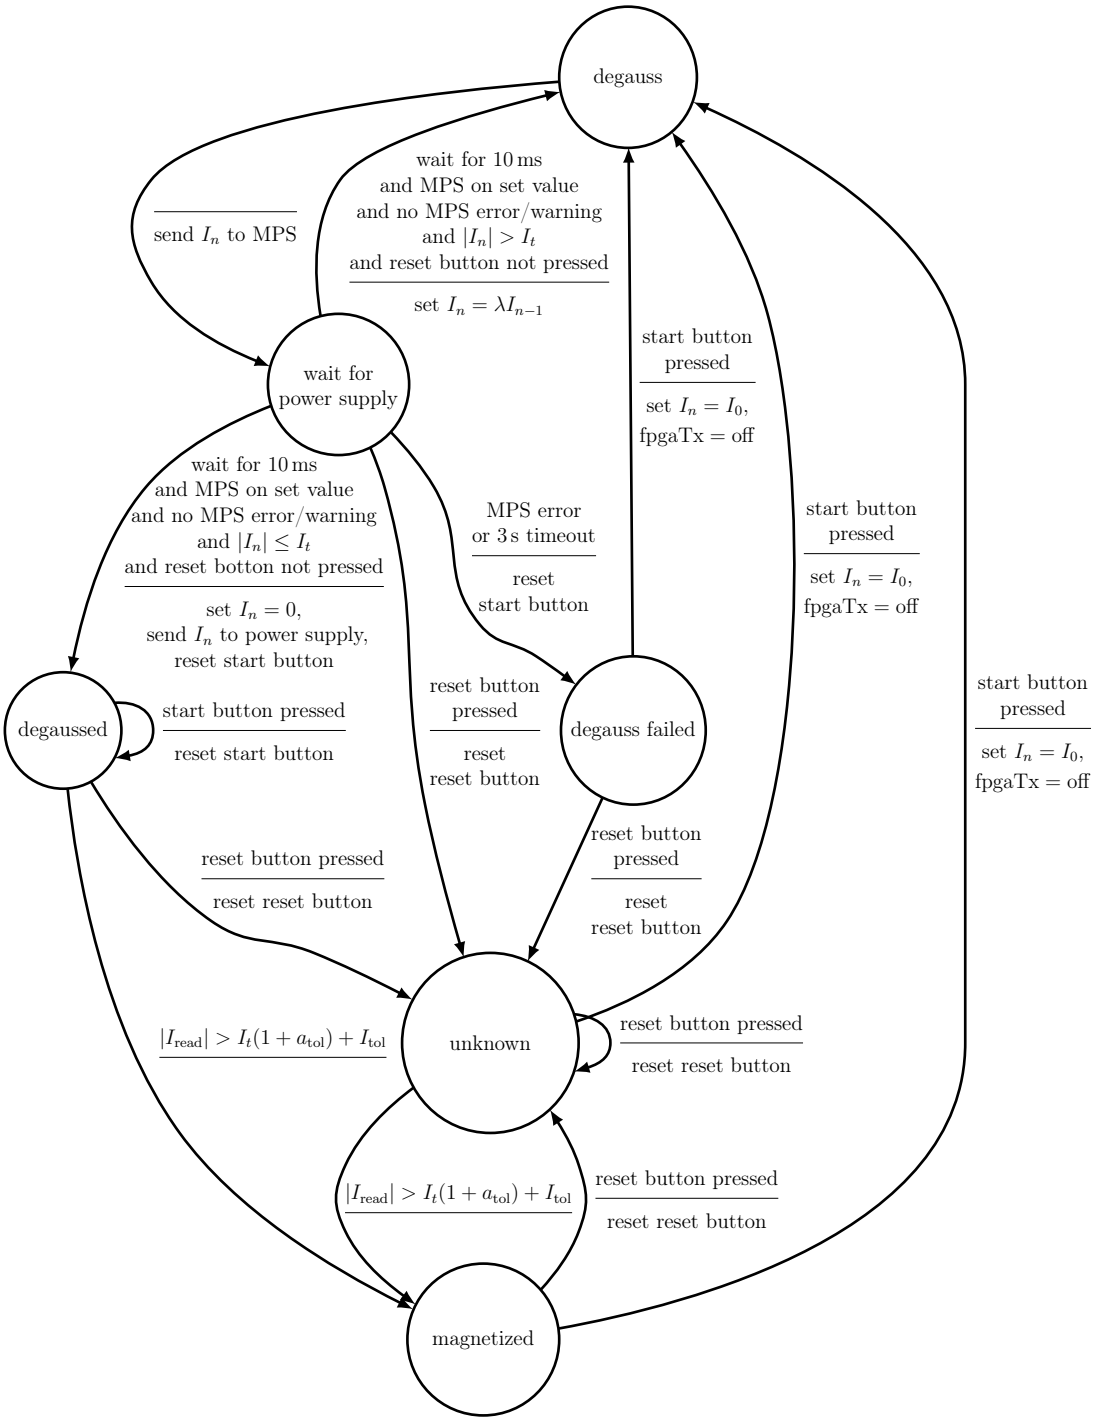

Figure 13.1: Finite state machine used for degaussing the magnetostrictive fine tuners. Therefore the current of the magnet power supply (MPS) is iteratively decreased with alternating sign.

SEQUENCER have been its big number of users and its actively developed code base. In addition to this the SEQUENCER can run state programs in a reentrant mode which allows to run more than one instance without any interference between the different copies.

In the context of the EPICS migration project several finite-state machines have been implemented. As an example for a state machine the process of degaussing the fine tuners will be described in detail. Each of the 12 fine tuners has its own instance of this state program.

13.1 Degaussing of the fine tuners

The process of degaussing a fine tuner requires applying a current of alternating sign and decreasing strength to the coil of the tuner (see section 4.6.1). Therefore in each step a new current

$$I_n = \lambda\, I_{n-1} \quad \text{with} \quad -1 < \lambda < 0 \tag{13.1}$$

has to be calculated and sent to the power supply. This is represented by the "degauss" state in fig. 13.1. Before a new command can be sent to the power supply, the finite-state machine has to wait until the power supply has successfully regulated the current to the requested value (state "wait for power supply"). This is represented by the condition connected to the transition back to the degauss state. The state machine uses a set of records provided by the power supply template that detects if the actual output current matches the current set by the state machine. This takes the accuracy of the `CPS05` power supplies of 0.2% plus 10 digits into account and is represented by the "MPS on set value" condition in the state diagram.

Before the finite-state machine enters the "degauss" state, transmission of tuner commands by the controller board has to be disabled. In addition I_0 is set to the start current provided by the operator. The process stops when I_n becomes smaller than the termination current I_t which can also be adjusted by the operator. Typical values are $I_0 = 2\,\mathrm{A}$ and $I_t = 100\,\mu\mathrm{A}$. These values are provided as default values.

After degaussing, the state machine enters the "degaussed" state where the output current is switched to zero and the start button is reset to signal the operator

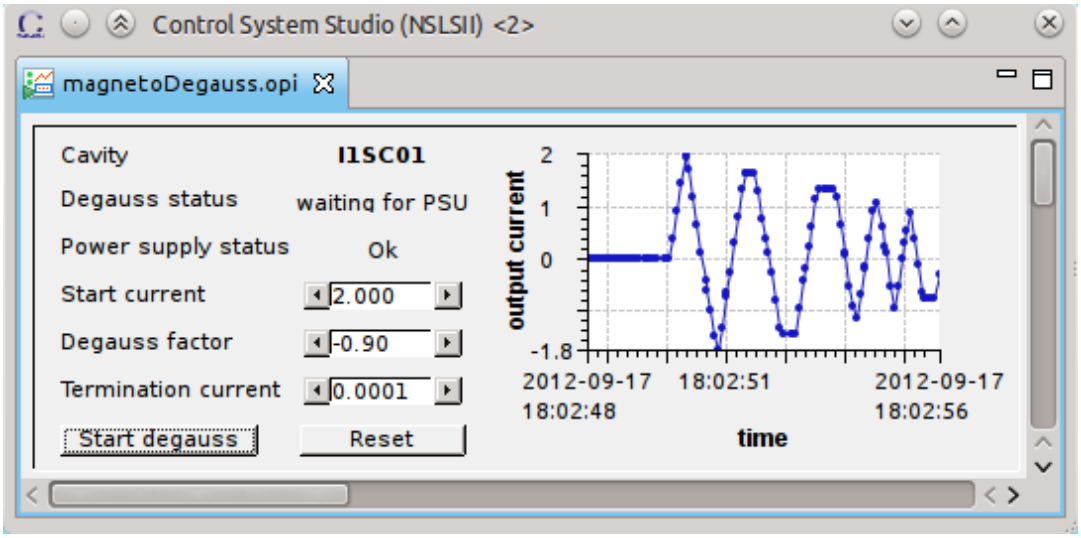

Figure 13.2: Graphical user interface of the magneto degauss finite-state machine.

that the degauss process has been completed. If the start button is pressed again, no further action is needed. The state machine only resets the start button to prevent an unexpected start of the degauss process as soon as the machine enters the "magnetized" state.

If the power supply is set to an output current that exceeds the termination current I_t, the state machine changes to the "magnetized" state from where the operator can start the degauss process again.

After starting the IOC the state machine begins in state "unknown" to signal the operator that the state of the fine tuner is unknown. As soon as an output current exceeding I_t is detected the machine changes to the "magnetized" state.

If the power supply does not apply the output current to the coil during a period of three seconds, the state machine switches to the "degauss failed" state. This can happen if the coil or the cables are damaged, the power supply detects an error, the power supply is broken, or some other process changes the current while the degauss process is running. The operator can confirm an error by pressing the reset button which switches the finite-state machine back to the "unknown" state.

The graphical user interface of the state machine shows the current state as well as the controls for entering the parameters and starting the degauss process (see fig. 13.2).

14 Archiver

Archiving systems that collect and store data from an accelerator control system play an important role in modern control systems environments. They can e. g. help to identify broken hardware, decreasing performance of components or wrong operation of systems by operators.

The previous accelerator control system provides an archiving software that collects process data from the helium liquefier and stores it in a relational database [70]. It reads the data directly from the hardware without using a control server. Thus it is not possibility to attach clients to monitor the current values or connect them with other PVs like it can be done in the EPICS database. Access is only provided via a web client [71] which retrieves data from the database.

For EPICS different implementations of archiving systems are available that all retrieve data via Channel Access from the IOCs. The traditional Channel Archiver [97] stores data in its own binary format. This provides high performance but requires special clients. Other archivers store the data in a relational database and thus allow access from a wide variety of clients. They also provide a higher flexibility with regard to data retrieval and analysis than the EPICS Channel Archiver. With the Control System Studio Archive Engine [98] one of these relational database archiving solutions has been set up for the S-DALINAC in the context of this thesis. It makes use of the Control System Studio Data Access Layer and thus can collect data from different control systems at the same time. With MySQL®, Oracle®, and PostgreSQL® the CSS Archive Engine supports different relational databases as back-end. PostgreSQL has been chosen as archive back-end for the S-DALINAC because it is a powerful open-source relational database management system that comes without license fees. Another advantage is that PostgreSQL is already used for the previous archiving system and for the configuration management system thus keeping the maintenance effort low.

The previous archiving system collects data at fixed intervals (usually once a minute). In contrast the CSS Archive Engine monitors PVs for changes. This makes sure short events like spikes are not missed. At the same time this saves memory for PVs that only change seldom or very slowly (e. g. power status and

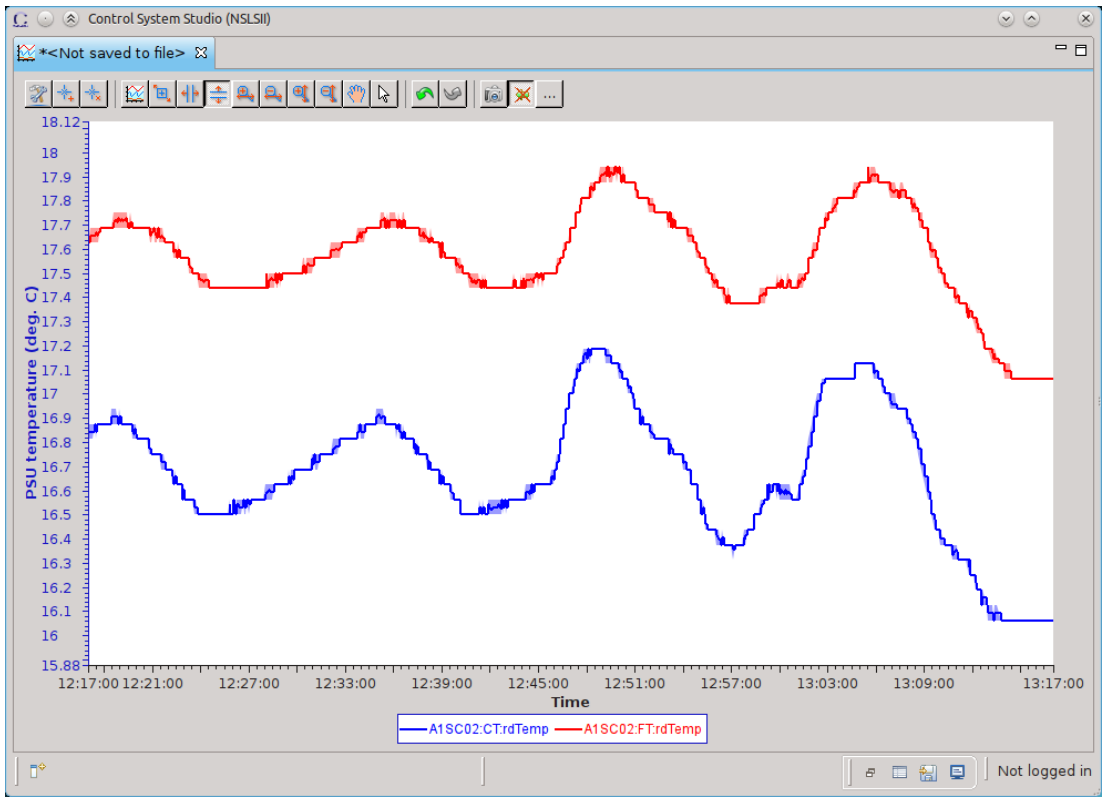

Figure 14.1: Screenshot of the Control System Studio Data Browser plug-in. It plots data from PVs and also retrieves archive data from the relational database. The visualization makes the correlation of the temperature drift of two tuner power supplies obvious.

supply voltages of components). For PVs that change continuously (like the current through the fine tuners) or that are very noisy an archive death band can be configured on the IOC. Only when the value of the PV has changed more than the archive death band a new value is transmitted to the Archive Engine. Archive death bands have been configured for all records that provide a high amount of data.

The Control System Studio Data Browser plug-in [99] is used to visualize data from the PVs and the archive (see fig. 14.1). It plots data directly from the IOC. If archive data is available, the history of the PV value is loaded from the database automatically. The Data Browser plug-in provides a very powerful and intuitive way to navigate through the data and thereby helps operators and engineers to analyze trends and track down failures efficiently.

Currently more than 5.000 PVs are archived at the S-DALINAC resulting in about 2.5 TB of data per year. The PostgreSQL server's hardware and software have been optimized for fast data retrieval to make sure data can be analyzed in a reasonable time [27]. Database indexes are used to speed up data retrieval. A redundant array of 36 independent disks (RAID 10) improves disk performance and reliability [100]. A non-volatile write-back cache improves write speed while at the same time maintaining data integrity in the case of power outages. The database back-end can answer queries a lot faster if it does not need to read data from disk. With 128 GB the size of the main memory has been chosen to be large enough to hold the data of the last days including the relevant indexes.

If the size of the main archive table grows larger than physical memory and even its index stops fitting into memory query times can escalate. One way to improve performance is to split data into several partitions that each fit into main memory. At the S-DALINAC most database queries ask for data from the last three days making partitioning over time with a partition size of a week a reasonable choice. When executing queries the database management system can skip partitions that are outside the requested range. For most queries only one or two partitions have to be considered.

In contrast to enterprise databases like Oracle or Microsoft SQL Server, PostgreSQL does not provide built-in functions for partitioning. Using PostgreSQL's extensive server-side programming features a partitioning solution tailored to the particular needs of the CSS archiver has been developed in the context of this thesis. It uses table inheritance to combine the data of weekly sub-tables into one table. A trigger function redirects `INSERT`s to the appropriate partition. New partitions are added automatically. The partitioning feature only affects the database side of the archiver and has been included into the Control System Studio distribution.

Partitioning speeds up data retrieval and simplifies maintenance. For example creating backups or deleting old data is much faster with partitioning. With typical retrieval times of a few seconds for the data of a whole day the system has proven to be fast enough for everyday usage. As soon as all devices have been migrated to EPICS the old archiver can be shut down.

15 Control system infrastructure

In the process of migrating the control system to an EPICS-based system the control system infrastructure has been modernized as well. This includes major improvements concerning computer security as well as maintainability.

15.1 Network

In the last years control systems have been increasingly threatened by cyber attacks [101, 102]. At the beginning of this thesis the accelerator network was not separated from the office network. A firewall prohibited access from outside the department's network but the accelerator's network was accessible from all office PCs and laptops. Thus e. g. a student's laptop infected with a virus could disrupt accelerator operation. In the context of this thesis network security has been improved by segregating the accelerator network from the office network. The segregation has been accomplished by configuring virtual local area networks (VLANs) on the involved network switches. Thus the existing network infrastructure (network switches and cables) can be used to host multiple isolated virtual networks. Each port of the network switches can be flexibly configured to be part of one of these networks.

A firewall connects the accelerator network to the department network. It prohibits access to the accelerator network from the outside. Several instances of the Channel Access Gateway [103, 104] act as proxies that allow read-only access of PVs from the office and the experiment network.

A second VLAN for accelerator services which need to be available from outside the accelerator core network is currently being implemented. It will accommodate hosts running services like the relational database containing the archive as well as accelerator related web services such as WebOPI. Access from the core accelerator network into the services network will be possible to allow processes like the Archive Engine to write into the database. Access from the services network into the core accelerator network will be prevented by a firewall. Isolating the vulnerable web servers from the core accelerator network in the described way provides an extra layer of security.

The described security measures help to maintain the high availability of the S-DALINAC's new accelerator control system while still providing convenient read access to all EPICS PVs from the office network.

15.2 Virtualization

The previous control system of the S-DALINAC concentrated all services on a few PCs to keep the probability of hardware failures low. This considerably complicated upgrades of the operating system, the libraries, and the applications. The EPICS-based control system introduced at the S-DALINAC in the context of this thesis instead uses an own machine for each service to keep each installation as simple as possible. This also allows a finer split-up of maintenance responsibilities and allows to conduct maintenance work on hosts that are not needed in the current mode of accelerator operation. Virtualization and software deployment techniques are used to assure a high availability of the control system while at the same time keeping the maintenance effort low.

Computer hardware virtualization uses a software running on a real computer (host system) to simulate virtual computers (virtual machines). The software controlling this virtualization (often called hypervisor) emulates the complete hardware of the virtual machines including CPUs, main memory, disks etc. The hypervisor can execute multiple virtual machines on the same real machine.

In the S-DALINAC's control system environment most machines are operated at low processor, memory, and disk load wasting a lot of computing resources. By running these machines on a hypervisor, hardware resources can be used more efficiently. At the same time the amount of hardware and thereby the risk of hardware failures is reduced. Even if the host system is equipped with more expensive high-reliability functions like redundant power supplies, fans, disks etc. virtualization can be more cost-efficient. The virtual machines of the S-DALINAC's control system are executed on a cluster of servers machines in the department's data center running the VMware ESXi hypervisor. The cluster uses redundancy and high availability techniques for power supply, storage and network.

Since the hardware of the virtual machines is emulated by the hypervisor it is not possible to virtualize control system PCs that need access to hardware directly

Host name	Description
`archiver`	CSS Archive Engine storing PV data to the relational database
`cagateway`	Channel Access Gateway providing access to PVs from outside the accelerator network
`debianbuild`	Build-server for accelerator control system software
`debianfai`	Server for automatic software deployment
`monitoring`	Nagios server that monitors the accelerator network, its hosts, and services
`serial2ethernetioc`	IOC for devices connected to serial device servers

Table 15.1: List of the most important virtual machines used in the S-DALINAC accelerator control system.

connected to the PC. However machines like the Channel Access Gateway, servers used for software development and deployment can be virtualized. Even IOCs that require only network access to communicate with devices like serial device servers can be virtualized. Table 15.1 lists the machines that are virtualized at the S-DALINAC.

An important advantage of virtualization of control system machines is the centralization of administrative tasks. Replacing the hardware a virtual machine runs on is very easy. Virtual machines can even be moved from one node of the cluster to another without shutting down (live migration). Virtual machines can easily be cloned which can be very handy to create a test machine. The state of a machine can be saved in snapshots which allows to return to a snapshot if e. g. something goes wrong during major maintenance tasks. Using the snapshot technology virtual machines can also be automatically backed up during operation. All these features have proven to make maintenance of control system machines much easier while at the same time increasing availability.

15.3 Monitoring

With an increasing number of control system hosts and services it becomes more and more time-consuming to make sure they are all up and running. This is particularly difficult for operators that are not familiar with the details of the control system infrastructure. To simplify this task, a monitoring service has been set up. Nagios [105] has been chosen as monitoring software because it provides a wide range of plug-ins for monitoring computer infrastructure like network switches, disk usage, and protocols like HTTP, SSH, and SNMP. Even a plug-in for monitoring EPICS channels is available [106]. Nagios provides extensive reports on its website (see fig. 15.1). For convenient everyday use by the control system administrators user-friendly browser extensions and applications for hand-held devices like smart-phones are available. In the context of this thesis monitoring for 32 hosts and 57 services has been set up. This covers all important sub-systems of the accelerator network.

15.4 Software development

In contrast to the previous control system, revision control software has been used to track changes of the source code from the first day of the project. The complete software developed for the S-DALINAC's control system is kept in a central Subversion repository [107, 108]. This allows to track changes to all control system components. Additionally a web-based issue tracker software is used to track bugs in control system components [109]. Both tools facilitate collaborative development and maintenance of the S-DALINAC accelerator control system.

For the installation of the control system PCs binary Debian packages provided by the NSLS-II project are used for all EPICS core components [110]. This keeps the maintenance effort low and makes the S-DALINAC's control system as compatible as possible to other EPICS installations. The most important S-DALINAC-specific control system components are automatically built and packaged by a local build server. This build server automatically generates binary packages for Debian and Ubuntu systems. New packages are pushed to a Debian package repository on the same machine. From there they can be installed to the control system computers using the standard Debian/Ubuntu package manager.

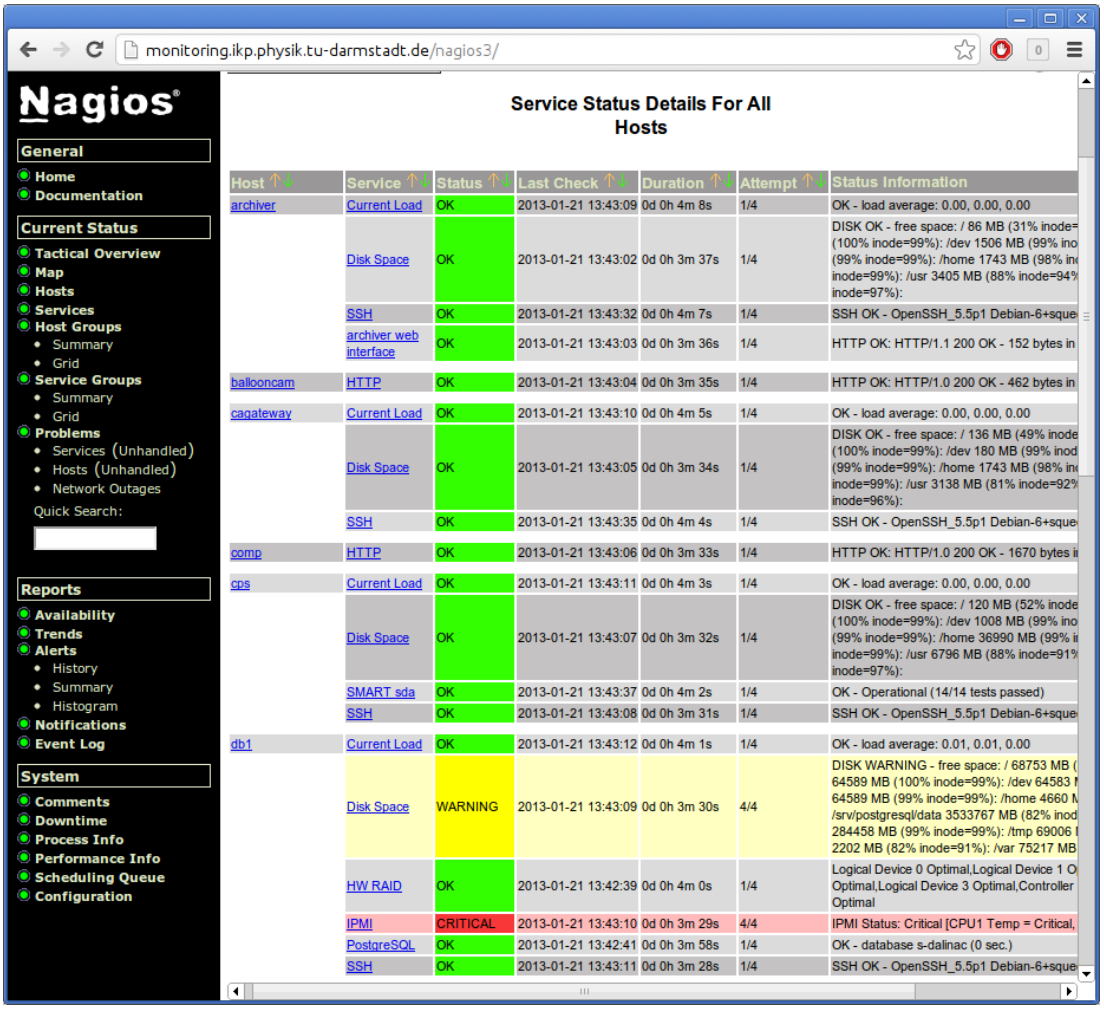

Figure 15.1: Screenshot of the web frontend of the Nagios monitoring service. The management controller of the database server `db1` reports critical CPU temperatures. In addition to this the machine is short on disk space.

15.5 Software deployment

The subsystems of the EPICS control system have been commissioned as soon as the core functions allowed stable operation. Within the concept of continuous integration, new features and corrections have been integrated through frequent deployment of the software to the control-system PCs. This needs to be done in a repeatable fashion and has been achieved by using the Debian Fully Automatic Installation (FAI, [111]) to setup the control system machines over the network. New software and its deployment can be tested with virtual machines before it is applied to the production machines at the accelerator. Since most software

components can be installed as binary Debian packages created by the build server, changes can be deployed to the test system or to the production system within minutes. A complete installation of one of the control system PCs takes about 10 minutes. Hence downtimes of the accelerator of as short as half an hour can be used to upgrade the control system software. During installation the IOC configuration is generated automatically from a relational database and a set of templates during installation. This way a quick recovery from a PC failure is warranted. A single spare machine is sufficient for all control system PCs running Linux.

16 Summary and outlook

The development of the digital rf control system with its re-programmable hardware and thousands of PVs triggered the migration of the accelerator control system from an in-house developed system to an EPICS-based system which has been presented in part II of this thesis.

The extensive basic functionality of the EPICS framework allows the integration of many new features via configuration instead of programming. This significantly improves the flexibility of the control system and speeds up the development process. Standard EPICS tools have been used wherever possible, e. g. to implement complex processes in form of finite-state machines. With Control System Studio a modern graphical user interface has been chosen that improves usability by integrating multiple plug-ins for operation as well as monitoring into a single application.

Since all in-house developed hardware communicates via CAN bus, a CAN device support has been implemented. While device support for commercial devices equipped with serial ports is readily available, a special device support had to be implemented for the high-speed USB read-out of the rf control system. The data is streamed directly to the clients for maximum performance. At the same time the data is written to EPICS records at a lower rate. A software oscilloscope, providing much higher accuracy than conventional digital oscilloscopes, as well as a spectrum analysis software offer different visualizations of the monitored data. Errors of amplitude and phase are calculated on-line which can be used to optimize the control parameters very precisely and to observe the status of the controller.

The Control System Studio archiver is used to keep the history of all relevant EPICS PVs. The PVs are monitored for changes, so that spikes are not missed but only necessary data is stored. A new database server has been installed and optimized for maximum performance. It allows fast data retrieval despite huge amounts of archived data.

State-of-the-art technology like virtual machines, centralized monitoring, and automatic deployment have been exploited to improve the maintainability of the

new accelerator control system. Due to automation of the install process, even recovery from a PC failure is a matter of minutes.

Until now, the rf control system, the vacuum controls, part of the beam diagnostics and the magnet power supplies have been implemented in EPICS. Currently the other power supplies, the monitoring of the cryogenic plant as well as the remaining part of the beam diagnostics are still operated using the previous accelerator control system. As soon as the remaining features have been migrated to EPICS, the old control servers and the previous archiving system can be shut down.

A Transfer function of the baseband filters

Third-order low-pass π-filters are used for filtering the baseband signals on the FPGA board. The same type of filter is used as anti-aliasing filter and as reconstruction filter, respectively. A schematic diagram of the filter is shown in fig. A.1. According to Kirchhoff's voltage law the sum of the voltages around each mesh is zero:

$$V_{\text{in}} = V_{R_1} + V_{C_1} \tag{A.1}$$
$$V_{C_1} = V_L + V_{C_2} \tag{A.2}$$
$$V_{C_2} = V_{R_2} \tag{A.3}$$
$$V_{R_2} = V_{\text{out}} \tag{A.4}$$

Kirchhoff's current law yields:

$$I_{R_1} = I_{C_1} + I_L \tag{A.5}$$
$$I_L = I_{C_2} + I_{R_2} \tag{A.6}$$

The currents through the components are:

$$I_{C_1} = C_1 \dot{V}_{C_1} \tag{A.7}$$
$$I_{C_2} = C_2 \dot{V}_{C_2} \tag{A.8}$$

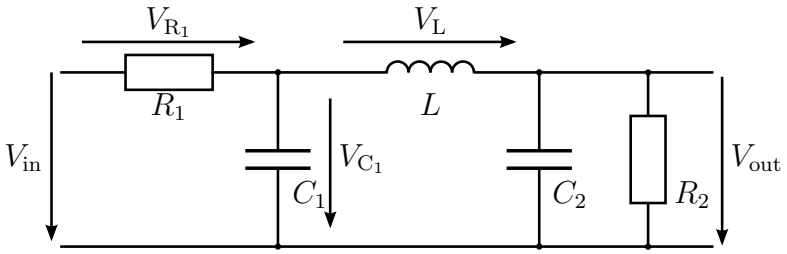

Figure A.1: Schematic diagram of the third-order low-pass filters used as anti-aliasing filters and reconstruction filters on the FPGA board.

$$\dot{I}_L = \frac{V_L}{L} \tag{A.9}$$

$$I_{R_1} = \frac{V_{R_1}}{R_1} \tag{A.10}$$

$$I_{R_2} = \frac{V_{R_2}}{R_2}. \tag{A.11}$$

Combining these eleven equations yields the differential equation of the filter:

$$\begin{aligned} V_{\text{in}} = {} & R_1 L C_1 C_2 \dddot{V}_{\text{out}} + \left(\frac{R_1 L C_1}{R_2} + L C_2 \right) \ddot{V}_{\text{out}} \\ & + \left(R_1 C_1 + R_1 C_2 + \frac{L}{R_2} \right) \dot{V}_{\text{out}} + \left(\frac{R_1}{R_2} + 1 \right) V_{\text{out}} \end{aligned} \tag{A.12}$$

Transforming the differential equation into Laplace space yields

$$\begin{aligned} V_{\text{in}}(s) = {} & R_1 L C_1 C_2 s^3 V_{\text{out}}(s) + \left(\frac{R_1 L C_1}{R_2} + L C_2 \right) s^2 V_{\text{out}}(s) \\ & + \left(R_1 C_1 + R_1 C_2 + \frac{L}{R_2} \right) s V_{\text{out}}(s) + \left(\frac{R_1}{R_2} + 1 \right) V_{\text{out}}(s) \quad . \end{aligned} \tag{A.13}$$

Solving the equation for $V_{\text{out}}(s)/V_{\text{in}}(s)$ results in the transfer function of the filter:

$$\begin{aligned} H(s) = \frac{V_{\text{out}}(s)}{V_{\text{in}}(s)} = {} & \Bigg[R_1 L C_1 C_2 s^3 + \left(\frac{R_1 L C_1}{R_2} + L C_2 \right) s^2 \\ & + \left(R_1 C_1 + R_1 C_2 + \frac{L}{R_2} \right) s + \left(\frac{R_1}{R_2} + 1 \right) \Bigg]^{-1} \end{aligned} \tag{A.14}$$

This transfer function has been used to determine the values of the components. The final set of values is given in table A.1.

Component	Value
R_1	$200\,\Omega$
C_1	$8.2\,\text{nF}$
L	$560\,\mu\text{H}$
C_2	$8.2\,\text{nF}$
R_2	$200\,\Omega$

Table A.1: Values of the baseband filter components.

B Soft-DSP instruction set and syntax

A program for the soft DSP has to comply to the following syntax:

```
<DSPProgram> ::= <Header> <Commands> <Footer>
<Header>     ::= 'HF_CMDS hf_cmds[] = {'
<Commands>   ::= <Command> | <Command> ',' <Commands>
<Command>    ::= '{' <CMD_ADDR> ',' <SRC> ',' <OP> ',' <TGT> ','
                 <ReadParam> ',' <ReadVar> ',' <WriteVar> ','
                 <START> ',' <INTEGRATOR_HOLD> ','
                 <LOOP_RESTART> ',' <COMMENT> '}'
<Footer>     ::= '};'
```

Each command must have a unique integer command address (<CMD_ADDR>) in the range of 0 to 127. Commands are executed in the order of their command addresses. If several commands use the same command address the last of these commands overwrites older ones.

For a detailed understanding of the DSP functions refer to fig. B.1.

Operations

ADD_V_K

Signed addition of <ReadVar> + <ReadParam> using saturation arithmetic according to eq. (6.4).

SUB_V_K

Signed subtraction of <ReadVar> − <ReadParam> using saturation arithmetic according to eq. (6.4).

SUB_K_V

Signed subtraction of <ReadParam> − <ReadVar> using saturation arithmetic according to eq. (6.4).

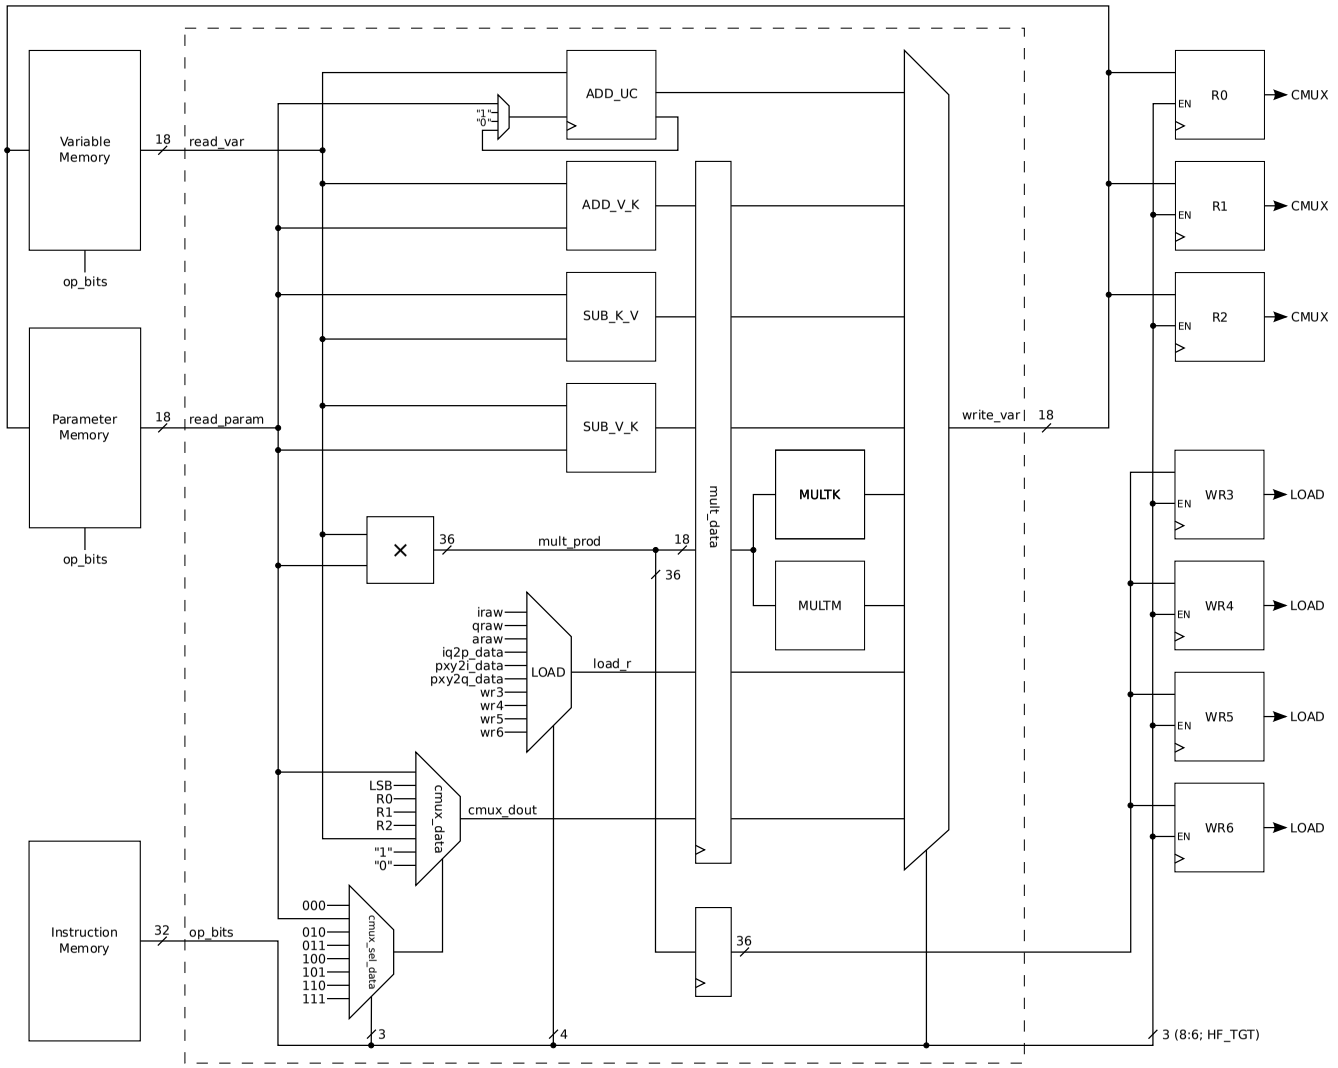

Figure B.1: Block diagram of the DSP including the data-paths for all commands.

<ADD_UC_SRC>	Output
ADD_K	ReadVar + parameter specified in <ReadParam>
SUB_K	ReadVar − parameter specified in <ReadParam>
ADD_OV	ReadVar + carry bit from last ADD_UC operation
ADD_LSB	ReadVar + LSB (0x00001)

Table B.1: Variants of the ADD_UC command.

ADD_UC

Add or subtract using non-saturation arithmetic. The variant of the command is specified by the ADD_UC_SRC field (see table B.1). The first operand is read from the variable specified by <ReadVar>. The source of the second operand depends on the variant of the command. The carry bit (overflow bit) is stored in OV and can be used until the next ADD_UC operation overwrites this register. The carry bit can be used as an operand by using ADD_OV as source. This e. g. allows to create counters with a multiple of 18 bits width.

MULTK

This command multiplies the 18 bit variable specified by <ReadVar> with an 18 bit parameter specified by <ReadParam>. The bits 34 to 17 of the product are used as the result. The most significant bit is dropped because it does not carry relevant information. This corresponds to a multiplication with a parameter in the interval $[-1.0, +1.0[$. This command uses saturation arithmetic.

MULTM

This command behaves like MULTK but it picks out the bits 26 to 9 of the product for the result. This corresponds to a multiplication with a parameter in the interval $[-256.0, +256.0[$. This command uses saturation arithmetic and is intended for the implementation of high-gain controllers.

<SRC>	Output
SRC_K	Value from parameter specified in <ReadParam>
SRC_KMUX	Select signal according to the three least significant bits of <ReadParam>:
	000: Value from parameter specified in <ReadParam>
	001: LSB (0x00001)
	010: Value from register R0
	011: Value from register R1
	100: Value from register R2
	101: Value from variable specified in <ReadVar>
	110: 1.0
	111: 0.0
SRC_R0	Value from register R0
SRC_R1	Value from register R1
SRC_R2	Value from register R2
SRC_V	Value from variable specified in <ReadVar>
SRC_EINS	1.0
SRC_NULL	0.0

Table B.2: Sources provided by the CMUX command.

CMUX

This command controls two multiplexers. Together they allow to select one of eight signals depending on the <SRC> field and (if SRC_KMUX is selected) on the value of a parameter (see table B.2).

OP_NOP

This command does nothing. It can be used to wait for some other operation to finish.

<HF_LOADSRC>	Output
IRAW	Most recent value of the I ADC
QRAW	Most recent value of the Q ADC
ARAW	Most recent value of the amplitude ADC
IQ2P	Phase result of the Cartesian-to-polar CORDIC
PXY2I	I result of the polar-to-Cartesian CORDIC
PXY2Q	Q result of the polar-to-Cartesian CORDIC
LOAD_WR3	The 18 most significant bits of the 36 bit register WR3
LOAD_WR4	The 18 most significant bits of the 36 bit register WR4
LOAD_WR5	The 18 most significant bits of the 36 bit register WR5
LOAD_WR6	The 18 most significant bits of the 36 bit register WR6

Table B.3: Sources supported by the LOAD command.

LOAD

The load command allows to load external values into a DSP variable. Table B.3 lists the available sources.

Load sources for parameters

For commands which operate on a parameter, <ReadParam> specifies which parameter to load. <ReadParam> can be the name of any parameter used in the control algorithm as well as one of the keywords listed in table B.4.

Note: After writing to register R0 to R2 it takes 3 clock cycles before the new value can be read back (insert two OP_NOPs between the commands).

Load sources and write targets for variables

For commands operating on a variable, <ReadVar> specifies which variable to load as an operand. <WriteVar> determines into which variable to write the result. The name of any variable used in the control algorithm as well as one of the special variables listed in table B.5 can be used.

<ReadParam>	Parameter to load
CNOP	Do not load a parameter (use 0 instead)
LOAD_R0	Load value of register R0
LOAD_R1	Load value of register R1
LOAD_R2	Load value of register R2
LOAD_R3	Load the 18 most significant bits of register WR3
LOAD_R4	Load the 18 most significant bits of register WR4
LOAD_R5	Load the 18 most significant bits of register WR5
LOAD_R6	Load the 18 most significant bits of register WR6

Table B.4: Sources provided by the CMUX command.

Special purpose registers

18 bit registers

If one of the targets SAVE_R0, SAVE_R1, and SAVE_R2 is specified, the computed result from the ALU is copied to a 18 bit register (R0/R1/R2) in addition to the variable specified by <WriteVar>. This is useful e. g. if two variables should be multiplied.

36 bit accumulating registers

If one of the targets SAVE_R3 to SAVE_R6 is specified, the computed result from the arithmetic logic unit is added to a 36 bit ("wide") register (WR3 to WR6). Additionally, the value is copied to the variable specified by <WriteVar>. Wide registers can be used to implement integrators or low-pass filters. The registers WR5 and WR6 use saturation arithmetic and can be used for amplitude or tuner controllers whereas the registers WR3 and WR4 use normal (overflow) arithmetic which is intended for phase controllers.

All wide registers stop accumulating if the controller is turned off by the operator or by the interlock. It is also possible to "freeze" their value by specifying INTEGRATOR_HOLD = HOLD_COND_1 instead of the default HOLD_DONT_HOLD.

<VarParam>	Variable to load
VNOP	Do not load a variable (use 0 instead)
IZ	Most recent value of the I ADC
QZ	Most recent value of the Q ADC
AZ	Most recent value of the amplitude ADC
IOUT	Value that is sent to the I DAC after the loop has finished
QOUT	Value that is sent to the Q DAC after the loop has finished
DG_OUT	Value that is sent to the demodulator gain DAC after the loop has finished
UG_OUT	Value that is sent to the modulator gain DAC after the loop has finished
LED_OUT	Value that is sent to the LEDs on the front panel after the loop has finished
XYP2IQ_X	First component of the input vector for the polar-to-Cartesian CORDIC
XYP2IQ_Y	Second component of the input vector for the polar-to-Cartesian CORDIC
XYP2IQ_P	Phase input value for the polar-to-Cartesian CORDIC
IQ2P_I	I component of the input vector for the Cartesian-to-polar CORDIC
IQ2P_Q	Q component of the input vector for the Cartesian-to-polar CORDIC

Table B.5: Special sources that can be used instead of normal variables.

Cartesian-to-polar CORDIC

This special block calculates the angle (phase) of an (I, Q) vector (VECTORING mode). The calculation is started with `<START>` = `S_IQ2P` and needs 27 clock cycles to complete[8]. The input values are taken from the variables `IQ2P_I` and `IQ2P_Q` whereas the result can be loaded by `<OP>` = `LOAD` and `<ReadParam>` = `IQ2P`.

Polar-to-Cartesian CORDIC

This block rotates a vector specified in Cartesian coordinates (X, Y) by a given angle φ (ROTATION mode). The input values are read from the variables `XYP2IQ_X`, `XYP2IQ_Y` and `XYP2IQ_P`. The calculation is started with `<START>` = `S_P2IQ` and needs 31 clock cycles to complete[9]. The output values I and Q can be loaded using the `LOAD` command with `<ReadParam>` = `<PXY2I>` or `<PXY2Q>`, respectively.

[8] 24 clock cycles for the CORDIC IP core and 3 additional clock cycles because of registers and pipelining.

[9] 28 clock cycles for the CORDIC IP core and 3 additional clock cycles because of registers and pipelining.

References

[1] A. Richter: Operational Experience at the S-DALINAC in *Proceedings of EPAC 1996* (Barcelona, Spain, 1996), p. 110–114.

[2] R. Eichhorn: *Optimierung des Strahltransportsystems und experimentelle Umsetzung verschiedener Methoden zur Gütemessung am S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 1999.

[3] C. Eckardt, T. Bahlo, P. Bangert, R. Barday, U. Bonnes, M. Brunken, C. Burandt, R. Eichhorn, J. Enders, M. Espig, C. Ingenhaag, J. Lindemann, M. Platz, Y. Poltoratska, M. Roth, F. Schneider, H. Schüßler, M. Wagner, A. Weber and B. Zwicker: The S-DALINAC Polarized Injector SPIN – Performance and Results in *Proceedings of PAC 2011* (New York City, New York, USA, 2011), p. 853–855.

[4] Y. Fritzsche: *Aufbau und Inbetriebnahme einer Quelle polarisierter Elektronen am supraleitenden Darmstädter Elektronenlinearbeschleuniger S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 2011.

[5] K. Sonnabend, D. Savran, J. Beller, M. Büssing, A. Constantinescu, M. Elvers, J. Endres, M. Fritzsche, J. Glorius, J. Hasper, J. Isaak, B. Löher, S. Müller, N. Pietralla, C. Romig, A. Sauerwein, L. Schnorrenberger, C. Wälzlein, A. Zilges and M. Zweidinger: *The Darmstadt High-Intensity Photon Setup (DHIPS) at the S-DALINAC*, Nucl. Inst. & Meth. A **640**, 6 (2011).

[6] D. Savran, M. Fritzsche., J. Hasper, K. Lindenberg, S. Müller, V. Y. Ponomarev, K. Sonnabend and A. Zilges: *Fine Structure of the Pygmy Dipole Resonance in* ^{136}Xe, Phys. Rev. Lett. **100**, 232501 (2008).

[7] O. Burda, N. Botha, J. Carter, R. W. Fearick, S. V. Förtsch, C. Fransen, H. Fujita, J. D. Holt, M. Kuhar, A. Lenhardt, P. von Neumann-Cosel, R. Neveling, N. Pietralla, V. Y. Ponomarev, A. Richter, O. Scholten, E. Sideras-Haddad, F. D. Smit and J. Wambach: *High-Energy-Resolution Inelastic Electron and Proton Scattering and the Multiphonon Nature of Mixed-Symmetry* 2^+ *States in* ^{94}Mo, Phys. Rev. Lett. **99**, 092503 (2007).

[8] M. Knirsch: *Konzeption, Aufbau und Erprobung eines hochauflösenden QCLAM-Elektronenspektrometers mit großem Raumwinkel und hoher Impulsakzeptanz am Elektronenbeschleuniger S-DALINAC*, dissertation D17, Technische Hochschule Darmstadt, 1991.

[9] P. von Neumann-Cosel, C. Rangacharyulu, A. Richter, G. Schrieder, A. Stascheck and S. Strauch: *Complete Spectroscopy of an Isobaric Analog Resonance with the (e,e'p) Reaction*, Phys. Rev. Lett. **78**, 2924 (1997).

[10] N. Ryezayeva, H. Arenhövel, O. Burda, A. Byelikov, M. Chernykh, J. Enders, H. W. Grießhammer, Y. Kalmykov, P. von Neumann-Cosel, B. Özel, I. Poltoratska, I. Pysmenetska, C. Rangacharyulu, S. Rathi, A. Richter, G. Schrieder, A. Shevchenko and O. Yevetska: *Measurement of the Reaction* ^{2}H(e,e') *at 180° close to the Deuteron Breakup Threshold*, Phys. Rev. Lett. **100**, 172501 (2008).

[11] K. Lindenberg: *Development and Construction of the Low-Energy Photon Tagger NEPTUN*, dissertation D17, Technische Universität Darmstadt, 2007.

[12] D. Savran, K. Lindenberg, J. Glorius, B. Löher, S. Müller, N. Pietralla, L. Schnorrenberger, V. Simon, K. Sonnabend, C. Wälzlein, M. Elvers, J. Endres, J. Hasper and A. Zilges: *The low-energy photon tagger NEPTUN*, Nucl. Inst. & Meth. A **613**, 232 (2010).

[13] P. Schardt: *Mikrowellenexperimente zum chaotischen Verhalten eines supraleitenden Stadionbillards und Entwicklung einer Einfangsektion am S-DALINAC*, dissertation D17, Technische Hochschule Darmstadt, 1995.

[14] T. Bahlo: *Konfiguration und Test eines Hochfrequenz-Chopper- und Buncher-systems für den S-DALINAC*, bachelor's thesis, Technische Universität Darmstadt, 2010 (unpublished).

[15] G. B. Steiner: *Strahldynamik-Simulation einer polarisierten Quelle für den S-DALINAC (SPIN)*, dissertation D17, Technische Universität Darmstadt, 2008.

[16] F. Hug: *Erhöhung der Energieschärfe des Elektronenstrahls am S-DALINAC durch nicht-isochrones Rezirkulieren*, dissertation D17, Technische Universität Darmstadt, in preparation.

[17] T. Rietdorf: *Entwurf und Realisierung einer variablen supraleitenden Hochfrequenz-Einkopplung für die Beschleunigungsstrukturen des supraleitenden Darmstädter Elektronenbeschleunigers S-DALINAC*, dissertation D17, Technische Hochschule Darmstadt, 1993.

[18] H.-D. Gräf: Experience with Control of Frequency, Amplitude and Phase in *Proc. of the 5th Workshop on RF Superconductivity*, (DESY, Hamburg, Germany, 1992), **1**, p. 317–333.

[19] D. Flasche: *Entwicklung und Aufbau eines modularen und rechnergesteuerten Kontrollsystems für die Hochfrequenzregelung des Darmstädter supraleitenden Elektronenbeschleunigers*, dissertation D17, Technische Hochschule Darmstadt, 1989.

[20] C. Burandt: *Aufbau eines 6-GHz-Systems für die Hochfrequenzregelung am S-DALINAC*, master's thesis, Technische Universität Darmstadt, 2010 (unpublished).

[21] M. Konrad, U. Bonnes, C. Burandt, R. Eichhorn and N. Pietralla: A Digital Low Level RF Control System for the S-DALINAC in *Proceedings of LINAC 2010* (Tsukuba, Japan, 2010), p. 268–270.

[22] M. Konrad, U. Bonnes, C. Burandt, J. Conrad, R. Eichhorn, J. Enders, P. Nonn and N. Pietralla: Design and Commissioning of a multi-frequency digital Low Level RF Control System in *Proceedings of IPAC 2011* (San Sebastián, Spain, 2011), p. 433–435.

[23] M. Konrad, U. Bonnes, C. Burandt, R. Eichhorn, J. Enders and N. Pietralla: A digital base-band RF Control System in *Proceedings of ICALEPCS 2011* (Grenoble, France, 2011), p. 82–85.

[24] M. Konrad, U. Bonnes, C. Burandt, R. Eichhorn, P. Nonn, J. Enders and N. Pietralla: *Digital base-band rf control system for the superconducting Darmstadt electron linear accelerator*, Phys. Rev. ST Accel. Beams **15**, 052802 (2012).

[25] M. Konrad, U. Bonnes, C. Burandt, R. Eichhorn, J. Enders, P. Nonn and N. Pietralla: Development of a Digital Low-level RF Control System for the

p-LINAC Test Stand at FAIR in *Proceedings of IPAC 2012* (New Orleans, Lousiana, USA, 2012), p. 3461–3463.

[26] O. Patalakha: *Design and implementation of a modular client/server control system for the S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 2006.

[27] M. Konrad, C. Burandt, J. Enders and N. Pietralla: Control System Studio Archiver with PostgreSQL Backend: Optimizing Performance and Reliability for a Production Environment in *Proceedings of PCaPAC 2012* (Kolkata, India, in press).

[28] C. Burandt, U. Bonnes, J. Enders, M. Konrad and N. Pietralla: SocketCAN Device Support for EPICS IOCs in *Proceedings of PCaPAC 2012* (Kolkata, India, in press).

[29] C. Burandt, U. Bonnes, J. Enders, F. Hug, M. Konrad and N. Pietralla: Status of the Migration of the S-DALINAC Accelerator Control System to EPICS in *Proceedings of PCaPAC 2012* (Kolkata, India, in press).

[30] J. Delayen: *Phase and Amplitude Stabilization of Superconducting Resonators*, Ph. D. thesis, California Institute of Technology, 1978.

[31] F. M. Gardner: *Phaselock techniques*, 3rd edn. (John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 2005).

[32] H. Padamsee, J. Knobloch and T. Hays: *RF Superconductivity for Accelerators* (Wiley-VCH, Weinheim, Germany, 2008).

[33] J. Lunze: *Regelungstechnik 1*, 8th edn. (Springer, Heidelberg, Germany, 2010).

[34] C. Ziomek and P. Corredoura: Digital I/Q Demodulator in *Proceedings of PAC 1995* (Dallas, Texas, USA, 1996), p. 2663–2665.

[35] L. Doolittle: Operational performance of the SNS LLRF interim system in *Proceedings of PAC 2003* (Portland, Oregon, USA, 2003), **3**, p. 1464–1466.

[36] S. N. Simrock: Digital low-level rf controls for future superconducting linear colliders in *Proceedings of PAC 2005* (Knoxville, Tennessee, USA, 2005), p. 515–519.

[37] J. Delayen, T. Allison, C. Hovater, J. Musson and T. Plawski: Development of a Digital Self-Excited Loop for Field Control in High-Q Superconducting Cavities in *Proceedings of SRF 2007* (Beijing, China, 2007), p. 230–234.

[38] S. N. Simrock: Considerations for the choice of the intermediate frequency and sampling rate for digital rf control in *Proceedings of EPAC 2006* (Edinburgh, Scotland, 2006), p. 1462–1464.

[39] U. Laier: *Güte der Beschleunigungsstrukturen und Entwicklung eines HF-Moduls am S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 2004.

[40] M. Platz: *Neuentwicklung und Aufbau einer Rechnerschnittstelle zur Ansteuerung der Geräte am S-DALINAC und Konzeption einer digitalen HF-Regelung*, dissertation D17, Technische Universität Darmstadt, 2004.

[41] A. Araz: *Aufbau und Erprobung einer digitalen HF-Regelung und Aufbau eines modularen Messsystems zur Energiestabilisierung für den S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 2009.

[42] A. Araz, U. Bonnes, R. Eichhorn, F. Hug, M. Konrad, M. Platz, A. Richter and R. Stassen: *3 GHz digital rf control at the superconducting Darmstadt electron linear accelerator: First results from the baseband approach and extensions for other frequencies*, Phys. Rev. ST Accel. Beams **13**, 082801 (2010).

[43] T. Kürzeder: *Neueinstellung der Feldglattheit der supraleitenden Beschleunigungsstrukturen für den S-DALINAC*, diploma thesis, Technische Universität Darmstadt, 2006.

[44] J. Serrano: FPGA technology in instrumentation and related tools in *Proceedings of DIPAC 2005* (Lyon, France, 2005), p. 132–136.

[45] V. Ayvazyan, G. Petrosyan, K. Rehlich, S. N. Simrock and P. Vertrov: Hardware and software design for the DSP based LLRF control in *Proceedings of PCaPAC 2005* (Tsukuba, Japan, 2005), p. WEP37.

[46] W. Cichalewski, W. Jalmuzna, D. Makowski and A. Napieralski: Multipurpose LLRF field controller for various superconducting cavity applications in *Proceedings of ICALEPCS 2009* (Kobe, Japan, 2009), p. 242–244.

[47] H.-S. Kim, H.-J. Kwon, K.-T. Seol and Y.-S. Cho: LLRF control system using a commercial board in *Proceedings of LINAC 2008* (Victoria, British Columbia, Canada, 2008), p. 1057–1059.

[48] The MathWorks Inc.: *MATLAB*, `http://www.mathworks.de/products/matlab/`.

[49] E. McCune: Taming the quadrature modulator in *Proceedings of the Wireless Symposium/Portable by Design Conference* (Penton Media, Inc., San José, California, USA, 2000).

[50] Analog Devices Inc.: *ADL5375 datasheet Rev. B*, 2011.

[51] X. Huang and M. Caron: Gain/phase imbalance and DC offset compensation in quadrature modulators in *IEEE International Symposium on Circuits and Systems, 2002. ISCAS 2002.* (Scottsdale, Arizona, USA, 2002), **4**, p. 811–814.

[52] S. Sievers: *Verbesserung der Güte von Beschleunigungsstrukturen im S-DALINAC*, dissertation D17, Technische Universität Darmstadt, in preparation.

[53] D. E. Thomas and P. R. Moorby: *The Verilog hardware description language (5^{th} ed.)* (Kluwer Academic Publishers, Dordrecht, Netherlands, 2002).

[54] Xilinx Inc.: *ISE WebPACK*, `http://www.xilinx.com/ise/logic_design_prod/webpack.htm`.

[55] Andrew Rogers et al.: *xc3sprog*, `http://sourceforge.net/projects/xc3sprog/`.

[56] C. Röder: *Entwicklung einer Software zur Fernprogrammierung der Geräte am S-DALINAC und Implementierung eines zugehörigen Anwendungsprogramms*, bachelor's thesis, Technische Universität Darmstadt, 2008 (unpublished).

[57] S. Williams: *Icarus Verilog*, `http://www.icarus.com/eda/verilog/`.

[58] Pragmatic C Software Corp.: *GPL Cver*, `http://www.pragmatic-c.com/gpl-cver/`.

[59] D. A. Patterson and J. L. Hennessy: *Computer Organization and Design*, 4th edn. (Morgan Kaufmann Publishers, Burlington, Massachusetts, USA, 2009).

[60] J. Volder: *The CORDIC Trigonometric Computing Technique*, IRE Transactions on Electronic Computers **EC-8**, 330 (1959).

[61] A. Neumann, W. Anders, O. Kugeler and J. Knobloch: *Analysis and active compensation of microphonics in continuous wave narrow-bandwidth superconducting cavities*, Phys. Rev. ST Accel. Beams **13**, 082001 (2010).

[62] T. Allison, J. Delayen, C. Hovater, J. Musson and T. Plawski: A digital self excited loop for accelerating cavity field control in *Proceedings of PAC 2007* (Albuquerque, New Mexiko, USA, 2007).

[63] C. W. Leemann, D. R. Douglas and G. A. Krafft: *The Continuous Electron Beam Accelerator Facility: CEBAF at the Jefferson Laboratory*, Annu. Rev. Nucl. Part. S. **51**, 413 (2001).

[64] T. Walcher, R. Frey, H.-D. Gräf, E. Spamer and H. Theissen: *High resolution electron scattering facility at the Darmstadt linear accelerator (DALINAC): II. Beam transport system and spectrometer (energy-loss system)*, Nucl. Inst. & Meth. A **153**, 17 (1978).

[65] A. W. Lenhardt, U. Bonnes, O. Burda, P. von Neumann-Cosel, M. Platz, A. Richter and S. Watzlawik: *A silicon microstrip detector in a magnetic spectrometer for high-resolution electron scattering experiments at the S-DALINAC*, Nucl. Inst. & Meth. A **562**, 320 (2006).

[66] *Geant4. Monte Carlo particle tracking software*, `http://geant4.org/`.

[67] F. Neumeyer: *Entwicklung eines interaktiven Rechenprogramms zur Untersuchung von elektromagnetischen Übergängen in Atomkernen am S-DALINAC*, diploma thesis, Technische Hochschule Darmstadt, 1993.

[68] P. Nonn: *Entwicklung einer gepulsten HF-Regelung für den FAIR p-Linac*, dissertation D17, Technische Universität Darmstadt, in preparation.

[69] A. Krugmann: *Entwicklung einer lokalen Target-Steuerung am S-DALINAC*, bachelor's thesis, Technische Universität Darmstadt, 2006 (unpublished).

[70] M. Hertling: *Aufbau einer zentralen relationalen Datenbank und Entwicklung einer graphischen Benutzeroberfläche für den S-DALINAC*, dissertation D17, Technische Universität Darmstadt, 2008.

[71] A. Kuhl: *Aufbau einer Kanonensteuerung auf HV-Potential und Entwicklung einer Prozessdatenerfassung mit dem QM07 für den S-DALINAC*, diploma thesis, Technische Universität Darmstadt, 2009.

[72] C. Burandt: *Entwicklung einer EPICS-basierten Rechneranbindung für die gepulste Hochfrequenzregelung am Teststand für den p-Linac*, dissertation D17, Technische Universität Darmstadt, in preparation.

[73] T. Schösser: *Entwicklung eines EPICS-basierten Steuerungskonzepts für den S-DALINAC*, dissertation D17, Technische Universität Darmstadt, in preparation.

[74] L. Dalesio, M. Kraimer and A. Kozubal: EPICS architecture in *Proceedings of ICALEPCS 1991* (Tsukuba, Japan, 1991).

[75] L. R. Dalesio, J. O. Hill, M. Kraimer, S. Lewis, D. Murray, S. Hunt, W. Watson, M. Clausen and J. Delesio: *The experimental physics and industrial control system architecture: past, present, and future*, Nucl. Inst. & Meth. A **352**, 179 (1994).

[76] *Experimental Physics and Industrial Control System*, `http://www.aps.anl.gov/epics/`.

[77] M. Sekoranja, S. Hunt and A. Luedeke: Visual DCT - Visual EPICS Database Configuration Tool in *Proceedings of ICALEPCS 2001* (San José, California, USA, 2001), p. 540–542.

[78] *The SocketCAN project*, `http://developer.berlios.de/projects/socketcan/`.

[79] D. Schirmer, E. Kasel, B. Keil and D. Zimoch: Standardization of the DELTA control system in *Proceedings of ICALEPCS 1999* (Trieste, Italy, 1999).

[80] *StreamDevice 2 device support*, `http://epics.web.psi.ch/software/streamdevice/`.

[81] *Telnet Com Port Control Option (RFC 2217)*, http://tools.ietf.org/html/rfc2217.

[82] J. Sinclair: *Extensible Display Manager*, http://ics-web.sns.ornl.gov/edm/.

[83] K. Evans: An overview of MEDM in *Proceedings of ICALEPCS 1999* (Trieste, Italy, 1999), p. 466–468.

[84] K. Kasemir: Control System Studio applications in *Proceedings of ICALEPCS 2007* (Knoxville, Tennessee, USA, 2007), p. 692–694.

[85] M. Clausen, J. Hatje, M. Moeller and H. Rickens: Control System Studio integrated operating, configuration and development in *Proceedings of ICALEPCS 2009* (Kobe, Japan, 2009), p. 667–669.

[86] *Control System Studio*, http://cs-studio.sourceforge.net/.

[87] *TAco Next Generation Objects (TANGO)*, http://www.tango-controls.org/.

[88] *Three-fold Integrated Networking Environment (TINE)*, http://adweb.desy.de/mcs/tine/.

[89] J. Hatje, M. Clausen, C. Gerke, M. Moeller and H. Rickens: Control System Studio (CSS) in *Proceedings of ICALEPCS 2007* (Knoxville, Tennessee, USA, 2007), p. 37–39.

[90] O. Gruber., B. J. Hargrave, J. McAffer, P. Rapicault and T. Watson: *The Eclipse 3.0 platform: Adopting OSGi technology*, IBM Systems Journal **44**, 289 (2005).

[91] X. Chen and K. Kasemir: BOY, a modern graphical operator interface editor and runtime in *Proceedings of PAC 2011* (New York City, New York, USA, 2011), p. 1404–1406.

[92] *Osqoop. The open oscilloscope*, http://gitorious.org/osqoop.

[93] SigBlips: *Baudline. A time-frequency browser*, http://www.baudline.com.

[94] A. J. Kozubal, L. R. Dalesio, J. O. Hill and D. M. Kerstiens: *A State Notation Language for Automatic Control*, Los Alamos National Laboratory report LA-UR-89-3564, 1989.

[95] A. J. Kozubal, D. M. Kerstiens and R. M. Wright: *Experience with the State Notation Language and Run-Time sequencer*, Nucl. Inst. & Meth. A **352**, 411 (1994).

[96] *State Notation Language and Sequencer*, `http://www-csr.bessy.de/control/SoftDist/sequencer/`.

[97] K. Kasemir and L. Dalesio: Data archiving in EPICS in *Proceedings of ICALEPCS 1999* (Trieste, Italy, 1999), p. 463–465.

[98] M. Giacchini, L. Giovannini, G. Bassato, M. Montis, J. Vásquez, G. Prete, A. Andrighetto, J. Jugo and M. Campo: *EPICS Latest Technology: Control System Studio and RDB Archiver in Production in the SPES Project*, LNL annual report, 2010.

[99] K. U. Kasemir: Control System Studio (CSS) Data Browser in *Proceedings of PCaPAC 2008* (Ljubljana, Slovenia, 2008), p. 99–101.

[100] D. A. Patterson, G. Gibson and R. H. Katz: A Case for Redundant Arrays of Inexpensive Disks (RAID) in *Proceedings of SIGMOD 1988* (Chicago, Illinois, USA, 1988), p. 109–116.

[101] S. Lüders: Securing control systems against cyber attacks in *Proceedings of PAC 2009* (Vancouver, British Columbia, Canada, 2009).

[102] S. Lüders: Summary of the 3$^{\text{rd}}$ control system cyber-security (CS)2/HEP workshop in *Proceedings of ICALEPCS 2011* (Grenoble, France, 2011).

[103] K. Evans: The EPICS process variable gateway – version 2 in *Proceedings of ICALEPCS 2005* (Geneva, Switzerland, 2005).

[104] K. Evans and M. Smith: Experience with the EPICS PV gateway at the APS in *Proceedings of PAC 2005* (Knoxville, Tennessee, USA, 2005), p. 3621–3623.

[105] *Nagios IT infrastructure monitoring software*, `http://www.nagios.org/`.

[106] M. Giacchini, N. Richtes and R. Lange: LivEPICS: An EPICS Linux Live CD Nagios equipped in *Proceedings of ICALEPCS 2007* (Knoxville, Tennessee, USA, 2007), p. 161–162.

[107] *Apache Subversion version control software*, `http://subversion.apache.org/`.

[108] *S-DALINAC Subversion source code repository*, `svn://b1.ikp.physik.tu-darmstadt.de/acs/`.

[109] *Mantis bug tracking software*, `http://www.mantisbt.org/`.

[110] *NSLS-II controls package repository*, `http://epics.nsls2.bnl.gov/debian/`.

[111] *FAI. The Debian Fully Automatic Installation project*, `http://fai-project.org`.

Danksagung

An dieser Stelle möchte ich mich bei allen bedanken, die zum Gelingen dieser Arbeit beigetragen haben.

Zuerst gilt mein Dank Herrn Professor Dr. Dr. h. c. Norbert Pietralla, der mir diese Arbeit in seiner Arbeitsgruppe sowie die Teilnahme an diversen internationalen Konferenzen ermöglicht hat. Das mir entgegengebrachte große Vertrauen sowie seine mutige Entscheidung für die Umstellung des Kontrollsystems auf EPICS haben die Vielfalt der in dieser Arbeit beschriebenen Softwarefunktionen erst ermöglicht.

Professor Dr. Harald Klingbeil danke ich für die Übernahme des Korreferats.

Ich danke Dr. Ralf Eichhorn und Florian Hug für ihre Unterstützung in allen Fragen der Beschleunigerphysik. Ihre Erfahrung im Umgang mit dem S-DALINAC hat die Gestaltung des Kontrollsystems wesentlich beeinflusst.

Dr. Markus Platz und Jonny Birkhan schulde ich Dank für ihr Vertrauen und die große Freiheit, die ich bei der Umgestaltung des Beschleunigernetzwerks und der Nutzung der Virtualisierungsinfrastruktur genießen durfte.

Herzlich bedanken möchte ich mich bei Ralph Lange (BESSY) für seine umfangreiche Hilfe beim Einstieg in EPICS.

Einen besonderen Dank möchte ich Uwe Bonnes und den Mitarbeitern der Elektronikwerkstatt aussprechen, ohne deren hervorragende Hardware und ständige Hilfsbereitschaft dieses Projekt nicht möglich gewesen wäre.

Allen Mitgliedern der Beschleunigergruppe danke ich für ihr Engagement und die angenehme Arbeitsatmosphäre. Vor allem die intensive und vertrauensvolle Zusammenarbeit mit Christoph Burandt war eine große Bereicherung.

Andreas Krugmann und den anderen Mitgliedern der Spektrometergruppe danke ich für ihre Hilfe bei der Durchführung der Messungen am 169°-Spektrometer.

Ein herzlicher Dank gilt Stefanie und Arne für ihr Verständnis und ihre Unterstützung während der Entstehung dieser Arbeit sowie meinen Eltern für ihre stete Hilfsbereitschaft.

Die vorliegende Arbeit wurde gefördert durch Mittel aus dem DFG Sonderforschungsbereich 634.

Der Lebenslauf ist in der Online-Version
aus Gründen des Datenschutzes nicht enthalten.

Eidesstattliche Erklärung:

Hiermit erkläre ich, dass ich die vorliegende Dissertation selbstständig verfasst, keine anderen als die angegebenen Hilfsmittel verwendet und bisher noch keinen Promotionsversuch unternommen habe.

Darmstadt, im Februar 2013