

CERN's Super Proton Synchrotron 200 MHz cavity regulation upgrade: Modeling, design optimization, and performance estimation

T. Mastoridis 

California Polytechnic State University, San Luis Obispo, California 93407, USA

P. Baudrenghien

CERN, Geneva 1211, Switzerland



(Received 4 October 2021; accepted 2 February 2022; published 17 February 2022)

CERN's Super Proton Synchrotron (SPS) accelerates protons to 450 GeV/ c and transfers them into the Large Hadron Collider (LHC). It is currently one of the limiting factors in increasing the beam intensity and thus the luminosity of the LHC. As part of the LHC injectors upgrade project, the SPS 200 MHz rf system has been modified during the CERN long shutdown 2 (January 2019–April 2021), resulting in a new layout containing two additional cavities. The goal is to improve longitudinal stability required for the planned doubling of the beam intensity for the high luminosity LHC. In parallel with the upgrade of the high-power rf, a new low-level rf (LLRF) system has been designed, including a new cavity field regulation system. This work presents a model of the beam-rf interaction which includes a detailed representation of the LLRF controlling the cavity. This model is used to determine the optimal LLRF design for maximum loop stability and beam loading compensation. Finally, the performance of the upgraded LLRF is estimated.

DOI: [10.1103/PhysRevAccelBeams.25.021002](https://doi.org/10.1103/PhysRevAccelBeams.25.021002)

I. MOTIVATION

The protons accelerated in the LHC first pass through a chain of accelerators, the last of which is the Super Proton Synchrotron (SPS). The SPS accelerates protons from 26 to 450 GeV/ c before they are injected into the LHC. Determining optimal operational parameters for the SPS low-level rf (LLRF) loops is not straightforward but can have a significant impact on the SPS and the LHC. Currently, the high uncompensated transient beam loading leads to significant phase variation along the SPS bunch train, which in turn causes high beam losses during both the beam injection into the SPS and the extraction to the LHC [1]. It is also a problem for the phasing of the harmonic cavities (800 MHz) required for beam stabilization through Landau damping [2]. Finally, it can drive longitudinal instabilities [3]. These issues will be amplified with the almost double beam intensity planned for the high-luminosity LHC (HL-LHC).

Until 2018 the SPS had four 200 MHz cavities, of the traveling wave type, among which two consisted of five sections (55 cells) and two consisted of four sections

(44 cells) [4]. This configuration was modified during long shutdown 2 (LS2, January 2019–April 2021), resulting in a new layout containing four three-section cavities and two four-section. The goal is to improve longitudinal stability required by the planned doubling of the beam intensity for the high luminosity LHC (HL-LHC) [5]. In parallel with the upgrade of the HPRF, a new LLRF system has been designed [6] with improved transient beam loading compensation. Figures 1 and 2 in [6] provide an overview of the system.

This work presents simulations implemented in MATLAB and SIMULINK, which model the beam-rf interaction. The simulations can then be used to determine the optimal design and operational parameters, to estimate the performance improvement with the proposed SPS HPRF and LLRF upgrades, to identify possible limitations so that the hardware and firmware designers can make the necessary adjustments, and to develop algorithms for the optimal configuration and operation of the upgraded SPS rf.

Section II describes the SPS rf and its LLRF cavity field regulation. Section III presents the cavity models in the frequency and time domains. Sections IV and V introduce the SPS one-turn delay feedback (OTFB) and feedforward systems. Section VI summarizes the MATLAB/SIMULINK models. Section VII presents the design studies and the performance estimates for the upgraded LLRF ignoring transmitter bandwidth and power limitations. The performance with the upgraded LLRF and the optimal configurations including hardware limitations, is presented in

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Sec. VIII. Finally, Sec. IX summarizes this work and presents the performance estimates for the HL-LHC beam.

II. SYSTEM DESCRIPTION

The SPS rf system is composed of two subsystems operating at 200 and 800 MHz, each including a LLRF system. The main rf system at 200 MHz is used for acceleration, whereas the 800 MHz (harmonic) system is used to stabilize the beam via bunch shortening or lengthening (Landau damping). Models of both systems have been developed, but since the transient beam loading compensation is dominated by the 200 MHz system, this work presents the 200 MHz system model and studies. In addition, the 200 MHz LLRF has just been upgraded and is being commissioned in 2021. A completely redesigned LLRF has been implemented.

Each cavity has a dedicated power amplifier (TX), either tube based or solid state, and a LLRF system. Each LLRF includes a feedback and a feedforward system, presented in Secs. IV and V, respectively. Both systems act on the TX drive, but the feedforward samples the beam via a pickup, whereas the feedback samples the cavity voltage, as shown in Fig. 1.

III. I/Q MODEL OF A TRAVELING WAVE CAVITY

The SPS traveling wave cavities (TWC) have been presented in several past publications [4] and their accelerating properties are well known. In this publication, we introduce the TWC response in the Cartesian representation to match the LLRF implementation. This formalism is used in the derivation of an optimal feedforward algorithm in Sec. V and Appendix B, as well as in the time-domain simulations of the SPS TWC and LLRF presented in Sec. VI.

The voltage in the TWC results from generator current I_g and rf component of beam current I_b ,

$$V(f) = Z_g(f)I_g(f) + Z_b(f)I_b(f). \quad (1)$$

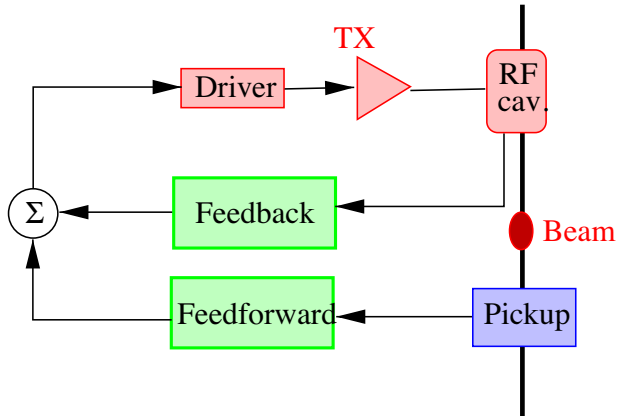


FIG. 1. SPS rf and LLRF block diagram.

We use the notation

$$\text{sinc}(x) \triangleq \frac{\sin(\pi x)}{\pi x} \quad (2)$$

In [4], the impedance Z_g is (using our notations)

$$Z_g(f) = R_1 \text{sinc}[\tau(f - f_0)].$$

In the time domain, this corresponds to a symmetric rectangular pulse lasting for time τ , modulated at f_0 , the cavity center frequency in Hz. Such a response is non-causal. To make it causal, we add the $\tau/2$ delay (τ is the cavity filling time compensated for the particle velocity [7]). We also add the image resonance at $-f_0$ to preserve Hermitian symmetry. These manipulations are important as we wish to derive real-valued causal impulse responses for use in the time-domain simulations. The two impedances are [7]

$$Z_g(f) = R_1 \{ \text{sinc}[\tau(f - f_0)] e^{-j\pi\tau(f - f_0)} + \text{sinc}[\tau(f + f_0)] e^{-j\pi\tau(f + f_0)} \} \quad (3)$$

$$Z_b(f) = -R_2 \left\{ \{ \text{sinc}[\tau(f - f_0)] \}^2 + \{ \text{sinc}[\tau(f + f_0)] \}^2 - j \frac{1 - \text{sinc}[2\tau(f - f_0)]}{\pi\tau(f - f_0)} - j \frac{1 - \text{sinc}[2\tau(f + f_0)]}{\pi\tau(f + f_0)} \right\}. \quad (4)$$

The parameters are

$$\begin{aligned} \tau &= \frac{NL}{|v_g|} \left(1 - \frac{v_g}{v} \right) \\ R_1 &= NL \sqrt{\frac{Z_0 r_2}{2}} \\ R_2 &= (NL)^2 \frac{r_2}{8}. \end{aligned} \quad (5)$$

In the above equations, L is the cell length, N is the number of interacting cells, v_g is the group velocity of the accelerating traveling wave in the cavity, which must be taken negative for a backward TWC as it is measured in the direction of the beam motion, r_2 is the series impedance [4,7], Z_0 is the characteristic impedance (50 ohm) and v is the particle velocity. The particle velocity varies only slightly during acceleration and in the simulation, we take the value at synchronism, which is when the rf frequency is equal to the cavity center frequency f_0 . Table I lists the basic cavity parameters.

Table II gives the main parameters of the three- and four-section cavities.

Figures 2 and 3 show the impedances for three-section and four-section cavities. Note that the impedance driven

TABLE I. TABLE I. Cavity data.

Cell length L in mm	374
Number of interacting cells N (three sections/four sections)	32/43
Group velocity v_g/c	-0.0946
Particle velocity at synchronism v	0.999138 c
Cavity center frequency f_0 in Hz (design value)	200.222 MHz
Series impedance r_2 (Ω/m^2)	27100

TABLE II. Cavity parameters.



Parameter	Three-sections	Four-sections
Cavity filling time τ (ns)	462	621
Resistance driven by generator R_1 (Ω)	9851	13237
Resistance driven by beam R_2 (Ω)	485202	876112

by the generator (Z_g) is always real, while the impedance driven by the beam (Z_b) has both a real and an imaginary part (using the accelerator physics definition of Z_b , the ratio of the beam induced voltage to the beam current) [8]. Due to the nonlinearity of the dispersion diagram, the measured zeros of the cavity response are not exactly equispaced as in

the above equations and figures. But we will see in Sec. VIII that compensation is not possible much beyond the first cavity zeros.

It is interesting to compare the impedance driven by the generator Z_g and the impedance driven by the beam Z_b . For a single-cell standing wave cavity (SWC) these two are identical (within a scaling factor), and it is, in theory, possible to compensate the beam induced voltage via a correction to the generator current. This is not the case for TWCs [8]. The impedance driven by the generator Z_g is zero on sidebands spaced by the inverse cavity filling time τ (in the linear dispersion diagram approximation). At these frequencies the real part of the impedance driven by the beam Z_b is also zero, but its imaginary part is not. The beam will induce some reactive voltage at these frequencies *and there is no way to compensate it via the generator*. On the other hand, perfect correction is possible for the resistive voltage induced by the beam, since the zeros are common for impedance driven by the generator and the beam.

The signals in the TWC chain are inherently band limited to a few MHz around the cavity center frequency. The LLRF therefore uses I/Q demodulation with a local oscillator (LO) at a frequency f_{LO} close to the cavity center frequency, followed by digitization of the narrow band I/Q pairs at 125 mega samples per second (MS/s) [6].

FIG. 2. Three-section cavity. Left: $Z_g(f)$ after compensation of the $\tau/2$ delay. Right: real (red) and imaginary parts of $Z_b(f)$.FIG. 3. Four-section cavity. Left: $Z_g(f)$ after compensation of the $\tau/2$ delay. Right: real (red) and imaginary parts of $Z_b(f)$.

In addition, the LLRF signal processing operates in the time domain on successive samples at the 125 MS/s rate. In [7], the impedances in Eqs. (3) and (4) are thus demodulated and inverse Fourier transformed to get the I/Q model of a TWC in the time domain. The impulse responses in the time domain are presented in Appendix A.

IV. THE FEEDBACK

The amplifiers are located on the surface while the cavities are in the accelerator tunnel, resulting in a loop delay of more than $2 \mu\text{s}$ (LLRF, TX, cabling to cavity, cavity response, and cabling back to LLRF). A direct rf feedback would therefore have limited bandwidth (less than 100 kHz, to be compared to the 43 kHz revolution frequency) and could only compensate the beam loading transient on the first two revolution frequency bands. The OTFB (Fig. 1) was designed to cope with this long loop delay. It is a feedback around the cavity/amplifier in which the loop delay has been intentionally extended to one exact turn (around $23 \mu\text{s}$), and the gain is limited to narrow frequency bands around the revolution harmonics [9]. It reduces the beam loading, including the transients caused by the gaps in beam current, thereby equalizing the bunch parameters (length) and increasing the longitudinal coupled-bunch instability threshold. First introduced in the early 1980s for the SPS [10] it has since been installed on many machines, sometimes as a complement to a direct rf feedback [11–13]. Three main elements are present in the processing performed in an OTFB: a filter $H_{\text{cav}}(f)$ to cope with the sign inversions in the impedance Z_g (explanations below), a comb filter $H_{\text{OTFB}}(f)$ matching the revolution frequency harmonics, and a delay element to properly match the corrective action in the cavity rf field with the

next passage of the same beam portion. The three blocks are shown in Fig. 4.

Since the cavity impedance driven by the generator has multiple zeros, the feedback response should also be zero at those frequencies, otherwise it would strongly drive the TX. In addition, the sign of the cavity impedance changes at those frequencies. The feedback sign should therefore also change for stability. This is achieved by the $H_{\text{cav}}(f)$ filter shown in Fig. 5. The first zero of this filter matches that of the cavity, so that the feedback sign tracks that of the cavity. The open loop gain is *significantly* lower than unity for higher frequencies though, so the rest of the zeros do not have to match. The bandwidth of the main lobe as well as the gain at the first sidelobe can be modified. Their effect on performance and stability are evaluated later in this paper.

The OTFB has a large gain only at the revolution frequency harmonics, but very low gain elsewhere. The z-transform of the simplest OTFB filter is given by

$$H_{\text{OTFB}}(z) = (1 - \alpha) \frac{1}{1 - \alpha z^{-N}}, \quad (6)$$

where N corresponds to the sampling periods in one turn and the parameter α controls the filter bandwidth around each revolution harmonic. The total loop delay is adjusted to be exactly equal to one turn; therefore, the phase shift is zero at all revolution harmonics. Details on the selection of parameters values for the OTFB can be found in [14].

V. THE FEEDFORWARD

The feedforward (Fig. 1) follows a very simple principle: We measure the rf component of the beam current with a pickup, and process it to generate a generator drive that




FIG. 4. SPS feedback block diagram.




FIG. 5. $H_{\text{cav}}(f)$ filter response and phase for three- and four-section cavity.

compensates the beam induced voltage. From Eq. (1), we get

$$Z_g(f)I_g(f) = -Z_b(f)I_b(f) \quad (7)$$

with the generator current being a processed version of the beam current

$$I_g(f) = H_{ff}(f)I_b(f). \quad (8)$$

The H_{ff} filter represents the processing in the LLRF. The limited TX bandwidth is not included. Merging the above two equations, we get the ideal feedforward transfer function

$$Z_g(f)H_{ff}(f) = -Z_b(f). \quad (9)$$

The real part of the beam loading impedance can be compensated exactly. We can ignore the $\tau/2$ delay since the synchrotron tune Q_s is smaller than 0.02 in the SPS. The beam current changes very little from one turn to the next. The feedforward therefore measures on a turn and corrects on the next turn. The delay $\tau/2$ can thus be subtracted from the full revolution period. From Eqs. (3) and (4) we get

$$\text{Re}[H_{ff}(f)] = \frac{R_2}{R_1} \{\text{sinc}[\tau(f - f_0)] + \text{sinc}[\tau(f + f_0)]\}. \quad (10)$$

In the time domain, and after I/Q demodulation, this corresponds to an even-symmetric impulse response, a scaled version of the responses shown on Figs. 25 and 26 (Appendix A). The imaginary part cannot be fully compensated. Z_g is zero at frequencies where the cavity response is zero, while the imaginary part of Z_b is nonzero at these frequencies. The design of the feedforward filter must therefore be optimized for *some* beam current signal. When used as an LHC injector, the SPS beam current consists of a series of 1.8 μs long batches separated by empty gaps. The optimization is therefore done for a *beam current step* and is described in detail in Appendix B.

VI. MATLAB/SIMULINK MODEL

A model of the SPS TWC 200 MHz cavities and the proposed LLRF was created in MATLAB and SIMULINK. The SIMULINK model reproduces the system's architecture, as shown in Fig. 1. All digital components in the feedback/feedforward are replicated exactly as presented in Secs. IV and V. Analog components, the transmitter (TX), and the cavity are modeled by digital filters whose frequency response is identical or very similar to the real elements.

In Secs. VI and VII an ideal TX is used. The finite TX bandwidth and saturation are introduced in Sec. VIII.

The finite bandwidth is modeled as a Butterworth filter with the same 3-dB and 15-dB bandwidth as the existing SPS transmitters. Two kinds of transmitters are used in the upgraded SPS: four existing tetrode transmitters on the three-section cavities, and two new solid-state amplifiers on the four-section cavities. The tetrode transmitters have a 3 dB bandwidth of 1.5 MHz. No measurements of the new transmitter frequency response at the required HL-LHC power level are currently available, but they will definitely have a wider bandwidth than the existing ones. Therefore, they have been modeled with the same narrower bandwidth, providing a conservative estimate of their performance. The TX saturation is modeled as a threshold at the peak TX power level.

The LHC beam consists of batches of up to 72 bunches, spaced by 25 ns. The simulation clock is set to the inverse of the 25 ns bunch spacing. Each bunch in the beam is modeled as a rigid macroparticle that tracks the cavity phase *exactly* at each turn to simplify the beam induced voltage comparison with various LLRF settings. As a result, the bunch tracks the *perturbed* bucket (due to beam loading, different for each bunch), but with the stable phase offset (identical for all bunches). Therefore, the simulation correctly estimates bunch phase variation due to transient beam loading but does not include longitudinal dynamics (synchrotron oscillation). The results apply to static conditions (stable circulating beams) but cannot be extrapolated to injection transients with mismatched buckets. This simplification is appropriate since we wish to apply the simulations to estimate the transient beam loading compensation when the beam is extracted to the LHC.

The model was validated by comparison with theoretical expressions, as well as with data from the SPS. The measurements were performed with the pre-LS2 SPS rf system, two four-section cavities set to 1 MV and the two five-section cavities at 1.25 MV. Data were acquired with both the feedback and feedforward systems on, or with one of them off. The operational gain settings were used for these systems. A 72-bunch batch was injected. The figure of merit is the beam phase variations along the batch, caused by the uncompensated beam loading. As mentioned above, such a variation creates losses at transfer to the LHC and causes bunch-by-bunch phase mismatch with the harmonic 800 MHz cavities. Figure 6 includes some of the benchmark measurements, showing good agreement with the model. The small differences are due to noise added from the demodulation of the cavity antenna signal. In addition, the measurement sampling is not synchronized with the rf buckets, and therefore, high frequency peaks are undersampled. Many measurements in various different feedback and feedforward settings were conducted to fully validate the simulations.

The validating measurements showed a small variation in the beam phase over the batch when both the feedback and feedforward systems were on, but a substantial variation




FIG. 6. Measured and simulated total cavity voltage phase, with feedforward off (left) and feedforward on (right). Feedback is on for both figures.

when one of them was off. This conclusion was confirmed with extensive simulation studies.

It should be noted that the old SPS feedforward designed in the late 1990s was injecting significant rf noise. The machine was therefore operated with the feedforward off in the last couple of years of LHC run 2 (2017, 2018), to reduce noise. Significant losses have been experienced as a result, both at the injection of protons into the SPS and at the extraction to the LHC. Therefore, it was essential to use these simulations to determine LLRF configurations that would include the feedforward and that would minimize uncompensated beam loading.

VII. LLRF OPTIMAL CONFIGURATION

The model was then used to study the beam loading compensation for various possible feedback schemes and determine the optimal configuration for the SPS upgrade. The beam phase variation over a batch of 72 bunches in the SPS is used as a performance metric in this work, since the transient beam loading compensation is of particular importance in these studies.

A. Baseline design

The baseline design combines elements of the existing LLRF architecture, upgrades based on the LHC experience, and design choices that would simplify the commissioning and operation of the system. In particular, the feedback $H_{\text{cav}}(f)$ filter has *very little* gain past the first cavity zero (more than 20 dB lower than the main band). The feedback gain is about 40 linear and the single-sided bandwidth of the peaks at the revolution harmonic is 107 Hz (OTFB $a = 63/64$). The feedforward filters were calculated as shown in Appendix B. Bunches of 2.3×10^{11} protons, 25 ns apart, were used corresponding to the HL-LHC intensity [15], for a peak rf component of beam current of 2.94 A (assuming a conserve bunching factor of 1).

It should be noted that the synchrotron frequency during the SPS beam acceleration exceeds 100 Hz and is thus outside the bandwidth of the feedback filter peaks. Therefore, the design is more complex, adding resonances on the synchrotron bands to achieve better cavity impedance reduction. This “triple-comb” filter implementation is presented in detail in [6]. In this note we focus on the transient beam-loading compensation and the modeling uses the simpler comb filter [Eq. (6)].

The simulations presented in this work correspond to the rf parameters right before the beam extraction from the SPS for LHC injection. Transient beam loading compensation is critical at this part of the SPS cycle since uncompensated beam loading leads to bunch phase distortion along the SPS batch which in turn leads to capture losses when transferred into the LHC 400 MHz buckets. The rf frequency is at 200.3945 MHz and the total rf voltage is 10 MV.

All simulations include the effects of the six cavities. Figure 7 shows the beam induced voltage and the beam loading compensation with just the feedforward, with just the feedback, or with both systems on. The induced voltage of 10.9 MV without compensation matches the theoretical value which can be calculated by multiplying the cavity impedances from Table II with the peak rf beam current. For comparison, the total accelerating voltage is just 5 MV at 26 GeV/c and 10 MV at 450 GeV/c.

With just the feedback on, the beam induced voltage is reduced to about 260 kV at the middle of the batch, a reduction of approximately a factor of 42, in agreement with the feedback gain. Similarly, with just the feedforward on, the *peak* beam induced voltage is reduced to about 540 kV, a reduction by a factor of 20, in agreement with Fig. 31 in Appendix B. The beam loading compensation is *not* additive: the resulting performance with both feedback and feedforward is comparable to that achieved with just the feedforward. This is expected, since the uncompensated beam loading in the presence of the feedforward is mostly

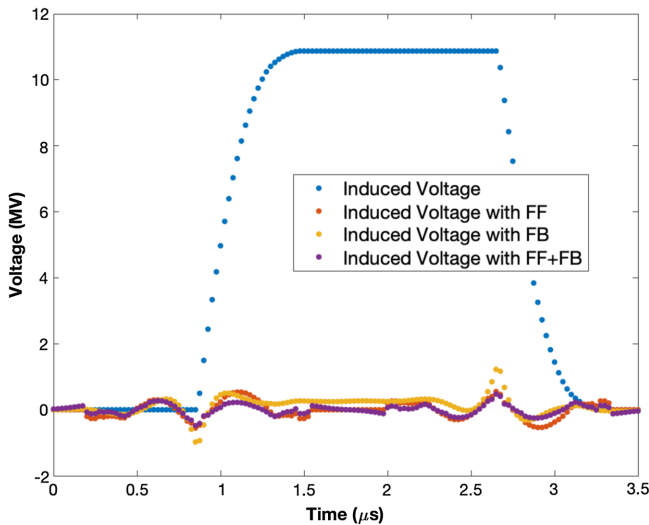


FIG. 7. Feedforward/feedback beam loading compensation. Total beam induced rf voltage from four three-section cavities and two four-section cavities.

at the cavity zeros and the TX cannot generate an accelerating field at these frequencies.

Figures 8 and 9 show the transmitter power for the three- and four-section cavities without compensation, with just feedforward on, with only feedback on, and with both systems together. Clearly, the presence of the feedforward system leads to very high-power transients, especially for the four-section cavities. This can be explained from the optimal feedforward responses derived in Sec. V (Fig. 29 in Appendix B): while the even-symmetric filter response is *smooth* (tapered pulse) and will not generate sharp transient in response to the beam step current, the odd-symmetric one has a few very large samples, and as such generates the huge response spike at the start of the batch.




FIG. 8. Transmitter power for three-section cavity.




FIG. 9. Transmitter power for four-section cavity.

The peak generator power at cavity input is 1 MW for the three-section cavities and 1.6 MW for the four-section cavities. We will impose these as hard limits in Sec. VIII.

B. Nonideal feedforward gain

Since the feedforward aims to fully cancel the beam current at the cavity, its performance will be limited by the unavoidable drifts in power amplifier gain and phase shift. The new SPS LLRF includes a regulation loop that keeps the gain and phase shift constant, from LLRF input to TX output, during the acceleration ramp (130 kHz frequency sweep at 200 MHz) and also assures long term phase and gain stability. Such a regulation called the klystron polar loop is in use in the LHC [12] and in Linac4 [16]. With such a system we expect only small changes in feedforward gain




FIG. 10. Uncompensated beam loading with nonideal feedforward gain (feedback off). Vertical markers indicate the extent of the batch.

TABLE III. rms error with feedforward gain (feedback off). Uncompensated beam loading.

Gain	0.9	0.95	1	1.05	1.1
rms error (kV)	583	373	246	323	520

and phase. To simulate this situation, the feedforward gain was changed by $\pm 5\%$ and $\pm 10\%$. The resulting uncompensated beam loading is shown in Fig. 10 (feedback is turned off). Table III shows the rms error for each case.

In the middle of the batch segment, the compensation comes from the even-symmetric part of the feedforward impulse response (Appendix B). We showed that this could provide exact compensation, and indeed the beam loading is zero for a perfectly calibrated feedforward (Fig. 10). As the calibration degrades, the match is less perfect, as expected. But in all cases, we cannot cancel the transient at the head of the batch during the cavity filling time (462 ns for a three-section cavity, 621 ns for a four-section, see Table II) and at the tail. Smaller transients are visible outside the beam segments, before the batch head and after the batch tail. These are of no importance for the beam.

The same procedure was repeated with feedback on. The results are shown in Fig. 11. The rms error is shown in Table IV. There is *significantly* less sensitivity to the feedforward gain when the feedback is on. It is also interesting to note the nonconstant uncompensated beam loading in the middle of the batch (between 1.5 and 2 μs in Fig. 11). The OTFB delay is set so that the loop stability is maximized (gain and phase margins). Figure 11 indicates that the OTFB compensates for the feedforward misalignments if these remain in the anticipated $\pm 10\%$ in gain. The feedforward calibration will also drift in phase.




FIG. 11. Uncompensated beam loading with nonideal feedforward gain (feedback on). Vertical markers indicate the extent of the batch.

TABLE IV. rms error with feedforward gain (feedback on).

Gain	0.9	0.95	1	1.05	1.1
rms error (kV)	161	158	155	153	151

Since the implementation is in Cartesian (I, Q) coordinates, a 0.1 radian phase error has the same effect as a $\pm 10\%$ gain error. So, we can draw identical conclusions concerning the effect of phase drifts if these remain within ± 0.1 radian (± 5.7 rf degrees).

C. Wider bandwidth and lower gain OTFB

The 107 Hz single-sided bandwidth on the revolution frequency harmonics will not fully cover the synchrotron sidebands as the synchrotron frequency ranges between 200 and 1000 Hz during the SPS HL-LHC cycle. While the transient beam loading is reduced by 40 (linear), the impedance on the synchrotron sidebands would be reduced much less, and that is a concern for the longitudinal coupled-bunch instability threshold. One LLRF solution is to have additional filter resonances on the synchrotron sidebands [6]. A simpler one is to operate the OTFB with increased bandwidth to cover the synchrotron sidebands with the simple response presented in Sec. IV. The OTFB bandwidth was thus adjusted to 428 Hz ($\alpha = 15/16$) and the gain was correspondingly decreased by a factor 4 (in order to keep the same regulation gain margin) to investigate any negative effects on the system performance. The rms error only slightly increased from 155 to 172 kV, as shown in Fig. 12 when the feedforward is perfectly calibrated.

When the feedforward is not perfectly calibrated, the rms error is still only slightly elevated, as shown in Table V.




FIG. 12. Uncompensated beam loading for different OTFB bandwidth/gain and ideal feedforward. The batch is indicated with vertical dashed lines.

TABLE V. rms error with feedforward gain (feedback on) and wider bandwidth, lower gain OTFB.

Gain	0.9	0.95	1	1.05	1.1
rms error (kV)	187	177	172	169	171

Therefore, there will only be a minimal effect on performance if operation with the higher bandwidth for the OTFB peaks at the revolution harmonics becomes necessary.

D. Coupled cavity feedback

As mentioned in Sec. III, the TX of a cavity cannot compensate for the beam induced voltage at frequencies where the cavity impedance Z_g is zero. As the frequencies of these zeros are different for the three- and four-section cavities we could however compensate, in the band of three-section cavity zeros, by injecting a corrective signal into the four-section cavities, and vice versa. The idea is attractive and was already proposed long ago [11] but was never implemented. With the flexibility of modern electronics, it could now be implemented, as long as it brings a significant improvement.

The proposed coupled cavity feedback could operate as follows. The six cavities are separated in two identical groups of two three-section cavities and one four-section cavity each. Each LLRF system will sample the voltage of its dedicated cavity, as well as the voltage of the other cavity type, as shown in Fig. 13. As such, the feedback systems will be coupled. The two voltages will be processed with different gains and filters. For the analysis below, the gain and filters used on the dedicated cavity voltage are the ones from the baseline design. The other cavity voltage is filtered by a bandpass filter centered around the cavity zero of the *other* cavity type, with the gain optimized to minimize transient beam loading. The combined output sets the TX drive.

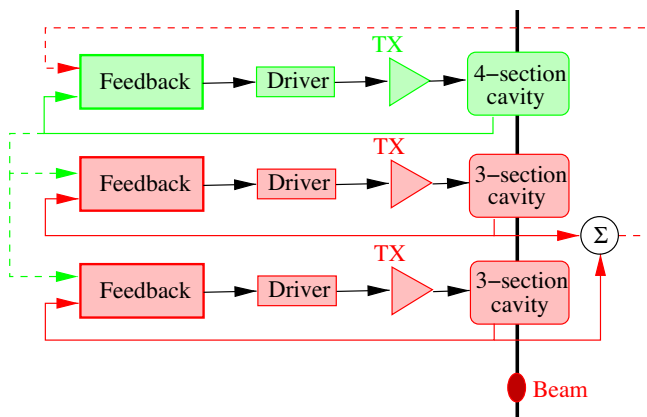


FIG. 13. Coupled cavity feedback block diagram. Solid lines show feedback inputs from cavities the LLRF system controls. Dashed lines show feedback inputs from other cavities.

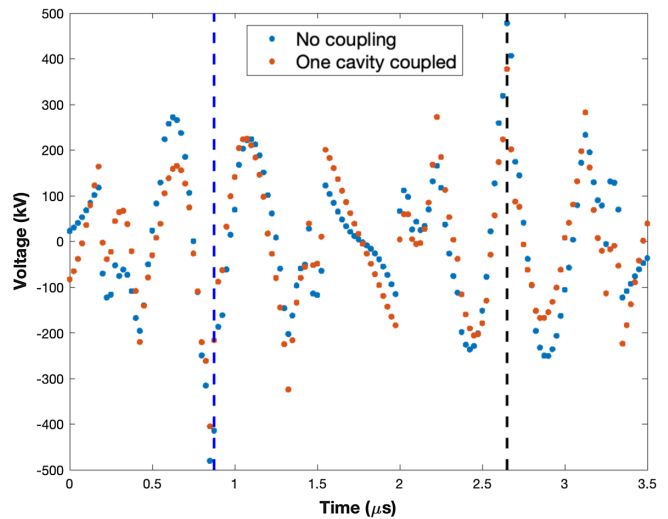


FIG. 14. Uncompensated beam loading with coupled feedback (three-section cavity only). The feedforward is also on. The batch is indicated with vertical dashed lines.

When the feedback $H_{cav}(f)$ filter has very little gain past the first cavity zero (as in the proposed baseline design), coupling is only used acting through the three-section cavity. The three-section cavity zero is at such a high frequency offset (2.16 MHz from the cavity center frequency) that the four-section cavity feedback has almost no gain at those frequencies to make a difference, as shown in Figs. 2 and 3 (left traces). Therefore, the gain on the total voltage is set to zero for the four-section cavity. Figure 14 shows the resulting induced voltage. In this case, the coupled feedback provides almost no improvement and does not justify the complicated implementation, commissioning, and operation.

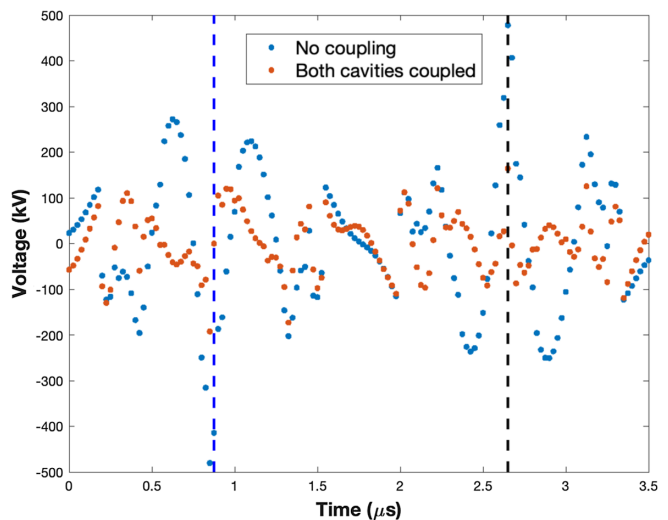


FIG. 15. Uncompensated beam loading with coupled feedback. The feedforward is also on. The batch is indicated with vertical dashed lines.




FIG. 16. Transmitter power for three-section cavity with coupled feedback.

If on the other hand the feedback $H_{\text{cav}}(f)$ filter sidelobe is at the same level as the main lobe as shown in Fig. 5 (the original LLRF design), then both coupling paths are used (gains are again optimized to minimize transient beam loading). Figure 15 shows the resulting induced voltage. A reduction of up to a factor of 2 is achieved. On the other hand, commissioning and operation of such a system will be significantly more difficult, so the trade-off of performance improvement with operational complexity should be evaluated. In addition, the demanded power increases with the coupling scheme, resulting in large power demands at the first zero of the companion cavity as shown in Figs. 16 and 17. This scheme will not be used for the initial operation of the updated SPS LLRF, but could be




FIG. 17. Transmitter power for four-section cavity with coupled feedback.

reconsidered if higher beam loading rejection is required *and* if the new TX response has a much wider bandwidth than the existing ones.

VIII. PERFORMANCE WITH UPGRADED LLRF AND OPTIMAL CONFIGURATION WITH HARDWARE LIMITATIONS

The studies presented in Sec. VII assume an ideal transmitter. In reality, the transmitter will limit the performance in two ways. First, the transmitter *input* signal will be clamped to protect the solid-state amplifier's driver. As a result, the transmitter output never exceeds 1 MW (three-section cavities) or 1.6 MW (four-section cavities). The TX saturation is modeled as a threshold at the peak TX power level. In addition, the transmitter's bandwidth will also influence the system performance. The finite bandwidth is modeled as a Butterworth filter with the same 3- and 15-dB bandwidth as the existing SPS transmitters (1.5 MHz 3-dB bandwidth and 2.5 or 4 MHz 15-dB bandwidth depending on the TX).

As shown in Figs. 8, 9, 16, and 17, the power averaged during the batch is very close to the transmitter limit. In addition, the limit is even exceeded during part of the batch.

The SPS cavities have no tuner. They were designed with a nominal center frequency at 200.222 MHz. The addition of higher order mode dampers during LS2 resulted in detuning the cavity [17], and, in high beam loading conditions, this helps to significantly reduce the power requirement and thus increase the operational margin at extraction. The rf frequency at 450 GeV/c (extraction energy) is 200.3945 MHz. The actual SPS cavity center frequency is at about 200.1 MHz (lower than the design value listed in Table I). The voltage is set to 1.3 MV for the three-section cavities and to 2.4 MV for the four-section cavities, to match their peak power margins at extraction, with a 2.94 A rf component of beam current (bunching factor of 1). The resulting *steady state* power required for the HL-LHC beam is then just 665 kW and 1.123 MW for the three- and four-section cavities respectively. We define steady state as the power required with a continuous beam of the same bunch intensity (no gap). Therefore, detuning reduces the *average* required transmitter power at extraction. The following simulations assume a 200.1 MHz cavity center frequency.

There are still very high-power transients though, especially at the beginning and end of the batch. These are smoothed out when the input clamping and transmitter response are introduced in the simulations.

Figure 18 shows the required transmitter power for three- and four-section cavities with detuned cavities, clamping the transmitter input stage (saturation at peak power level), and including the transmitter frequency response. Clearly, the demanded power at 450 GeV/c is now significantly reduced and below the transmitter limits.




FIG. 18. Transmitter power with detuning, clamping, and band-limited transmitter response. Steady state required power with (continuous) beam shown for comparison. The batch is indicated with vertical dashed lines.

A. H_{cav} filter study

The $H_{cav}(f)$ filters shown in Fig. 5 (the original LLRF design) have significant gain near the cavity zero. We designed smoother filters and investigated their effect on the beam phase regulation and required transmitter power. Figure 19 shows the original and four new filters. The 3 dB passband of the new filters was set to 20%, 30%, 40% or 50% of the frequency corresponding to the first zero of the cavity impedance Z_g . Their gain past the first zero is very low. The 50% case corresponds to the filter used in the proposed baseline design.

Figure 20 shows the bunch per bunch phase for these five filters. The beam batch contains 72 bunches spaced by five

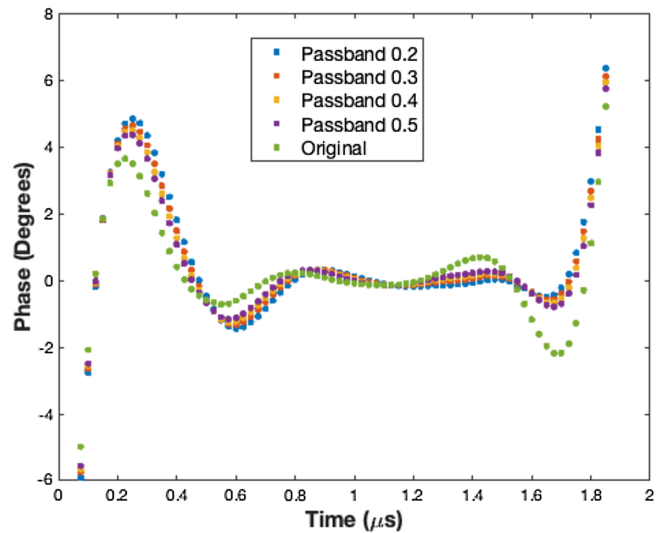


FIG. 20. Beam phase with original and new filters.

200 MHz rf periods (about 25 ns). Both feedback and feedforward systems are on. Clearly, there is very little loss of performance with the new filters. There is also minimal difference among them. The rms phase is 1.53 rf degrees for the original filter and 1.70–1.94 degrees for the new filters. The peak-to-peak phase is 10.2 degrees for the original filter and 11.3–12.3 degrees for the new filters. Most of this small performance loss is due to the first and last bunches in the batch.

This small loss of performance leads to some power savings though. The peak power for the original filter is 928 and 1514 kW for the three- and four-section cavities respectively. With the new filters, the peak power is 910–915 kW and 1366–1412 kW for the three- and four-section cavities, as shown in Fig. 21. Even though the power savings are just 1%–2% for the three-section

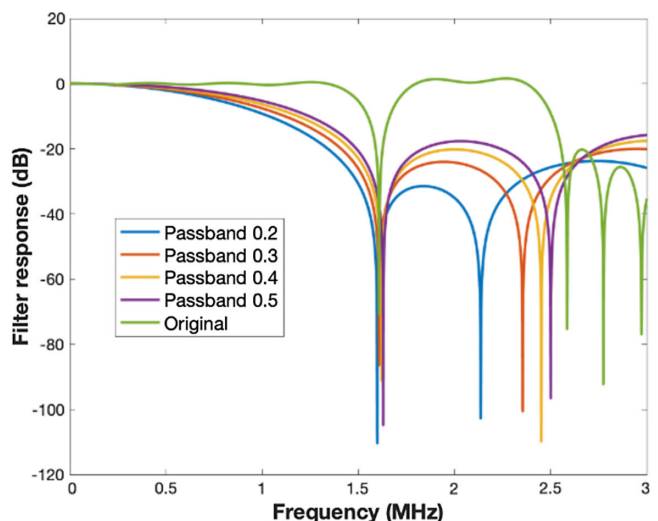



FIG. 19. $H_{cav}(f)$ filter comparison. Left: three-section cavity. Right: four-section cavity.




FIG. 21. Transmitter power comparison Left: three-section cavity. Right: four-section cavity.

cavities and 3%–10% for the four-section cavities, they might be significant given the hard TX power limit of 1 and 1.6 MW respectively.

The less aggressive filters also increase the rf loop stability limits and thus ease commissioning and operation of the SPS LLRF. The system is being commissioned in 2021 with a beam intensity lower than the HL-LHC, so a small performance reduction for a significant operational simplification is even more reasonable. As a result, the passband 0.5 filter was selected for the SPS design. This choice can of course be revisited if even higher beam performance is desired or lower power limits are established. The third LHC run starting in early 2022 will allow us to fully evaluate this design and reconsider the possible additional functionalities presented in Secs. VII and VIII.

IX. SUMMARY AND ESTIMATED PERFORMANCE WITH HL-LHC BEAM PATTERN

The studies presented in Secs. VII and VIII suggest that the design summarized below achieves the best tradeoff between performance, demanded power, and operational complexity. The planned feedforward system already achieves high beam loading compensation. The feedback does not improve the performance much in the presence of the feedforward, but it is needed to reduce the effect of possible gain and phase shift drifts in the feedforward. The feedback filter is selected to reduce peak power requirements, with a very small impact on performance. A wider bandwidth OTFB and a coupled feedback were explored but will not be included in the original design as they do not bring significant improvement.

The model verification and studies shown so far in this work used a 72-bunch (25 ns spacing) beam pattern. The

HL-LHC beam pattern in the SPS, consists of four such batches, spaced 200 ns apart. The estimated performance with the proposed design summarized above and with the HL-LHC beam pattern is presented in this section.

Figure 22 presents the resulting induced voltage. The induced voltage has an rms value of just 284 kV, to be compared with the 10.9 MV of uncompensated beam loading. The peak induced voltage for filled buckets is 1.61 MV. Values higher than 1.1 MV are only for the first bunch of the second, third, and fourth batches, so for only three out of 288 bunches.

Figure 23 shows the resulting 200 MHz beam phase. The rms phase is 2.0 rf degrees and the peak-to-peak variation is 15.6 degrees (dominated by the first and last bunches of each batch). The pre-LS2 SPS LLRF system achieved an




FIG. 22. Induced voltage with proposed design.




FIG. 23. Beam phase with baseline design.

rms phase of ~ 5 degrees and a peak-to-peak variation of ~ 33 degrees, at almost *half* the beam intensity (1.3×10^{11} instead of 2.3×10^{11} protons per bunch). The upgraded SPS LLRF will thus have an almost fivefold higher transient beam loading compensation than the pre-LS2 system in rms value, and an almost fourfold compensation in peak to peak.

Finally, Fig. 24 presents the transmitter power with the proposed design. The peak power is 885 kW for the three-section cavities and 1.39 MW for the four-section cavities, leading to a 13% and 15.4% power margin respectively with respect to the 1 and 1.6 MW transmitter limits. Comparing these peaks to the calculated 665/1123 kW steady state values, we conclude that the regulation requires 33% (three-section) and 23% (four-section) power overshoot.





FIG. 24. Transmitter power with proposed design. First and last bunch for each batch are indicated with vertical dashed lines. Left: three-section cavity. Right: four-section cavity.

X. CONCLUSIONS AND FUTURE PLANS

As part of the LHC injectors upgrade project (LIU), the SPS rf system has been modified during long shutdown 2 (LS2, January 2019–April 2021), resulting in a new layout containing four three-section and two four-section 200 MHz traveling wave cavities. The goal is to improve longitudinal stability required by the planned doubling of the beam intensity for the high luminosity LHC (HL-LHC). In parallel with the upgrade of the high-power rf, a new LLRF system has been designed and is presently commissioned. This work presents models and simulations of the SPS rf and LLRF, as well as the feedback and feedforward algorithms. Simulation studies were conducted and summarized. Through these studies, a design for the upgraded SPS LLRF system is presented in this work.

As the proposed design is commissioned in 2021, more data and technical details will be published. The predictions presented in this work will be compared with actual measurements. The simulation studies will be repeated when the response of the new SPS transmitter is fully documented. Optimization will take place over the coming years, as the SPS beam current is increased during the LHC run 3 starting in early 2022.

Finally, if higher beam loading rejection is required, the following options can be considered: (i) more aggressive feedback filters, with higher gain past the first cavity zero (this solution will probably be limited by the TX power threshold); (ii) feedback coupling among cavities, which achieves moderate performance improvement, but significantly complicates the design, setting up, and operation (the effectiveness of this possible solution will also be limited by the TX power threshold); (iii) the feedforward coefficients are currently optimized for each cavity [given that the performance limitations and power transients are largely related to the cavity zeros, more sophisticated

optimization of the feedforward coefficients might be possible: (a) for *all rf stations at the same time*, including the distinct cavity responses, (b) including TX power and frequency response limitations, (c) including the peak TX power constraint]; and (iv) design of more sophisticated feedforward and feedback algorithms that regulate the voltage in *populated buckets only* and extrapolate in the no-beam segment so that peak power is minimized. The current LLRF tries to reduce transient beam loading for all buckets, often leading to high power transients. All simulations show that the performances are worst for the few bunches following a gap. This is the result of unnecessary TX drive adjustments during the gaps. It is possible to adjust the cavity voltage setpoint to reduce the transients corresponding to empty buckets. The setpoint adjustment is automated and adapts to changes in bunch population. Effectively, the LLRF is trained to smoothly compensate for the induced voltage during the empty buckets. This scheme would follow our experience with the LHC full-detuning adaptive algorithm [18].

The models and formalism developed for this work have been instrumental in designing and are assisting with the commissioning of the upgraded SPS LLRF. As the beam current is increased during the LHC run 3 (2022–2025),

they will also be used to design and evaluate alternative operational algorithms and LLRF layouts as needed.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award No. DE-SC-0019287. The authors would like to thank Thomas Bohl (CERN) for many useful discussions, assistance with data gathering, and all his contributions to SPS. We thank Ivan Karpov (CERN) for providing data on the power required in static conditions, confirming our numbers. We are grateful to Eric Montesinos (CERN), in charge of the SPS power amplifiers, for his help in modeling these. Finally, we thank the CERN rf team for inviting us to present these results in internal meetings and for the subsequent fruitful discussions, comments, and suggestions.

APPENDIX A: MODEL OF A TRAVELING WAVE CAVITY

The cavity impedance driven by the generator [Eq. (3)], is transformed to the impulse response in the I/Q model [7]:

$$\begin{aligned} \begin{bmatrix} v_{gI} \\ v_{gQ} \end{bmatrix} &= \begin{bmatrix} h_{gs}(t) & -h_{gc}(t) \\ h_{gc}(t) & h_{gs}(t) \end{bmatrix} * \begin{bmatrix} i_{gI} \\ i_{gQ} \end{bmatrix} \\ &= R_1 \begin{bmatrix} \frac{1}{\tau} \Pi(\frac{t}{\tau} - \frac{1}{2}) \cos(2\pi\Delta f_0 t) & \frac{1}{\tau} \Pi(\frac{t}{\tau} - \frac{1}{2}) \sin(2\pi\Delta f_0 t) \\ -\frac{1}{\tau} \Pi(\frac{t}{\tau} - \frac{1}{2}) \sin(2\pi\Delta f_0 t) & \frac{1}{\tau} \Pi(\frac{t}{\tau} - \frac{1}{2}) \cos(2\pi\Delta f_0 t) \end{bmatrix} * \begin{bmatrix} i_{gI} \\ i_{gQ} \end{bmatrix} \end{aligned} \quad (\text{A1})$$

In the above equation, all variables are in the time domain, the notation $*$ represents the convolution operator and the $\Pi(x)$ function is a unit-amplitude rectangular symmetric pulse of width 1. The demodulator frequency f_{LO} is a free parameter. In the implementation we will choose it as a fixed frequency, close to but not exactly at the center frequency of each individual cavity.

$$\Delta f_0 = f_{LO} - f_0.$$

For understanding the above equation, it is easiest to assume that the demodulator is at the cavity center frequency. Then Δf_0 is zero and the off-diagonal terms vanish. According to Eq. (A1) the response to a generator current impulse is then a rectangle lasting for a time equal




FIG. 25. Impulse response to generator current. Diagonal term $h_{gs}(t)$ normalized. For $\Delta f_0\tau = 0$ (left) and $\Delta f_0\tau = 0.2$ (right).




FIG. 26. Impulse response to generator current. Off-diagonal term $-h_{gc}(t)$ normalized. For $\Delta f_0 \tau = 0$ (left) and $\Delta f_0 \tau = 0.2$ (right).

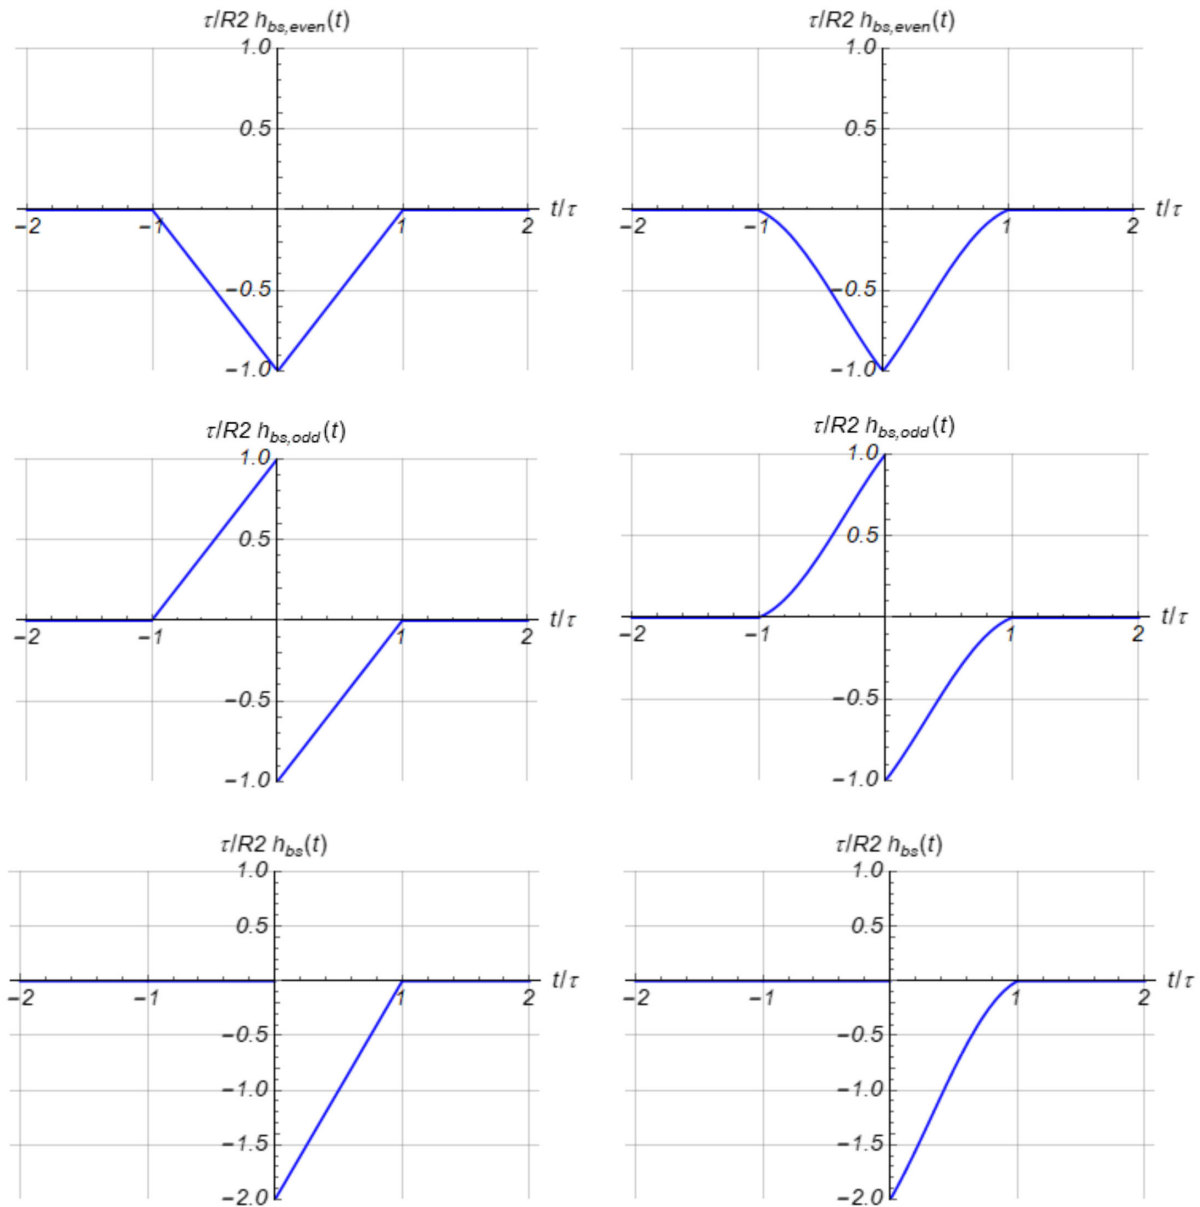


FIG. 27. Impulse response to beam current. Diagonal term $h_{bs}(t)$ normalized. For $\Delta f_0 \tau = 0$ (left) and $\Delta f_0 \tau = 0.2$ (right). The top trace is the even-symmetric part, the middle trace shows the odd-symmetric and the bottom trace is the sum that is causal.

to the cavity filling time, in agreement with previous publications [8]. Figure 25 shows the diagonal term $h_{gs}(t)$ after normalization by R_1/τ for two values of Δf_0 . Both axes have dimensionless units. Figure 26 shows the off-diagonal term $-h_{gc}(t)$.

$$\begin{aligned} \begin{bmatrix} v_{bI} \\ v_{bQ} \end{bmatrix} &= \begin{bmatrix} h_{bs}(t) & -h_{bc}(t) \\ h_{bc}(t) & h_{bs}(t) \end{bmatrix} * \begin{bmatrix} i_{bI} \\ i_{bQ} \end{bmatrix} \\ &= -R_2 \begin{bmatrix} \{\frac{1}{\tau} \Lambda(\frac{t}{\tau}) \cos(2\pi\Delta f_0 t) + \frac{1}{\tau} \Lambda(\frac{t}{\tau}) \text{sgn}(\frac{t}{\tau}) \cos(2\pi\Delta f_0 t)\} & \{\frac{1}{\tau} \Lambda(\frac{t}{\tau}) \sin(2\pi\Delta f_0 t) + \frac{1}{\tau} \Lambda(\frac{t}{\tau}) \text{sgn}(\frac{t}{\tau}) \sin(2\pi\Delta f_0 t)\} \\ \{-\frac{1}{\tau} \Lambda(\frac{t}{\tau}) \sin(2\pi\Delta f_0 t) - \frac{1}{\tau} \Lambda(\frac{t}{\tau}) \text{sgn}(\frac{t}{\tau}) \sin(2\pi\Delta f_0 t)\} & \{\frac{1}{\tau} \Lambda(\frac{t}{\tau}) \cos(2\pi\Delta f_0 t) + \frac{1}{\tau} \Lambda(\frac{t}{\tau}) \text{sgn}(\frac{t}{\tau}) \cos(2\pi\Delta f_0 t)\} \end{bmatrix} \\ &\quad * \begin{bmatrix} i_{bI} \\ i_{bQ} \end{bmatrix}, \end{aligned} \quad (\text{A2})$$

where $\Lambda(x)$ is the unit peak amplitude symmetric triangular function of full base width 2, and $\text{sgn}(x)$ is the sign function. Each of the four impulse responses in Eq. (A2) consists in two terms: The first one is an even-symmetric impulse response [$h(-n) = h(n)$] corresponding to the real part of the impedance. The second term is an odd-symmetric impulse corresponding to the imaginary part of the impedance. The two impulse responses cancel exactly for negative time samples, resulting in a causal response (Fig. 27).

The impulse response expressions derived in this section [Eqs. (A1) and (A2)] will be used for the estimation of an

Similarly, the cavity impedance driven by the beam [Eq. (4)], is transformed to the impulse response (wake-field) in the I/Q model in [7]:

optimal feedforward algorithm (Appendix B). They are also implemented in the time-domain simulations presented in Sec. VI.

APPENDIX B: OPTIMIZATION OF FEEDFORWARD COEFFICIENTS

The SPS beam consists of a series of batches with fairly uniform bunch intensity, separated by empty gaps. The largest beam loading transients are observed at the transitions between gaps and batches. Optimization will therefore be done for a *beam current step*. To keep the derivation




FIG. 28. Step response to beam current (normalized). The time index refers to the 31.25 MHz rate. The top left trace is the response to the real part of the impedance (even-symmetric impulse response), the top right trace shows the odd-symmetric part and the bottom trace is the sum.




FIG. 29. Even-symmetric (left) and odd-symmetric impulse responses of the feedforward FIR filter.




FIG. 30. Compensation voltage due to the feedforward (blue) and beam induced voltage (orange). Left: voltage caused by the even-symmetric impulse response compared to the one caused by the real part of cavity impedance (perfect superposition). Right: voltage caused by the odd-symmetric impulse response compared to the one caused by the imaginary part of cavity impedance.

simple, we will assume that the demodulation frequency is equal to the cavity center frequency ($\Delta f_0 = 0$). Then the impulse responses are pure rectangular or triangular shapes [Eqs. (A1) and (A2)]. The off-diagonal elements of the matrices are all zeros, meaning that the I and Q components can be treated separately, and with the same processing. In the planned implementation the I/Q signals are digitized at a frequency set to 31.25 MHz. With the shorter three-section cavity (462 ns filling time), we have fourteen samples ($L_{filling} = 14+1$) during cavity filling. This section is concerned with the processing algorithm. To ease the notations all signals are normalized to: A static unit beam current induces a unit cavity voltage, and so does a static unit generator current. The time index refers to the 31.25 MHz processing rate. The feedforward processing H_{ff} consists of

a finite impulse response (FIR) filter with N_{tap} taps (assumed odd). The generator driven voltage will come from processing the beam current with filter H_{ff} (N_{tap} samples), followed by the cavity response ($L_{filling} = 14$ samples). The fit should thus be done on a time record lasting for $N_{tap} + L_{filling} = P_{fit}$ samples. The figures consider a $N_{tap} = 31$ taps FIR and the three-section cavity ($L_{filling} = 14$). The fit is therefore done on $P_{fit} = 31 + 14 = 45$ samples. Figure 28 shows the beam induced voltage caused by a unit step in beam current. This is in agreement with past publications [8].

We have N_{tap} unknowns (the FIR coefficients) that we can choose so that the compensation voltage matches the beam induced voltage. We will split the FIR filter into an even-symmetric part that must match the even-symmetric

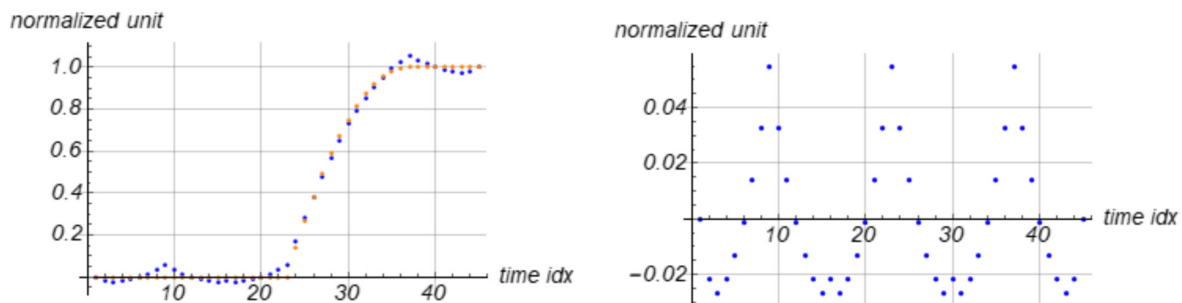


FIG. 31. Left: Compensation voltage (blue) and beam induced voltage (orange). Right: Compensation error.

part of the beam loading, and an odd-symmetric FIR to match the odd-symmetric part of beam loading. Our criterion is to minimize the mean square error over the $Pfit$ window. As the error depends linearly on the FIR filter coefficients, the minimization is a classic linear least squares problem [19]. The derivation is shown in detail in [7]. Figure 29 shows the even-symmetric and odd-symmetric parts of the feedforward FIR filter. As expected, the even-symmetric part consists of a rectangular response lasting for the cavity filling time (14 samples).

Figure 30 shows the voltage created by the feedforward as a response to a beam current step. It is compared to the beam induced voltage: The even-symmetric part of the response is an exact match to the corresponding part of the beam induced voltage. The odd-symmetric part is an approximation.

Finally, Fig. 31 compares the total beam induced voltage to the total compensation. We observe small remaining ripples in the uncompensated beam loading, but their peak value has been reduced below 0.06, from an initial value of 1. The scheme is very promising.

-
- [1] L. Medina, T. Argyropoulos, R. Calaga, and H. Timko, Studies of longitudinal beam losses at LHC injection, in *Proceedings of the 12th International Particle Accelerator Conference, Campinas, Brazil* (JaCoW Publishing, Geneva, 2021).
- [2] T. Argyropoulos *et al.*, Controlled longitudinal emittance blowup in a double rf system at CERN SPS, in *Proceedings of the ICFA Workshop HB2010, Morschach, Switzerland* (2010).
- [3] M. Schwarz, K. Iliakis, A. Lasheen, G. Papotti, J. Repond, E. Shaposhnikova, and H. Timko, Flat-bottom instabilities in the CERN SPS, in *Proceedings of the 10th International Particle Accelerator Conference, Melbourne, Australia* (JaCoW Publishing, Geneva, 2019).
- [4] J. Dome, The SPS acceleration system traveling wave drift-tube structure for the CERN SPS, Report No. CERN-SPS/ARF/77-11, 1976.
- [5] J. Coupard, H. Damerou, A. Funken *et al.*, LHC injectors upgrade, Technical Design Report No. CERN-ACC-2014-0337, CERN, 2014.
- [6] G. Hagemann, P. Baudrenghien, J. D. Betz, J. Egli, G. Kotzian, M. Rizzi, L. Schmid, A. Spierer, T. Wlostowski, and F. J. G. Guarch, The CERN SPS low level rf upgrade project, in *Proceedings of the 10th International Particle Accelerator Conference, Melbourne, Australia* (JaCoW Publishing, Geneva, 2019).
- [7] P. Baudrenghien and T. Mastoridis, I/Q model of the SPS 200 MHz traveling wave cavity and feedforward design, Report No. CERN-ACC-NOTE-2020-0032, 2020.
- [8] D. Boussard, Traveling-wave structures, Joint U.S.–Cern–Japan International School, Tsukuba, 1996.
- [9] D. Boussard, Control of cavities with high beam loading, *IEEE Trans. Nucl. Sci.*, 1852 (1985), <https://cds.cern.ch/record/160201?ln=en>.
- [10] D. Boussard and G. Lambert, Reduction of the apparent impedance of wideband accelerating cavities by rf feedback, *IEEE Trans. Nucl. Sci.* **30**, 2239 (1983), <https://cds.cern.ch/record/143505?ln=en>.
- [11] P. Baudrenghien and G. A. Lambert, Reducing the impedance of the traveling wave cavities: Feed-forward and one turn delay feedback, in the *10th Workshop on LEP-SPS Performance, Chamonix, France* (CERN, Geneva, 2000), pp. 94–101.
- [12] P. Baudrenghien, G. Hagemann, J. C. Molendijk, R. Olsen, T. Rohlev, V. Rossi, D. Stellfeld, D. Valuch, and U. Wehrle, The LHC low level rf, in *Proceedings of the 10th European Particle Accelerator Conference, Edinburgh, Scotland, 2006* (EPS-AG, Edinburgh, Scotland, 2006).
- [13] F. Pedersen, RF cavity feedback, in *Proceedings of the Conference on B Factories, Stanford, CA, USA* (Stanford University, Stanford, 1992).
- [14] P. Baudrenghien, Low level rf systems for synchrotrons: Part II, *CERN Accelerator School: Radio Frequency Engineering, Seeheim, Germany* (CERN, Geneva, 2000), pp. 185–188.
- [15] S. Antipov *et al.*, Update of the HL-LHC operational scenarios for proton operation, Report No. CERN-ACC-NOTE-2018-0002.
- [16] P. Baudrenghien, J. Galindo, G. Hagemann, J. Noirjean, D. Stellfeld, and D. Valuch, Commissioning of the LINAC4 low level rf and future plans, in *Proceedings of LINAC2014 Conference, Geneva, Switzerland* (JaCoW Publishing, Geneva, 2014).
- [17] P. Kramer and C. Vollinger, Measurement of rf voltage in long traveling wave accelerating structures with application to the CERN Super Proton Synchrotron, *Phys. Rev. Accel. Beams* **22**, 101002 (2019).
- [18] T. Mastoridis, P. Baudrenghien, and J. Molendijk, Cavity voltage phase modulation to reduce the high-luminosity Large Hadron Collider rf power requirements, *Phys. Rev. Accel. Beams* **20**, 101003 (2017).
- [19] G. Strang, *Linear Algebra and its Applications* (2006), ISBN-13: 978-0030105678.