ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

LOW LEVEL RF CONTROL ALGORITHMS FOR THE CERN PROTON LINAC4

P. Baudrenghien†, B. Bielawski, R. Borner, CERN, Geneva, Switzerland

Abstract

The CERN Linac4 Low Level RF (LLRF) uses a Linear Quadratic Gaussian regulator and an Adaptive Feed Forward to stabilize the accelerating field in the cavities in the presence of strong beam loading. A Klystron Polar Loop is also implemented to compensate the RF perturbations caused by the ripples and droop in the klystron High Voltage supply. The paper presents the important parts of the regulation, shows results as the system has evolved from first prototype (2013) to operational beams (2020), and mentions some important issues encountered during the commissioning and the first years of operation, with their mitigations.

LINAC4

The Linac4 machine accelerates H- ions to 160 MeV Kinetic Energy and injects these into the PSB synchrotron, through a stripping foil. The RF operates at 352.2 MHz. It includes a chopper at 3 MeV removing the bunches that would fall outside the 1 MHz PSB bucket (h=1). The chopper also creates empty 2 µs long beam gaps to cope with the switching time of the distributing magnet that routes the Linac beam to the four superposed PSB rings. As a consequence, the cavities see strong transient beam loading as the beam intensity changes from zero to maximum beam current (presently 25 mA DC) in just 3 ns. The machine operates at a 1.2 s repetition time. At each pulse it can accelerate up to 600 µs of beam consisting of four batches (one per PSB ring) spaced by 2 µs [1,2]. Linac4 has been producing protons for the CERN complex (PSB, PS, SPS, LHC) since Dec. 2020 [3]. Its first years of RF operation are presented in a companion paper [4]. The LLRF consists of a tuning system keeping the cavity at the tune that minimizes the required power [5], and a field regulation that modifies the generator drive to keep the cavity field at the desired value.

THE NEED FOR FIELD REGULATION

End-to-end beam dynamics simulations had been carried out, early in the machine design, to define (among other tolerances) the maximum level of RF phase and amplitude jitter that the system can tolerate before beam quality at injection in the PSB is compromised [6]. This study resulted in a one RF degree, one percent amplitude pk-pk budget during the beam pulse, for the (then) nominal 40 mA DC intensity. Although the source cannot presently give the target intensity, the RF performances presented here are scaled for the 40 mA.

† Philippe.Baudrenghien@cern.ch.

The causes for field fluctuations are numerous:

- Environmental causes (temperature and humidity) will affect the tunnel floor changing the cavity spacing, and the cable length, thereby introducing phase shifts.
- The cavities are subject to vibrations and microphonics that may be too fast to be corrected by the mechanical tuners.
- The power amplifiers suffer from RF gain/phase ripples caused by the noise in their High Voltage (HV) DC supply. This is very severe with klystrons: The changes in the HV modulate the velocity of the klystron electron beam, resulting in a change in the delay between input and output cavities. The LEP tubes (reused in Linac4) show 0.1 dB and 8 RF degree per precent HV drift. The slow drift during the pulse is called klystron *droop*. It will lead to a reduction of the field along the batch. The higher frequency (10 kHz) was traced to the switching frequency of the HV power converters. It is important to note that this affects the RF system as a *multiplicative* noise.
- Another source of perturbation is the beam current. An RF cavity is a resonant circuit excited by two currents, the RF amplifier output and the beam [7]. This is by far the largest perturbation in the Linac4 cavities. The 25 mA DC beam current induces almost 1 MV in the CCDTL1 cavity for example, to be compared to the 8 MV accelerating voltage. For nominal 40 mA, the beam induced voltage will be 1.6 MV, that is 20% amplitude variation. To respect the specifications, our LLRF must (and does) reduce the beam loading peak by twenty linear minimum. The beam loading is an *additive* perturbation in the RF system.

LLRF HARDWARE

This paper is focused on the algorithms. Yet a short presentation of the hardware will help the understanding.

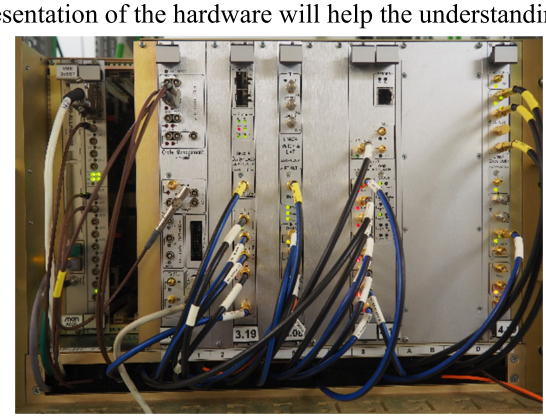

Figure 1: VME crate housing the LLRF modules for one power amplifier and its cavity.

ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

We have one VME crate per *line* (consisting of one power amplifier powering one or two cavities). A commercial front-end computer is housed in slot zero. The other cards are custom designed [5]. The field regulation is implemented in the Cavity Loop module (Fig. 1). We use the classic I-Q Digital demodulation scheme first proposed in 1995 [8]. The RF signal (antenna, coupler or reference line) at f_{RF} (352.2 MHz) is first down-mixed with an LO at 15/16 f_{RF} (330.2 MHz) to produce an IF at 1/16 f_{RF} (22.0125 MHz). The analog IF signal is sampled by an ADC whose clock runs at ¼ f_{RF} (88.05 MHz). After time de-multiplexing and sign inversion, a stream of baseband (I,Q) pairs is generated at 44.025 MSps [5]. Linac4 has a coaxial reference line, powered with a 100 W signal, and running in the tunnel, with -30 dB coupler adjacent to each cavity. The signals from these couplers are routed to the LLRF on cables (~100 m long) running together with the corresponding antenna signals. This layout is intended to minimize phase variations caused by temperature changes. On the surface, the reference line signal is used to generate the demodulations clocks (LO and ADC clock). A reference phase, obtained from the demodulation of the reference signal, is subtracted from the antenna demodulation, so that the scheme is not sensitive to drift in the generation of the LO and ADC clocks. After processing in baseband (field regulation loops implemented in an FPGA clocked at 88.05 MHz), the output is mixed to the 22.0125 MHz IF frequency, converted to analog and mixed up to generate a 352.2 MHz generator drive. More details and a block diagram can be found in [5].

THE LQG REGULATOR

The prototype LLRF (2013-2018) included a simple RF feedback (Proportional-Integral, PI controller). Such systems have been used in accelerators since the 1970s [9]. Their gain (and therefore efficiency in reducing perturbances) is limited by the overall delay in the feedback path (including amplifier group delay, waveguides, cables and latency of LLRF electronics). This delay could not be reduced below 2.5 µs in the Linac4 layout. The performances with PI controller therefore did not fulfil the specifications. The obvious next step was to search for a method that would give an *estimate* of the cavity voltage *without delay*. Then we could use this estimate, in place of the instantaneous cavity voltage, in the feedback. In this estimation quest we can take advantage of our knowledge of the generator current input to the cavity (or, at least the LLRF drive sent to the generator). We also know the system's dynamics, that is the cavity Q and detuning. And finally, we have access to the measurement of the *delayed* cavity voltage.

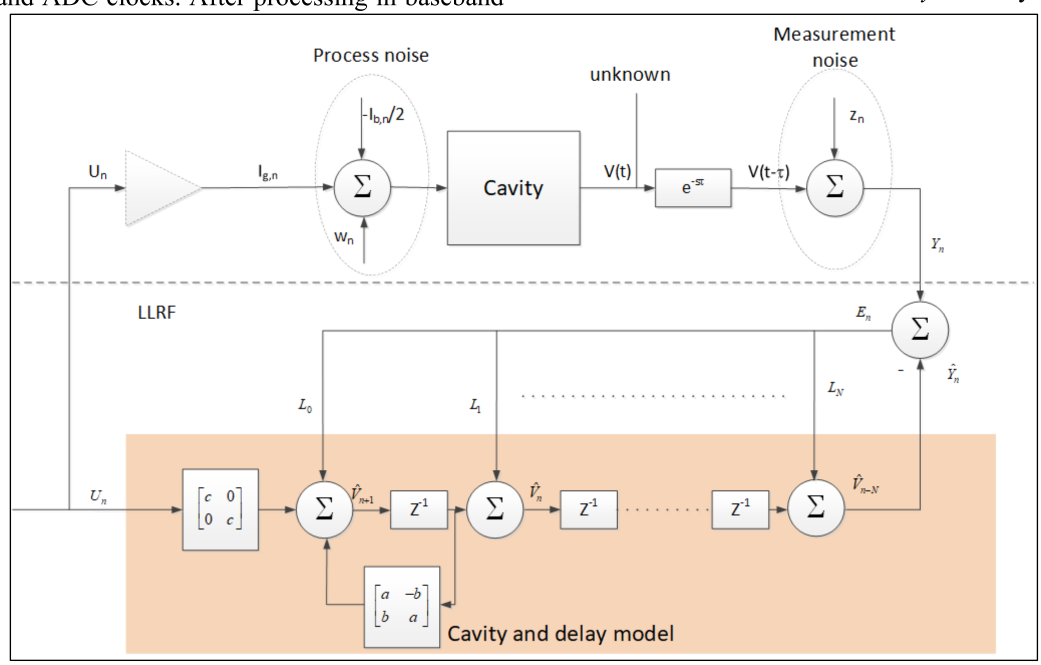

Figure 2: The Kalman Predictor

This is a classic problem in Controls, and the solution is the *Kalman Predictor*. Refer to Fig.2. In the LLRF firmware we implement a model of the *cavity plus delay* (orange box). All variables are in cartesian coordinates (I,Q). The time index is noted n at the 22.0125 MHz rate. The 2x2 cavity state-transition matrix is diagonal if the cavity is on-tune (b=0). This allows to separate the processing of the I and Q components, saving firmware resources. The chain of z^{-1} represents successive one sample delays. The model receives, as input, the same drive as the cavity amplifier (U_n). Of course, the model output ($\hat{Y}_n$) will not be exactly equal to the cavity output (Y_n). The difference comes from the perturbations injected into the cavity and absent in the model. This is called *Process noise*. In our case that is mainly the beam loading and klystron noise. At each iteration, the model state is corrected by comparing model and cavity outputs and weighting the error by the matrix L. There is noise at the cavity output, the so-called *Measurement noise* (small in our case). In the Kalman formalism the optimum matrix L depends on the relative strength of process noise and measurement noise. If the process noise is high compared to the measurement noise

ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

(our case), the correction will be applied strongly as we have much confidence in the measurements, while expecting that the model output deviate much from these as beam loading is not included in the model. After some fine trimming, we adjusted the Linac4 Kalman Predictor to react with a time constant below five iterations (< 200 ns).

We can now implement a proportional-integral (PI) feedback using the cavity voltage *estimate* (Fig.3). We have optimized the gains using the Linear Quadratic Regulator method (LQR). The feedback must obviously react much slower than the predictor. After some trials its reaction time was adjusted to 20 samples (900 ns). The combination of the Kalman Estimator and the LQR is called Linear Quadratic Gaussian regulator (LQG) in literature. See for example [10] for a tutorial.

Figure 4 shows the beam loading compensation in CCDTL1 with the PI controller, 12 mA DC (left) and with the LQG, 25 mA DC (right). The LQG has much improved the regulation. The reaction at the head of the batch is much faster. But we still see a significant peak deviation for the first few µs after beam arrival. That is easily understandable referring to Fig.3. The beam loading is seen by the regulation *after* propagating through the cavity and the ~ 1 µs long cable from antenna to the electronics. Only then can the LQG react, resulting in this unavoidable transient.

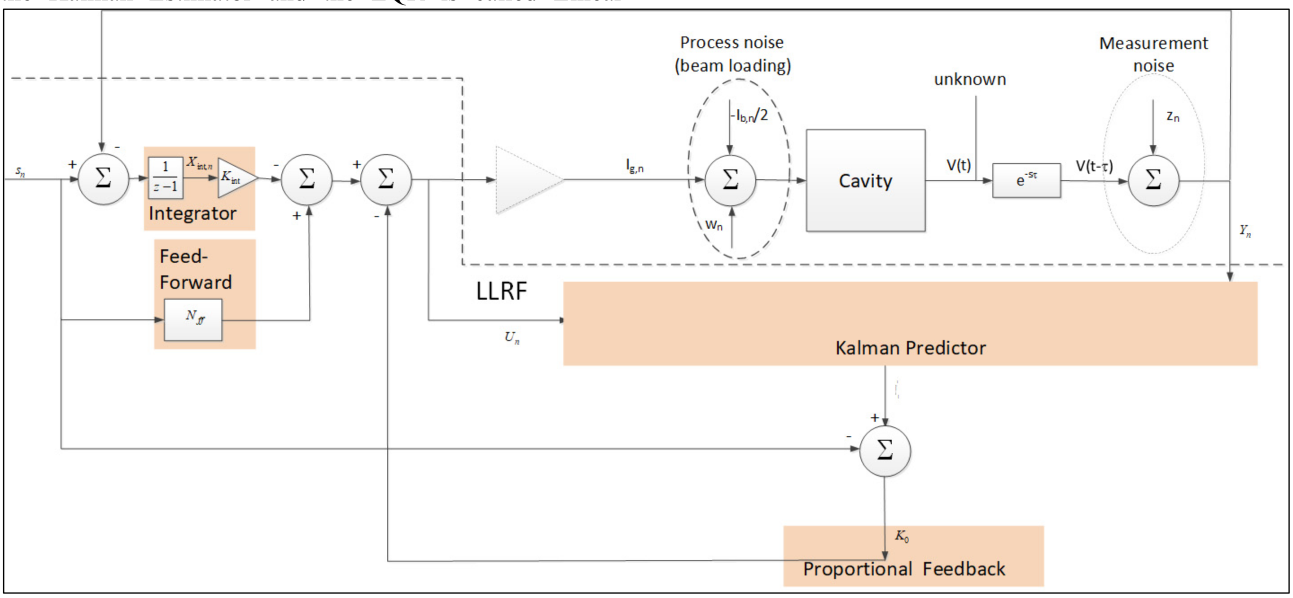

Figure 3: The LQG. Kalman Predictor and LQR using the predicted cavity voltage.

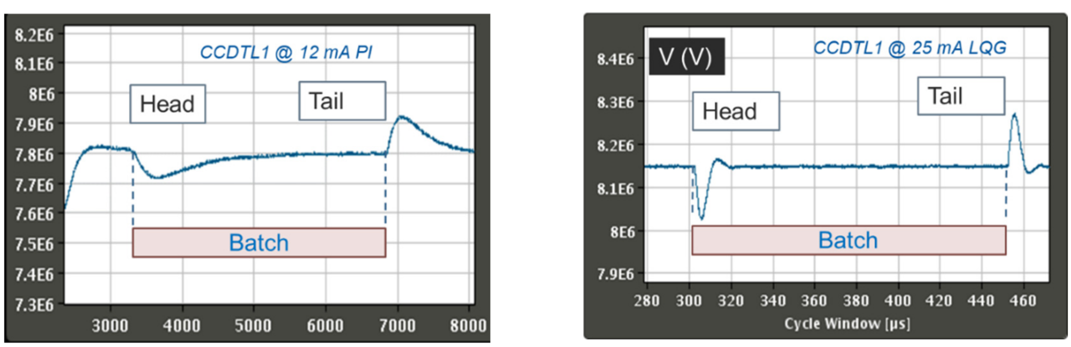

Figure 4: Cavity voltage. Left. PI controller (2017), 12 mA DC (45.42 ns/sample). Right. LQG regulation (2019), 25 mA DC.

THE ADAPTIVE FEED FORWARD

As long as we react pulse per pulse, we cannot completely cancel the transient at the head of the batch due to causality; we cannot anticipate the required increase of klystron power. However, assuming that this transient is reproducible from pulse to pulse, we could base our correction on the observation of previous pulses and anticipate for the coming pulse. This Adaptive Feed Forward (AFF) works very well for repetitive perturbations such as beam loading (since the source current does not change much from pulse to pulse). It has no effect on non-reproducible noises such as RF noise caused by klystron HV ripples and microphonics. In the Linac4 design these are mitigated by the LQG. The AFF correction, based on the observation of the voltage error over the few previous pulses (with more weight on the more recent ones), is added to the klystron drive (Fig. 5). It is computed off-line, and the correction (waveform covering the entire beam batch) is loaded in the firmware at each pulse. We use an algorithm similar to the one developed at TRIUMF [11]; we filter the voltage error with an impulse response that is the LQG closed-loop response, time-inverted, and *advanced* correctly in time. As

ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

noted at TRIUMF, fine adjustment of this *negative* delay is critical. It is close to 1.8 µs in our applications.

Figure 6 shows the performances with LQG and AFF. The feed-forward reduces the uncompensated beam loading to 30 kV pk-pk at the head of the batch, that is 0.375 % for 25 mA DC. Scaling to 40 mA DC current, the performances are factor two better than specifications. The beam becomes barely visible in the overall pulse. The AFF is switched OFF at the end of the batch, therefore the transient. The AFF *reduces* the peak of demanded power at the beginning of the batch. Figure 7 shows the power without (left) and with AFF. The LQG regulation cannot anticipate the arrival of the head of the batch. It is caught by surprise and reacts violently causing a peak of demanded power. With the AFF active the system anticipates and there is *no power surge* at the batch head.

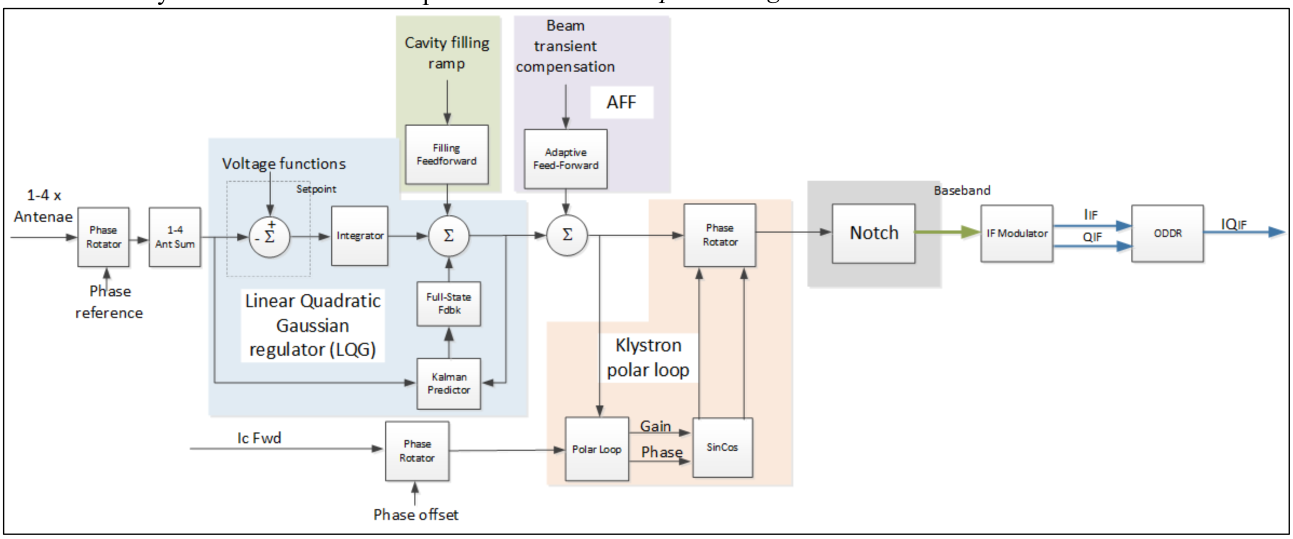

Figure 5: The complete regulation with LQG, AFF and Klystron Polar Loop.

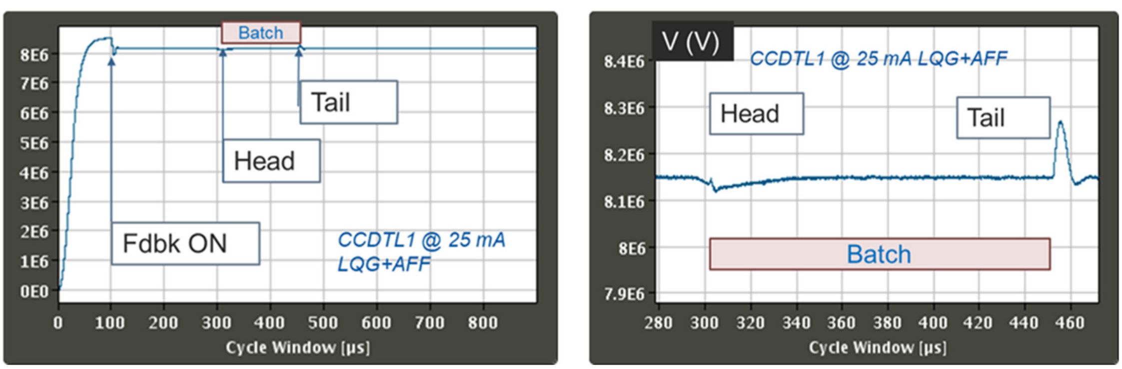

Figure 6: CCDTL1 cavity voltage with LQG and AFF. Full RF pulse (left) and enlargement during the beam passage.

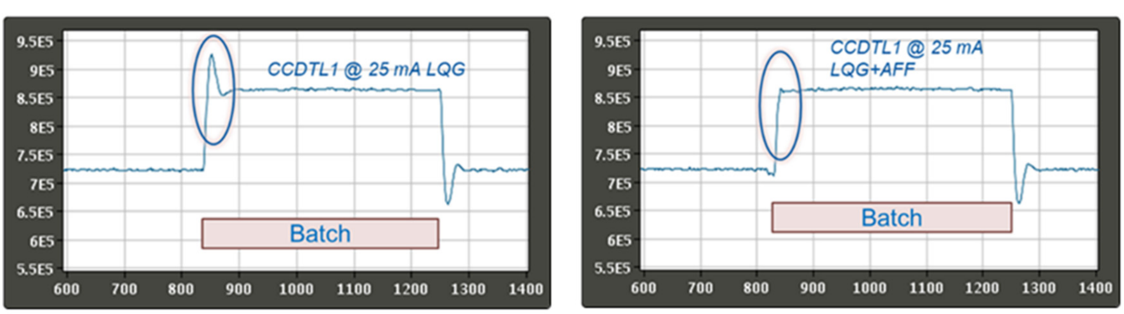

Figure 7: Power required (W) during the beam passage without (left) and with AFF.

THE KLYSTRON POLAR LOOP

So far, we have focused on the compensation of beam loading that is an *additive* perturbation to the cavity field. Another major noise is the modulation of amplifier gain and phase, consequence of the ripples on the High Voltage supply. This noise is *multiplicative*, and therefore best compensated by acting on the LLRF gain and phase shift. RF feedbacks (such as the LQG) are typically adjusted for 50-60 degrees phase margin. They will start oscillating if

ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

the amplifier's phase rotates by that amount. As mentioned above the LEP klystron phase drifts by eight RF degree for one percent drift in the HV. To mitigate this, a solution was developed and implemented in the LHC RF, the Klystron Polar loop (KPL) [12]. It has been included in the Linac4 design. Refer to Fig.5. We compare the forward current at cavity input $I_{c,fwd}$ (measured with a coupler in the waveguide before the cavity main coupler) to the sum of LQG and AFF outputs, and apply gain and rotation to keep the overall gain and phase shift constant, including klystron and circulator. Its reaction time is chosen to be much slower than the LQG, but sufficient to cover slow HV drifts and the rectifiers ripples measured around 10 kHz. As the $I_{c,fwd}$ measurement is taken after the circulator, the loop also compensates for its phase drifts.

ISSUES

The CCDTL and PIMS are multi-cell cavities and have resonances close to the accelerating mode [13,14]. Without mitigation, these will limit the gain at which the LQG remains stable. For the CCDTL the closest two parasitic resonances are at ±1 MHz offset from the 352.2 MHz accelerating mode. A digital notch filter, whose frequency and bandwidth can be adjusted for each cavity, is inserted in the LLRF (Fig. 5). To identify the offending frequency(ies), we drive the klystron with rectangular 352.2 MHz pulses and observe the cavity field, first with Notch filter bypassed (fine blue trace on Fig. 8). That allows for an easy identification of the location of the closest disturbing modes (-0.930 MHz and + 1.02 MHz for CCDTL5). As the LLRF operates on (I,Q) coordinates it will generate symmetric notches. The result of inserting the notch is the black trace on Fig. 8. The parasitic resonances have disappeared.

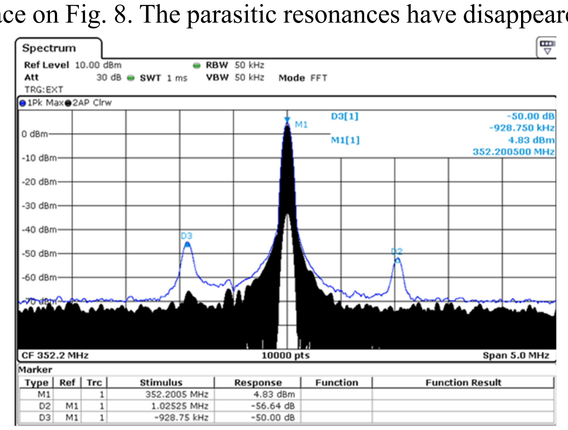

Figure 8: Spectrum of the voltage in CCDTL5 when driving the klystron with 352.2 MHz RF pulses. Without notch (blue trace) and with notch (black background).

Recovering stable operation after a fault was very lengthy until a sequencer was designed for the correct restart of the whole line: interlocks, klystron, HV, LLRF including tuner and regulation loops. If the RF trips during the pulse (for example due to arcing in the cavity), the compensations calculated by AFF and KPL will be very large and will keep tripping the RF on the subsequent pulses. Therefore these regulations are reset in the sequence.

Linac4 must supply beam with different characteristics, achieved by a set of dedicated settings (such as number of batches, batch length, distribution to various PSB rings or dump) that are called *user*. The user changes from pulse to pulse and all LLRF settings are correctly uploaded before the start of the user's pulse. However, the operation crew can change the beam intensity for a *given* user. If we reduce the intensity significantly the AFF correction may trip the RF on the next cycle by causing an arc in the cavity (overvoltage). To avoid this the AFF corrections are automatically reset if the user's intensity is changed. Similarly, the AFF learning is reset if a user is requested while it has not been played for a long time.

CONCLUSIONS

Figure 9 shows the energy at the Linac4 output, with four batches 75 µs long. Let us examine the deviation, along the batch, from the design 160.7 MeV value: we see ±50 keV transient in the first few µs at the head of the first beam batch. That is the result of the uncompensated transient beam loading (Fig. 6, right). The rest of the beam pulse deviates by less than 20 keV. When compared to the 250 keV rms energy spread (design value) of the *individual* 352.2 MHz bunches, the above ±50 keV deviation causes absolutely *no degradation on the transfer to the PSB*.

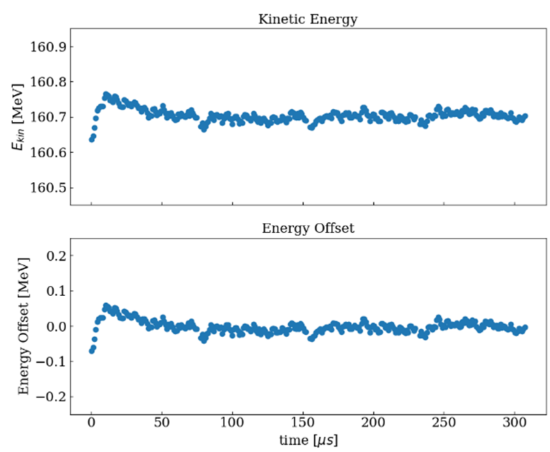

Figure 9: Energy measured at Linac4 output (four 75 µs long batches spaced by 2 µs, 21 mA DC). Each measurement covers 1 µs. Courtesy of P. Skowronski (CERN).

AKNOWLEDGEMENTS

The design of the LLRF started in 2009. At CERN several persons have participated to the LLRF electronics and software developments. In order of appearance in the project (staff, fellow and student): A.K. Bhattacharrya, D. Stellfeld, J. Noirjean, G. Hagmann, D. Valuch, J. Galindo, M. Ojeda Sandonis, V. Costa, B. Kremel, J. Simonin, K. Adrianek, M. Andersen. We thank M. Crofford, L. Doolittle and T. Hardek for their tutoring and hospitality during our visit to SNS in 2009. Thanks also to K. Fung from TRIUMF for extremely helpful discussions on the Adaptive Feed Forward at the LLRF17 workshop. We have also much appreciated the enthusiasm and motivation of the Linac4 operation crew during commissioning.

ISBN: 978-3-95450-215-8 ISSN: 2226-0366 doi:10.18429/JACoW-LINAC2022-TH1AA06

REFERENCES

[1] Linac4 Technical Design Report, CERN, Geneva, Switzerland, CERN–AB–2006–084 ABP/RF, Dec. 2006 https://cds.cern.ch/record/1004186/files/ab-2006-084.pdf

[2] C. Carli, R. Garoby, "Active longitudinal painting for the h-charge exchange injection of the linac4 beam into the PS booster", CERN, Geneva, Switzerland, AB-Note-2008-011 ABP, March 2008 http://cds.cern.ch/record/1092704/files/AB-Note-2008-011.pdf

[3] P. K. Skowro *et al.*, "Summary of the First Fully Operational Run of LINAC4 at CERN", *in Proc. IPAC'22*, Bangkok, Thailand, Jun. 2022, pp. 58-61. doi:10.18429/JACoW-IPAC2022-MOPOST007

[4] S. Ramberger, "First Years of LINAC4 RF operation", presented at the LINAC'22, Liverpool, UK, Sept. 2022, paper MO1PA03, this conference

[5] P. Baudrenghien, J. Galindo, G. Hagmann, J. Noirjean, D. Stellfeld, and D. Valuch, "Commissioning of the Linac4 Low Level RF and Future Plans", *in Proc. LINAC'14*, Geneva, Switzerland, Aug.-Sep. 2014, paper THPP027, pp. 892-894

[6] G. Bellodi, M. Eshraqi, M. Garcia Tudela, L. Hein, J.B. Lallement, S. Lanzone, A.M. Lombardi, P. Posocco, E. Sargsyan, "Alignment and Field Error Tolerance in Linac4", CERN, Geneva, Switzerland, CERN-ATS-Note-2011-021, April 2011. https://cds.cern.ch/record/1342092/files/CERN-ATS-Note-2011-021.pdf

[7] J.Tuckmantel, "Cavity-Beam-Transmitter Interaction Formula Collection with Derivation", CERN, Geneva, Switzerland, CERN-ATS-Note-2011-002 TECH, Jan 2011. https://cds.cern.ch/record/1323893/files/CERN-ATS-Note-2011-002%20TECH.pdf

[8] C. Ziomek and P. Corredoura, "Digital I/Q Demodulator*", in Proc. PAC'95*, Dallas, TX, USA, May 1995, paper RPQ02

[9] D. Boussard, "Control of cavities with high beam loading", IEEE Transactions on Nuclear Science, Vol. NS-32, No.5, Oct. 1985.

[10] G.F. Franklin, J.D. Powell, M. Workman, "Digital Control of Dynamic Systems", Ellis-Kagle Press, 3rd edition, section 9.4.3, pp. 394-396.

[11] M. P. Laverty and K. Fong, "An Iterative Learning Feedforward Controller for the TRIUMF e-linac", in *Proc. LINAC'16*, East Lansing, MI, USA, Sep. 2016, pp. 485-487. doi:10.18429/JACoW-LINAC2016-TUPLR009M

[12] P. Baudrenghien *et al*., "The LHC Low Level RF", *in Proc. EPAC'06*, Edinburgh, UK, Jun. 2006, paper TUPCH195, pp. 1471-1473.

[13] M. Vretenar *et al.,* "Development of a Cell-Coupled Drift Tube Linac (CCDTL) for Linac4", *in Proc. LINAC'08*, Victoria, Canada, Sep.-Oct. 2008, paper MOP008, pp. 67-69.

[14] M. Vretenar *et al.,* "Development Status of the Pi-Mode Accelerating Structure (PIMS) for Linac4", *in Proc. LINAC'08*, Victoria, Canada, Sep.-Oct. 2008, paper THP064, pp. 939-941