

ADAPTIVE FEED FORWARD BEAM-LOADING COMPENSATION EXPERIENCE AT THE SPALLATION NEUTRON SOURCE LINAC *

Kay-Uwe Kasemir, Mark Champion, Mark Crofford, Hengjie Ma,
ORNL, Oak Ridge, TN 37831, USA.

Abstract

When initial beam studies at the Spallation Neutron Source (SNS) indicated a need for better compensation of the effects of beam-loading, a succession of rapid-prototyping and experimentation lead to the development of a simple yet successful adaptive feed forward (AFF) technique within a few weeks. We describe the process and first results.

INTRODUCTION

In order to meet specifications, both feedback and feed-forward methods were part of the initial design for the SNS Low-Level Radio Frequency (LLRF) control system [1]. Specifically, the LLRF hardware included five Programmable Logic Devices (PLDs) and dual Digital Signal Processors (DSPs) with direct access to the RF data sampled at 40 MHz [2]. The current SNS LLRF hardware is an extension of the system originally developed for the SNS front-end [3]: A single Xilinx Virtex II Field Programmable Gate Array (FPGA) implements a 40 MHz control loop with proportional and integral gain. It provides history buffers for the digital output as well as the measured forward, reflected, and cavity fields. Each buffer holds 512 pairs of in-phase and quadrature $\{I, Q\}$ data samples, i.e. complex number waveforms. In normal operation, the resolution is configured to about $2\mu\text{s}$ per sample pair in order to capture the full RF pulse. A Motorola MVME2100 CPU board accesses the history buffers via VXIbus. Software based on the Experimental Physics and Industrial Control System (EPICS) toolkit[4] can present data to network clients via the ChannelAccess network protocol (CA) at up to 60 Hz, both the original $\{I, Q\}$ data as well as computed amplitudes and phases. In practice, history buffer updates are usually reduced to 2 Hz and aligned with beam pulses.

When beam pulses with a length of $40\mu\text{s}$ were accelerated by the cavity, the PID loop barely managed to recover from the beam loading effect by the end of the beam pulse. The need for a feed-forward based approach became especially obvious when trying to exceed beam currents of 20 mA.

*SNS is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy. SNS is a collaboration of six US National Laboratories: Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Thomas Jefferson National Accelerator Facility (TJNAF), Los Alamos National Laboratory (LANL), Lawrence Berkeley National Laboratory (LBNL), and Oak Ridge National Laboratory (ORNL).

FEED-FORWARD CAPABILITIES

The feed-forward buffer in the current SNS LLRF hardware allows control over the first 1.6ms of the RF output with a granularity of $0.4\mu\text{s}$. The feed-forward buffer accepts arbitrary $\{I, Q\}$ waveforms, except that consecutive entries must not vary by more than 0.4% of the full digital output range.

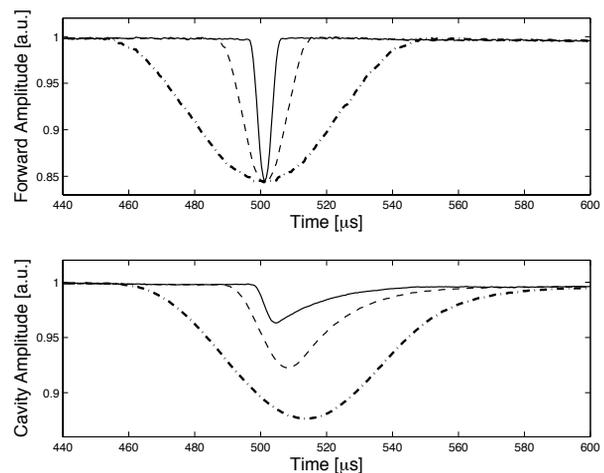


Figure 1: Variations in the SNS DTL4 forward RF amplitude and resulting cavity amplitude.

The feed-forward buffer provides excellent control of the RF drive waveform sent to the klystron. The forward waveform, measured via directional couplers in the waveguide to the cavity, also matches very well with the programmed feed-forward shape. The cavity, however, exhibits a delayed response due to its limited bandwidth. Fig. 1 shows how symmetric “dips” in the forward RF sent to the SNS Drift-Tube-Linac (DTL) cavity number 4 result in variations of the cavity amplitude that are distorted and delayed.

Assuming that the cavity amplitude over time, $c(t)$, follows this differential equation:

$$\frac{d}{dt}c(t) = \frac{1}{\tau}(f(t) - c(t)), \quad (1)$$

where $f(t)$ is the forward RF amplitude and τ is the time constant of the cavity, it is possible to simulate the observed cavity behavior with good agreement, compare Fig. 2. It is worth noting that in both experiment and simulation, only the width of the “dip” in the forward amplitude was changed. Its location and depth stayed constant. The cavity field, however, exhibits a “dip” that changes in width, amplitude and location.

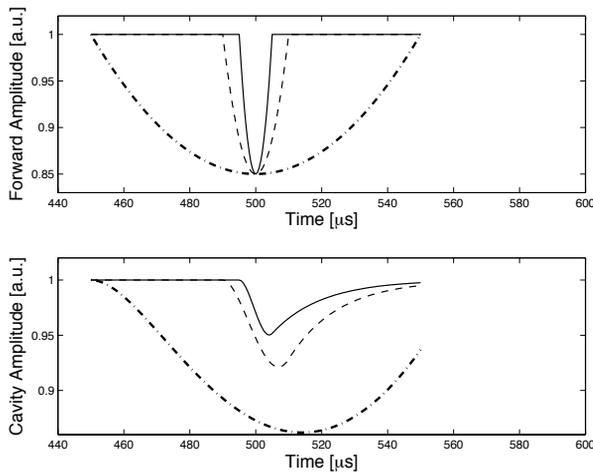


Figure 2: Simulated variations in the forward RF amplitude and the resulting cavity amplitude, obtained by numerically integrating eq. 1 for $\tau = 15\mu s$.

In a preliminary attempt, we reduced the beam loading effect significantly by manually placing a simple “step” of appropriate amplitude and phase in the feed-forward buffer, thereby increasing the RF drive in anticipation of the beam pulse. This, however, was impractical, because the exact position of the step proved to be critical, yet there was no obvious way to determine this position from available data without trial and error. Once the beam chopper becomes operational, the LLRF will have to compensate known but otherwise arbitrarily shaped beam profiles, which are unlikely to be handled by a single “step”-type feed-forward.

ADAPTIVE FEED-FORWARD

We investigated an adaptive approach similar to the one described in [5]: Assuming good reproducibility of the system response to a given RF drive, we adjust the feed-forward buffer for the next pulse based on errors observed in preceding pulses.

A first prototype, developed in MATLAB [6], read the cavity amplitude via CA, determined the beam-induced error by subtracting the cavity amplitude during the beam from the average cavity amplitude during a configurable time window before the start of beam. Instead of determining the optimal feed-forward shape for the desired error correction from eq. 1, the error waveform itself was used as the feed-forward waveform, only shifted in time by a configurable offset to compensate for the delayed cavity response.

The phase of this correction needs to be chosen appropriately in order to actually *reduce* the error. The correct phase is influenced by time delays due to cable length as well as the klystron gain curve. Software used to close the control loop determines this “gain rotation”, the phase relation between the programmed output and the resulting cavity field. The control loop then adds minor corrections within the

duration of one RF pulse. The MATLAB script assumed that these mechanisms have determined the correct output phase up until the start of the beam, and continued to use that same output phase for the error correction.

This approach was used manually and on-demand. About every 10 seconds, a fraction of the error was added to the feed-forward buffer, shifted by a few microseconds. It resulted in sufficient amplitude correction, but failed to correct beam-induced phase deviations of the cavity field.

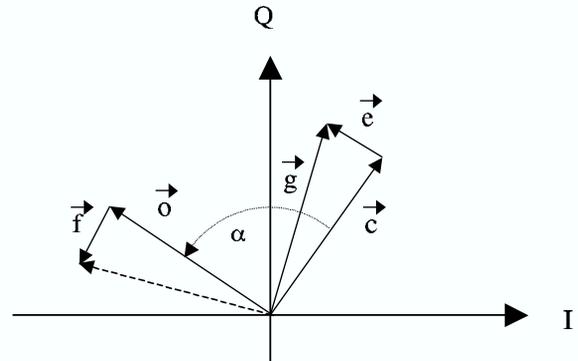


Figure 3: Illustration of vectors involved in adaptive feed-forward calculation, refer to text.

A second MATLAB prototype determined the feed-forward correction for both amplitude and phase based on the complex cavity waveforms. The AFF goal, \vec{g} in Fig. 3, is defined as the average cavity $\{I, Q\}$ during a configurable time window before the start of beam. Ideally this should match the amplitude and phase set point of the LLRF control loop, but by using measured cavity data, adaptive feed-forward can also be tested on cavity disturbances while operating in open loop. The complex error \vec{e} is calculated by subtracting the cavity vector during the beam, \vec{c} , from the AFF goal. This error within the general $\{I, Q\}$ coordinates of all measured waveforms needs to be rotated into the coordinates of the current output, \vec{o} , in order to compensate cable and control loop delays. This rotation (angle α in Fig. 3) would ideally match the aforementioned gain rotation, but it is determined dynamically from the averaged cavity and output vectors before the start of beam, allowing independent studies of the control loop gain rotation without interfering with AFF operation.

The intermediate result is a waveform of feed-forward vectors $\vec{f}(t)$ with a resolution of about $2\mu s$. Linear interpolation is used to update the feed-forward buffer at a granularity of $0.4\mu s$, where a fraction $K \cdot \vec{f}(t + t_s)$ is added to the current feed-forward buffer content at time t , again using a time shift t_s to compensate for the delayed cavity response. The proportional gain factor K , typically around 0.05, determines how fast the AFF algorithm learns.

This approach was first tested by manually triggering updates of the feed-forward buffer after visual inspection of the calculated correction, resulting in updates of the feed-forward buffer about once every 10 seconds. Since

the results were satisfactory, a C-code implementation was loaded onto the MVME2100 CPU board. It runs at about 2 Hz, triggered by the SNS timing system, which also provides information on whether the last RF pulse should be ignored by AFF because beam was not available for the full duration of the pulse.

When this version of the code ran in a laboratory test setup, no instabilities like self-oscillations were observed. On the accelerator, however, noise in the measured cavity data often was amplified by the AFF method. After experimenting with several finite impulse response filters, a simple nine-element sliding average was applied to the feed-forward buffer after each update, reducing the amplitude of higher-frequency components that are introduced by noise.

Fig. 4 shows how the application of adaptive feed-forward to DTL4 reduced the beam-induced cavity amplitude error from about 7% down to 1%, while at the same time the error in phase was reduced from 2° to less than 1° .

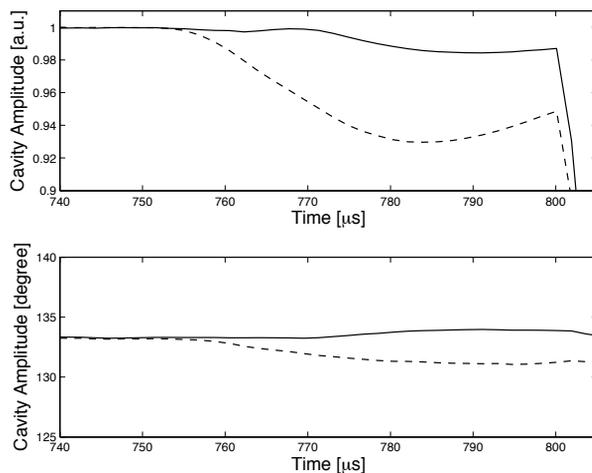


Figure 4: Example of SNS DTL4 cavity amplitude and phase under the influence of beam loading without (dashed) and with adaptive feed-forward (solid lines).

POSSIBLE IMPROVEMENT

Our method of shifting the error waveform in time to compensate for the delayed cavity response is clearly a simplification. The proper choice of a time shift value is not obvious. It is currently determined experimentally and needs to be adjusted in response to different beam parameters. When we convert eq. 1 to

$$f(t) = c(t) + \tau \Delta c(t) / \Delta t, \quad (2)$$

we can compute a forward waveform $f(t)$ for a desired cavity correction $c(t)$ that should be very close to the ideal correction, requiring only the known cavity time constant τ . In addition to a proportional fraction of the most recent error, that cavity correction can also include a fraction of the cavity error sum, accumulated over previous RF pulses. We are currently investigating such a nonlinear AFF method with proportional and integral gain on a test cavity.

CONCLUSION

Within a very short time frame, we have implemented an adaptive feed-forward algorithm which proved instrumental in allowing SNS beam studies beyond currents of 20 mA. By utilizing the repetitive nature of the accelerator and the good pulse-to-pulse reproducibility through suitable averaging and filtering, we can develop and test ideas within a comparably slow MATLAB architecture. We achieved acceptable results on the existing hardware, which is simpler than the one originally planned for the SNS LLRF.

Further tests are required to determine how well the methods work for different beam shapes or for the superconducting SNS cavities.

REFERENCES

- [1] D. Rees, M. Lynch, P. Tallerico, W. Roybal, K. Cummings, T. Hardek, J. Bradley, W. Reass, A. Regan, "The RF System Design for the Spallation Neutron Source", PAC01, Chicago, June 2001.
- [2] A.H. Regan, S.I. Kwon, T.S. Rohlev, Y.M. Wang, M.S. Prokop, D.W. Thomson, "The SNS Linac RF Control System", PAC01, Chicago, June 2001.
- [3] L. Doolittle, C. C. Lo, M. Monroy, A. Ratti, "The SNS Front End LLRF System", LINAC2002, Gyeongju, Korea, 2002.
- [4] "Experimental Physics and Industrial Control System", [Http://www.aps.anl.gov/epics](http://www.aps.anl.gov/epics).
- [5] Dazhi Li, Yanshan Wang, Yonggui Li, Jialin Xie, "Beam-Loading Compensation for RF Gun by Feed-Forward Control System", Intl. Journal of Infrared and Millimeter Waves, Vol. 22, No. 8, August 2001.
- [6] For Info on Matlab see <http://www.mathworks.com>.