A Digital Long Pulse Integrator*

J.D. Broesch, E.J. Strait, R.T. Snider, M.L. Walker
General Atomics
P.O. Box 85608, San Diego, California 92186-9784

ABSTRACT

A prototype digital integrator with very long integration capabilities has been developed and field tested on an inductive magnetic sensor on the DIII–D Tokamak. The integrator is being developed for use on ITER with a pulse length of 1000 s, and has direct applications for other long pulse Tokamaks. Inductive magnetic sensors are routinely used on existing Tokamaks, are well understood, and are extremely robust, however, they require integration of the signal to determine the magnetic field strength. The next generation of Tokamaks, will have pulse lengths of 1000 s or longer, require integrators with drift and noise characteristics compatible with the very long pulse lengths. This paper will discuss the architecture, algorithms, and programming of the Long Pulse Integrator (LPI). Of particular interest are the noise control and the built-in offset correction techniques used in this application.

PROBLEM DEFINITION

Integrators in tokamak inductive sensor systems are used to convert acquired voltages V from magnetic probes to field strength B using the relationship $B = (1/NA) \int V\,dt$ where N is the number of turns and A is the cross-sectional area of the probe. Long pulse operation simultaneously requires a small integrated error due to integrator drift and sufficient resolution to detect slow physical drifts of the plasma. The integration circuit must be able to accurately integrate signals with a very large dynamic range from ≈0 V for plasma equilibrium operation to on the order of 1000 V for disruptions.

The major electronic problem is to distinguish small physical drifts of the plasma from the inherent drift of acquired signals due to small uncorrected offsets in the acquisition circuit. In ITER, it is expected that during normal plasma operation a typical integrated magnetic probe signal will have values on the order of 0.1 to 1 Vs [1]. A perturbation analysis with equilibrium reconstruction codes indicates that ITER reconstructed equilibria are accurate to within 1 to 2 cm for 3% rms noise on integrated diagnostics (including magnetic probes) [2]. From these estimates we synthesize an approximate requirement on integrator performance. We require that the total error in the integrated signal due to drift at any time during a 1000 s pulse be just 3% of the smaller 0.1 Vs value. That is, drift $<(0.03)(0.1\text{Vs}) = 3$ mVs.

At the same time, we wish to be able to detect small "constant" plasma drifts down to the error level of the integrator. We state this in the following way: Define $V_{min} = 0.03(01\ \text{Vs}/1000\ \text{s}) = 3\ \mu\text{V}$. For any signal V(t) such that $V(t) > V_{min}$, for $0<t<T$, the integration process acting on the signal V(t), $0<t<T$, should give an answer which is greater than $V_{min}T$. The significance of this requirement on integrated signal resolution becomes clear later when we discuss digital integration schemes.

Obviously, uncompensated drift can be made a smaller fraction of the sensor voltage signal by amplifying the signal voltage level using larger number of turns N and/or area A subject to certain physical constraints. Very simple assumptions on the values of N and A were made in deriving the above estimated requirements.

CURRENT APPROACHES TO INTEGRATION

Over short periods of time, on the order of seconds to a few minutes, simple analog circuits can be used to efficiently perform the integration. Beyond this time frame the inherent drift of the analog integrator causes a much larger output than the small signal being measured. While techniques exist for minimizing this offset, the problem cannot be totally eliminated in analog circuits. Careful trimming of component values combined with efficient temperature control can achieve good short term results. Integration on the order of minutes is about the best that can be achieved using these methods in cost sensitive applications, however.

Chopper stabilized operational amplifiers, on the other hand, have good stability but a narrow bandwidth. This limits their use for measuring the rapidly changing signals found during strong ramp-ups and disruptions. Further, it can be difficult to efficiently incorporate these amplifiers into traditional instrumentation amplifier configurations. Thus it appears that the requirements for magnetic sensor signal integration cannot be satisfied with purely analog instrumentation.

LPI ARCHITECTURE

A number of approaches were initially evaluated to achieve the required stability for the Long Pulse Integrator. Analog, digital, and hybrids of analog and digital techniques were evaluated. It was decided that the approach that offered the best chance of success was a digital technique based on a modified classical auto-calibration scheme. This approach has been discussed previously [3].

The system architecture of the magnetic inductive sensor and the LPI is shown in Fig. 1. The system is divided into three main components: the magnetic probe, the signal conditioning module, and the digital signal processing sub-system. For a detailed discussion of the inductive magnetic probe see [4].

*Work supported by the U.S. Department of Energy under Contract No. DE-AC03-89ER52153.

The signal conditioning circuit consists of an RC filter, an electronic switch, an instrumentation amplifier, and an operational amplifier. The RC filter may be viewed in two ways: first, it may be thought of as performing a role similar to a conventional anti-aliasing filter; alternately, it may be viewed as a passive integrator that continues to acquire the signal during the sampling interval and during the period in which the baseline measurements are being made (see below).

The purpose of the switch following the RC filter is to provide periodic measurement of the baseline offset voltage of the electronic circuit which a DSP algorithm can use for removing offset voltage from the diagnostic signal. The analog switch was fabricated using two high-performance MAX303 [5] analog switches in series. The switches were required to meet the high switching rates used by the integration algorithm. This switch configuration was used to achieve high isolation between signal and baseline channels.

The instrumentation and operational amplifiers provide dual low gain (x1) and high gain (x10) signal paths, respectively. The high gain(x10) path is used by the DSP algorithm to obtain higher resolution (more digitizer counts/signal volt) measurement of low level signals caused by small variations in plasma position while the low gain (x1) path provides sufficient dynamic range to handle data during disruptions. The instrumentation amplifier is an AMP-05 [6] chosen for its low drift characteristics and its flexible design attributes.

In this architecture, the RC filter serves the dual purpose of noise anti-aliasing filter and reducing the dynamic range of large narrow (≈1000 V) data spikes characteristic of disruptions to approximately ±10 V. It was assumed that the filter produces maximum signal levels on order of 1 V or less for normal plasma operations. These numbers were used to define the full scale voltages for the design, so that the dynamic range of data at input to the ADC is always ±10 V if the integration algorithm uses the high gain (x10) amplifier data during normal operations and low gain (x1) amplifier data during disruptions.

The DSP processing system is a dual processor configuration. The actual signal processing is done by a TMS320C30 [7]. The applications development and man/machine interface is handled by an industrial IBM compatible PC. The system operates in three possible modes:

Mode 1: The data are acquired and processed in real time by the TMS320C30 DSP.

Mode 2: The data are acquired and transferred in real time to the PC. The PC stores the data to the hard disk for post acquisition analysis. Operation in this mode is limited to approximately 10,000 samples per second. Normally, two channels of data are transferred, with each channel having a maximum of 5000 samples per second throughput.

Mode 3: This is a combination of the first two modes, with some real time processing done by the DSP, and the results logged by the PC.

For this testbed, the integrated signal was converted back to the analog domain and output via the DAC. Performance of the LPI was then compared to the existing DIII–D analog integrators. These comparisons showed that the data from the LPI closely matched the output of the analog integrators for individual plasma discharges. It is anticipated that production versions of the LPI would eliminate the DAC stage. Instead, the output would be sent to the instrumentation and control computers in a high speed digital format.

The objective of the LPI is to provide real time operation. The ability to save the raw data for later analysis proved critical for testing, however. By having the raw data available, it is possible to reconstruct how the integrated signal was obtained. More importantly, however, having the raw data set available made it possible to try a number of different algorithms and compare the results in a consistent manner.

LPI ALGORITHM

The flexibility of the architecture shown in Fig. 1 was used to evaluate a number of integration algorithms. The algorithm which produced the best results is the following. The baseline and the input signal were integrated separately and the difference taken on a point-by-point basis:

$$y_k = \frac{1}{R}\sum_{n=0}^{k}\left(x_n - X\right) - \frac{1}{R}\sum_{n=0}^{k}\left(b_n - B\right) \qquad (1)$$

where x_n is the digitized input signal, b_n is the digitized baseline measurement, offset X and B are averages of x and b for some number of samples prior to starting the integration, and R is the sampling rate.

Integer arithmetic was used. The calculated offset is typically noninteger, so a technique analogous to conventional pulse width modulation (PWM) techniques was used to produce a result that could reliably correct for fractional bit offset errors. Two correction coefficients c_f and c_c are defined as the greatest integer less than, and smallest integer greater than the calculated offset, respectively. The appropriate correction coefficient is then used in place of the offset in Eq. (1) a number of times which is inversely proportional to its distance from the calculated floating point offset. For example, assume the calculated offset is 33.25 counts. Two correction coefficients would then be generated: c_c = 34.0 and c_f = 33. Assuming the correction is applied over cyclic intervals of 100 samples, the corrected signal would be $y_n = x_n - c_f, n = 0\ldots74,\ y_n = x_n - c_c,\ n = 75\ldots99$, where y_n is the value to be integrated and x_n is the input signal value. This technique produced superior results to either a floating point or fixed point computation — even though the latter two techniques would appear to have an advantage in the degree of resolution obtainable. Why this is true is not entirely clear at this point in our study.

A typical integration of a DIII–D magnetic probe signal using the low gain (x1) amplifier is shown in Fig. 2. The data rate on each of the input and baseline signals was 5 kHz. Each point on the plot in Fig. 2(a) represents a block average of 2500 data points. Only one of every 2500 points of the integrated signals is plotted in Figs. 2(b–d). The signal in Fig. 2(d) (when correctly scaled) represents measured field which should return to zero after each plasma discharge. There are two plasma discharges during the 1300 s – one at

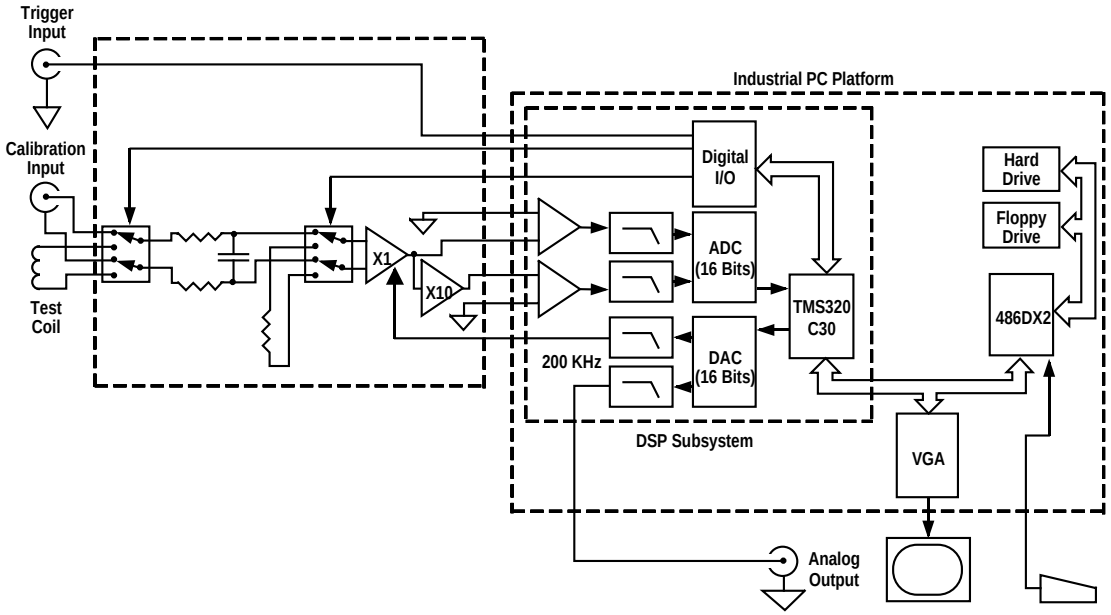

Fig. 1. Long Pulse Integrator (LPI) architecture.

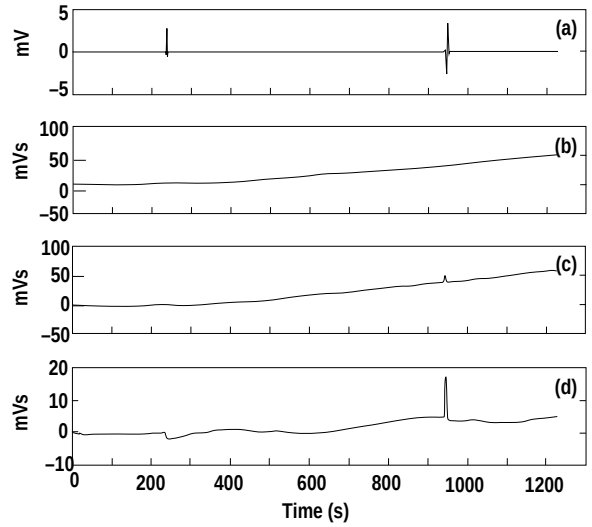

Fig. 2. Example analysis of data acquired on DIII–D: (a) raw input signal from magnetic probe, (b) integrated baseline signal, (c) integrated input signal, (d) integrated signal minus integrated baseline signal.

about t=235 s and the other at about t=940 s. The first discharge resulted in a disruption and so can barely be seen in the integrated signal plots (Figs. 2(c) and (d)). Fig. 2(d) shows a drift error of approximately 5 mVs after 1200 s and also at about 900 s. This value exceeds the stated requirement of 3 mVs but the high gain amplifier path was not exploited for these tests. Note that the magnitude of the "blip" representing an ITER plasma discharge would be about 10 times as large as the maximum signal size which represents a DIII–D plasma discharge in Fig. 2.

Fig. 3 shows results of bench tests using the high gain (x10) path to integrate a very small signal level. The input signal was nominally constant zero with the usual environmental noise giving typical ADC output values in the range of ±8 counts. The data rate on each of the input and baseline signals was 5 kHz. Only one of every 2500 points of the integrated signals is shown in each plot. The worst case drift error of approximately 1 mVs clearly satisfies the required 3 mVs limit. Even for the worst of these bench tests, drift error never exceeded 2 mVs. However, our data set is not large enough at this point to completely characterize the drift error performance.

Preliminary tests have been conducted to determine whether the resolution requirement is satisfied. A simple analysis shows that nominal resolution of the ADC/high gain amplifier channel is given by 2 V/(2^{16} counts) = 30.5 μV/count. Jitter on the signal due to noise provides an improvement in effective resolution to approximately 1/4 of that number.

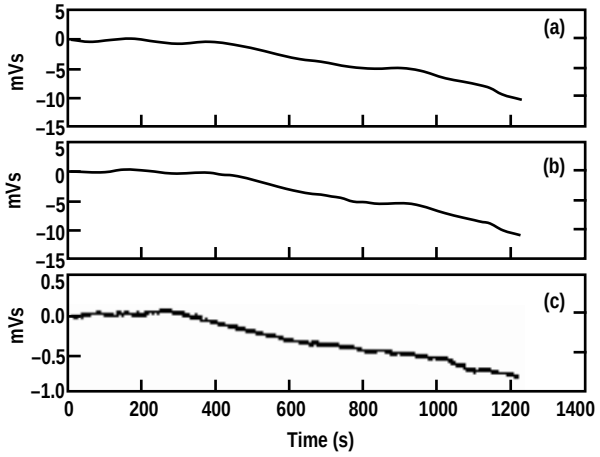

Fig. 3. Bench test of small signal integration: (a) integrated baseline signal, (b) integrated input signal, (c) integrated signal minus integrated baseline signal.

This still does not quite meet the stated resolution requirement of 3 μV. However, it is not unreasonable to assume that could be used for the operational amplifier. This implies a data during normal operations can be taken to have a maximum value on the order of 100 mV so that a gain of 100 3.05 μV nominal resolution, even before resolution improvement by noise jitter. Current work uses this higher gain setting.

ALTERNATIVE ALGORITHMS

Variations on the algorithm discussed above were also attempted. Feedback from the DAC to the instrumentation amplifier shown in Fig. 1 is often used to execute an auto-zeroing cycle where the input offset is automatically nulled by applying a correcting offset voltage to the output via the DAC. Auto zeroing the circuit produced an output from the instrumentation amplifier that had several bits more noise than when simply grounding the ADC reference pin.

Ideally, this auto-zeroing should not be necessary. The offset is common to both the signal and the baseline value. By simply subtracting the baseline value b_k from the input signal value x_k on a point-to-point basis and integrating

$$y_n = \frac{1}{R} \sum_{k=0}^{n} \left(x_k - b_k \right) \tag{2}$$

the initial offset as well as drift should both be corrected for. Here R is the sample rate in samples per second. In practice, however, this ideal was not achieved. The difference x_k - b_k typically produced noise values in the range ±8 counts. This occasionally increased up to ±16 counts. Even for an ideal input of zero volts a real world ADC will typically produce an output that varies by several counts above and below zero. A common practice where high precision is required is to oversample the signal, then average samples down to produce a more stable reading. This technique works well over a relatively large range of sampling rates; factors from three to several hundred are common.

When applied to a zero input signal with simple circuit self-noise, the difference (x_k - b_k) in Eq. (2) resulted in a signal with a non-zero-mean distribution. Statistical analysis indicated that the noise statistics were apparently non-stationary. For example, test sample sets taken over 100 s at sample rates of 10,000 samples per second result in an average value between one and three counts. This produced an effective offset which caused the output of the digital integrator to drift. The drift was small in comparison to typical analog integrators, on the order of 50 μV/s, but was too high to meet the requirements of the anticipated ITER application. It was observed, however, that the drift resulting from separate integration of the baseline and the input signal tended to follow the same trend line. It was this observation which led to the algorithm described in the previous section.

CONCLUSION AND FUTURE WORK

Our work clearly demonstrates that a digital integrator capable of accurate integration over 1000 s periods is realizable. Additional work is required to more precisely quantify the requirements for integration errors and resolution on ITER, to develop and quantitatively evaluate the optimum integration algorithm, and to evaluate the LPI performance under a wider range of conditions. In the near term a high neutron flux experiment is planned at the Brookhaven National Laboratory fission facility that will increase the LPI database.

In a sampled data (digital) system, the integrated error requirement directly conflicts with the simultaneous demand for sufficient resolution to detect slow physical drifts of the plasma. An apparently simple solution to the problem of drift errors is to reduce the number of bits used to represent the signal so that any uncompensated voltage offset is masked by the discarded low order bits. However, this also masks any small constant voltage representing a slow drift of the plasma.

In contrast to analog integrators, we can compensate very well for intrinsic offset in the electronic circuit using a digital scheme. The remaining difficulty comes from noise in the signal which has a nonzero mean. This noise mean provides an *effective* offset which integrates to a drift error in a similar manner as the classical electronics offset. This noise mean is not drift in the classical sense because it appears even when data is acquired from the ADC with all conditioning electronics disconnected and inputs to the ADC shorted. The standard trick of averaging down an oversampled signal produces a signal with lower noise variance, but does not affect the mean. It is desireable not to average too much in order to maintain sufficient noise variance to improve resolution.

Specific areas of additional investigation include understanding the distribution characteristics of the noise, the effect of this noise on floating point and fixed point calculations, and optimizing the hardware design to realize the most cost effective implementation. A particularly interesting area for investigation is found in the selection of the ADC. The LPI's standard SAR instrumentation ADC could be replaced with a delta-sigma converter Our work to this point suggests the increased latency would be acceptable, while the drift rejection characteristics would be improved.

REFERENCES

[1] E.J. Strait, private communication.
[2] J. Leuer, et. al., The Influence of Blanket Magnetic Probes on Equilibrium Reconstruction in a Full Current ITER Plasma, to be published in Nuclear Fusion
[3] DIII–D Physics Memo D3DPM No. 9012, A Hybrid Digital-Analog Integrator Concept.
[4] T. Hoddap, et al., *Magnetic Diagnostics for Future Tokamaks*, this conference.
[5] Maxim Semiconductor.
[6] Burr Brown.
[7] Texas Instruments.