Fermilab FERMILAB-Conf-02/290-E November 2002

The DO Solenoid NMR Magnetometer

R. Yamada, S. Hansen, T. Kiper, J. Lofgren, T. Regan,
Fermi National Accelerator Laboratory, Batavia, Illinois, USA*
G. Boero
Institute of Microsystems, EPFL-Swiss Federal Institute of Technology CH-1015 Lausanne, Switzerland

***Abstract*-- A field monitoring system for the 2 Tesla Solenoid of the D0 detector is described. It is comprised of a very small NMR probe cabled to a DSP based signal processing board. The design magnetic field range is from 1.0 to 2.2 Tesla, corresponding to an RF frequency range of 42.57 to 93.67 MHz. The desired an accuracy is one part in 10^5. To minimize material in the interaction region of the D0 detector, the overall thickness of the NMR probe is 4 mm, including its mounting plate, and its width is 10 mm. To minimize cable mass, 4mm diameter IMR-100A cables are used for transmitting the RF signals from a nearby patch panel 25 meters to each of four probes mounted within the bore of the solenoid. RG213U cables 45 meters long are used to send the RF from the movable counting house to the patch panel. With this setup, the detector signal voltage at the moving counting room is in the range of 250-400 mV.**

I. Introduction

As part of the Run II upgrade, a superconducting solenoid magnet has been installed in D0 detector [1]. For this magnet, we have developed a low cost NMR magnetometer system. Four of these magnetometers are used for continuous monitoring of the absolute field at four points within its bore. The nominal magnetic field is two Tesla, within a cylindrical volume 1.2 meters in diameter and 2.6 meters in length [2]. The solenoid is mounted physically and magnetically separated from a conventional toroidal magnet. With this layout, the field decays rapidly beyond the end of the active volume, but the field distribution is quite uniform within the bore in the vicinity of the probes. Commercially available NMR probes have a thickness of typically greater then one cm, with an equally bulky cable. We therefore chose to use a small probe originally developed at The Swiss Federal Institute of Technology, Lausanne, Switzerland [3].

We have designed and built a VME module, which performs a frequency domain search for an NMR signal and locks to any signal found. If lock is achieved, the NMR frequency is measured and the result can be read over the VME bus. The hardware design and software algorithms written for automatic scanning and locking are described. We use a proton sample for the NMR signal, which has a frequency to field strength relation of 42.577340 MHz/Tesla.

*This work is supported by the U.S. Department of Energy under contract DE-AC02-76H0300.

II. System Description

A. Miniature Probe

The probe is assembled on a printed circuit board, 12 mm wide and 83 mm long, and has been tested in static magnetic fields between one and two Tesla. The measured precision is better than 5 ppm [3]. The probe assembly is surrounded by a copper shield and has a total thickness of 4 mm. This assembly is mounted on a 1.6 mm thick 2.5 cm wide, 1.3 m long G10 wand installed in the interior bore of the D0 solenoid.

The NMR sample element of this probe is a small cylinder of natural rubber 4.5 mm in diameter and 2.5 mm high, placed on a 4.6 mm diameter planar coil of 7 turns. The coil is placed in parallel with a variable capacitance (vari-cap) diode to create a tunable resonant circuit. By varying the bias voltage of this diode over a range of 2 to 12 volts, we can achieve resonance over a span of 50 MHz to 100 MHz in our system. The planar coil is placed with its axis perpendicular to the direction of the magnetic field.

When the coil is driven at the NMR frequency of the sample, some of the proton nuclei in the sample begin to change the orientation of the their nuclear spins. These changes absorb energy from the field generated by the coil, which appears as a small reduction in the magnitude of the voltage developed across the resonant circuit. The frequency of the RF signal can be modulated around this inflection point to accurately determine the point of NMR resonance. Between each sweep of the RF source across the NMR frequency, the proton nuclei need some time to return to their nominal state. We have found that 20KHz of FM modulation with a sweep rate of 25Hz gives satisfactory results. Fig.1 shows the schematic for the NMR probe. The Planar coil and the vari-cap diode form a parallel resonant circuit, while the two Schottky diodes form a voltage doubler type AM detector. The operational amplifier is configured as bandpass filter to exclude high frequency noise and to differentiate the NMR signal, which is thus changed from a valley into a bipolar inflection.

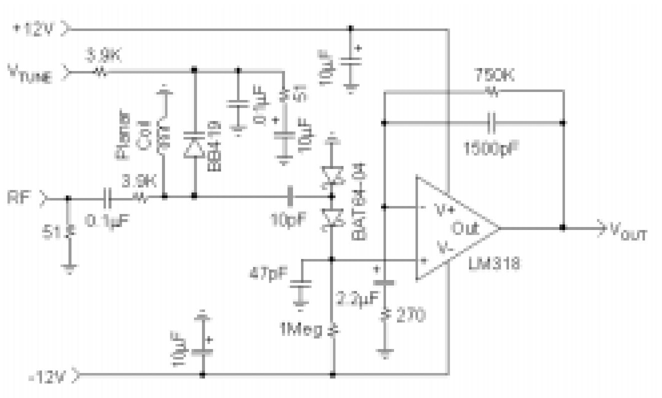

Fig. 1. Schematic of the probe assembly

B. *Control Board*

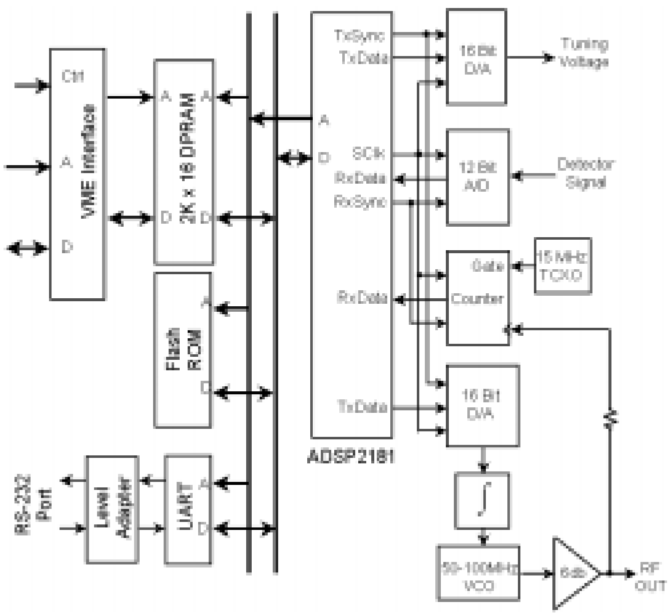

Fig. 2. Block Diagram of NMR Control Board

Figure 2 shows a block diagram of the control board cabled to the NMR probe. This board supplies the RF drive for the resonant circuit in the probe, the tuning voltage for the var-cap diode, the power for the probe amplifier and digitizes the signal returning from the probe. Setup parameters and NMR data are exchanged over the VME bus. Two serial 16 bit D to A converters and a 12 bit serial A to D converter are used. A counter (also with a serial interface) with its input attached to the VCO output driver is gated by a temperature compensated crystal oscillator (TCXO) with a specified stablity of 2 ppm. The crystal has a frequency trim adjustment which we set using a spectrum analyzer equipped with an ovenized crystal as its frequency standard. One of the DACs drives an integrator, which in turn drives the tuning input of a voltage-controlled oscillator (VCO). An RF driver whose output is about 2V rms buffers the VCO output. The integrator extends the dynamic range of the DAC by removing its constant term. The board is a 6U size VME slave with a 16 bit data path using standard 24 bit addressing. The information accessible from VME includes the board status (e.g. locked on resonance or scanning), the resonant frequency if lock has been achieved and the field strength in units of micro Tesla. Optionally, a 256-point record of the ADC values taken during one RF modulation can be read out to give an idea of the signal quality coming from the probe. VME was chosen because the D0 control and monitoring system is VME based. The DSP program is stored in a flash ROM chip that can be reprogrammed with code sent over the RS-232 connection.

III. Software Algorithms

A. *Vari-cap diode biasing*

The bias voltage on the capacitance diode must be adjusted to maintain resonance even as the RF frequency is being changed. A square wave modulation with in the range of 0.5-1 mV at a frequency of about 2kHz is superimposed onto the nominal tuning voltage. The DSP processor does this modulation by writing to the D/A converter that drives the tuning diode. The feedback voltage from the tank circuit during this modulation can then be used to determine whether the circuit is above or below resonance. The digitized values of the signal returning from the probe from one half of the modulation cycle are summed and compared to the sum returned from the other half. This difference of sums is appropriately scaled and added to the nominal DAC value for subsequent modulation cycles. Using this algorithm, the circuit will always move toward resonance. When the algorithm has set the bias at resonance, the difference between the sums from each half of the modulation will be zero, resulting in a constant voltage. This is the digital form of a lock-in amplifier.

B. *Inflection finding algorithm*

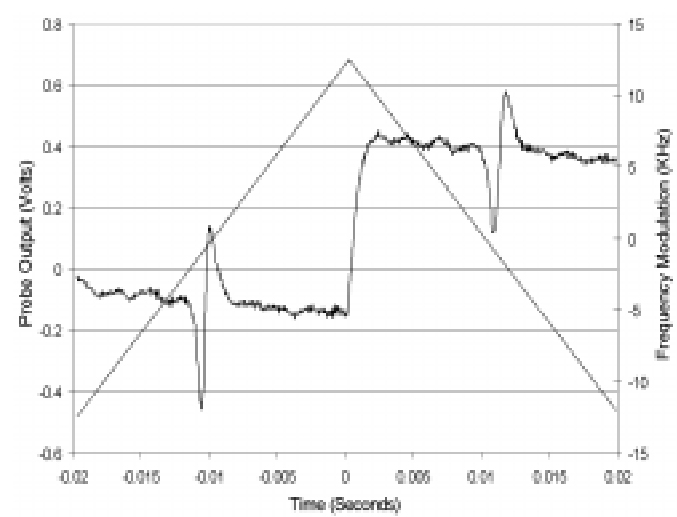

Fig. 3. Plot of the NMR Signal at the ADC input

The first step in finding, tracking, and calculating the NMR frequency is to develop a method to determine what sort of signal qualifies as an NMR resonance. Fig. 3 shows a good quality NMR signal with a graphical representation of the frequency modulation. As can be seen in the figure, regardless of the direction of the frequency modulation, the NMR signal is characterized by a sharp downward inflection of "greater than usual" magnitude, followed immediately by an approximately equal upward inflection. These criteria are used for triggering on the NMR signal. The DSP is updating a number of calculations at a 50 KHz rate. For this reason, the simplest possible algorithm is required in order not to overload the processor. An algorithm that calculates the difference between the most recent ADC samples and the longer-term average level of the signal is used.

TABLE 1.
64-WORD ADC SAMPLE BUFFER

	Sample #	Calculation
Earliest	0	Averaged
	:	to Establish
	31	Baseline
	32	
	:	Ignored
	59	
	60	Averaged
	:	for latest
Most Recent	63	Trend

Table 1 shows the circular buffer used for storage of the 64 most recent ADC samples. The earliest 32 samples in the buffer are averaged to determine the long-term average baseline voltage of the signal, and the 4 most recent samples are averaged to determine the latest trend in the signal. The recent trend value is then subtracted from the baseline value to obtain the differential filtered value. Fig. 4 shows the filtered version of the input signal shown in Fig. 3. A threshold

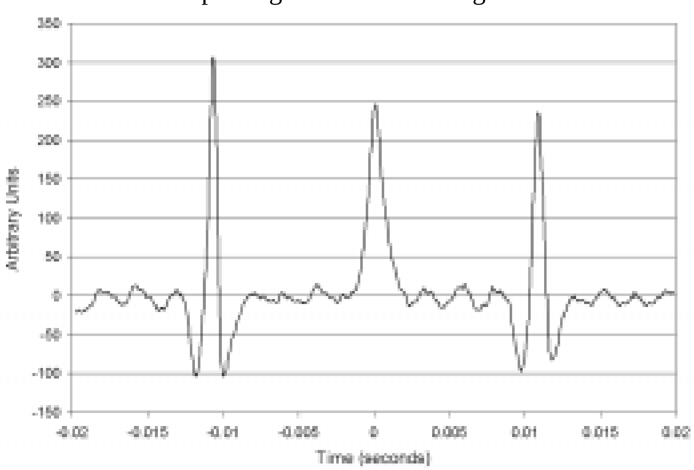

Fig. 4. Plot of the filtered NMR Signal

which is a software controllable parameter is applied to the filtered signal to determine what is an inflection of "greater than usual" magnitude. We have found empirically a setting of 20% of the maximum inflections gives acceptable results. Figure 5 shows a flowchart of the algorithm used to find an inflection. The processor waits for a new ADC sample, and then runs the filter on the current contents of the ADC buffer. If the new filtered value passes the negative threshold, a counter is incremented. If it does not, the same counter is reset. If five consecutive negative excursions are found, then the processor proceeds to look for five consecutive positive excursions. If the positive excursions are not found within 50 sampling intervals (the "Neg. to Pos. Timeout" decision box) then the counters are reset and the search starts over.

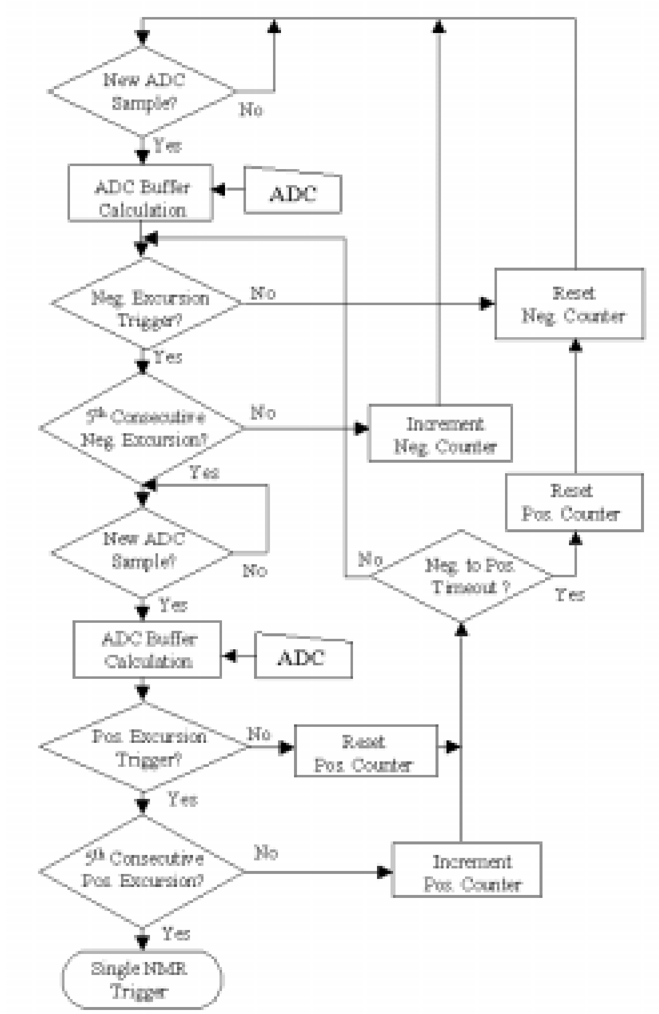

Fig. 5. Flow chart for inflection finding algorithm

C. *Scanning algorithm*

To search for an NMR resonance, the RF frequency is ramped up and down using a "resonance-terminated frequency modulation" scheme. At startup, the frequency generator is set to ramp up until an inflection is detected, or the maximum frequency has been reached. If an inflection is detected, a short delay timer is set. At the end of the delay, the frequency generator is set to ramp down and the search for an NMR trigger begins again. Each time an inflection is encountered, the delay timer is run and the frequency modulation changes direction. In this manner, assuming a sweep through the resonant frequency causes a trigger on every pass, the frequency modulation automatically tracks the resonance, even if it slowly moves. Also, assuming the frequency modulation rate is the same going up as it is going down, The NMR signal will always be at the center of the modulation range. The NMR frequency, therefore, is equal to the average frequency over an integral number of FM cycles. Sweeping through the resonant frequency does not always cause a trigger. For this reason, it is necessary to set an upper and lower limit on the frequency output. If the upper or lower limits are reached, the frequency generator is simply switched to ramp in the opposite direction. If a lock on the NMR signal is lost because of a single missed trigger, it will be restored once the frequency has swept to the limit and back again.

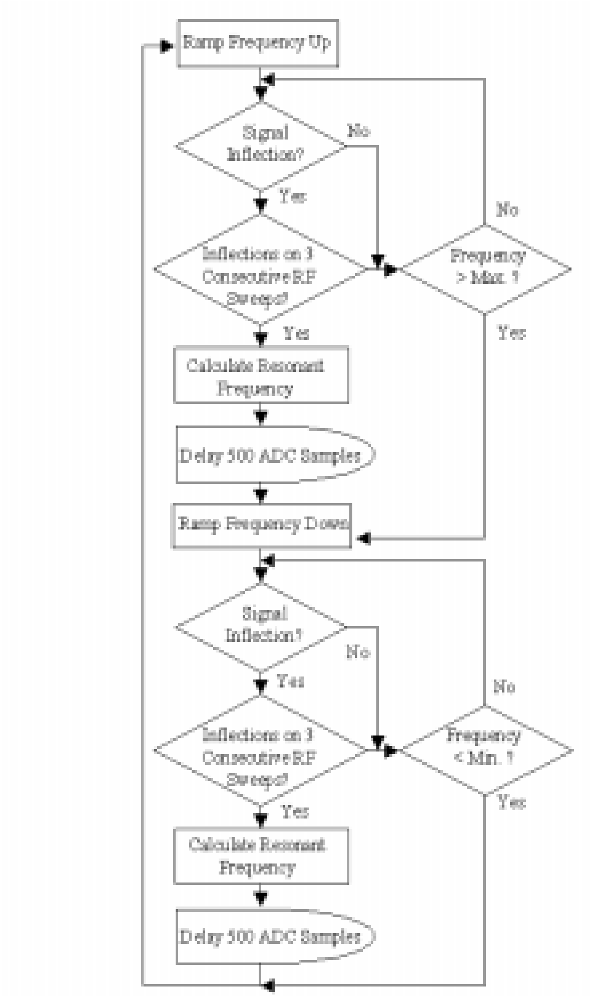

Fig. 6. Flow chart of frequency scanning algorithm

Fig. 6 shows a flowchart of the resonance tracking algorithm. For every new ADC sample, the processor checks if an NMR trigger has occurred. If not, the processor checks the frequency bounds. If a trigger has occurred, the processor checks if it is the 3rd consecutive trigger before a timeout occurred. (A full FM cycle occurs in the time between the first trigger and the third trigger.) A timeout occurs when there is too much time between the first and third trigger. If this timeout occurs, it is likely that the system has lost lock on the NMR signal. A timeout also occurs when the frequency has reached either the upper or lower bounds. This prevents any false triggers near the frequency bounds from causing the processor to calculate false resonant frequency values. When three NMR triggers have occurred within the time limit, the resonant frequency can then be obtained by simply finding the average frequency over the full FM cycle from the first trigger to the third. The frequency is determined using the following relation:

$$F = \frac{Count_{acc} \bullet Rate_{samp}}{N_{samples}} \quad (1)$$

Where: F is the calculated frequency, $Count_{acc}$ is the accumulated frequency counter value for three RF sweeps, and $Rate_{samp}$ is the sampling rate. For example, if the frequency counter has a total count of 3,485,000 from the first trigger to the third trigger, the ADC sampling rate is 50 kHz, and there have been 2050 samples between the first and third triggers, then the NMR frequency is 85 MHz.

D. Key Parameters

1) Nominal Frequency Generator DAC Setting

The accuracy of the resonant frequency calculation is dependent on the frequency modulation sweeping up and down symmetrically. For this reason, it is very important to establish an accurate zero DAC setting for the frequency generator. Because an integrator is used in the frequency generator between the DAC and the VCO, modulating up and down in a triangle wave becomes a simple matter of adding a constant value to or subtracting the same value from the DAC zero setting. Determining this setting is done by setting up a frequency servo loop in which a certain goal frequency is set somewhere near the expected NMR resonance, and incremental changes are made in the DAC setting to set and hold the frequency at that frequency. While the frequency is held at the specified value, the DAC settings are buffered for some period of time and averaged to establish an accurate zero setting.

The software for the frequency servo is a standard proportional, integral and differential (PID) loop. Since there is an integrator between the DAC and the VCO, the feedback signal is the I term. Differentiating once provides the P term and a second differentiation provides the D term.

2) Triggering Threshold Levels

The choice of the positive and negative excursion threshold levels for the NMR signal triggering is perhaps the most important parameter in making the entire resonance tracking algorithm work well. If the thresholds are set too high, the algorithm will either never find the resonance or find it now and then, but not be able to keep lock. If the thresholds are set too low, the algorithm will trigger on noise and won’t be able to obtain or keep lock on the real resonance. If set low enough the algorithm will trigger enough times consecutively on noise to think it has locked on resonance and it will report a false resonant frequency.

IV. Preinstallation Test and Calibration

We tested all four probes inside the D0 solenoid at 2 Tesla, using a commercial RF signal generator and an oscilloscope. The voltage signals from the probes were ± 200mV peak to peak, at the frequency of 85.565 MHz, corresponding to 2.00964 Tesla, while a Metrolab model PT2025 teslameter with its probe placed nearby gave a reading of 2.00907 Tesla. Their difference is 2.8 x 10-4 or 22kHz (very near to the modulation width) which we suspect is due to a DC offset in the signal we use to drive the modulation input of the signal generator. We will verify this with a spectrum analyzer. The noise level was less than 10% of the signal value, averaging over four frequency modulation cycles.

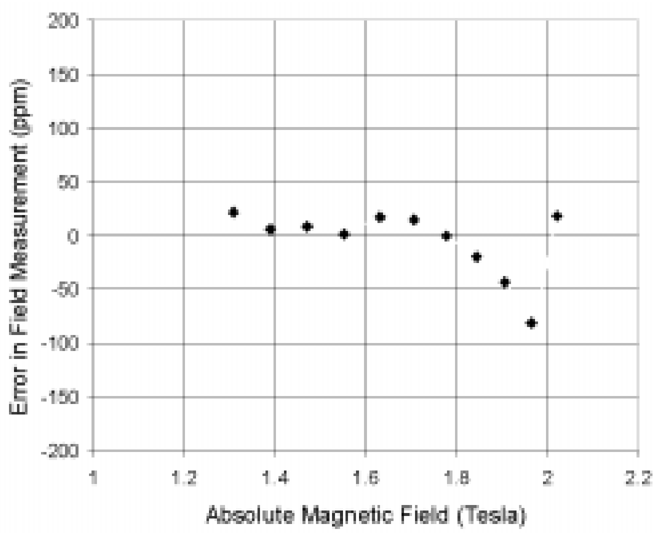

Fig. 7. Comparison between D0 and commercial magnetometer

The calibration was done using a conventional magent and the same Metrolab teslameter. The result of the comparison between the Metrolab and D0 magnetometers is shown in Fig.7. We believe the major portion of the difference is due to the different positions of these probes with respect to the pole faces of the test magnet. We consider the overall operation range of the D0 NMR is about from 1.1 to 2.1 Tesla with an accuracy of better than 10^{-4}. This accuracy is sufficient for our application so there is little incentive to pursue this issue further. We were able to get data points in a magnetic field range from1 0.3 to 2.02 Tesla. We believe the upper limit is set by field non-uniformities of the test magnet that start to appear when the pole tips are driven hard into saturation. The lower limit is set by the size of the NMR signal which diminishes at frequencies below 60MHz.

V. Installation and Signal Observation

A cable run of 70 meters runs between the VME crates containing the readout boards and the probes in the solenoid. Each module has a coaxial RF cable and a multi-conductor cable containing the power, tuning voltage and probe signal return linking it to its probe.

We were able to find NMR resonance in the D0 installation, but encountered some difficulties. One unexpected result is RF crosstalk between measurement channels, which forces us to energize one channel at a time. We also see other types interference on our signal cable, most likely due to the very long run. This interference level varies over time depending on the operating mode of the D0 detector. The difference in the field measurements between channels is of order 5 to 10 gauss. Again we do not know whether this is due to actual field differences around the periphery of the solenoid or to measurement error of the magnetometers. However, based on our comparisons to the Metrolab magnetometer we are reasonably confident that the solenoid field varies by that amount around its perimeter.

VI. Ideas for Improving the Design

An open loop frequency synthesizer would simplify the DSP code significantly. With advances in direct digital synthesizer chips, it would be interesting to try to use one of these devices. We did some tests using the Analog Devices AD9851, but found a very high spurious spectral content near (within 100Hz) the center frequency in the configuration that we tried. It would be of interest to pursue these studies further. Our major difficulties arise from the need to use long cables. The idea of mounting the frequency source, tuning DAC and detector ADC directly on the probe with an entirely digital link would be attractive.

VII. Present Status

Four channels of NMR magnetometer have been installed and demonstrated to work with the D0 superconducting solenoid. Further work needs to be done in order to fully integrate the data from the system into the overall D0 control and monitoring system.

VIII. References

[1] S. Abachi, et al, "The D0 upgrade", Nucl. Inst. & Meth. A408, 103-109, 1998.

[2] R.P.Smith, et al., "Preliminary testing of the D0 superconducting Solenoid", Advances in Cryogenic Engineering Vol 43a, 229 (1998).

[3] G. Boero, C. de Raad Iseli, P.A. Bess, R.S. Popovic, "An NMR magnetometer with planar microcoils and integrated electronics for signal detection and amplification." Institute of Microsystems, EPFL-Swiss Federal Institute of Technology.

[4] K. Borer, G. Fremont, "The nuclear magnetic resonance magnetometer type 9298", Nucl. Instrum. Methods 154, 61-82 (1978)