A Low-Power FPGA-Based Time-Domain NMR Relaxometry System for Field Deployment

David P. Wamai[a,b], Gabriel Howe[a,b], Josh McGuire[a,b], M. Usman Khan[a], Suyash Vardhan Singh[b], Winford Janvrin[a], Md Asifuzzaman Khan[a], Austin R.J. Downey[a,c], and Jason D. Bakos[b]

[a]Department of Mechanical Engineering, University of South Carolina, Columbia, South Carolina 29208, USA
[b]Department of Computer Science and Engineering, University of South Carolina, Columbia, South Carolina 29208, USA
[c]Department of Civil and Environmental Engineering, University of South Carolina, Columbia, South Carolina 29208, USA

ABSTRACT

Nuclear magnetic resonance (NMR) systems have long been appreciated for their ability to rapidly and non-destructively characterize materials with minimal need for extensive sample preparation. Standard NMR techniques, like NMR spectroscopy, have been shown to be remarkably effective in providing detailed insights into the molecular structures of various materials. However, the inherent requirements for high magnetic field strength and homogeneity, coupled with the reliance on high-power, high-speed signal processing platforms, render these systems both expensive and challenging to miniaturize. Time-domain systems, instead of leveraging NMR relaxometry, have been proposed for the in situ real-time monitoring of environmental conditions such as water quality, soil composition, anomalous substances, or other environmental factors. Efforts toward optimizing the power demands of these systems, as well as making them economically viable, provide a foundation for the possibility of conducting NMR experiments outside the laboratory. This work reports on the development of a custom low-power and low-cost digital signal processing platform for time-domain NMR based on field-programmable gate arrays (FPGAs). FPGA development boards demonstrate a capability to match, and in some cases, surpass the performance of traditional digital signal processing platforms typically used in proprietary systems at a fraction of the cost while consuming significantly less power. The system focuses primarily on the generation of high-fidelity Larmor sinusoids and Carr-Purcell-Meiboom-Gill pulse trains, ensuring the implementation of an effective NMR system and maximizing the yield of impactful data. A profiling of the FPGA system reveals exceptional levels of performance, a 95.6% reduction in power consumption, and a 97.3% reduction in cost when compared to a compact NMR system developed using off-the-shelf control systems and laboratory-standard data acquisition modules.

Keywords: Time-domain NMR, FPGA, Relaxometry, Environmental, Field-deployable, CPMG, Low-power

1. INTRODUCTION

Nuclear Magnetic Resonance (NMR) devices employing NMR spectroscopy allow for the characterization and analysis of systems in a time-efficient and non-destructive fashion, often requiring little to no sample preparation. Despite its performance and reliability, the high cost and complexity of NMR spectroscopy, as well as its demand for high magnetic field strength and homogeneity, pose a significant challenge for discretization in a practical manner. NMR relaxometry, however, provides a functionally similar analog to Fourier-based NMR spectroscopy,[1] as it primarily focuses on the relaxation process in a sample from the time domain, and requires comparatively

Further author information:
Austin R.J. Downey.: E-mail: austindowney@sc.edu

lower system sensitivity in order to acquire useful results. Such techniques, combined with the precision of digital signal processing systems, have allowed for the ability to improve upon simple, lab-built time-domain NMR (TD-NMR) systems that can be reconfigured for any range of applications.[2]

In situ water quality monitoring represents one such application for which TD-NMR relaxometry has shown considerable promise. Traditional approaches to water quality assessment rely on laboratory-based chemical analysis, which introduces meaningful delays between field sample collection and the availability of results.[3] Portable NMR instruments, in particular, have demonstrated the ability to operate outside the laboratory and perform measurements on arbitrary-sized samples without the need for reagents or extensive sample preparation.[4] Low-field NMR spectroscopy has further been shown to be applicable in industrial and process monitoring environments, supporting its use as a practical field-deployable sensing modality.[5] Prior work has demonstrated that the T_2 relaxation time of water changes measurably in the presence of organic contaminants such as hydrocarbons, enabling quantification of contaminant concentration directly from the NMR signal.[6] This relationship between T_2 and contaminant content has since been extended to a broader range of environmental matrices, suggesting that relaxometry-based platforms can serve as viable sensing modalities for detecting dissolved or suspended contaminants in surface water and groundwater.[7] The barrier to field adoption of such systems, however, remains the cost and power consumption of the digital signal processing hardware on which they depend, which motivates the development of the low-power, FPGA-based alternative described in this work.

While TD-NMR systems leveraging relaxometry have been shown to provide an effective means of identifying molecular structures attached to hydrogen atoms, they are often quite large in regard to both the magnet and digital signal processing (DSP) instruments involved. Laboratory-standard measurement platforms offered by companies like National Instruments and Keysight Technologies are often used in the development of these systems, but they are typically expensive and require a significant amount of power to operate. Their benefits are attributed to their resilience and reliability; however, the cost of these systems can be prohibitive for many applications, and the power consumption can be a limiting factor for field deployment or other applications where power is a concern. Due to their size, pricing, power efficiency, and level of accessibility, these systems leave the user with a limited capacity to perform in-situ measurements. This work considers an alternative to the components of such systems involved with DSP, and describes the design and analysis of a portable, modular FPGA platform capable of driving the operation of a compact TD-NMR system for field-ready molecular characterization. The code and design artifacts for the FPGA-based controller, referred to as SPIN (Signal Processing Instrument for NMR), are available in a public repository,[8] and the associated NMR hardware design is available in[9] as described in detail by Janvrin et al.[1]

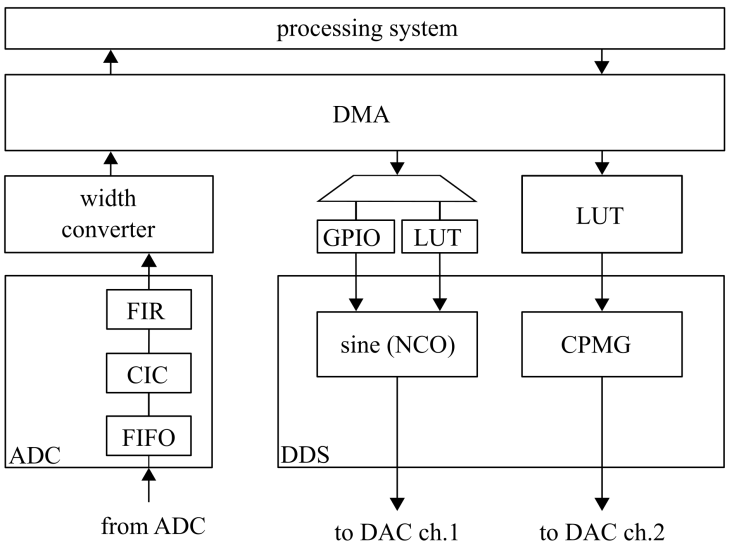

Figure 1. A block diagram of the deployed hardware description.

Figure 1 illustrates the overall architecture of the FPGA-based NMR system, including the signal generation, acquisition, and data transfer components. The FPGA platform consists of a custom hardware description deployed on an FPGA evaluation board. The evaluation board highlighted in this work, the Eclypse-Z7 from Digilent,[10] was chosen for its modest cost and contextually appropriate selection of peripheral modules, which include a digital-to-analog converter (DAC) and an analog-to-digital converter (ADC). These modules, particularly the ADC, offer maximum sampling rates of 125 MS/s, which are more than adequate for the TD-NMR system. The FPGA's hardware description is designed to generate variable RF signals and pulse-train sequences for the TD-NMR system, as well as acquire and transfer the collected data to a host PC or (signal board computer) for post-processing. While it is often in one's best interest to consider the potential for high power consumption due to a dense or otherwise irresponsibly designed hardware description, the Eclypse-Z7 evaluation board will typically default to a dramatically lower rate of power consumption as compared to other systems, making it an ideal candidate for field deployment or generally low-power cyberphysical system designs.

We proceed first with an overview of the FPGA platform's design, followed by a detailed discussion of the hardware description and the implementation of the system's components, beginning with digital-to-analog conversions, then analog-to-digital conversions, and finally direct memory access. We then discuss the post-processing of the acquired data, and conclude with a discussion of the results and a comparison of the FPGA platform to a commercially available DSP system.

2. SYSTEM DESIGN

The FPGA platform can be generalized to a signal generation and acquisition system comprised of a hardware description implementing controllers for both functionalities. The board on which the hardware description is deployed is equipped with a Zynq-7020 SoC, which integrates a dual-core ARM Cortex-A9 processor and an Xilinx 7-series FPGA. The programmable logic (PL) inside the FPGA is responsible for hosting the hardware description, hence providing logic to drive the ZMOD AWG-1411 and ZMOD Scope-1410 daughter cards and providing the board with interfaceable DAC and ADC modules supporting SMA or similar connections. The ARM processor is responsible for the configuration and control of the FPGA, as well as any in-hardware post-processing of the acquired signal. The system is modular, with the PL and ARM processors being able to operate independently of each other or in conjunction, allowing for any post-processing operations the ARM processor otherwise would have performed to be accelerated by the PL hardware. The system is portable; due to its modest power requirements and small form factor, one can exploit its potential to be powered by a battery pack or small solar cell setup. The system is low-cost, with the Eclypse Z7 and its complementary daughter cards (the ZMOD AWG-1411 and ZMOD Scope-1410) totaling approximately $900 at the time of writing. The system is also reconfigurable and scalable, allowing for the replacement of either daughter card with modules compatible with the SYZYGY connector standard, or the integration of additional sensor modules interfaceable through general-purpose I/O pins on the board.

The platform was designed with the intention of utilizing the PL's potential for parallelization and accelerating computationally intensive tasks like digital filtering and data streaming, so any further discussion of the system's design will primarily focus on the configuration of the PL. The PL contains the three principal components of the system that allow for the generation and acquisition of signals; digital to analog conversion, analog to digital conversion, and direct memory access (DMA) are driven exclusively by the PL, with any synchronization or choreographic operations left to be handled by the processing system (PS) or ARM processor. We continue, in order, with the discussion of the implementation of the DAC, ADC, and finally the DMA.

3. DIGITAL TO ANALOG CONVERSION

Figure 2 shows the physical implementation of the compact NMR system, including the FPGA platform, control electronics, and experimental setup. The ZMOD AWG-1411 is a 14-bit, 125 MS/s DAC module that provides the Eclypse Z7 with two channels of high-speed analog output. Since this platform was made to complement an already existing TD-NMR PCB, the DAC was configured to output both a radio frequency (RF) sinusoidal signal alongside a gating square wave sequence from channels one and two, respectively. The referred PCB includes a

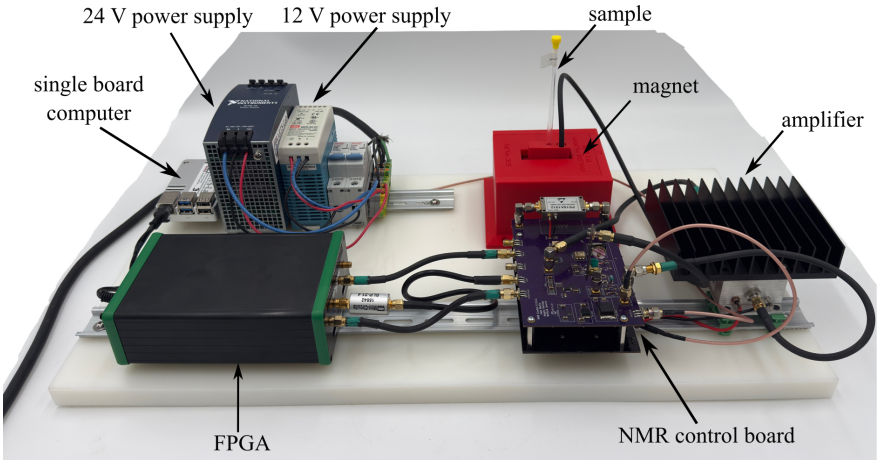

Figure 2. Compact NMR configuration with Eclypse-Z7 FPGA.

switch that permits the sinusoidal pass through to the sample governed by the gating sequence. One can think of channel one of the DAC as a direct replacement of the waveform generator denoted in,[?] and channel two as the replacement of the pulse generator.

3.1 Direct Digital Synthesis

There exists a variety of methods to generate an analog waveform digitally, the most common of which is via a numerically controlled oscillator (NCO). The NCO generates a waveform by updating the phase θ of the signal at each clock cycle according to the equation

$$\theta_{n+1} = \theta_n + \Delta\theta, \tag{1}$$

where θ_{n+1} and θ_n are the phase values at the current and next clock cycles, respectively, and $\Delta\theta$ is the phase increment, which is proportional to the desired output frequency. The phase increment or frequency tuning word (FTW) is calculated as

$$\Delta\theta = \frac{f_{\text{out}} \cdot 2^N}{f_{\text{clk}}}, \tag{2}$$

where N is the bit width of the phase accumulator, and f_{out} is the desired output frequency.

The output waveform, typically a sine wave, is obtained either by directly providing the phase increment value $\Delta\theta$ to the NCO's phase increment input or mapping the phase θ to a corresponding amplitude value using a waveform lookup table (LUT). By adjusting the FTW, the output of the NCO can be configured dynamically if interfaced with a control register or general purpose I/O. For the sake of simplicity and replicability, the NCO is implemented in the form of the Direct Digital Synthesis (DDS) IP core provided by Xilinx. The DDS core is a feature-rich NCO that supports waveform generation dictated either by a data stream or an FTW. Both data streaming and FTW methods are implemented here for the sake of demonstration, but the data streaming method is preferred specifically for the ease of generating the gating square wave sequence in exchange for a slight loss in timing precision. Figure 3 provides a conceptual representation of a numerically controlled oscillator used for digital waveform generation.

3.1.1 Sinusoidal DDS

The compact-NMR system is designed to operate at a given input frequency determined by the sample being measured and the field strength of the magnet. This frequency, known as the Larmor frequency, typically falls between 20–30 MHz for our particular configuration.[1]

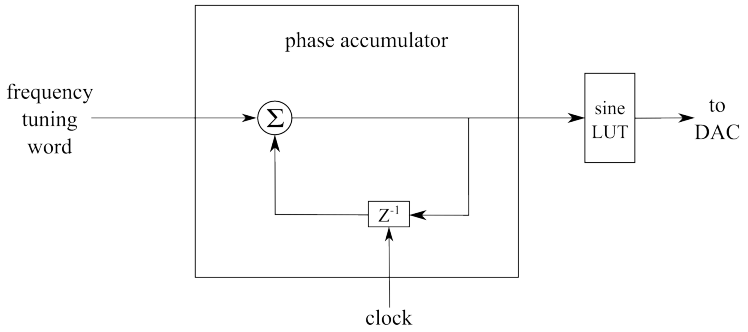

Figure 3. Block diagram of a generic numerically controlled oscillator.

The Larmor frequency (ω_L) can be approximated as

$$\omega_L = \gamma B_0, \tag{3}$$

where ω_L is the Larmor frequency in radians per second (rad/s), γ is the gyromagnetic ratio of the nucleus, which is known for hydrogen (^{1}H), $\gamma \approx 42.57748\,\text{MHz/T}$ or $267.5221 \times 10^6\,\text{rad/s/T}$, and B_0 is the magnetic field strength in Teslas (T).

For example, given the modest magnetic field strength 0.5 T, the Larmor frequency for Hydrogen can be defined as:

$$\omega_L = 42.577\,\text{MHz/T} \times 0.5\,\text{T} = 21.288\,\text{MHz}. \tag{4}$$

This yields

$$\omega_L = 21.288 \times 10^6\,\text{MHz}. \tag{5}$$

Thus, for 0.5 T, the Larmor frequency or required RF pulses for the hydrogen nucleus excitation in the compact-NMR system would be within $\pm 10\,\text{kHz}$ of 21.3 MHz.

To determine the $\Delta\theta$ for the DDS compiler given our found Larmor frequency, we can observe that the desired output frequency (f_{out}) is ω_L and consider again the equation for $\Delta\theta$: Assuming a DDS system clock frequency (f_{clk}) of 100 MHz and a phase accumulator width of 32 bits ($N = 32$), the FTW to generate a 21.3 MHz signal is

$$\Delta\theta = \frac{21.288 \times 10^6 \cdot 2^{32}}{100 \times 10^6}. \tag{6}$$

The calculated $\Delta\theta$ value, $914{,}342{,}702_{10}$, can then be provided to the DDS compiler in unsigned or two's complement binary form to configure the NCO for generating an RF signal with a frequency within $\pm 10\,\text{kHz}$ of the determined Larmor frequency for the compact-NMR system.

3.1.2 Gating DDS

The Carr-Purcell-Meiboom-Gill (CPMG) sequence is widely used in Nuclear Magnetic Resonance (NMR) spectroscopy to measure the transverse relaxation time T_2 of spins in a sample. This sequence is particularly effective in reducing the effects of inhomogeneous broadening on decay signals. The CPMG sequence consists of a 90° pulse to flip the magnetization into the transverse plane, followed by a series of 180° refocusing pulses.[11] These pulses refocus the dephasing spins caused by inhomogeneities in the magnetic field, leading to the formation of spin echoes. The timing of pulses in a CPMG sequence is critical for accurately measuring T_2. The sequence begins with a 90° pulse, followed by the first 180° pulse at a time τ after the initial pulse. Subsequent 180° pulses are applied at intervals of 2τ, and the echo signals are observed at these intervals. The structure and timing of this pulse sequence are illustrated in Figure 4.

Algorithm 1 Carr-Purcell-Meiboom-Gill pulse-train

```
pulse_counter ← 0
period_counter ← 0
pulse_state ← 0
tau_done ← 0
TAU, TWO_TAU ← 0xDEAD                                   ▷ Set to desired τ and 2τ
TAU_LOW, TWO_TAU_LOW ← 0xBEEF
while True do
    if pulse_state = HIGH then
        if (NOT tau_done AND pulse_counter < TAU) OR (tau_done AND pulse_counter < TWO_TAU) then
            pulse_counter ← pulse_counter + 1
            data ← HIGH_VALUE
        else
            pulse_state ← LOW
            pulse_counter ← 0
            period_counter ← 1
            data ← LOW_VALUE
        end if
    else                                                ▷ pulse_state is LOW
        if (NOT tau_done AND period_counter < TAU_LOW) OR (tau_done AND period_counter <
TWO_TAU_LOW) then
            period_counter ← period_counter + 1
            data ← LOW_VALUE
        else
            pulse_state ← HIGH
            period_counter ← 0
            pulse_counter ← 1
            data ← HIGH_VALUE
            if NOT tau_done then
                tau_done ← TRUE
            end if
        end if
    end if
end while
```

The echo signals in a CPMG sequence are observed at times given by

$$t = n(2\tau), \quad n = 1, 2, 3, \ldots, N, \tag{7}$$

where n is the echo number and N is the total number of echoes observed.

The pulse timing for the CPMG sequence was often, for this work in particular, determined experimentally. The reference pulse interval used by the PXI-based compact-NMR system was $2\tau = 14\,\mu$s (i.e., $\tau = 7\,\mu$s), so this value was also used in the generation of the pulse-train.

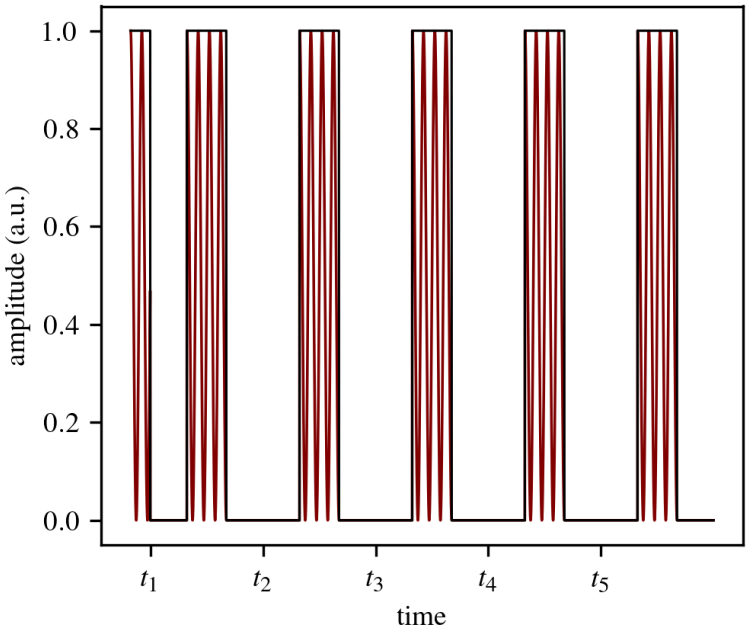

Figure 4. Visualization of the Carr-Purcell-Meiboom-Gill pulse-train.

One method used to generate the pulse train was the utilization of a LUT to stream DAC codes to the DDS compiler in the PL from the PS using the AXI Direct Memory Access (DMA) IP core. The LUT was initialized in a software application running within the PS written in C, then populated with the exact pulse train pattern as a function of the desired τ value.

One can also find more consistency in the CPMG signal by generating it logically or in combination in hardware. Rather than using the DDS compiler or similar NCO IP cores, the gating square wave can be generated with a finite state machine (FSM) in the PL. Here, we develop the FSM logically, shown in Algorithm 1. The use of an FSM instantiated in hardware allows for greatly enhanced flexibility as compared to the LUT method. Since the FSM resides on the register transfer level (RTL), and access to related hardware signals or RTL cores, which are otherwise unreachable from the PS,[1] is granted. Such examples of enhanced flexibility include, but are not limited to, the ability to count clock cycles to generate the pulse-train at the desired interval, and resetting after the desired number of pulses have been generated, or to continue generating pulses indefinitely.

4. ANALOG TO DIGITAL CONVERSION

The ZMOD Scope-1410 is a 14-bit, 100 MS/s ADC module that provides the Eclypse Z7 with two channels of high-speed analog input. The ADC was configured to sample the signal from the TD-NMR PCB at 1 MS/s, and stream the acquired data to DDR memory using the AXI DMA IP core. The native sampling rate of 100 MS/s was not used in this configuration, as the implementation of a platform-to-host data transfer method suitable for a rate as high as 100 MS/s without driving the entire platform from a Linux application had not yet been determined. The data transfer method used here is simply a serial bus between the FPGA board and a host PC, initiated by the user in the software application running in the PS, where access to familiar GPIO interfaces like `printf` is available. The data transfer method is not ideal, but it does serve as a reliable method of transferring data to a host PC for analysis without implementing software-coherent cache maintenance.

[1] Addressable control registers can be used to interface with the PL from the PS; however, doing so in excess can lead to a loss of performance and increased power consumption.

4.1 Downsampling

The conversion from the native 100 MS/s sampling rate to 1 MS/s was performed in the PL using both the Xilinx Cascaded Integrator-Comb (CIC) Compiler and the Xilinx Finite Impulse Response (FIR) Compiler IP cores. The data stream is initiated with a hardware trigger monitoring the values of the samples entering the ADC. The trigger is set to initiate the data stream when any given sample in the stream exceeds a certain value, and the stream is terminated when the number of samples determined by the user has been acquired. The stream of data is then downsampled using the CIC filter, followed by the FIR filter to correct for any attenuation in the pass band of the CIC filter. The CIC filter is configured to decimate samples in the data stream by a factor of 100, resulting in a 1 MS/s data stream after being passed through the single-rate FIR filter.

CIC filters are commonly used in digital signal processing for decimation and interpolation. A CIC filter consists of a series of integrators followed by a series of combs, or difference operators. This filter is particularly effective due to its simple structure, which only requires addition and subtraction operations. Decimation is the process of reducing the sampling rate of a signal by an integer factor R. In the context of CIC filters, decimation involves the reduction of the sampling rate from an input frequency f_{in} to an output frequency f_{out} where

$$f_{\text{out}} = \frac{f_{\text{s,in}}}{R}. \tag{8}$$

Here, $f_{\text{s,in}}$ is the input sampling frequency.

A basic CIC filter structure includes N integrators and N combs. The transfer function of a single integrator stage in the z-domain is given by

$$H_{\text{int}}(z) = \frac{1}{1 - z^{-1}}. \tag{9}$$

The transfer function of a single comb stage is

$$H_{\text{comb}}(z) = 1 - z^{-D}, \tag{10}$$

where D is the delay element in the comb stage, and z^{-1} represents a delay of one sample.

When cascading N integrators and N combs, the overall transfer function $H(z)$ of the CIC filter becomes the product of all individual stages:

$$H(z) = \left(\frac{1}{1 - z^{-1}}\right)^N \left(1 - z^{-D}\right)^N. \tag{11}$$

The CIC filter architecture allows for effective decimation by accumulating multiple samples (integrators) over several clock cycles, and differencing over a delay of D samples, effectively combing the signal as well as removing high-frequency components. The decimation factor R typically equals the delay D, and the rate at which the signal is decimated impacts the frequency components that are preserved or attenuated. The output signal $y(n)$ is given by

$$y(n) = x(nR), \tag{12}$$

where $x(n)$ is the input signal. This expression shows how every R^{th} sample is taken from the input, effectively reducing the sampling rate by the factor R. The behavior and frequency response of the CIC filter used in this system are illustrated in Figure 5.

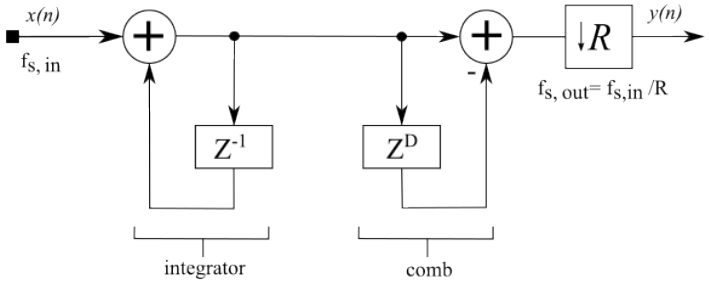

Figure 5. CIC filter transfer function diagram demonstrating decimation.

5. DIRECT MEMORY ACCESS

Direct Memory Access (DMA) is a feature of many generic computer systems that allows certain hardware subsystems to access the system's main memory independently of the PS. The compact-NMR platform utilizes scatter-gather DMA to write sample data into system memory. Scatter-gather DMA allows for the collection of data into non-contiguous memory regions within a single DMA transaction, eliminating the need for performing multiple DMA transactions in sequence.[12]

The DMA transaction is initiated once the sample data buffer defined and allocated by the software application running on the PS is filled with the desired number of samples. The scatter-gather DMA engine in the PS first defines a descriptor table, which contains the addresses of the non-contiguous memory blocks where the data will be written. Each entry specifies a memory address and the length of the data to be written to that address. The DMA controller in the PL is then provided with the descriptor table's base address, enabling it to index through the table entries correctly. Once a stream of data becomes available, the DMA controller begins the process of writing the data to memory by first referencing the base address of the descriptor table, using the data width to determine the offset that it will use to populate that particular table entry. The DMA controller then writes the data to the memory address specified in the table entry, and continues to the next index in the descriptor table.

6. POST PROCESSING

To generate the free induction decay (FID) plot, a post-processing algorithm is implemented on a host PC that collected the data from the FPGA over a serial connection. The overall experiment is configured such that a 4-second scan is performed. At a 1 MS/s sampling rate, this results in 4,000,000 samples. The FID plot is characterized by transverse relaxation from spins going out of phase. Repetition of the 180° pulse in the CPMG sequence causes the spins to decrease in magnitude when they refocus, and the FID plot is a representation of this decrease in magnitude over time. The FID plot is generated by performing a peak-finding algorithm on the time-domain signal acquired by the FPGA.[13] This algorithm first eliminates the CPMG pulses from the acquired signal, then finds the local maximum of a window in the data defined by Eq. (7). The local maxima are then plotted in sequence to generate the FID plot shown in Fig. 7.

Further characterization of the properties of the FID signal involves determining the transverse (spin-spin) relaxation time T_2 of the sample. The transverse relaxation time is the decay constant of the component of the magnetization vector perpendicular to the static magnetic field. The transverse relaxation time is determined by fitting the FID plot to the exponential decay function

$$M_{xy}(t) = M_0 \exp\left(\frac{-t}{T_2}\right), \tag{13}$$

where $M(t)$ is the magnetization at time t, M_0 is the initial magnetization, and T_2 is the transverse relaxation time. The presence of the exponential component in the expression suggests that the transverse magnetization vector decays to 37% of its initial value after a time T_2. The measured transverse relaxation time is calculated by

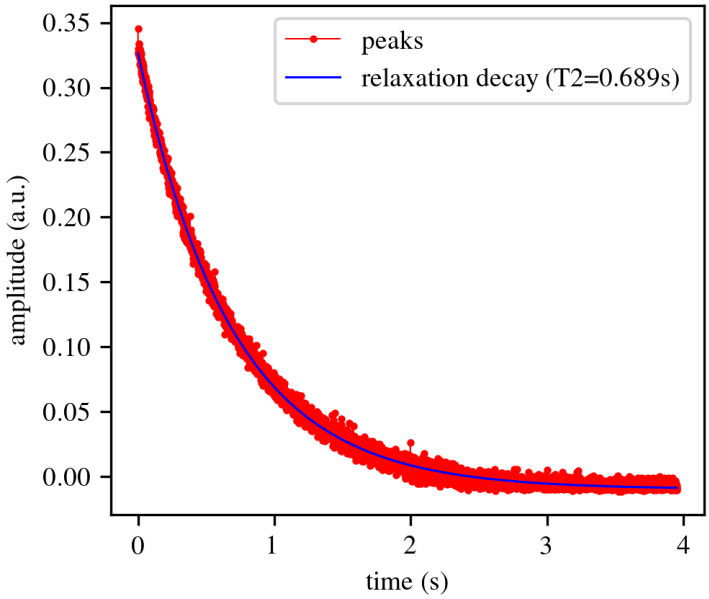

Figure 6. n-decane FID peaks and relaxation decay time constant calculated from the PXI-1062Q platform's acquisition.

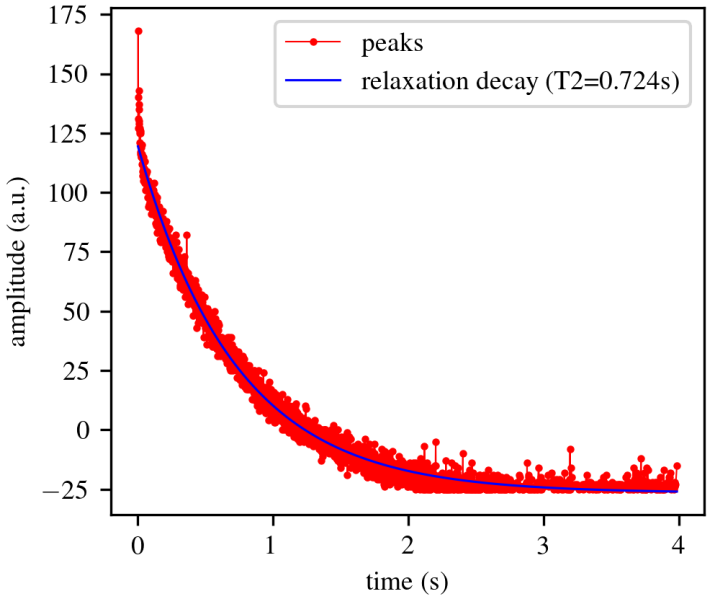

Figure 7. n-decane FID peaks and relaxation decay time constant calculated from the FPGA platform's acquisition.

fitting the peaks of the FID plot to an exponential decay function using a Python script, providing a quantitative basis for comparison between the PXI-based and FPGA-based systems. The resulting fits, shown in Fig. 6 and Fig. 7, demonstrate consistent decay behavior across both platforms.

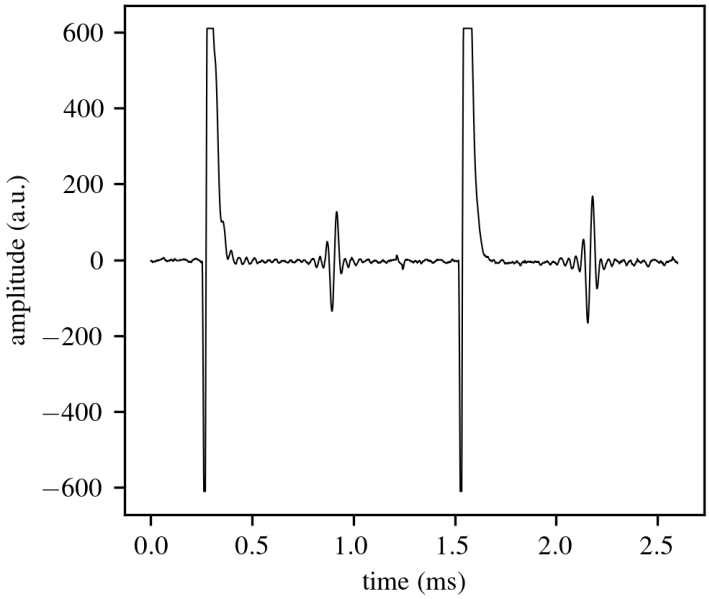

Figure 8. The first few spin echo response signals captured by the FPGA platform.

7. RESULTS AND DISCUSSION

Fig. 8 reports the first two 180˚ spin echo response signals of n-decane, a hydrocarbon compound, acquired by the FPGA platform. The experiment was performed using the same control-circuit configuration as the PXI-based compact-NMR system. The FPGA configuration of the experiment yielding the result in Fig. 7 involved a 4-second scan with a Larmor frequency of 27.608 MHz and a logically generated CPMG sequence with a pulse interval τ of 7 μs. The FID plot was generated by the post-processing routine described in the previous section, detailing a representation of the decrease in magnitude of the spin echo response signals over time, which can be used to estimate the hydrogen composition in the sample solution.

A qualitative control experiment was performed using the PXIe-1062Q platform to compare the performance of the FPGA platform against it. The PXI chassis was configured almost identically to the FPGA platform, the only exception being the sampling rate of the ADC module, which was fixed well below that of the FPGA at 300 kS/s. The experiment was performed using the same n-decane sample as used in the FPGA experiment. Fig. 6 reports only the T_2 relaxation time of the sample, as the LabVIEW software used to control the PXI platform undergoes a functionally similar but technically different post-processing routine than the FPGA platform.

A power profiling using a generic BL6523GX-based power meter was performed on both the FPGA platform and the NI PXIe-1062Q. Table 1 details the idle and active power characteristics of both platforms. The values in the table reflect the average of each metric taken over 10 total tests for each platform. We describe the idle state for each platform to be powered on, but running no processes related to the NMR experiment. We describe the active state for each platform to be powered on and generate the RF and gating signals for the NMR experiment. Post-processing and the transfer of the data to a host PC are functional aspects not included in the power profiling. A single test run involves the platform being powered on and left idle for one minute, then recording the power value. The platform is then prompted to execute one 4-second scan, and the maximum power value recorded during the scan is taken as the active power value.

Table 1. Power profiling results for the NI PXIe-1062Q and the Digilent Eclypse Z7.

Metric	**NI PXIe-1062Q**	**Eclypse Z7**
Idle Power (W)	109.0	2.9
Active Power (W)	110.6	4.9

Table 2. Comparison of the technical specifications of the NI PXIe-1062Q and the Digilent Eclypse Z7.[10,14] PXIe-1062Q cost includes the chassis and two PXI-5421 AWG modules and one PXI-5122 digitizer. Eclypse Z7 cost includes the board, ZMOD AWG-1411, and ZMOD Scope-1410 modules.

Platform	**NI PXIe-1062Q**	**Digilent Eclypse Z7**
Cost	∼$37,000	∼$900
Input Voltage Range	120 VAC	9–12 VDC
Input Current	8 A	5 A
Width (mm)	271.4	99
Length (mm)	396.5	160
Weight (kg)	8.8	0.612

The reduction in cost of the FPGA platform as compared to the PXIe-1062Q is observably significant, as summarized in Table 2. When configured with the modules required to perform the NMR experiment—two PXI-5421 arbitrary waveform generators and one PXI-5122 digitizer for the PXI system, and the ZMOD AWG-1411 and ZMOD Scope-1410 daughter cards for the FPGA platform—the total system cost is approximately $37,000 and $900, respectively. This represents a reduction of approximately 97.5% in cost, with the FPGA platform costing roughly 2.5% of the equivalent PXI-based configuration.

Limitations of the system involve the rate at which data is transferred between the FPGA platform and a host PC. Transferring data over serial at a rate equal to the sampling rate of the ADC is impossible, since the maximum possible baud rate for serial is 921,600 bits per second, assuming an 8-bit word and no multiplexing between multiple serial channels. The data transfer portion of the experiment is easily the most time-consuming, taking approximately 90 minutes to transfer all 4,000,000 samples to a host PC. Improvements to the data transfer method are being considered, primarily the introduction of an Ethernet connection between the FPGA platform and the host PC.

8. CONCLUSION

This work demonstrates the feasibility of a low-power, FPGA-based platform for time-domain NMR that maintains performance comparable to conventional laboratory systems while significantly improving portability and efficiency. The system successfully generates and acquires NMR signals, producing results consistent with those obtained from established PXI-based instrumentation.

By leveraging FPGA-based digital signal processing, the platform achieves substantial reductions in power consumption and system cost, enabling operation in environments where traditional systems are impractical. These improvements support the broader goal of transitioning NMR technologies from laboratory settings to field-deployable applications, such as environmental monitoring and in situ analysis. While the current implementation is limited by data transfer speeds between the FPGA and host system, the core architecture provides a flexible foundation for future enhancements. Continued development of communication interfaces and on-board processing capabilities will further improve system performance and usability.

ACKNOWLEDGMENTS

This material is based upon work supported by the AQSEL project funded by the South Carolina Quantum Association/State of South Carolina. Additional funding for this work comes from the National Science Foundation, United States, under Grant Nos. CMMI-2152896, CPS-2237696, CISE-2324052, and ITE-2344357. The views and conclusions contained within this document are those of the authors and should not be interpreted as representing the official policies of the NSF or the U.S. Government.

REFERENCES

1. W. Janvrin, J. Martin, D. Hancock, A. Varillas, A. R. Downey, P. J. Pellechia, J. Satme, and S. H. Won, "Open-source compact time-domain hydrogen (1H) NMR system for field deployment," *HardwareX* **22**, p. e00651, 2025.
2. J. S. Martin, A. R. J. Downey, M. Baalousha, and S. H. Won, "Rapid measurement of magnetic particle concentrations in wildland–urban interface fire ashes and runoff using compact NMR," *IEEE Sensors Journal* **24**(6), pp. 7355–7363, 2024.
3. I. Yaroshenko, D. Kirsanov, M. Marjanovic, P. A. Lieberzeit, O. Korostynska, A. Mason, I. Frau, and A. Legin, "Real-time water quality monitoring with chemical sensors," *Sensors* **20**(12), p. 3432, 2020.
4. B. Blümich, J. Perlo, and F. Casanova, "Mobile single-sided NMR," *Progress in Nuclear Magnetic Resonance Spectroscopy* **52**(4), pp. 197–269, 2008.
5. F. Dalitz, M. Cudaj, M. Maiwald, and G. Guthausen, "Process and reaction monitoring by low-field NMR spectroscopy," *Progress in Nuclear Magnetic Resonance Spectroscopy* **60**, pp. 52–70, 2012.
6. T. R. Bryar and R. J. Knight, "NMR relaxation measurements to quantify immiscible organic contaminants in sediments," *Water Resources Research* **44**(4), p. W04421, 2008.
7. E. L. Fay and R. J. Knight, "Detecting and quantifying organic contaminants in sediments with nuclear magnetic resonance," *Geophysics* **81**(6), pp. EN87–EN97, 2016.
8. ARTS-Lab, "Signal processing instrument for NMR (SPIN)." GitHub. https://github.com/ARTS-Laboratory/signal-processing-instrument-for-NMR.
9. ARTS-Lab, "Compact-NMR." GitHub. https://github.com/ARTS-Laboratory/Compact-NMR.
10. A. Brown, "Eclypse Z7," 2024. Accessed March 7, 2024.
11. Michigan State University, "Pulsed Nuclear Magnetic Resonance," 2024. Accessed March 7, 2024.
12. A. Brown, *Eclypse DDR Streaming Project User Manual*. Digilent Reference, 2024. Accessed March 7, 2024.
13. P. Huggins, J. S. Martin, A. R. Downey, and S. H. Won, "Interpretable machine learning for predicting the derived cetane number of jet fuels using compact TD-NMR," *Sensors and Actuators B: Chemical* **426**, p. 137018, 2025.
14. National Instruments, "NI PXIe-1062Q specifications," 2024. Accessed March 7, 2024.