A single-board NMR spectrometer based on a software defined radio architecture

2011 Meas. Sci. Technol. 22 015902

(http://iopscience.iop.org/0957-0233/22/1/015902)

Meas. Sci. Technol. **22** (2011) 015902 (8pp) doi:10.1088/0957-0233/22/1/015902

A single-board NMR spectrometer based on a software defined radio architecture

Weinan Tang and Weimin Wang

Institute of Quantum Electronics, School of Electronic Engineering and Computer Science, Peking University, Beijing, People's Republic of China

E-mail: wmwang@ele.pku.edu.cn

Received 4 September 2010, in final form 29 October 2010
Published 13 December 2010
Online at stacks.iop.org/MST/22/015902

Abstract
A single-board software defined radio (SDR) spectrometer for nuclear magnetic resonance (NMR) is presented. The SDR-based architecture, realized by combining a single field programmable gate array (FPGA) and a digital signal processor (DSP) with peripheral radio frequency (RF) front-end circuits, makes the spectrometer compact and reconfigurable. The DSP, working as a pulse programmer, communicates with a personal computer via a USB interface and controls the FPGA through a parallel port. The FPGA accomplishes digital processing tasks such as a numerically controlled oscillator (NCO), digital down converter (DDC) and gradient waveform generator. The NCO, with agile control of phase, frequency and amplitude, is part of a direct digital synthesizer that is used to generate an RF pulse. The DDC performs quadrature demodulation, multistage low-pass filtering and gain adjustment to produce a bandpass signal (receiver bandwidth from 3.9 kHz to 10 MHz). The gradient waveform generator is capable of outputting shaped gradient pulse waveforms and supports eddy-current compensation. The spectrometer directly acquires an NMR signal up to 30 MHz in the case of baseband sampling and is suitable for low-field (<0.7 T) application. Due to the featured SDR architecture, this prototype has flexible add-on ability and is expected to be suitable for portable NMR systems.

Keywords: nuclear magnetic resonance, spectrometer, software defined radio

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Recently, with the development and application of portable nuclear magnetic resonance (NMR) systems, it was desirable to build a highly integrated and cost-effective NMR spectrometer that satisfies modern pulse NMR measurement requirements. Even though there are several commercial spectrometers available such as the Terranova (Magritek, Inc.), the demand for designing and developing application-specific portable spectrometers has not decreased. This is because researchers often have unique requirements for complexity, reconfigurability and cost. With development in the field of electronics and computer science, a number of groups are working to design home-built NMR spectrometers [1–9]. Direct digital synthesizer (DDS) chips are used to replace highly dedicated radio frequency (RF) sources [10, 11], and digital signal processing chips are used in building the receiver [12].

For constructing a portable NMR spectrometer, some schemes mentioned above seem complex and bulky (many circuit boards or chips are used) [5, 7–9], and hardware is not easy to update [6]. As a result, these designs are difficult to use, and cost is generally high. In this work, we present a complete laptop-sized NMR spectrometer based on a configurable software defined radio (SDR) architecture that mainly combines a single field programmable gate array (FPGA) chip and a digital signal processor (DSP). Compared to the heterodyne architecture that has been applied in NMR spectrometer design for many years, the SDR architecture allows the signal sampling to be placed as close to the receiver coil as possible [14]. Therefore the noise and distortion associated with the analog mixing stage can be reduced. With our spectrometer, the key parameters and

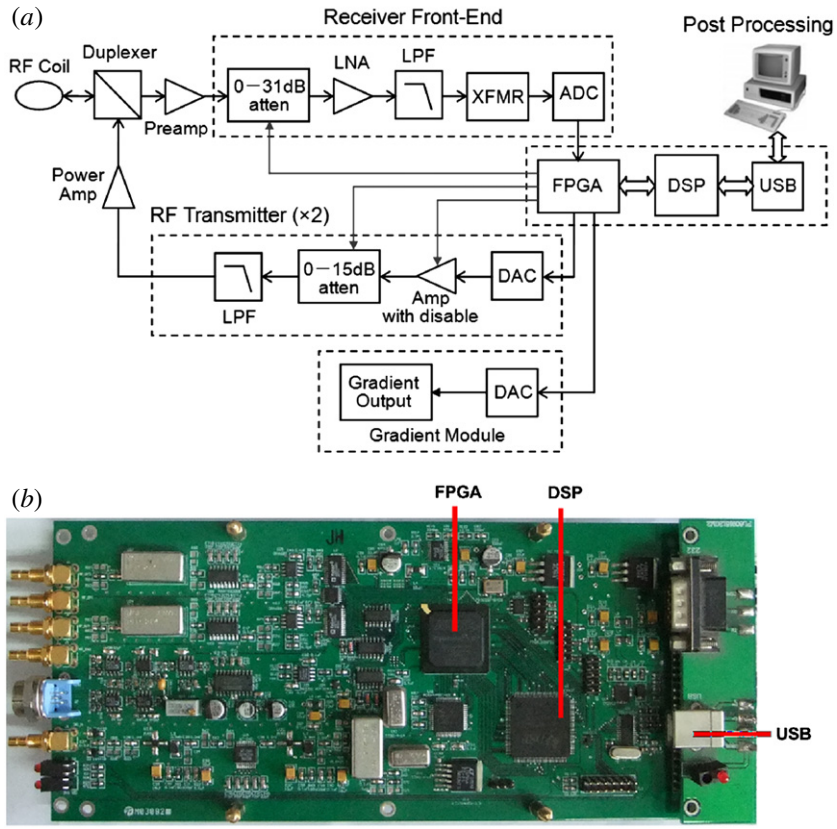

Figure 1. (*a*) The block diagram of the single-board spectrometer. The glossary is defined as follows: Preamp, pre-amplifier; atten, attenuator; LNA, low-noise amplifier; LPF, low-pass filter; XFMR, transformer; ADC, analog-to-digital converter; Amp, amplifier; DAC, digital-to-analog converter. (*b*) Photograph of the spectrometer, which is constructed on a single 100 mm × 220 mm printed circuit board.

operations are allowed to be determined and reconfigured by upgrading software. In SDR application, the FPGA features reconfigurability and parallelism and is hence a highly flexible design platform and signal processing engine [15]. So we employ a single FPGA to perform digital processing tasks such as a numerically controlled oscillator (NCO), digital down converter (DDC) and gradient waveform generator. The design of the FPGA with similar functions has been reported by Takeda [2, 3], but differently, we adopt an NCO-based quadrature demodulation, optimize the digital filters and additionally provide gradient pulse generation. Most significantly, Takeda built a pulse programmer by a highly specialized method, which is complex as the translation from pulse programming language into hexadecimal codes is accomplished manually. Herein we use a low-cost ($8.85), easy-to-use DSP as the pulse programmer in which the sequence interpretation can be simply and automatically completed with the assembly language tool.

The DSP is a good choice as a digital pulser in spectrometer design [6], but the number of control bits is generally constrained as its digital outputs are limited. In our spectrometer, with a parallel interface between the FPGA and DSP, the FPGA receives command words and decodes them into various control lines. As the FPGA has abundant programmable I/Os, the number of external control lines that can be configured are as many as hundreds. This is desirable in highly demanding NMR environments, e.g. the experiment with multiple RF channels and many external devices [3, 13]. The DSP we chose is more a digital signal controller (DSC) than a general purpose DSP. The chip combines the power of the DSP and the function of a microcontroller, and is a suitable device for sequence control.

In addition to the featured flexibility and low cost (the total price of the design is less than $300), all components are constructed on a single printed circuit board as small as 100 mm × 220 mm in dimensions. The described spectrometer has been demonstrated successfully to incorporate in benchtop NMR systems (UNIQ-PMR series, Peking University) in a low magnetic field (<0.7 T) by performing time-domain NMR (TD-NMR) experiments such as solid fat content and oil seed analysis, and a two-dimensional NMR measurement in petrophysics. The details of the spectrometer and results are presented in the following sections.

2. The spectrometer architecture

The block diagram of the spectrometer is shown in figure 1(*a*), and the photograph of the spectrometer board is shown in figure 1(*b*).

The spectrometer contains one RF receiver and an associated front end, two RF transmitters, one gradient waveform generator, a USB interface and a pulse programmer. The DSP, TMS320LF2407A (Texas Instruments, Inc., Dallas, TX), acts as the pulse programmer, controlling the timing of

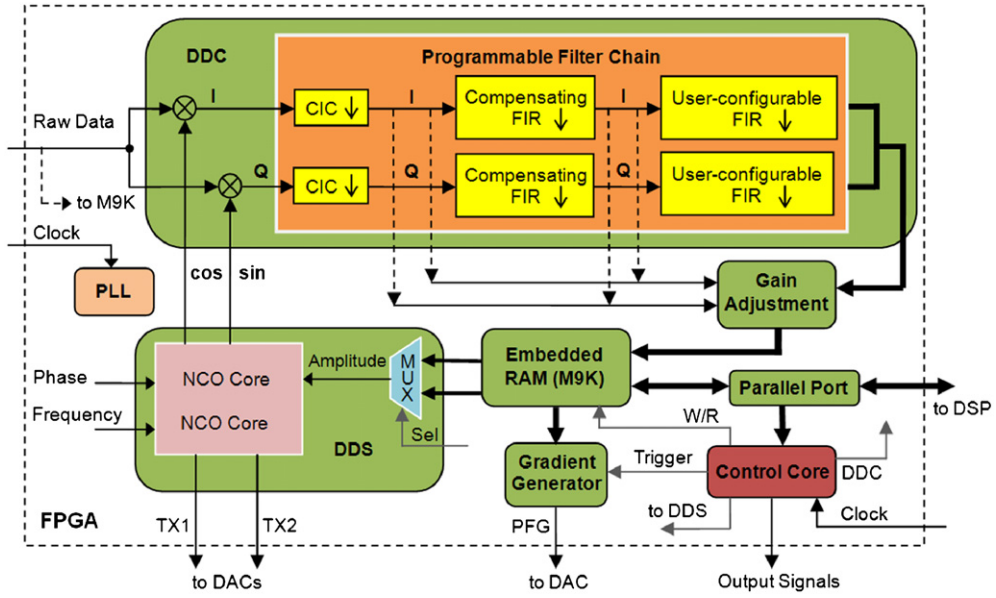

Figure 2. The block diagram of the FPGA. The glossary is defined as follows: DDC, direct digital converter; CIC, cascaded integrator-comb; FIR, finite impulse response; PLL, phase-locked loop; NCO, numerically controlled oscillator; DDS, direct digital synthesizer; MUX, multiplexer; PFG, pulsed field gradient; W/R, write/read; TX, transmit.

pulse sequences and configuring the registers inside the FPGA via a 16 bit parallel port.

The FPGA chip in the present design is EP3C55F484, which belongs to the Cyclone III device family (Altera, Inc.). It consists of 55 856 logic elements, 2396 160 RAM bits, 156 multipliers (18 bit × 18 bit), four phase-locked loops (PLLs) and 377 user I/Os with a package size of 23 × 23 mm^2 [16]. The high memory-to-logic and multiplier-to-logic ratios make the chip inherently suitable for operations such as finite impulse response (FIR), DDC and NCO functions. Figure 2 shows the block diagram of the FPGA.

2.1. The pulse programmer

The pulse programmer runs at 30 MHz with a time resolution of 33.3 ns. It has a 16 bit external memory interface connected with the FPGA for register configuration and data transfer. The main loop counting in the pulse sequence is achieved using four 16 bit timers in the event manager block, and thus loop cycle times of 2^{64} can be achieved. When the event manager's interrupt for the timer is encountered, the corresponding interrupt service routine is called, and it branches into the next instruction. The control program of each event, for example, RF pulse generation or data acquisition, is packaged into a subroutine. The cycle time of each subroutine is carefully calculated and demonstrated to provide precise duration of the event.

The FPGA receives 16 bit command words from the pulse programmer and decodes them into different control signals. There are two possible modes of control signals: the internal and external signals. The internal signals are related to digital functions inside the FPGA, such as acquisition and transmission triggering. The external signals comprise controls for attenuation, RF switch and power-amplifier unblanking.

In addition, we built user-friendly software in the host computer. The software invokes a specified dynamic link library to interpret the pulse sequence programming. The sequence is programmed using assembly language and converted into hexadecimal codes by the Hex Conversion Utility of the assembly language tool. The codes are then downloaded into the DSP internal program memory through a 2 Mbyte s^{-1} USB interface controlled by a USB interface device (PDIUSBD12, Philips). Before running the sequence, a bootload program in the DSP is used to configure registers and arguments (e.g. soft pulse and gradient waveforms) in the FPGA. After that, the DSP enters the idle phase, ready for the execution of the pulse sequence.

2.2. The receiver

The receiver consists of an analog front-end circuit and a DDC inside the FPGA. The former is responsible for achieving appropriate receiver gain before the RF signal reaches the analog-to-digital converter (ADC), while the latter performs frequency translation, decimation, filtering and data transfer in the digital domain.

Two low-noise gain block amplifiers (SGA-4586(Z), RFMD, Inc.) and a digital attenuator (AT-260, M/A-COM, Inc.) provide a controllable receiver gain of 9 to 40 dB with a resolution of 1 dB. Generally, for a fixed gain of 30 dB in the preamplifier at a bandwidth of 20 MHz, 20 dB of gain is implemented in the receiver to provide a total dynamic signal range of 60 dB for the maximum input level (10 dBm) of the ADC (Analog Devices AD6644) with the resolution of 14 bits and sampling rate of 60 MHz.

The sampled signal is passed into the FPGA's DDC core. The DDC process contains the following stages: (1) quadrature demodulation, (2) multi-stage programmable filtering including a cascaded integrator-comb (CIC) decimator, CIC compensator and a user-configurable FIR filter and (3) gain adjustment. The multiplexer in

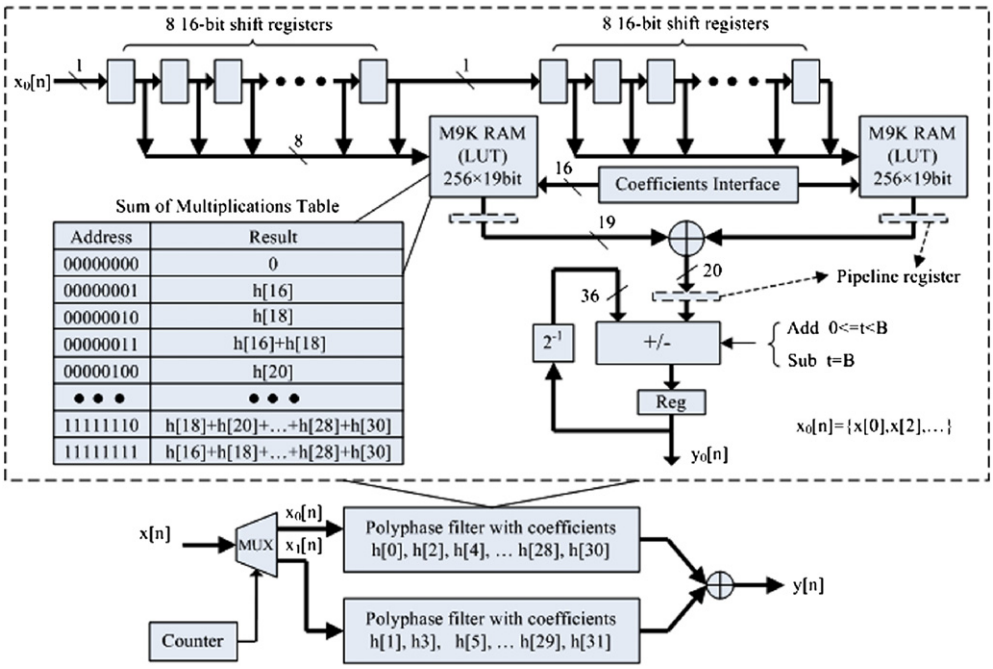

Figure 3. The block diagram of a 32-tap polyphase CIC compensating FIR filter based on a soft multiplier-based distributed arithmetic algorithm.

the control core allows each filtering stage to be bypassed depending on the application. Due to the DDC, the baseband bandwidth of the receiver is adjustable from 3.9 kHz to 10 MHz.

During the first stage, the input is multiplied by two quadrature reference signals; then it is separated into in-phase and quadrature parts. To ensure phase coherence between the RF pulse and received signal, we developed a scheme similar to that proposed in [17]: the NCO that is utilized to generate reference signals synchronizes with the one used in the transmitters. For direct sampling, the reference frequency f_0 should be determined as follows:

$$f_0 = \begin{cases} f_c \bmod f_{\mathrm{ADC}}, & \text{if } (f_c \bmod f_{\mathrm{ADC}}) \leqslant f_{\mathrm{ADC}}/2 \\ f_{\mathrm{ADC}} - (f_c \bmod f_{\mathrm{ADC}}) & \text{if } (f_c \bmod f_{\mathrm{ADC}}) > f_{\mathrm{ADC}}/2 \end{cases}, \tag{1}$$

where f_{ADC} is the frequency at which the input signal is sampled. For synchronous processing, the AD6644's data-ready output, equivalent to f_{ADC}, is used as the operation clock in the DDC.

The baseband signal is then low-pass filtered through a CIC decimator followed by two FIR filters to achieve a low output sample rate. CIC decimators are widely used in digital receivers because they require no multipliers and limited storage. These characteristics make CICs especially useful in high sample rate change. In this work, the CIC decimator is specified with programmable parameters including the decimation factor R_{CIC} from 2 to 240, differential delay of 1 and number of integrator-comb pairs of 5. Herein the Hogenauer pruning method is applied to the CIC decimator considering effective resource utilization.

A lax transition-band CIC compensator with inverse sinc response follows the CIC decimator to improve the flatness of the CIC pass band. This decimate-by-two compensator has programmable 16 bit coefficients with selectable taps of 32, 16 and 8. The filter is decomposed into two polyphase subfilters with even and odd parts of the original coefficients. Figure 3 depicts an example of designing a 32-tap CIC compensator. The demultiplexer controlled by the counter delivers even or odd input samples ($x_0[n]$, $x_1[n]$) to only one subfilter at a rate of $f_{\mathrm{ADC}}/2R_{\mathrm{CIC}}$. A single output is taken by adding the outputs of both subfilters after two input samples are applied. Each polyphase subfilter is realized using a soft multiplier-based distributed arithmetic algorithm [18]. The combinations of the coefficient summation are preprogrammed and stored in the FPGA embedded RAM block (M9K, 256 × 19 bit) as a partial look-up table (LUT). The 16 bit input samples are first shifted into a bit shift register. At each clock cycle, the shift register's taps drive the address of the LUT to read out the multiplication result for a specific bit position. In each iteration, the sum of the multiplication result is accumulated and shifted one bit to the right. In the end, the sum of the results from two polyphase subfilters gives the output of the CIC compensator.

A monorate FIR filter with up to 256 programmable 16 bit coefficients, followed by a sample rate converter, is used in the final stage for further signal processing. The FIR filter is featured with a sharp transition band to eliminate out-of-band noise. The sample rate converter has an integral rate change factor between 1 and 32. We designed the final stage filter's response using the filter design toolbox of Matlab (the Mathworks, Inc.) and generated corresponding VHDL code using Filter Design HDL Coder.

Figure 4 shows the frequency response of the filter chain in which all three-stage filters are involved.

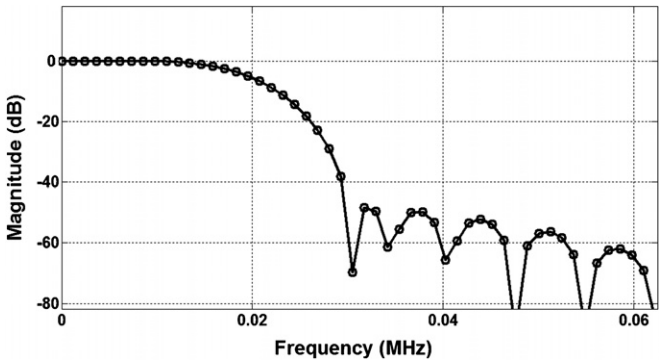

Figure 4. The overall frequency response of the three-stage multirate filter chain with an output bandpass of 10.5 kHz and an output sample period of 8 μs. The maximum ripple in the passband is 0.02 dB and the attenuation above the stopband frequency (30 kHz) is larger than 50 dB. The total group delay of the filter chain is 89 μs.

2.3. The transmitter

In addition to high resolution and spectral purity, an RF transmitter with rapid and flexible changes in its phase, frequency and amplitude modulation is preferable in a modern NMR spectrometer. We applied the direct digital synthesis technique to construct a high performance RF source. Instead of using a dedicated DDS chip, the spectrometer generates the RF pulse based on the NCO inside the FPGA. The NCO is realized based upon the coordinate rotation digital computer (CORDIC) algorithm [19]. Compared to other NCO generation methods, i.e. LUT-based or multiplier-based algorithms, the main advantage of the CORDIC-based NCO is that it achieves high phase, frequency precision, and high operating speed with fewer hardware resources. A block diagram of the one-channel DDS with the NCO core and the digital-to-analog converter (AD9742, Analog Devices, Inc.) is shown in figure 5. The NCO has precision-programmable components consisting of a phase accumulator, frequency modulator, phase modulator and magnitude register. The desired RF waveform is written as follows:

$$s(nT) = KX_0 \times \cos[2\pi(f_0 + f_{\mathrm{FM}})nT + \phi_P], \qquad (2)$$

where T is the reference clock derived from the PLL, KX_0 is the magnitude defined in the magnitude register (AMW) of N bit width, f_0 is the unmodulated output frequency based on frequency tuning word (FTW) of M bit precision, f_{FM} is the modulating frequency based on frequency modulation word (FMW) of M bit precision, and ϕ_P is the modulating phase based on phase modulation word (PMW) of P bit precision.

The relationships between related parameters are defined as follows:

$$\begin{aligned} &\mathrm{AMW} = KX_0 \times 2^N \qquad 8 \leqslant N \leqslant 32, \\ &\mathrm{FTW} = (f_0/f_{\mathrm{clk}}) \times 2^M, \\ &\mathrm{FMW} = (f_{\mathrm{FM}}/f_{\mathrm{clk}}) \times 2^M \quad 4 \leqslant M \leqslant 64, \\ &\mathrm{PMW} = (\phi_P/2\pi) \times 2^P \quad 10 \leqslant P \leqslant 32, \end{aligned} \qquad (3)$$

where the units of ϕ_P are in radians, and N, M and P are integers. For example, a high frequency resolution of 6.4×10^{-6} Hz and 0.000 02° phase offset resolution can be achieved when a 180 MHz clock drives the DDS with a frequency precision of 48 bits and phase precision of 24 bits. If a shaped pulse is desired, amplitude modulation is performed by choosing X_0 as $A[n]/K$, in which $A[n]$ is the amplitude of the modulation waveforms of the shaped pulse stored in the FPGA embedded RAM.

To demonstrate the phase, frequency and amplitude-switching agility, an example with a 10 MHz hard pulse with 180° phase switching was recorded (see figure 6(*a*)). Figure 6(*b*) shows a frequency and amplitude shift from a 20 to 10 MHz hard pulse.

The output driver of the AD9742 is a wideband operational amplifier with enable function (OPA680, Texas Instruments, Inc.), which is used as a current-to-voltage converter and RF switch. The output frequency range of the DDS is from DC to approximately 105 MHz.

2.4. Gradient waveform generator

The pulsed field gradient is usually required in NMR studies of diffusion and spatially selective spectroscopy, and is known as an alternative to phase cycling in coherence selection. Herein a gradient waveform generator is built with functions of arbitrary waveform generation and pre-emphasis inside the FPGA. The waveform is then converted into a voltage pulse and differentially amplified to drive the gradient amplifier. In NMR systems, the pre-emphasis technique could be effective in compensating for eddy currents induced by time-varying magnetic fields. Adhering to the eddy-current model [20], we used a recursive filter algorithm [21] to realize pre-emphasis

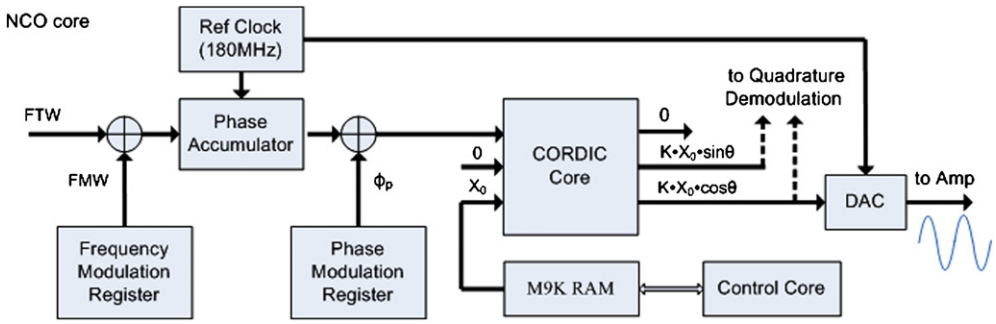

Figure 5. Block diagram of a complete NCO-based DDS

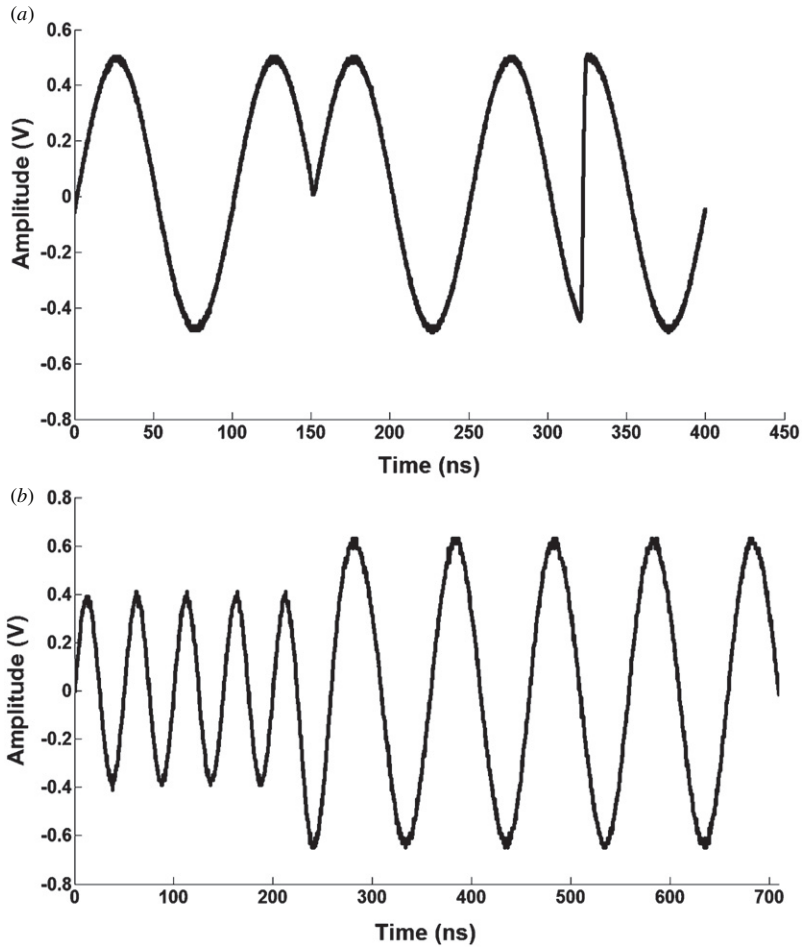

Figure 6. (*a*) 180° phase switching of a hard pulse with a frequency of 10 MHz. (*b*) Frequency and amplitude switching from 20 to 10 MHz.

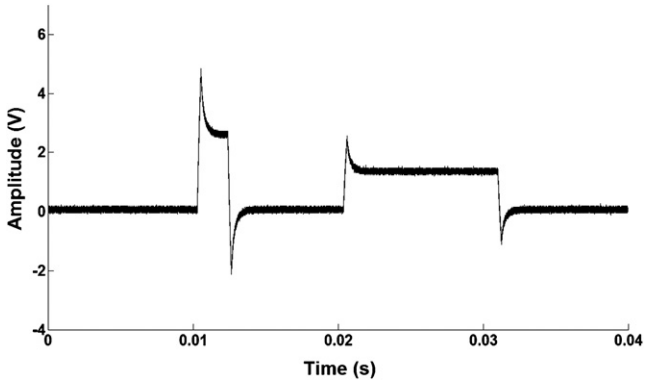

Figure 7. One-channel trapezoid-shaped gradient waveform with pre-emphasis. Two of four compensation parameters are used. Time constants are 100 and 300 μs, respectively; weights are both 80%.

calculation in which up to four adjustable compensation parameters (time constants and weights) are involved. Figure 7 shows the one-channel gradient waveform output from a 24 bit DAC (PCM1704, Texas Instruments, Inc.) with pre-emphasis setting.

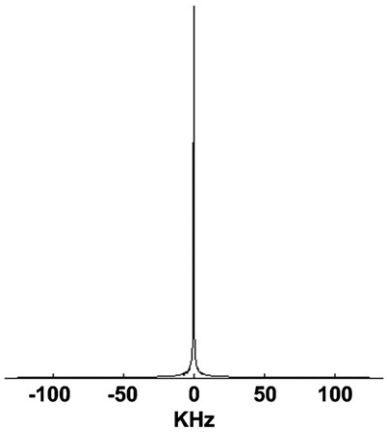

Figure 8. ^{1}H spectrum of the FID signal obtained from baby oil with the following experimental parameters: spectrometer frequency SF = 18.9 MHz, the 90° pulse width P90 = 16 μs, the dead time is 60 μs and 140 μs for probe and digital filters, respectively, the dwell time DW = 4 μs, the bandwidth of digital filters FW = 10 kHz, the number of sampling NP = 1024, and the number of excitations NEX = 1.

3. Experimental results

Several NMR experiments have been conducted to demonstrate the performance of the spectrometer in conjunction with different permanent magnetic systems (UNIQ-PMR series, Peking University). Figure 8 shows the typical ^{1}H spectrum of the free induction decay (FID)

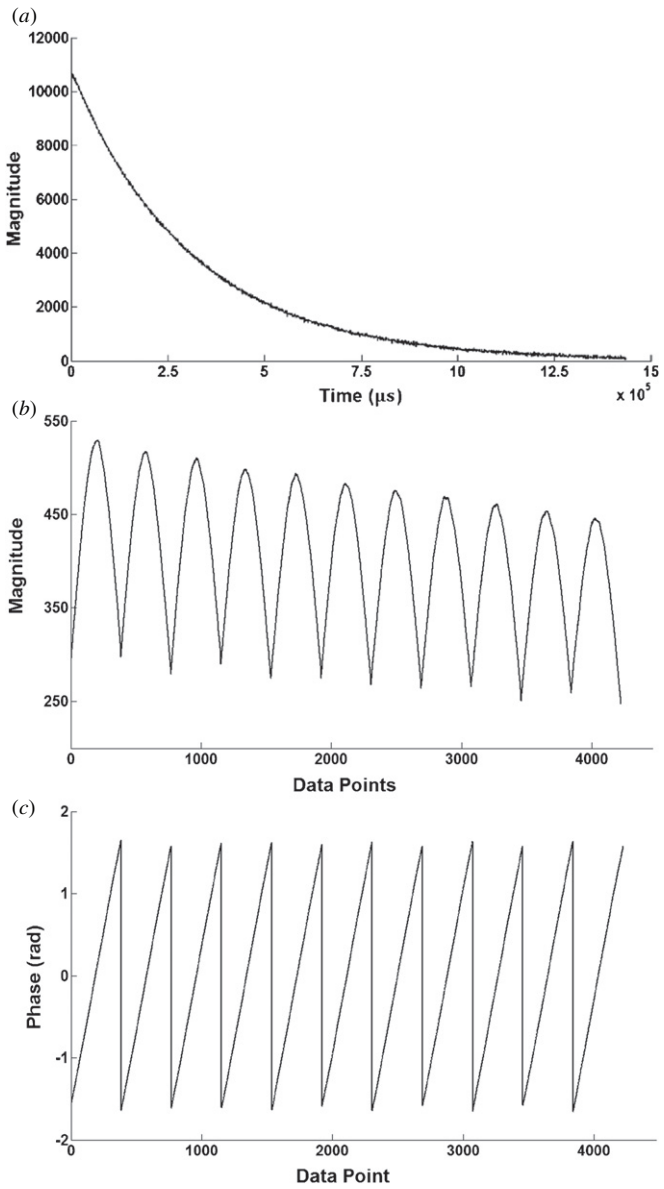

Figure 9. (*a*) Signal captured in the CPMG experiment for pure water with 0.12% copper sulfate, using a 5.44 MHz permanent magnet. The magnitude of the signal is the modulus calculated from quadrature and in-phase data with digital quadrature demodulation. (*b*) Signal captured in the CPMG experiment with the same magnet and samples as in (*a*) but 11 echoes are acquired with 384 points per echo. (*c*) The phase of the signal acquired in (*b*).

signal obtained from baby oil. The linewidth at half-height is 10 ppm. This broadening of the linewidth is mainly due to the inhomogeneity of the magnetic field in which the signal decays rapidly with a short T_2 relaxation time.

Figure 9(*a*) shows the experimental result when the Carr–Purcell–Meiboom–Gill (CPMG) sequence is manipulated and only the tops of the echoes are captured. The related experimental parameters are listed as follows: 2048 echoes with one point per echo, SF = 5.44 MHz, P90 = 34 μs, 180° pulse width P180 = 65 μs, the time between 180° pulses is 700 μs, the dead time is 50 and 60 μs for probe and digital filters, respectively, DW = 2 μs, FW = 10.5 kHz, the repetition

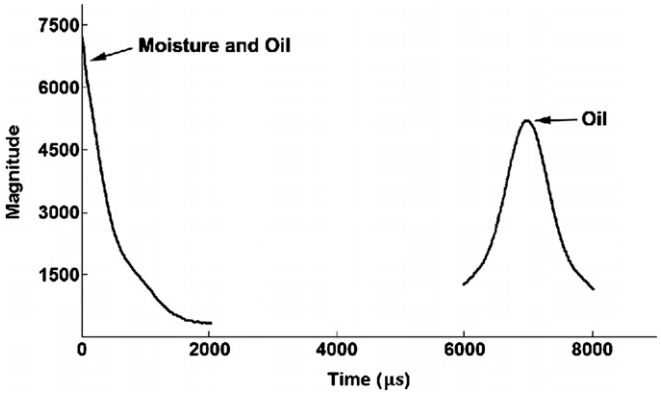

Figure 10. Hahn echo NMR signal of canola at 5.55 MHz. The acquisition points are 1024 for both FID and echo. P90 = 21 μs, P180 = 41.69 μs, the time between 90° and 180° pulses is 3.5 ms, the dead time is 20 μs for both probe and digital filters, DW = 2 μs, and NEX = 16.

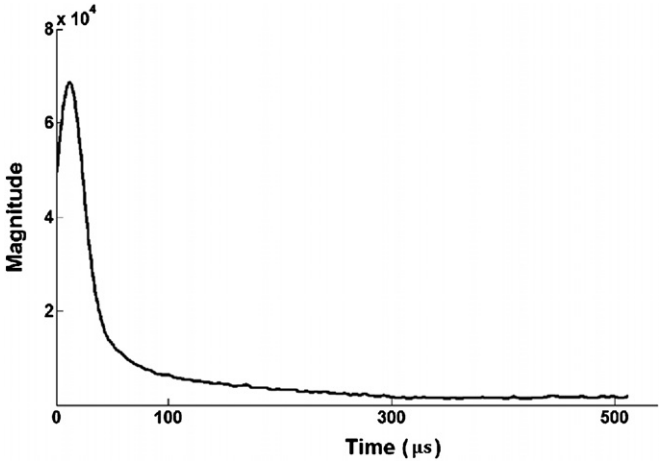

Figure 11. Solid-echo signal recorded in the 19.64 MHz permanent magnetic system from a sample of wax.

delay time RD = 1000 ms and NEX = 8. Figures 9(*b*) and (*c*) depict the magnitude and phase of 11 CPMG echoes with 384 points captured per echo, respectively.

We also demonstrated simultaneous determination of moisture and oil content in oilseeds using the Hahn echo pulse sequence method. In figure 10, the amplitude of the FID (at 50 μs) and echo (at 7 ms) is proportional to the sum of the oil and moisture signals and the oil component from the sample (canola), respectively. Therefore, the moisture content could be determined from the amplitude difference.

Figure 11 shows the solid-echo signal that is recorded in the 19.64 MHz UNIQ-PMR system from the sample of wax. Herein only the CIC decimator is used in the filter chain to give a short dead time of 8 μs for the digital filter. The related experimental parameters are defined as follows: SF = 19.64 MHz, P90 = 4.3 μs, the time between two 90° pulses is 30 μs, the dead time is 5 and 8 μs for probe and digital filters, respectively, DW = 0.5 μs, NP = 1024, and NEX = 32.

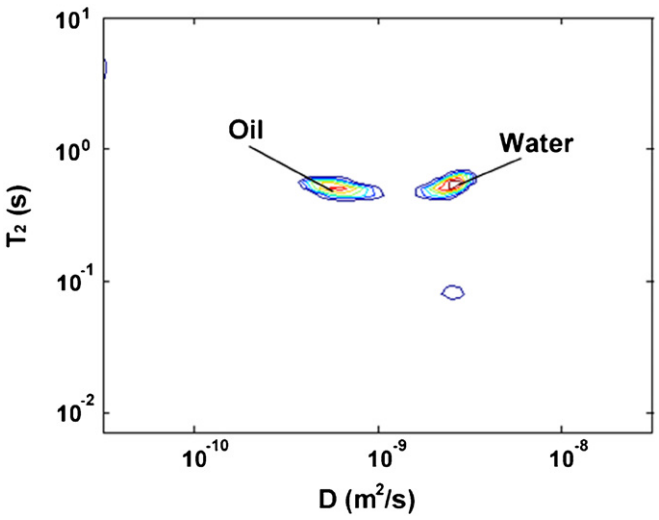

Figure 12. Two-dimensional diffusion–T_2 contour plot for samples of cores saturated with a mixture of water and oil using a 5.5 MHz permanent magnet. The vertical axis indicates the T_2 relaxation time, and the horizontal axis indicates the diffusion coefficient.

Finally we present a 2D NMR result that applied to core analysis using the CPMG sequence with field pulse gradient during the diffusion encoding time. Figure 12 shows that the water and oil peaks are identified by their diffusion coefficients being around 2.5×10^{-9} and 6.3×10^{-10} m^2 s^{-1}, respectively.

4. Conclusion

In conclusion, we presented a configurable single-board spectrometer design and its implementation based on an SDR architecture. The spectrometer was constructed mainly by combining a single FPGA chip and a DSP, thus achieving a small size and low cost. The DSP, which acts as a pulse programmer, offered a high accuracy and easy-to-program solution for the pulse sequence timing control. Inside the FPGA, one-channel DDC, two RF sources and one gradient waveform generator were implemented. The device was validated successfully in a series of NMR experiments and 2D NMR measurement, and it is expected to be suitable for portable NMR systems.

Acknowledgment

This work was supported by the Ministry of Science and Technology of China under project no 2008ZX05008-004-015.

References

[1] Michal C A 2010 *Meas. Sci. Technol.* **21** 105902
[2] Takeda K 2007 *Rev. Sci. Instrum.* **78** 033103
[3] Takeda K 2008 *J. Magn. Reson.* **192** 218–29
[4] Dykstra R, Adams M, Callaghan P T, Eccles C D, Hunter M W, Southern T and Ward R L 2007 *Proc. Int. Conf. on Sensing Technology (Palmerston North, New Zealand)*
[5] Shen J, Xu Q, Liu Y and Li G 2005 *Rev. Sci. Instrum.* **76** 105101
[6] Allodi G, Banderini A and De Renzi R 2005 *Rev. Sci. Instrum.* **76** 083911
[7] Ambrosetti R, Ranieri G A and Ricci D 1998 *Meas. Sci. Technol.* **9** 1303–10
[8] Job C, Pearson R M and Brown M F 1994 *Rev. Sci. Instrum.* **65** 3354–62
[9] Wu X, Patterson D A and Butler L G 1993 *Rev. Sci. Instrum.* **64** 1235–8
[10] Xiao L and Wang W 2009 *Rev. Sci. Instrum.* **80** 124703
[11] Jiang Y, Jiang Y, Tao H and Li G 2002 *Rev. Sci. Instrum.* **73** 3329
[12] Michal C A, Broughton K and Hansen E 2002 *Rev. Sci. Instrum.* **73** 453–8
[13] Toyoda T, Yoshida H, Oishi O and Miyajima S 1997 *Rev. Sci. Instrum.* **68** 3140–2
[14] Kenington P B 2005 *RF and Baseband Techniques for Software Defined Radio* (Norwood, MA: Artech House)
[15] Cummings M and Haruyama S 1999 *IEEE Commun. Mag.* 108–12
[16] 2009 *Cyclone III Device Handbook* vol 1 (San Jose, CA: Altera)
[17] Buszko M L, Grimm R and Griend B V 2005 The NMR processor: integrating the excitation and acquisition architectures on a single chip, *46th ENC (Providence, RI)*
[18] 2004 *Implementing Multipliers in FPGA Devices* (San Jose, CA: Altera) ver. 3.0
[19] Andraka R 1998 A survey of CORDIC algorithms for FPGA based computers *Proc. 1998 ACM/SIGDA 6th Int. Symp. on Field Programmable Gate Arrays* pp 191–200
[20] Jehenson P, Westphal M and Schuff N 1990 *J. Magn. Reson.* **90** 264–78
[21] Parhi K K and Messerschmitt D G 1989 *IEEE Trans. Acoust. Speech Signal Process.* **37** 1099–117