

# A Low-Cost Scalable Multichannel Digital Receiver for Magnetic Resonance Imaging

Ishaan L. Dalal, Ashwin L. Kirpalani, Fred L. Fontaine  
Department of Electrical Engineering, The Cooper Union, New York, NY 10003.  
{*ishaan, kirpal2, fred*}@cooper.edu

**Abstract**— Commercial receivers used for parallel MR imaging often present researchers with hurdles such as high cost-per-channel, low scalability for multiple coils and non-accessibility to intermediate data for research. A novel low-cost multichannel digital receiver for use with MR scanners has been developed to alleviate these concerns. MR signals from up to 16 coils are bandpass sampled at RF, with all subsequent downconversion performed on a single-chip Field-Programmable Gate Array (FPGA). Downconverted information is buffered and can be downloaded over a network or onto flash memory for image reconstruction. Arrays with more than 16 coils can easily scale by using more than one of these economical receivers. A 2-channel prototype has been designed, built and tested successfully with a combination of real-world signals and simulated MR data.

## I. INTRODUCTION

Today, magnetic resonance imaging (MRI) is one of the most common techniques for non-invasive medical diagnosis. Hydrogen nuclei in the body are excited by RF magnetic fields; upon subsequent relaxation of these fields, the spins of the nuclei generate a complex amplitude-modulated signal that is picked up by one or more coils. An MRI receiver amplifies, quadrature demodulates and digitizes these signals for subsequent image reconstruction.

Conventional MRI imaging of fast motion such as a beating heart used to be difficult because of the long scan times necessary. The current solution is to use parallel imaging, i.e. techniques such as SENSE and SMASH. These employ multiple coils with reduced phase encoding steps to compensate for the artifacts caused by shorter scan times. Each coil is a channel that requires its own receiver; multiple receivers may be integrated into one physical device.

Fig. 1 shows a conventional multichannel receiver that performs analog RF/IF downconversion before baseband digital sampling. Such receivers have a number of problems. First, imperfections inherent to analog stages, such as channel-to-channel mismatch, quadrature phase imbalance and DC offsets manifest as degrading artifacts in the final image. Second, available commercial receivers have eight channels at most and cannot easily scale to larger coil arrays that are becoming increasingly prevalent. Last, apart from non-scalability, the high cost-per-channel as well as the lack of access to intermediate data in proprietary receivers are significant impediments for

The design and results presented here were part of the first two authors' senior project[1], advised by F. Fontaine, at The Cooper Union.




Fig. 1. A Conventional MRI Receiver




Fig. 2. The FPGA-Based Multichannel Digital Receiver

research into image acquisition and reconstruction, especially in academic environments.

Digital receivers that eliminate the issues with analog processing have been built; they move sampling higher up the chain to IF [2] and/or compensate for hardware inaccuracies after digitization [3]. However, they invariably use custom hardware for data acquisition and multiple DSPs for digital downconversion. Multichannel versions usually duplicate single-channel circuitry; such multichannel receivers sampling at RF [4] have recently been demonstrated. Thus, while these experimental approaches resolve image quality issues, they are still quite expensive, not easily scalable and may not offer access to intermediate data from all points in the processing chain via a standardized interface.

Our goal was to design a multichannel digital receiver that sampled directly at RF (eliminating analog downconversion) and could handle 8-16 channels, scaling modularly beyond that. Hardware redundancies between multiple channels would be minimized through integration, and data from any point in the chain would be accessible over a network. In an academic environment, the primary concern was cost; the system had to consist of only off-the-shelf components. All of these objectives were met by bandpass sampling 16 channels at RF and digitally downconverting/filtering all of them on a single-chip Field-Programmable Gate Array (FPGA). Fig. 2 shows an outline of our receiver; in the remainder of the paper, its design methodology, implementation, prototype test results




Fig. 3. Analog Conditioning Subcircuit

and analysis are presented.

## II. DESIGN METHODOLOGY

The digital receiver is divided into four stages: (a) analog conditioning, (b) data acquisition, (c) digital downconversion and (d) data transfer. The external power supply is +5V; this is converted internally to 3.3V and 1.5V. All of the analog circuitry employs differential I/O because of the noise/harmonic rejection advantages of differential transmission.

### A. Analog Conditioning

This stage consists of input-protection and amplification as shown in Fig. 3. Periodic high-powered exciter spikes ( $\approx 0$  dBm) are picked up by the coils and must be blocked to avoid saturating the amplifiers. The MR scanner provides a TTL ‘unblank’ signal synchronized to the exciter; a comparator (Texas Instruments, TLV3501) gates an active RF switch (Analog Devices, ADG902) based on the unblank level. The switch provides an off-isolation of  $\geq 60$  dB, while the comparator has around 100 mV of hysteresis to counteract noise or slow-rising unblank signals. The response time of the protection circuit is  $\leq 15$  ns.

Internal pre-amplifiers directly following the coils amplify the RF signal to  $\approx -30$  dBm before the receiver input. As shown in Fig. 3, this is amplified further to 2 Vpp ( $\approx +20$  dBm) by a combined LNA/variable-gain amplifier (VGA; Analog Devices, AD8331) and an op-amp (Texas Instruments, THS4511) to fully exploit the dynamic range of the ADC. User-adjustable gains allow for individual coil weighting or other sensitivity compensations. The LNA converts the single-ended MR signal to a differential signal for the rest of the chain.

### B. Data Acquisition

For a 1.5T General Electric *Signa LX* MR scanner, the MR information (which has a maximum bandwidth of 250 KHz) has been amplitude-modulated onto a  $\approx 64$  MHz<sup>1</sup> carrier. Instead of Nyquist-sampling at  $> 128$  MHz, such a narrowband signal can be bandpass-sampled at lower rates. Fig. 4 illustrates how bandpass-sampling the bandlimited 64 MHz ( $f$ ) RF signal at a rate of 20 MSPS ( $f_s$ ) creates a baseband alias between 0 to  $f_s/2$ , centered around 4 MHz. The MR information is thus downshifted without any analog RF mixing.

<sup>1</sup>Actually 63.87 MHz. The carrier frequency is invariant and is given by  $\frac{\gamma}{2\pi} B_0$  where  $\gamma$  is the gyromagnetic ratio for the nuclei being imaged.




Fig. 4. Aliasing due to undersampling. The dashed arrows show center frequencies; content around these would also alias back into baseband at 4 MHz if the 64 MHz signal were not bandlimited.

Bandlimiting the original RF signal is key, since any spurious signals (including noise) centered around all

$$nf_s \pm (f \bmod f_s), n = 0, 1, 2, \dots \quad (1)$$

will also alias back into baseband (0 to  $f_s/2$ ) and corrupt the MR information. The ADC’s analog input bandwidth must still be  $\geq 64$  MHz, since its sample-and-hold circuitry has to run at the actual RF frequency ( $f$ ).

The data acquisition stage thus consists of a bandpass (BPF) filter followed by the ADC. Appropriate off-the-shelf BPFs are not available (previous receivers employing bandpass-sampling have used custom-built passive BPFs [4]). That approach was not financially viable, so an active filter was constructed from multiple op-amps. Building such a high-Q active BPF at 64 MHz presents certain challenges. Common topologies such as Sallen-Key do not work well because of high sensitivity to component values and impractical gain-bandwidth (GBW) requirements upon the constituent op-amps. Also, regular voltage-feedback op-amps are bandwidth-limited by the pole created by their compensation capacitors; applications needing high-GBW, high slew-rate op-amps (e.g., video) have generally used uncompensated current-feedback op-amps. These, however, are not well suited for use in active filters because a frequency-dependent impedance (i.e. capacitor) in their feedback path easily renders them unstable.

The solutions are to use the Akerberg-Mossberg topology [5] that makes filter Q independent from op-amp GBW and is relatively tolerant of component variations, along with recently released high-bandwidth voltage-feedback op-amps (Texas Instruments, THS4511). The resulting BPF is an 8th-order Butterworth, with a passband of 63.5-64.5 MHz and stopband attenuation of  $\geq 40$  dB below 62 and above 66 MHz.

The BPF outputs are connected to a 12-bit, 65 MSPS ADC (Texas Instruments, ADS5272) that integrates 8 channels in one physical package. The ADC has an analog input bandwidth of 300 MHz and an effective-number-of-bits (ENOB) of 11.5. Sampled data are output serially to save on I/O pins, and are converted back to parallel data by high-speed deserializers on-board the FPGA. ADCs with higher resolutions (14/16-bit) were investigated, but found unsuitable because they had similar ENOBs or much higher nonlinearity. The 20 MSPS ADC sampling clock must be synchronized to a 10 MHz reference provided by the MR scanner for accurate demodulation; a delay-locked-loop on the FPGA multiplies this reference to generate the necessary 20 MHz.

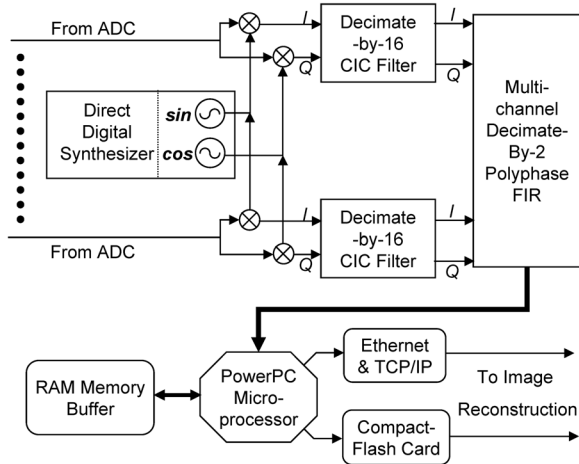


Fig. 5. Digital Downconversion and Data Transfer on the FPGA

### C. Digital Downconversion

The bandpass-sampled data is demodulated and filtered by up to 16 digital downconverters (DDC) implemented on a Xilinx *Virtex-II Pro* XC2VP30 FPGA. The FPGA has around 13,500 logic cells, two embedded PowerPC microprocessors and 136  $18 \times 18$ -bit hardware multipliers which are used whenever possible to speed up algorithms and save on logic resources.

Fig. 5 shows the digital downconversion process. A direct digital synthesizer generates 16-bit, 4 MHz I/Q LOs that are multiplied with the sampled 20 MSPS data stream. The DDC outputs are decimated by multiplier-less Cascaded-Integrated-Comb (CIC) [6] filters, which are essentially linear-phase FIR with a sinc-like response given by

$$|H(f)| = \left| \frac{\sin \pi f}{\sin \pi f/R} \right|^N \quad (2)$$

where  $R$  is the decimation ratio and  $N$  is the number of CIC stages. Time-multiplexed, four-stage CICs decimate both I and Q streams for each channel by 16 and also perform some alias rejection ( $\approx 65$  dB).

The CICs are followed by a 16-tap polyphase FIR filter that compensates for the  $\approx 1$  dB passband attenuation of the CIC and also performs a final  $/2$  decimation. Since the 16 channels share the same filter coefficients, this FIR block is a fully-pipelined, time-multiplexed implementation that handles all 32 I/Q streams, producing 625 KSPS baseband outputs. As with commercial receivers, the FIR coefficients can be configured by the user for custom lowpass factors. All processing within the DDC is done at 24-to-32-bit fixed-point precision with the final output rounded to 16 bits; thus, the design can easily use higher-resolution ADCs if desired, e.g. for clinical purposes.

### D. Data Transfer

Xilinx's reference design for a DDR RAM controller was adapted to run at 200 MHz with an inexpensive 1 gigabyte PC RAM module. The memory buffers data from the DDCs




Fig. 6. Bandpass-sampled spectrum at ADC output. The peaks at  $4 \pm 0.125$  MHz are due to the 250 KHz burst-repetition frequency of the signal generator.

and also serves as program memory for the PowerPC micro-processor. A custom-written high-speed interface allows the PowerPC to arbitrate I/O between the FPGA logic and RAM at data rates of up to 320 megabits/second.

The receiver is networked through a Fast Ethernet controller (MAC layer) built from FPGA logic and an external (PHY layer) interface. A subset of the lightweight *lwIP* TCP/IP stack [7] (written in C) was ported to the PowerPC and a basic FTP client coded to transfer buffered  $k$ -space data to a PC for image reconstruction. A CompactFlash controller and socket are also provided for non-volatile storage of data onto flash memory.

## III. TEST RESULTS AND DISCUSSION

A complete two-channel prototype of the FPGA-based receiver has been realized. The analog circuitry is surface-mount, soldered onto PCBs that were designed in-house; manufacturer-provided evaluation boards are used for high-pin-count chips such as ADC and the FPGA for rapid prototyping. Most blocks on-board the FPGA are custom-coded and optimized in VHDL; the PowerPC microprocessor is programmed with C and some assembly (for time-critical portions).

So far, the channels have been tested individually in two stages. The analog conditioning stage and the BPF were interfaced with an MR scanner. For a gain of 49 dB, the amplification chain has a noise figure of 3.2 dB. The BPF's response rolls off at around 25 dB/decade instead of the expected 40 dB/decade; this problem has been traced to parasitics and a slight component mismatch in the differential feedback loops of the filter op-amps.

The remaining stages, i.e. the ADC/FPGA have been tested with simulated MR signals in the lab. Real-world baseband  $k$ -space data of a phantom scan are extracted from GE's receiver and AM-modulated onto a 64 MHz carrier by an Agilent 4432B signal generator to create a simulated MR source.

Fig. 6 shows the bandpass-sampled spectrum at the output of the ADC. The noise floor is  $\approx 35$  dBc; this will improve with the sharper response of a revised BPF. Fig. 7 shows the baseband spectrum after being digitally downconverted and filtered on the FPGA. Data from simulated scans of a fat/water phantom are buffered to memory and transferred to a PC over the network. Fig. 8 shows a sample  $256 \times 256$  phantom reconstructed in MATLAB.

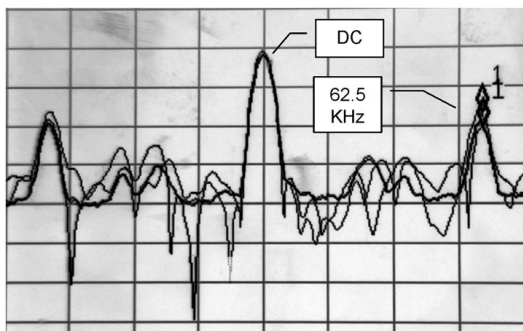


Fig. 7. Spectrum of downconverted baseband data (photographed)




Fig. 8. Reconstructed Fat-Water Phantom

Overall system performance is assessed by comparing the reconstructed images ( $M_r$ ) to those obtained from GE's receiver ( $M_0$ ), using the PSNR metric, defined as

$$20 \log_{10} \left( \frac{255}{\sqrt{\frac{1}{256^2} \sum_{i=0}^{255} \sum_{j=0}^{255} \|M_r(i, j) - M_0(i, j)\|^2}} \right) \quad (3)$$

PSNR for the reconstructed images varies from  $\simeq 30 - 33$  dB. TCP/IP network transfer was also benchmarked (1.6 megabits/sec).

Based on manufacturer-provided pricing and including PCB fabrication (but not labor), a 16-channel receiver costs approximately \$200/channel. This is more than fifty times cheaper than previous experimental efforts [2], [4]; although pricing for commercial receivers is not publicly available, it is obvious the savings would be even more substantial.

A 16-channel receiver implemented on the FPGA takes up 70% of the total logic cells available (Table I), and uses only one of the two embedded PowerPCs. Ways to utilize the remaining capacity are being considered; one concrete idea involves implementing a 1-D 256-point complex FFT and fast matrix transpose for transforming the baseband  $k$ -space data as well as an efficient CORDIC algorithm for calculating the magnitudes of the FFT outputs. This allows on-board image reconstruction, turning the receiver into a complete front-end; with the addition of a video-controller block and an external video DAC, images could be displayed in real-time on any attached monitor or TV. Preliminary simulations indicate that a 16-channel receiver/reconstruction/display system

TABLE I  
RESOURCE UTILIZATION ON THE FPGA

Block	Logic Cells	Number	Total
1-Ch DDC & Filtering	3.5%	16	56.0%
DDR RAM Controller	7.2%	1	7.2%
Ethernet & CF	6.5%	1	6.5%
Total			69.7%

implemented on the FPGA would be capable of handling real-time video imaging (up to 17 frames/sec, using sum-of-squares interpolation between coil images).

Also being investigated is the possibility of regridding spiral trajectories on the FPGA, as well as hardware implementations of the arithmetic for SENSE and SMASH. A user-interface will allow run-time reconfiguration, trading off the number of processed channels to accommodate advanced imaging/reconstruction techniques, resulting in an extremely flexible receiver-reconstruction engine.

#### IV. CONCLUSION

A design and working prototype for a novel FPGA-based multichannel digital receiver for MRI has been presented. Bandpass sampling up to 16 channels directly at RF and integrating downconversion/data transfer on a single FPGA leads to a system that is more than an order of magnitude cheaper than available receivers, is modularly scalable in 8- or 16-channel blocks and allows easy access to intermediate data from any stage over a network. Image fidelity for the prototype is comparable to commercial receivers and will improve as the implementation evolves.

#### ACKNOWLEDGMENT

The authors express their gratitude to Yi Wang and the Weill Medical College of Cornell University for use of their facilities and other cooperation.

#### REFERENCES

- [1] I. L. Dalal, M. F. R. Malik, S. M. Ahmad, J. C. Salomon, and A. L. Kirpalani, *A Low-Cost Scalable Multichannel Digital MRI Receiver*, Senior Thesis, EE Dept., The Cooper Union, New York, NY, May 2006.
- [2] P. N. Morgan, R. J. Iannuzzeli, F. H. Epstein, and R. S. Balaban, "Real-time cardiac MRI using DSPs," *IEEE Trans. Med. Imaging*, vol. 18, no. 7, pp. 649–653, 1999.
- [3] J. C. Hoenninger, L. E. Crooks, and M. Arakawa, "A floating-point digital receiver for MRI," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 7, pp. 689–693, 2002.
- [4] H. D. Morris *et al.*, "A wide-bandwidth multi-channel digital receiver and real-time reconstruction engine for use with a clinical MR scanner," *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 10, no. 7, p. 61, 2002.
- [5] D. Akerberg and K. Mossberg, "A versatile active RC building block with inherent compensation for the finite bandwidth of the amplifier," *IEEE Trans. Circuits and Systems*, vol. CAS-21, no. 1, pp. 75–78, 1974.
- [6] E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, no. 2, pp. 155–162, 1981.
- [7] A. Dunkels, "Full TCP/IP for 8-bit architectures," *Proc. First ACM/Usenix International Conference on Mobile Systems, Applications and Services*, 2003.