A 100 MHZ – 2.5 GHZ CMOS TRANSCEIVER IN AN EXPERIMENTAL COGNITIVE RADIO SYSTEM

Gio Cafaro, Neiyer Correal, Dave Taubenheim (Motorola, Plantation, FL, USA; gio.cafaro@motorola.com); John Orlando (Motorola, Schaumburg, IL, USA)

ABSTRACT

This paper describes a fundamentally flexible low power 90 nm RFIC transceiver chip that has been used in a cognitive radio demonstration. Novel circuit architectures have been applied to the 90 nm CMOS RFIC to overcome problems that have encumbered wideband transceivers in the past. Flexible programming allows the RFIC to process signals of multiple wireless protocols from 100 MHz – 2.5 GHz with channel bandwidths from 8 KHz to 20 MHz.

The cognitive radio demonstration unit integrates the RFIC, digital signal processing, a cognitive communication stack, and an embedded Linux operating system. A PC controls the unit and provides a graphical user interface, including visualization of the spectral sensing, signal detection, and dynamic frequency allocation that is taking place in the units.

1. INTRODUCTION

Radio spectrum has been historically managed via centralized planning and the allocation of static licenses. This command-and-control approach proved effective at addressing the problem of mutual interference with the unfortunate byproduct of spectrum underutilization. This is supported by spectrum measurements which indicate that at a given time a large fraction of the available spectrum lies fallow [1]. Cognitive radio promises improvements in spectral access and utilization by enabling radios to take advantage of spectral opportunities in the spatial, temporal and frequency domains and to rapidly adapt to the dynamic conditions. The potential of this technology has been recognized and acknowledged by regulatory agencies such as the U.S. Federal Communications Commission. Frequency agility, waveform flexibility, spectral awareness and intelligence are some of the key capabilities that are evolving to make cognitive radio systems a reality. This paper presents a Software Defined Radio (SDR) platform that exhibits the above capabilities. The paper is organized as follows: Section 2 describes the unique features of the RFIC transceiver, Section 3 presents one implementation of the cognitive radio platform, and concluding remarks are found in Section 4.

2. RFIC

Challenges in implementing SDR stem from its broadband nature. Receiver architectures must be chosen that minimize the need for filtering since low-loss, broadband tunable filters are not practical in today's technology [2]-[4]. Direct conversion is preferred for this reason, but it has some disadvantages that must be addressed through circuit techniques as discussed in Section 2.2. Quadrature Local Oscillator (LO) generation over a wide frequency range is another challenge and is typically addressed with multiple high frequency VCOs and complex divider schemes [5]. Unfortunately, this approach often leads to spot coverage that targets some subset of predefined frequencies. Direct digital synthesis offers some attractive features like fast switching and continuous coverage, but again, it has disadvantages that require particular attention. Our solutions to these challenges and others will be discussed in the remainder of this section.

Table I

Summary of Transceiver Performance

Freq. Range	100 MHz – 2.5 GHz
Rx NF	7 dB
Rx Gain	48 dB
Rx IIP2	+60 dBm
Rx IIP3	-6 dBm
Rx Current Drain	40 mA
Tx Output Power	+6 dBm
Tx Sideband Suppression	35 dBc
Tx Current Drain	40 - 90 mA
EVM π/4 DPQSK 3.5 MS/s	1% @ 800 MHz
LO Phase Noise	-123 dBc/Hz @ 25 KHz
LO Frequency Resolution	15 Hz
LO Current Drain per DDS	80 mA

2.1 Architecture

Referring to Fig. 1, three independent direct digital synthesizers (DDS) use a single 1 GHz PLL reference to provide differential quadrature LO signals to the receiver, transmitter, and transmitter linearization Cartesian feedback mixers. Direct conversion is used in the receiver and a

direct launch quadrature modulator is used in the transmitter. One of five receiver paths is selected to drive a common analog baseband low pass filter section with programmable corner from 4 KHz to 10 MHz. There are provisions for receiver Automatic Gain Control (AGC), DC offset correction, and in band and out of band Receiver Signal Strength Indicator (RSSI). Dynamic matching is used in the direct conversion mixers for improved second order intermodulation intercept point (IP2), flicker noise, and DC offset [6]. Differential baseband analog in-phase and quadrature receiver signal outputs are provided for external connection to an ADC and digital processing.

Differential baseband in-phase and quadrature inputs from an external transmitter DAC are applied to programmable low pass filters similar to the receiver with 10% bandwidth steps from 4 KHz to 10 MHz bandwidth. There is one of three selectable transmitter paths with up to 80 dB of on chip programmable gain available (for power control). A transmitter feedback network is provided for closed loop narrow band linearization or open loop alternative transmit signal analysis and processing. A summary of transceiver performance is shown in Table I.

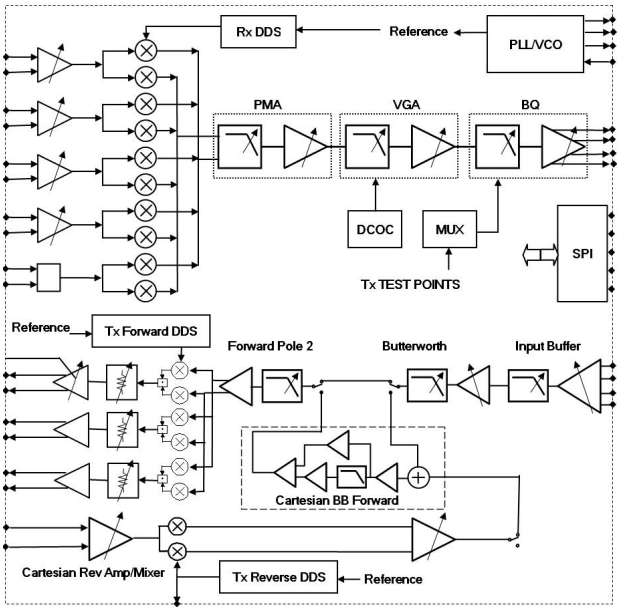

Figure 1. Transceiver Block Diagram

2.2 Receiver

The RFIC contains 5 fully differential receiver inputs that drive fully differential, dynamically matched (chopping), passive, quadrature (I/Q) mixers. Four of the receiver inputs each include an inductorless LNA. Baseband filters that support multiple bandwidths are also implemented along with gain control and DC offset correction.

The quadrature mixers on Rx inputs 1, 3, and 5 are non-chopped "passive" mixers built with a quad ring of CMOS transmission gates. The quadrature mixers on Rx inputs 2 and 4 use dynamic matching (chopping) to improve IP2, flicker noise and DC offset. The chopping is implemented with three mixers in series, where each mixer is built with a quad ring of CMOS transmission gates.

Because the mixer design is passive (with active CMOS devices acting as switches), excellent power drain, linearity and noise figure are achieved. Current drain from the 1.2V supply of the LO buffers for the non-chopped I/Q mixers is 3.7 mA at 1 GHz, while IIP3 of the mixers is +17 dBm. Noise figure of the mixers is 5 dB (essentially equal to the conversion loss). All of the mixer current drain comes from the LO buffers and multiplexers, since there is no DC current drain in the actual switching mixer CMOS devices.

The baseband filter architecture has four poles of filtering with two real poles and one complex pole pair in the Sallen-Key BiQuad. Filter bandwidth is programmable from 4 KHz to 10 MHz in 6.25% steps or less. Bandwidth selection is implemented by adjusting the resistor and capacitor values in the filter design. The user has independent control of the pole locations of the post mixer amplifier (PMA), voltage gain amplifier (VGA) and BiQuad as well as control of the BiQuad filter Q. This gives the user the flexibility to trade off filter shape and attenuation for pass-band amplitude and phase distortion to suit his application.

Baseband filter gain control is accomplished at three points – at the PMA input, in the VGA, and in the output buffer. The entire baseband filter lineup has a maximum gain of 64 dB and a minimum gain of –4 dB.

A notable feature of the baseband filter is the use of chopper stabilization to mitigate the undesirable effects that occur in direct conversion receivers when designed in a CMOS process. This is of particular concern in narrowband FM applications where CMOS flicker noise can degrade sensitivity and noise figure. The measured improvement of chopper stabilization is shown in Fig. 2.

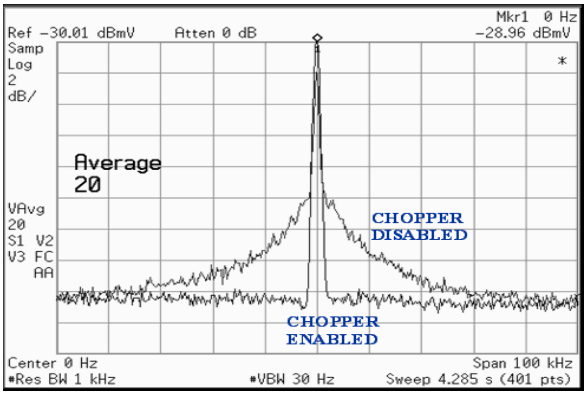

Figure 2. Received signal spectrum with and without chopping.

DC offset correction circuitry (DCOC) is implemented as a complete control loop that automatically corrects DC offsets at the output of the baseband filter. DCOC consists of a 1-bit ADC (comparator), control logic, and a 5-bit current mode DAC that injects current into the feedback resistors of the VGA to adjust the offset voltage. The control logic implements a successive approximation algorithm that converges on the correct 5-bit word that compensates for the filter's DC offset.

2.3 Transmitter

The transmit direct launch quadrature modulator will support both linear and constant envelope modulation formats to cover standards with baseband bandwidths of 4 KHz to 10 MHz (channel bandwidths of 8 KHz to 20 MHz) and RF carrier frequencies from 100 MHz to 2.5 GHz. There are two fully differential modes of operation – classical I/Q or polar (with an external PA). In either case, baseband bandwidths and output power are programmable to meet the spectral mask requirements of a given protocol. For narrow and medium band protocols, a Cartesian feedback system provides the necessary linearization. This system provides a downmix path (receiver) that samples the output of the power amplifier and uses that sampled signal to correct for any non-linearity induced errors in the forward transmission path.

The baseband transmit block provides filtering, programmable attenuation, level shifting and buffering for the DAC inputs and drives the forward RF section and/or the Cartesian baseband forward path. The input buffers provide stepped attenuation for the incoming baseband signals. Programmable active RC reconstruction filters limit the amount of far out quantization noise and images due to aliasing. An RC tracking network provides automatic filter pole adjustment for the reconstruction filters. Closed loop correction of baseband DC offsets and I/Q phase gain imbalance is done in the pre-transmit warm-up period.

The forward RF chain contains three separate RF mixer/driver paths along with the associated biasing and gain control. Each path is independently programmable to trade off bandwidth, power control range, and linearity according to the signal protocol being processed.

The Cartesian feedback path is essentially a highly linear direct conversion receiver with low sensitivity. The output of the external power amplifier is coupled into the RF inputs of this block where the signal is either amplified or attenuated and then down converted to baseband using the down-mixer and feedback LO. Two baseband amplifiers provide programmable gain for the feedback signal before being fed into the Cartesian forward path. With its programmable gain control, this block provides the feedback gain in the Cartesian system that ultimately controls the output power of the transmitter. Two 6-bit DC offset DACs tune out any offset errors at 2.5 mV per step.

2.4 Frequency Synthesizer

Direct digital synthesis is used to supply LO signals for the receiver and transmitter (forward and reverse paths) as well as the clock for the chopping mixers. DDS has a feature set that makes it attractive in SDR applications. A very wide tuning range can be achieved with a single VCO. In the RFIC, a single integrated VCO provides the 1 GHz clock to the digital processing blocks of multiple independent DDSs, each of which can be tuned independently and with phase coherent properties. This arrangement is immune VCO pulling (or transmit re-modulation) [7] since the VCO is not operating at the DDS output frequency. Each DDS generates square wave outputs with fast rise and fall times that are ideal to drive the switching mixers in the receiver and transmitter. Finally, the cycle-to-cycle frequency switching enables unique transceiver capabilities that are not possible with traditional phase locked loops with their associated lock times.

Traditional direct digital synthesizers have two disadvantages relative to PLL-based synthesizers – power consumption and spurious frequency content. The DDS architecture developed in Motorola Labs [8]-[10] uses a ROM-less architecture to achieve typical power consumption below 120 mW and non-zero-mean dither [11] to keep spurious frequency components below -35 dBc.

The differential I and Q outputs can switch frequency on a (glitch free) cycle-to-cycle basis anywhere from 100 MHz to 2.5 GHz. With 15 Hz resolution and measured phase noise of -123 dBc/Hz at 25 KHz offset and -150 dBc/Hz at 20 MHz offset, the DDS has the flexibility and the noise and switching time performance that is needed for SDR applications.

3. EXPERIMENTAL COGNITIVE RADIO SYSTEM

In this section, we describe an experimental cognitive radio platform that leverages the flexibility of the RFIC described above. We demonstrated an early prototype of this platform at DySPAN in 2007, capable of sensing its signal environment and determining what spectrum to use for a simple live video application.

The remainder of this paper is organized into sections that describe each major function of the radio system. Section 3.1 begins with a description of the system's physical layer, illustrating the radio's modulation and spectrum utilization. In Section 3.2 we cover the radio's spectral sensing module, which provides vital frequency usage information to enable dynamic use of the band. We introduce neighbor discovery and link maintenance

functions in Sections 3.3 and conclude with system implementation details in Section 3.4.

3.1 Communications Physical Layer

The PHY described in this section serves as an example of the flexibility of the platform. It demonstrates the rapid prototyping capabilities of the system and was chosen to meet the aggressive timeline needed to participate in DySpan.

The premise behind the PHY is that communication between cognitive radios takes place using arbitrary spans of spectrum that can be changed on-the-fly. We call these fundamental spans "spectrum blocks", or SBs. The larger spectrum pool is partitioned into N non-overlapping SBs. Each SB is further divided into M subcarriers (SCs), as illustrated in Fig. 3. In general, the sizes of the SBs and SCs are system dependent, but in our implementation we define an SB as the granularity of the spectrum sensing decision blocks, and an SC as the smallest communication utilization unit. While several PHY modulation techniques could be applied to this spectral scheme, we use OFDM. With OFDM it is possible to leave any subset of SBs unmodulated and thus provide a flexible spectral shape that can fill spectrum gaps with low interference to other users.

The OFDM communication system operates in the 2 GHz experimental band that was allocated for DySPAN, that is, two 50 MHz wide channels centered at 2.056 GHz and 2.231 GHz. We use 128 subcarriers in a data symbol, distributed over a 20 MHz bandwidth, resulting in a subcarrier frequency spacing of 156.25 kHz, centered about a programmable channel frequency. Thirty-two of the 128 subcarriers are reserved as guard tones with 16 at each end of the 20 MHz range. We further specify that an SB contains 8 subcarriers, each modulated with DQPSK. In total, the usable 96 data subcarriers comprise 12 spectrum blocks, each having bandwidth of 1.25 MHz and carrying 2.0 Mbps, giving a maximum raw data rate of 24 Mbps.

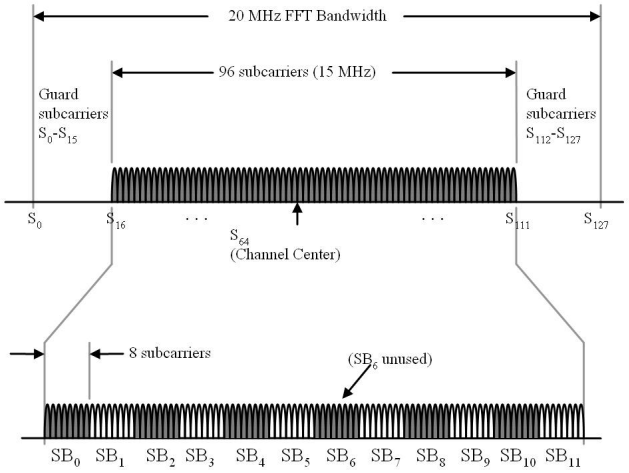

Figure 3. Allocation of subcarriers and blocks

For a single OFDM data symbol, each SB carries 16 bits of data, mapped onto the 8 complex data subcarriers within the SB. Data are simply taken from the source in a bit-wise fashion and mapped onto the SBs that constitute the logical channel. The lowest-frequency SBs are filled first, and then higher-frequency SBs are filled sequentially, until all selected SBs are full. We allot a guard time of 1.6 μs between symbols to provide tolerance for delay spread and time for pulse shape rolloff. Given the 6.4 μs FFT duration, the system's symbol period is 8 μs. Each OFDM symbol contains modulated subcarriers from one or more SBs. Only the subcarriers within the used SBs are modulated; all others are set to zero. Table II summarizes the characteristics of the experimental PHY.

Table II

Summary of Modulation Performance

Bandwidth (MHz), B	20
FFT Bins, N	128
# Guard Subcarriers, N_G	32
# Data Subcarriers, N_D	96
Subcarrier Spacing (kHz), $\Delta_F = B/N$	156.25
FFT Duration (μs), $T_{FFT} = 1/\Delta_F$	6.4
Guard Interval (μs), T_G	1.6
Symbol Length (μs), $T_S = T_G + T_{FFT}$	8
Symbol Rate (ksymb/s), $R_S = 1/T_S$	125
Subcarrier Data Modulation	DQPSK
Coding Efficiency, η_C	1
Channel Data Rate (Mbps) = $2*N_D*R_S*\eta_C$	24

The PHY protocol data unit consists of a preamble, a header, and a payload. The four-symbol-long (32 μs) preamble is designed to facilitate correlation for timing recovery and AGC settling, with its last symbol as a starting symbol for the DQPSK modulation. The preamble is defined as though all 96 subcarriers will be transmitted, but in practice, only a few of the SBs are enabled. As a result, the preamble waveforms need to be created dynamically, based on which SBs are active, in a manner similar to the data symbols. The PHY header is a 12-bit field within the first data symbol that specifies the length of the payload. The payload is variable in length, up to 4095 bytes, and contains the actual data to be transferred. The payload also contains a 32-bit CRC for error detection.

What is the appropriate coding and interleaving strategy when only a fraction of the subcarriers are used? Since we are already dynamically selecting spectrum blocks to avoid interference, we could choose the blocks to avoid fades. However, this would require active sounding of the spectrum in addition to passive sensing for interference, and its effectiveness would depend on the mobility of the device

and its surroundings. Another option is to use a single FEC code and then interleave across the active subcarriers. The coding and interleaving would be more effective as more subcarriers are used, so a likely strategy would be to use as many spectrum blocks as possible on a given transmission and occupy fewer time slots.

A third option is to employ simple frequency diversity. Here, the available spectrum blocks are divided into two sets, and data are duplicated and sent simultaneously on both sets. This is equivalent to a half-rate repetition code with interleaving across frequency to provide diversity. Due to its low implementation complexity and ease of adaptation for dynamic spectrum use, frequency diversity was the most attractive choice for our initial PHY protocol.

3.2 Spectrum Sensing and Detection

For a radio to be considered cognitive, it must monitor the local RF environment and identify available spectrum that can be used for communications. Our approach to spectral monitoring is based on spectral correlation analysis. We measure the correlation of spectral bands at particular frequency-shifted versions of the input signal. The spectral correlation density function is defined as the time averaged product of narrowband spectral components (with bandwidths ~ *1/T*) of the input signal separated in frequency by α where the time averaging in ‹·› goes to infinity and the bandwidth goes to zero [12]:

$$S_x^{\alpha}(f) \equiv \lim_{T\to\infty} \frac{1}{T} \left\langle X(t, f+\alpha/2) X(t, f-\alpha/2)^* \right\rangle$$

The sensing process begins with analog-to-digital conversion and acquisition of the in-phase (I) and quadrature (Q) samples. These samples are processed with a *K*-point complex FFT to obtain the 24-bit resolution spectrum of the input signal. The resulting bin samples are then processed frame by frame in real time. At DySPAN, the processing interval was 5 ms over a 10 MHz bandwidth, and K=1024.

Signal detection divides the observed spectrum into small sub-bands of interest and outputs a decision that corresponds to the analysis of the α=0 space of the spectral correlation density (SCD), which in turn corresponds to the power spectral density (PSD). The decision is used to classify whether or not the sub-bands are occupied; detection can be achieved across the whole band or over specific portions of the band. For the DySPAN system, the detectors measured the energy in each sub-band and decided if a signal was present, based on adjustable thresholds.

3.3 Neighbor Discovery

For peer to peer communication among secondary users in a cognitive system, the first task for a radio is to discover its communication partners; a secondary user "Tx" tries to discover another secondary user "Rx". At DySPAN we achieved discovery through channel scanning and beacon broadcasting [13] using the frequency-time pattern illustrated in Fig. 4. After spectrum scanning, both Tx and Rx identify one another and find *N* available SBs with which to communicate.

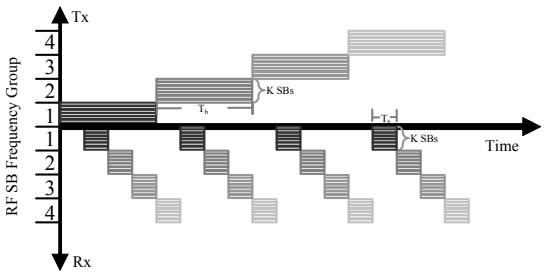

Figure 4. Frequency-time diagram for neighbor discovery process

Future research will consider dynamic cognitive radio environments in which channel status may change quickly. These are environments where, unless users take proactive action, the changes will likely result in loss of the RF link. Once communicating radios lose track of each other, they need to go through the costly neighbor discovery process again. A reliable link maintenance protocol is thus crucial for the success of this experimental cognitive radio system, and further details of our thinking for the next version of the system are provided in [14].

3.4 System Implementation

The radio unit in Fig. 5 contains the RFIC described in Section 2 as well as a Xilinx FPGA and a PowerPC microprocessor. The complex input and output of the transceiver interface via data converters to a Xilinx XC4VSX35 FPGA, which perform the digital signal processing for the cognitive radio's spectrum sensing, signal detection, modulation, and demodulation. In the FPGA a soft-core 32-bit processor, MicroBlaze [15], is instantiated to handle certain physical layer tasks, such as packet formatting and parsing, and to perform the neighbor discovery functions. The FPGA is interfaced via a compact form-factor PCI bus to a PowerPC microprocessor running ElinOS [16] embedded Linux. The PowerPC processor binds the wireless link to a TCP/IP Ethernet connection. The TCP/IP connection serves two purposes in the demonstration platform: to provide a means to observe the workings of the cognitive radio on a PC GUI, and as a port for data to be sent or received over the air.

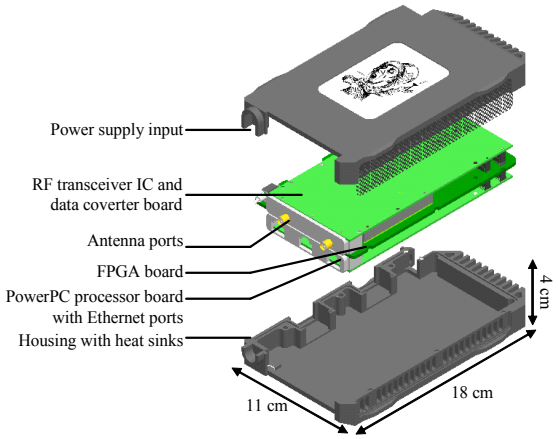

Figure 5. Cut-away view of IF- and signal-processing module hardware

The neighbor discovery protocol is implemented in C code, running in the FPGA's MicroBlaze microprocessor. While either of the radio's processors - MicroBlaze or PowerPC - afford sufficient computational resources, such as memory, speed, and timers, to handle the protocol code efficiently, there is an additional motivation for choosing to execute the protocol in the MicroBlaze; the collocation of the processor and the digital signal processing functions allow for each function to become a memory-mapped peripheral in the processor. As peripherals, spectral sensing and transceiver functions are available to the link management protocol with little latency or overhead, which in turn expedites the actuation of the real-time decisions made by the protocol.

The digital signal processing of the experimental radio was implemented using a suite of FPGA design tools. For example, Xilinx System Generator, a Simulink-based FPGA design tool, was used for the design of the radio's spectral sensing block and the digital signal processing blocks of the OFDM physical layer. These signal processing blocks make up approximately 80% of the FPGA's occupied logic.

Finally, a graphical user interface provides visualization and control of the radio. The GUI is written in Matlab and connects to the radio via Ethernet link. Part of the visualization suite includes a display of spectral analysis results over the allotted frequencies and SBs allocation.

4. CONCLUSIONS

We have presented a fundamentally flexible RFIC transceiver that enables SDR functionality in a compact form factor. We also presented an SDR platform, built around the RFIC, that has proven extremely effective as a rapid prototyping platform. In various projects, we have used it to implement UMTS, GSM, Tetra, and proprietary waveforms.

The SDR platform is also valuable as a testbed for cognitive radio system development. As an example, it allowed a small team to quickly develop the system described in this paper for a demo at DySpan.

5. REFERENCES

[1] M. McHenry, D. McCloskey, "Spectrum Occupancy Report for New York City during the Republican Convention August 30 - September 1, 2004" Shared Spectrum Company. (http://www.sharedspectrum.com/inc/content/measurements/nsf/NYC_report.pdf)

[2] M. Steyaert, B.D. Muer, P. Leroux, M. Borremans, K. Mertens, "Low-voltage low-power CMOS-RF transceiver design," IEEE Trans. on Microwave Theory and Tech., vol. 50, no. 1, pp. 281-287, January 2002.

[3] P. Zhang, et. al., "A single-chip dual-band direct-conversion IEEE 802.11a/b/g WLAN transceiver in 0.18 micron CMOS," IEEE Journal of Solid-State Circuits, vol. 40, no. 9, pp. 1932-1939, September 2005.

[4] T. Maeda, et. Al., "Low-power-consumption direct-conversion CMOS transceiver for multi-standard 5-GHz wireless LAN systems with channel bandwidths of 5-20 MHz," IEEE Journal of Solid-State Circuits, vol. 41, pp. 375-383, February 2006.

[5] R. Bagheri, et. al., "Software-defined radio: dream to reality," IEEE Communications Magazine, pp. 111-118, August 2006.

[6] E. Bautista, B. Bastani, J. Heck, "A high IIP2 downconversion mixer using dynamic matching," IEEE Journal of Solid-State Circuits, vol. 35, pp. 1934-1941, December 2000.

[7] O. Charlon, et. al., "A low-power high-performance SiGe BiCMOS 802.11a/b/g transceiver IC for cellular and Bluetooth co-existence applications," IEEE Jounal of Solid-State Circuits, vol. 44, no.7, pp. 1503-1512, July 2006.

[8] J. Juan, R. Stengel, F. Martin, D. Bockelman, "Cascaded delay locked loop circuit," US Patent 6891420, Dec. 2006.

[9] D. Bockelman, J. Juan, "Time interpolating direct digital synthesizer," US Patent 6353649, March 2002.

[10] F. Martin, R. Stengel, J. Juan, "Method and apparatus for digital frequency synthesis," US Patent 7154978, May 2005.

[11] T. Gradishar, R. Stengel, "Method and apparatus for noise shaping in direct digital synthesis circuits," US Patent 7143125, November 2006.

[12] W. A. Gardner, "Cyclostationarity in communications and signal processing", IEEE Press, 1994.

[13] J. Zhao, H. Zheng and G. Yang, "Distributed coordination in dynamic spectrum allocation networks", Proc. IEEE Dyspan 2005, pp. 259-268, Nov. 2005.

[14] Q. Shi, D. Taubenheim, S. Kyperountas, P. Gorday, and N. Correal, "Link maintenance protocol for cognitive radio system with OFDM PHY", presented at DySPAN 2007.

[15] http://www.xilinx.com/ipcenter/processor_central/microblaze

[16] Sysgo AG. http://www.sysgo.com/en/products/elinos