

A 10 MHZ – 4 GHZ DIRECT CONVERSION CMOS TRANSCEIVER FOR SDR APPLICATIONS

¹Gio Cafaro (gio.cafaro@motorola.com), ¹Tom Gradishar, ²Homero Guimaraes, ¹Joe Heck, ²Jeff Isbell, ¹Geetha Nagaraj, ²Andy Pagones, ¹Raul Salvi, ²Tony Schooler, ¹Bob Stengel, ¹Sumit Talwalkar, ²Poojan Wagh, ²Curt Williams, ²Jack Zvonik, ²Terry Blake, ²Greg Burton, ²Robert Crosswell, and ²Bill Ziemer

¹Motorola, Plantation, FL
²Motorola, Schaumburg, IL

ABSTRACT

In 2007 we reported a flexible integrated circuit transceiver operating from 100 MHz to 2.5 GHz with continuous coverage over that range [1-2]. A new version of the chip is now available with improved performance and significant new features. Notable performance improvements include operation up to 4 GHz using the on-chip direct digital synthesizers (DDS) and up to 6 GHz using an external LO source. Noise figure, phase noise, linearity, and Tx output power are much improved in the 2-4 GHz range. The most significant new features are on-chip data converters and a high speed digital interface based on the digRF standard. The new chip is 5.0 x 5.4 mm in 90 nm CMOS and is housed in a 10 x 10 mm 132-pin dual row MFL package. The system development kit includes GUI programming software and an evaluation board with a Virtex5 FPGA and dual 1 Gb Ethernet ports.

1. INTRODUCTION

A Software Defined Radio is capable of operating over a broad continuous frequency range, regardless of modulation type and channel bandwidth, where the hardware is reconfigurable via software. The operating ranges implied in that definition represent a significant challenge to hardware designers attempting to implement a radio within the size and power dissipation constraints of a portable device.

The practical motivation behind SDR is the ability to use a single transceiver for as many communication protocols as possible. Beyond the stated frequency and bandwidth goals lies consideration of varying system requirements and RF operating environments. Transceivers targeting a specific application will make engineering tradeoffs that most efficiently meet those requirements. Transceivers targeting multiple protocols need to meet the superset of requirements and therefore have less freedom to make engineering tradeoffs. If a concurrent goal is a high level of integration, i.e. on-chip, then the degrees of

freedom are further constrained. This paper describes a radio frequency integrated circuit (RFIC) that is the result of the engineering tradeoffs described above. Other examples can be found in [3]-[7].

This RFIC was developed with flexibility as the most important requirement. In general, only those circuit functions that could be made completely flexible were included on-chip. Where performance is limited by either the process technology or by design choices, the transceiver may be augmented with external components to extend the performance. For example, the linearity of the Rx front end is set by the inductorless broadband LNA that is on-chip. For applications that require very high linearity, an external LNA can be coupled into an alternate Rx port that couples directly into a very linear mixer. For applications above 4 GHz or those requiring lower phase noise than can be achieved with the on-chip DDS, an external LO source can be used. For applications that require real bandwidths greater than the 20 MHz bandwidth provided by the on-chip baseband circuits, external I/Q filters and amplifiers can be used and routed back on-chip through the sigma delta ADC.

The preceding examples are by no means a complete listing of the modes that can leverage high quality discrete components to improve overall system performance. However, we will limit the discussion here to the typical cases where all major functions are performed on chip.

The rest of the paper is organized as follows. Section 2 outlines the overall architecture of the RFIC. Section 3 describes the frequency synthesizers. Sections 4 and 5 detail the Rx and Tx chains, respectively, along with their data converters. Section 6 introduces the high speed digital interface and Section 7 the systems development kit.

2. ARCHITECTURE

Referring to Figure 1, four independent direct digital synthesizers (DDS) use a single 1 GHz PLL reference to provide differential quadrature LO/clock signals to the receiver, transmitter, Cartesian feedback mixers, data converters, and high speed digital interface. Direct

A programmable high speed digital interface uses low voltage differential signaling (LVDS) to transfer transmit and receive samples between the RFIC and a baseband processor. The interface clock, supplied by the RFIC, determines the data transfer rate of the serial interface.

[illegible]

The chip can be configured such that only the analog portion is used. The analog baseband inputs and output can be sent off-chip, thereby bypassing the data converters and high speed digital interface.

A single 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 completely immune to VCO pulling and transmit re-modulation since the VCO is not operating at the DDS output frequency [8]. Each DDS

Traditional direct digital synthesizers have two disadvantages relative to PLL-based synthesizers – power consumption and spurious frequency content. The DDS architecture developed in Motorola [9]-[11] uses a ROM-less architecture to achieve typical power consumption below 120 mW and non-zero-mean dither [12] to keep spurious frequency components below -35 dBc.

4. RX CHAIN

The RFIC contains 5 fully differential receiver inputs that drive fully differential, passive, quadrature (I/Q) mixers. Four of the five Rx inputs have on-chip LNAs with 50 Ohm differential input impedance and the last is designed with 200 Ohm differential input impedance for use with an external LNA. The Rx input path is selectable through software. The LNA is a common gate architecture with cascoding and RF AGC provided by the current steering pair. Source followers are used to drive the capacitive load of the mixer stage as well as the significant parasitic of the caps used for AC coupling.

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 for the non-chopped I/Q mixers is 3.7 mA typical, 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). The drain with the chopping mixers is somewhat higher than the non-chopped mixers due to the additional chopping clock buffers, and depends on the chopping frequency.

Baseband filters that support multiple bandwidths are implemented along with gain control and DC offset correction. The filter architecture has four poles of filtering with two real poles and one complex pole pair in the Sallen-Key BiQuad. The shape was designed by taking a fourth order Butterworth prototype and replacing the low Q complex pole pair with two real poles. Filter bandwidth is programmable from 4 kHz to 10 MHz in 6.25% steps or less. Sufficient margin is built into the design to allow for a 20% change in RC tolerance and still maintain the bandwidth range of 4 kHz to 10 MHz. 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 flexibility to trade off filter shape and attenuation for pass-band amplitude and phase distortion.

Baseband filter gain control is accomplished at three points. A programmable resistor divider at the input of the PMA allows attenuation in four 6 dB steps while maintaining an input impedance of 2 K Ω , differential. The PMA has a maximum gain of 32 dB and a minimum gain of -10 dB. The VGA has a gain range of 8 dB and the output buffer has a programmable gain control of 0 to 18 dB in 6 dB steps. 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 - even order distortion, flicker noise, and DC offset - 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.

Chopper stabilization is implemented around the first stage amplifier of each two-stage op-amp. This implementation was chosen due to the fact that the op-amp's input referred voltage offset and flicker noise performance are heavily dependent on the first stage of the op-amp. The chopping clock is produced by dividing the PLL VCO frequency by 16, 32, 64, or 128. Flicker noise is essentially eliminated when chopping is enabled and thus narrowband protocols will see improvement in receiver sensitivity.

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.

4.1 Analog to Digital Converters

Two sets of analog to digital converters (ADCs) exist on the RFIC. Each set consists of an I- and a Q- ADC core. One set is a pulse-width-modulation ADC (PWMADC), which is an advanced form of a sigma-delta modulator. The other is a 1-bit 2nd order continuous time sigma-delta modulator designed to leverage the high speed available from 90 nm CMOS. Due to multiplexing on the decimators, only one of the ADC's can output data at a given time.

In this section, a pair refers to a differential pair and a quad refers to two pairs, one pair for the in-phase (I) baseband signal and one for the quadrature (Q) baseband signal.

The ADC section has an internal routing bus that can route between multiple analog quad inputs and multiple analog quad outputs. There are also numerous clocking options for the ADC's. The ADC's can take any of the following as a clock source:

1. Direct PLL (~ 2 GHz)
2. On-chip DDS
3. Off-chip clock source

The PWMADC employs pulse-width modulation (PWM) to encode an analog signal as a binary signal. By varying the pulse width rather than amplitude, the PWMADC produces a binary signal that still has a multi-level attribute. The amplitude of the analog signal is encoded as the width rather than the height of the pulse. The pulse-width modulator is embedded in a Sigma-Delta loop which feeds back the pulse and rejects distortion and quantization noise. The pulse-width trait of the feedback offers several advantages for low distortion and high SNR.

Both ADCs feed into a common decimation structure. The PWMADC, since it runs at a higher sample rate (~ 2 GHz), drives a specialized down-sample-by-8 structure. This down-sampling brings the sample rate to ~ 250 MHz for easier routing to the more flexible decimation structures. The non-PWM ADC, running at 667 MHz drives a specialized down-sample-by-4 structure, reducing the sample rate to 166.5 MHz for the same reason. These specialized fixed down-samplers are then multiplexed into a single down-sample chain which has programmable decimation ratios.

The decimator architecture consists of one fixed-rate decimate-by-8 stage followed by multiple programmable stages of decimation. The total downsampling is given by:

$$K = 8 * [1...8] * [1...8] * [1...8]$$

The notation [1...8] represents an integer from 1 through 8 inclusive. As a result, there are 729 possible combinations of these integers, but only 80 result in unique values for K.

Note that fine control of the sample rate cannot be affected through this architecture. For example as K changes

from 1 to 2, the output sample rate changes from 250 MHz to 128 MHz. As K changes from 3584 to 4096 (its highest value), the output sample rate changes from 558.04 kHz to 488.28 kHz.

Varying sample clock can give much finer control of the sample rate. The intent is to reduce the sample rate to where transmitting over the serial link is feasible. The fine control of the sample clock simply allows this output sample rate to be a multiple of some symbol rate required by the DSP and/or avoid interference with a received RF signal.

The RxDSP module will decimate the high speed 2GHz A/D data down to a multiple of the symbol rate that is low enough to serialize over the digital interface.

The frequency response of the decimation filters is evaluated for different candidate systems and the sampling rates are summarized in Table 1.

Table 1. Sample rate summary for various protocols.

Standard	Nominal Sample Rate (kSa/s)	f_0 (GHz)	K
GSM/EDGE	541.666	1.941333	3584
	1083.3332	1.941333	1792
WCDMA	3840	1.966080	512
WiMax 2.5 MHz	2856	1.827840	640
WiMax 20 MHz	22856	1.828480	80
LTE 10 MHz	7680	1.966080	256
LTE 20 MHz	30720	1.966080	64

In a multi-stage decimator system, the comb filters filter out the unwanted signals around the zeros in the frequency response, so that aliasing can be controlled below a certain level. However, the comb filters do not have sharp transition between the passband and stopband. Therefore, the FIR filter at the final stage of the RxDSP is necessary to clean up the out-of-band components. In particular, when the nominal sample rate does not provide enough over sampling ratio when compared with the bandwidth of the signal, we can not use comb filters only to achieve the decimation.

5. TX CHAIN

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 10 MHz to 6 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 requires 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 oscillator provides automatic filter pole adjustment for the reconstruction filters. Closed loop correction of baseband DC offsets and I/Q phase and 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.

5.1 Digital to Analog Converter

The sigma delta DAC block diagram is shown below. The fully digital portion of the design consists of the three top blocks while the analog portion is the bottom three blocks. The buffer shown at the end of the signal chain is actually the input buffer of the TX baseband reference path in the TX lineup.

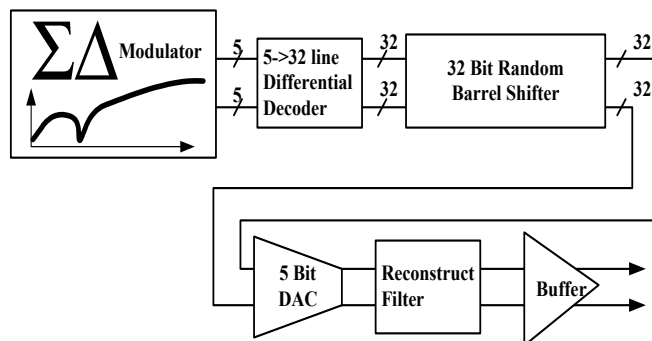


Figure 2. Transmit DAC architecture

By design, the sigma-delta DAC core is so much faster than the modulation, that its speed is not a significant contributor to signal distortion. Rise/settling times on the DAC core are set by the passive reconstruction poles, and further modified at the output by subsequent filtering.

The noise response of the forward path circuits is an important aspect of the analog design. Specifically, the thermal and 1/f noise contribution of the DAC core and buffer must be well below the quantization noise of the complete DAC modulator (ENOB). Thermal and 1/f noise are minimized by using large MOS devices and small resistive loads, respectively. The design tradeoff is that large devices translate into lower bandwidths and more area, while lower load resistors translate into larger load currents.

The number of differential pairs used in the DAC core, and therefore the effective number of bits out of the quantizer, is programmable. Furthermore, the user can programmatically modify the resistor size and current drain of the DAC core in order to take advantage of the tradeoffs discussed earlier for different communication protocols.

One of the most important aspects of any DAC design is device matching. Great care was taken in the design of the SDR DAC to ensure this accuracy. The effective number of bits (ENOB) of any DAC is tied to controlling the inband noise floor and spurious response of the DAC output. This also affects the linearity of its signal processing capabilities which ultimately determines how accurately a digital signal can be represented in its analog form. Three separate techniques were implemented to ensure this accuracy: 1) large devices for both the tail currents and differential pairs were utilized 2) common centroid placement ensured good gradient matching and 3) random digital selection of the tail and diff pair groups was implemented.

6. DIGITAL INTERFACE

The ADC and DAC digital interface module for the RFIC provides communications between the RFIC and the baseband processor, transmit/receive control timing, and receive decimation filters.

The transmit/receive control is preformed in the Bus Interface which contains a serial-to-parallel interface (SPI) control interface that provides communication from the host processor to all internal registers, as well as a sequence manager to automatically update internal registers based off of internal RFIC events.

The digital interface consists of one or two streaming data pipes for the transmit data, as well as one or two streaming data pipes for the receive data. This module also has queuing to hold the data until it can either be sent to the DAC for the transmit case, or sent to the baseband processor for the receive case.

The Rx DSP contains a flexible decimation filter line-up to decimate the approximately 2 GHz 2-bit I/Q ADC input down to the desired sample rate to transfer data over to the baseband IC. This module also contains the master timer, which produces time-stamps at the decimated ADC rate that can trigger RFIC events.

The following table describes some use cases for different systems that the digital interface module could support.

Table 2. Example of system requirements.

System	Transmission Rate	Over-Sample Rate	Data Width	Minimum High Speed Serial Clock
GSM GPRS EDGE	270.833 kBps	2x	24-bits	26 Mbps
		4x	24-bits	52 Mbps
W-CDMA	3.84 Mcps	2x	18-bits	276.5 Mbps
IS-95 IS-2000	1.2288 Mcps	4x	18-bits	177 Mbps
WiMax 2.5 MHz	2.856 MHz		18-bits	102.8 Mbps
WiMax 20 MHz	22.856 MHz		12-bits	548.5 Mbps
LTE 5 MHz	7.68 MHz		16-bits	245.8 Mbps
LTE 20 MHz	30.72 MHz		12-bits	737.3 Mbps

6.1 Digital Interface Signals

Transmit and receive samples are transferred between the baseband processor and the RFIC via the digIF serial high speed digital interface. The digital interface uses a low voltage differential signaling (LVDS) interface and an interface clock to transfer the data between the two ICs. The interface clock – also LVDS – is sent from the RFIC to the Baseband IC to establish the transfer rate of the serial interface. Internally both ICs will multiply the interface clock by 8 to derive the serial clock rate that will be used to transmit and receive the serial data. When data is transferred over the serial pipe, it will be randomized to reduce the number of spurs that are created by this high speed serial link.

The connections between the RFIC and the Baseband IC and are shown in Figure 3. The extra txData1P/N and rxData1P/N signals are optionally enabled when higher throughput is required.

The baseband processor determines when to enable the serial pipe using the sysClkEn signal. When the sysClkEn signal is low, the RFIC will disable the sysClk output, synchronously reset the digIf block, clear all the FIFO's, and will set the LVDS signals in a low power mode.




Figure 3. Streaming data pipe signals

When the sysClkEn signal is high, the RFIC drives the sysClk on the board, and enables the LVDS outputs. The txData0/1P/N is the LVDS signal pair that carries the randomized serial transmit data from the baseband processor to the RFIC. The rxData0/1P/N is the LVDS signal pair that carries the randomized serial receive data from the RFIC to the baseband processor.

7. SYSTEM DEVELOPMENT KIT

The RFIC has been tested using a system development kit comprised of the RFIC and it's supporting circuitry, voltage regulators, and a Virtex-5 FPGA. The I/O range from high frequency RF connectors to low-speed SPI interface to Gigabit Ethernet ports.

The purpose of the FPGA is to deliver I-Q samples to the RFIC (which passes them to its internal DAC), and also to receive I/Q samples from the RFIC (which passes them from its internal ADC's). The FPGA is charged with:

1. Handling the transfer of high-speed digital information via a proprietary protocol that is similar to the digRF standard protocol.
2. Performing minor DSP operations like DC offset correction, amplitude imbalance correction, and phase imbalance correction in both the TX and RX paths.
3. Sending and receiving samples from an external source to the SDK board. There are two backend options that the SDK board was specifically designed to accommodate: a UDP/IP interface to a PC via dual Gigabit Ethernet ports and a proprietary interface to specific Rohde and Schwartz test equipment that is capable of extensive generation and analysis of complex signals.

Each of these tasks occurs in its own clock domain, with FIFO boundaries between them as shown in Figure 4.

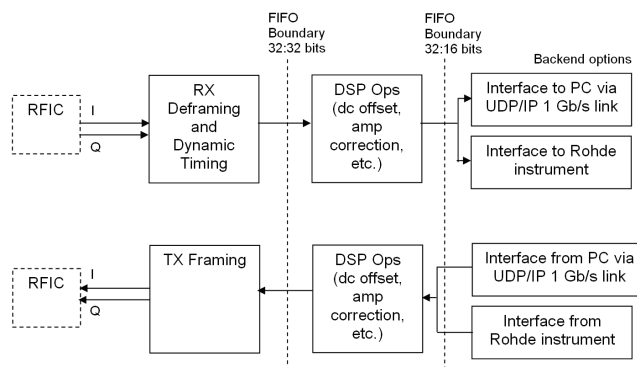


Figure 4. FPGA and RFIC interface.

8. CONCLUSION

The SDR transceiver described here was implemented in 90 nm technology. The chip measures 5 x 5.4 mm and is housed in a 10 x 10 mm dual-row MLF package.

The divergent requirements of the many standards this RFIC can process, makes a comprehensive performance report beyond the scope of this paper. While certain parameters may be optimized based on a particular standard, measured data based on typical conditions are listed in Table 3.

Table 3. RFIC performance summary

Freq. Range	
Internal DDS	10 MHz – 4 GHz
External LO	10 MHz – 6 GHz
Baseband Bandwidth	10 kHz – 10 MHz
Rx NF (with on-chip LNA)	
100 – 500 MHz	6 dB
500 – 1000 MHz	7 dB
1 GHz – 2 GHz	8 dB
2 GHz – 4 GHz	9 dB
Rx Gain	
LNA/Mixer	20 dB max
Baseband	62 dB max
Rx IIP2	+60 dBm
Rx IIP3	-6 dBm
Rx Current Drain	40 mA
Tx Output Power	
1 GHz	+7.5 dBm
2 GHz	+3.6 dBm
3 GHz	-3.2 dBm
4 GHz	-5.9 dBm
Tx Sideband Suppression	35 dBc
Tx Current Drain	40 - 90 mA
EVM $\pi/4$ DPQSK 3.5 MS/s	1% @ 800 MHz
LO Phase Noise	-113 dBc/Hz @ 25 KHz
LO Frequency Resolution	15 Hz
LO Current Drain per DDS	80 mA

9. REFERENCES

- [1] G. Cafaro et al., "A 100 MHz – 2.5 GHz direct conversion CMOS transceiver for SDR applications," IEEE RFIC Symposium, June 2007, pp. 189-192.
- [2] G. Cafaro et al., "A 100 MHz – 2.5 GHz CMOS transceiver in an experimental cognitive radio system," SDR Conference, November 2007
- [3] A. Bourdoux, J. Craninckx, A. Dejonghe, L. Van der Perre, "Receiver Architectures for Software-defined Radios in Mobile Terminals: the Path to Cognitive Radios," IEEE Radio and Wireless Symposium, Jan 2007, pp. 535 – 538.
- [4] A. Geis, et al., "A Compact low power SDR receiver with 0.5 - 20 MHz baseband sampled filter," IEEE Radio Frequency Integrated Circuits Symposium, June 2009 pp. 285-288.
- [5] V. Giannini, et al., "A 2mm² 0.1-to-5GHz receiver in 45nm digital CMOS," IEEE International Solid-State Circuits Conference, Feb 2009, pp. 408-409.
- [6] J. Craninckx, et al., "A fully reconfigurable Software-Defined Radios Transceiver in 0.13um CMOS," IEEE International Solid-State Circuit Conference, Feb 2007, pp. 346-607.
- [7] D. Kaczman, et al., "A Single-Chip 10 Band WCDMA/HSDPA 4-Band GSM/EDGE Saw-less CMOS Receiver with DigRF 3G Interface and +90 dBm IIP2," IEEE Journal Solid State Circuits, volume 44 issue 3, March 2009 pp. 718-739.
- [8] D. Miyashita, et al., "A Low-IF Single-chip Bluetooth EDR Transmitter with Digital I/Q Mismatch Trimming Circuit," 2005 Symposium on VLSI Circuits, June 2005, pp. 298-301.
- [9] J. Juan, R. Stengel, F. Martin, D. Bockelman, "Cascaded delay locked loop circuit," US Patent 6891420, December 26, 2006.
- [10] D. Bockelman, J. Juan, "Time interpolating direct digital synthesizer," US Patent 6353649, March 5, 2002.
- [11] F. Martin, R. Stengel, J. Juan, "Method and apparatus for digital frequency synthesis," US Patent 7154978, May 10, 2005.
- [12] T. Gradishar, R. Stengel, "Method and apparatus for noise shaping in direct digital synthesis circuits," US Patent 7143125, November 28, 2006.
- [13] E. Bautista, B. Bastani, J. Heck, "A high IIP2 downconversion mixer using dynamic matching," IEEE Journal of Solid-State Circuits, vol. 35, December 2000, pp. 1934-1941.
- [14] E. Bautista, B. Bastani, J. Heck, "Method and apparatus providing improved intermodulation distortion protection," US Patent 6125272, September 26, 2000.

Copyright Transfer Agreement: The following Copyright Transfer Agreement must be included on the cover sheet for the paper (either email or fax)—not on the paper itself.

"The authors represent that the work is original and they are the author or authors of the work, except for material quoted and referenced as text passages. Authors acknowledge that they are willing to transfer the copyright of the abstract and the completed paper to the SDR Forum for purposes of publication in the SDR Forum Conference Proceedings, on associated CD ROMS, on SDR Forum Web pages, and compilations and derivative works related to this conference, should the paper be accepted for the conference. Authors are permitted to reproduce their work, and to reuse material in whole or in part from their work; for derivative works, however, such authors may not grant third party requests for reprints or republishing."

Government employees whose work is not subject to copyright should so certify. For work performed under a U.S. Government contract, the U.S. Government has royalty-free permission to reproduce the author's work for official U.S. Government purposes.