PART I
Transceiver Concepts and Design

COPYRIGHTED MATERIAL

1 Software-Defined Radio Front Ends

JAN CRANINCKX
IMEC, Leuven, Belgium

1.1 INTRODUCTION

The ultimate dream of every software-defined radio (SDR) front-end architect is to deliver a radio-frequency (RF) transceiver that can be reconfigured into every imaginable operating mode, in order to comply with the requirements of all existing and even upcoming communication standards. These include a large range of modes for cellular (2G–2.5G–3G and further), WLAN (802.11a/b/g/n), WPAN (Bluetooth, Zigbee, etc.), broadcasting (DAB, DVB, DMB, etc.), and positioning (GPS, Galileo) functionalities. Obviously, each of them has different center frequency, channel bandwidth, noise levels, interference requirements, transmit spectral mask, and so on. As a consequence, the performances of all building blocks in the transceiver must be reconfigurable over an extremely wide range, requiring ultimate creativity from the SDR designer.

Reconfigurability is a requirement for SDR functionality, but often one forgets that it can also be an enabler for low power consumption. Indeed, once flexibility is built into a transceiver, it can be used to adapt the performance of a radio to the actual circumstances instead of those implied by the worst-case situation of the standard. Since linearity, filtering, noise, bandwidth, and so on, can be traded for power consumption in the SDR, a smart controller is able to adapt the radio at runtime to the actual performance required, and hence can reduce the average power consumption of the SDR.

In this chapter, several important innovations and concepts are presented that bring this ultimate dream closer to reality. These include circuits for wideband local oscillator (LO) synthesis, multifunctional receiver and transmitter blocks, and novel ADC (analog-to-digital converter) implementations. The result of all this is integrated in the world's first SDR transceiver covering the frequency range from 174 MHz to 6 GHz, implemented in a 1.2-V 0.13-μm CMOS technology.

Multi-Mode/Multi-Band RF Transceivers for Wireless Communications, Edited by Gernot Hueber and Robert Bogdan Staszewski
Copyright © 2011 John Wiley & Sons, Inc.

1.2 SYSTEM-LEVEL CONSIDERATIONS

A first choice to be made is the radio architecture to be used. In past decades, lots of studies and examples have been presented on heterodyne, homodyne, low-IF (intermediate frequency), wideband-IF, and other architectures, all having certain benefits and problems for a certain application. Which one to choose? In view of SDR, this question perhaps becomes a little easier to answer. Indeed, when the characteristics of all possible standards are taken into account, not a single intermediate frequency can be found that suits them all. And having multiple IFs and the associated (external) filtering stages increases the hardware cost of the SDR, which cannot be tolerated. So direct-conversion architectures are the right choice for the job. All of the well-known problems, such as dc offsets, I/Q mismatch, $1/f$ noise, and power amplifier (PA) pulling, that have limited the proliferation of zero-IF CMOS radios into mainstream products have been better understood in recent years, and it will enable the design of a low-cost front end.

A schematic vision of what the final SDR will look like is represented in Fig. 1.1. For a low cost in a large-volume consumer market, the active transceiver core is implemented in a plain CMOS technology. It includes a fully reconfigurable direct-conversion receiver, transmitter, and two synthesizers [for frequency-domain duplex (FDD) operation]. The functions that cannot be implemented in CMOS are included on the package substrate. These are related primarily to the interface between the active core and the antenna. They must provide high-Q bandpass filtering or even duplexing, impedance-matching circuits, and power amplification. In the remainder of the chapter we focus primarily on the transceiver implementation.

The hard works starts with determining performance specifications for each block in the chain. The total budget for gain, noise, linearity, and so on, must be divided

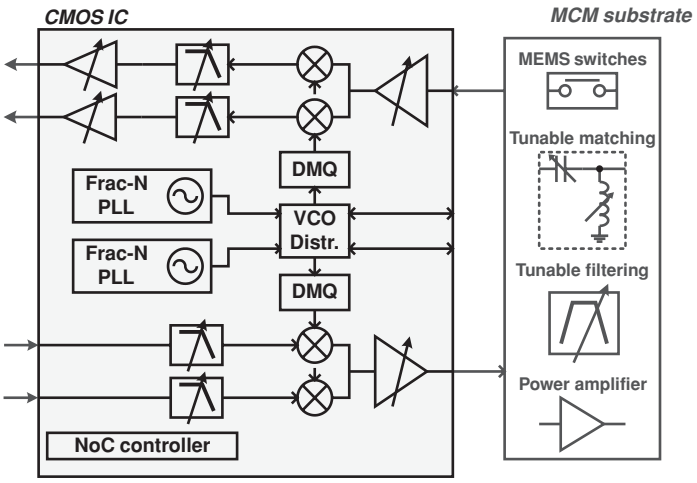

FIGURE 1.1 Conceptual view of an SDR transceiver front end.

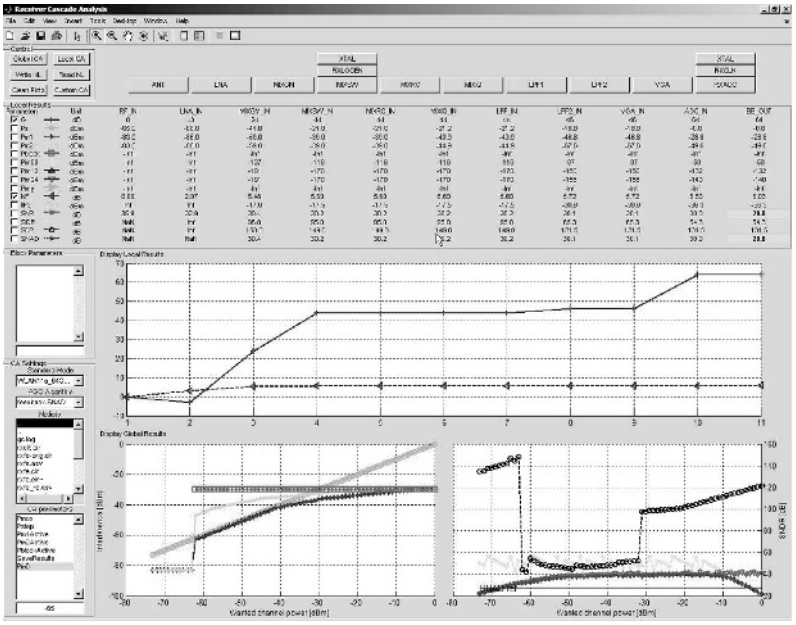

FIGURE 1.2 System-level analysis tool.

over all blocks, ensuring that all possible test cases are covered, and this must be done for every standard. Having very flexible building blocks helps a great deal, of course, but making a smart system analysis at this point is crucial to obtaining an optimal SDR solution.

A custom MATLAB tool has been developed to do this exercise [1]. It takes in a netlist that describes all building blocks, with the performance characteristics and gain ranges, and simulates on a behavioral level the complete chain for a list of different test cases. Figure 1.2 shows a screenshot. The performance under all circumstances can thus be evaluated, and the building block performance can be tuned to fulfill all requirements. Gain ranges and signal filtering must be set such that the signal levels are an optimal trade-off between noise and distortion. Although being a difficult exercise, the analysis can show that with the built-in flexibility, a software-defined radio can achieve state-of-the-art performance very close to that of dedicated single-mode solutions. In the next sections we go deeper into the design of some crucial building blocks.

1.3 WIDEBAND LO SYNTHESIS

To generate all required LO signals in the range 0.1 to 6 GHz, several frequency-generation techniques have been proposed to relax the tuning range specifications of a voltage-controlled oscillator (VCO). They use division, mixing, multiplication,

or a combination of these [2]. However, to make these systems efficient in terms of phase noise and power consumption, the VCO tuning range still has to be maximized. In the following section we discuss the design of such a wideband VCO, and the architecture required to generate all LO signals is discussed in Section 1.3.2. The target frequency band of the VCO is around 4 GHz, so that it does not coincide with any of the major RF frequency bands used. The actual LO frequency will be obtained by further division and mixing. Since the VCO frequency differs from the RF frequency, most direct-conversion problems will be relaxed or avoided.

1.3.1 3 to 5-GHz Voltage-Controlled Oscillator

To reach the stringent phase noise specifications for today's mobile communication systems, most RF transceiver integrated circuits (ICs) use LC-VCOs. Frequency tuning of LC VCOs is commonly done by changing the capacitance value of the resonant tank using varactors and/or an array of switched capacitors [3]. Switched or controlled inductor designs have been reported [4], but it remains difficult to cover the desired wideband continuously and to limit the deterioration of the phase noise performance caused by the insertion of these switches.

Instead of using a single large varactor to tune the frequency, a mixed discrete/continuous tuning scheme is usually chosen [3]. A small varactor is used for fine continuous tuning, and larger steps are realized by digitally switching capacitors in and out of the resonant tank. This has two advantages: The VCO gain is lower, allowing easier phase-locked loop (PLL) design, and digitally switched varactors have a higher ratio between the capacitance in the on-state (C_{on}) and the capacitance in the off-state (C_{off}). A higher $C_{\text{on}}/C_{\text{off}}$ ratio allows a larger VCO frequency tuning range. However, as the tuning range of a VCO is increased and exceeds the typical 20% range obtained in many designs, new problems and trade-offs appear that need a solution. In this design we have tackled the two main problems encountered in wideband LC-VCOs [5]. First, the negative resistance required to maintain oscillation varies a lot over the frequency range, leading to significant overhead when a fixed active core is used. Second, the large variation of the VCO gain (K_{VCO}) across the entire tuning range creates problems for optimal and stable PLL design. Solutions are proposed for both problems.

1.3.1.1 Tank Loss Variations In the target frequency range (< 5 GHz), the losses in the oscillator tank are usually dominated by the inductor. It can be modeled by an inductor series resistance R_S, which in this simple example we consider to be frequency independent. This simplification is, of course, not completely valid, since extra losses due to the skin effect, for examples will increase the resistance at higher frequencies, but that does not change the general conclusion we will make.

The negative resistance needed to compensate for the inductor losses is given by $G_m = R_S(\omega C)^2$, where C is the total tank capacitance and ω is the oscillation frequency, which is, of course, given by the simple equation $\omega = 1/\sqrt{LC}$, with L the inductance value [6]. If we want the oscillation frequency to change by a factor of 2, for example, the total capacitance of the resonant tank has to be changed by a factor of 4, and hence the required negative resistance must also change by a factor of

4. The transconductance required for the active core is four times higher at the lower end of the frequency tuning range than at the higher end.

Recent phase noise theory based on the impulse sensitivity function (ISF) theory of phase noise, together with a linear-time-variant circuit analysis [7], has shown that not only the small-signal transconductance must be considered. The VCO phase noise depends on the large-signal oscillation amplitude, and that is proportional to the bias current and the parallel tank resistance, which varies with frequency.

In a traditional design, the active core will be designed for the toughest case (i.e., for the lowest frequency). For the highest frequency, the active core is largely overdimensioned, which is obviously a waste of power. Changing the VCO operating point with frequency is beneficial. Another argument to take into account is the tuning range achievable. The key to a wideband VCO is, of course, to have a tank capacitance that consists as much as possible of varactors and as little as possible of parasitics. The smaller the active transistors, the better. Here it is obvious that if we could eliminate some of the parasitics at high frequencies, the tuning range could be extended considerably.

The basic idea behind the solution presented here is thus not only to scale the biasing current of the active core, but simultaneously to change the size of the transistors as well, to keep parasitics at a minimum, which is beneficial for both the phase noise performance and the tuning range achievable. Therefore, the active core will be constructed from an array of core units, which can be turned on or off when necessary. In each of these core units, switches must be added to turn the active transistors on or off. The position and size of those switches has to be considered carefully, to avoid degrading the oscillator phase noise performance as well as to ensure that additional parasitic capacitances are small. The circuit diagram is depicted in Fig. 1.3.

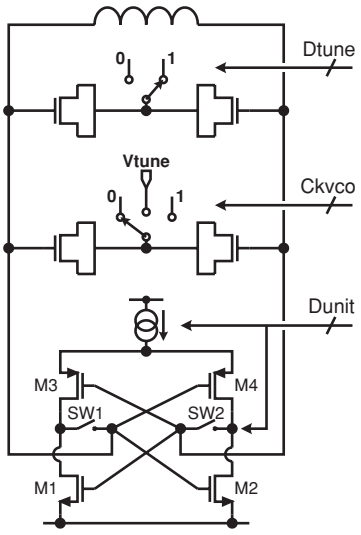

FIGURE 1.3 Wideband VCO architecture.

As is clear from the analysis above, it is of utmost importance that together with the negative resistance, the parasitic capacitance is removed from the oscillator tank to ensure a large tuning range. In the on-state the switch is closed, and the parasitic capacitance is then determined by the drain–gate and source–gate capacitance of the switch, plus the drain and gate capacitance of the active transistors. This is obviously larger than the parasitics of a simple negative resistance because of the added switch parasitics, but that is not an issue. Indeed, the core units are activated only when the oscillation frequency is lowered, and hence a larger capacitance is tolerated.

In the off-state, the active NMOS transistors (M1, M2) keep their gate connected to the LC tank and their parasitic remains. But the active PMOS transistors (M3, M4) are turned off by the positive gate–source voltage, and their gate capacitance drops considerably. Finally, the switch transistors (SW1, SW2) also turn off and only the drain parasitics stay attached to the tank, which results in a large drop is capacitance. This information is, of course, used in the sizing the various transistors. The NMOS is made small ($W = 2.6$ μm per unit) and the PMOS is approximately three times larger ($W = 8.5$ μm). The greatest width is given to the transistor that has the largest on/off ratio, so the switch size is set to $W = 18$ μm. With this structure a $C_{\text{on}}/C_{\text{off}}$ ratio close to 3 is obtained, without any significant contribution of the switches' series resistance to the overall phase noise. So, in fact, we have been able to use the negative resistance core as a varactor. For high oscillation frequency, the capacitance is low and there is no negative resistance. For lower frequencies, more and more core units are gradually activated, and the total bias current increases to keep the oscillation amplitude steady and the parasitic capacitance increases, helping the "normal" varactors in their goal to increase the total tank capacitance.

1.3.1.2 Sensitivity Variations The second problem solved in the design presented is the variation in VCO sensitivity for wide-tuning-range VCOs. A change in the control voltage V_{tune} results in a change ΔC in the analog varactor capacitance C_{var}. This causes a change in frequency Δf. The size of this frequency change depends on the relative weight of the analog capacitance change with respect to the total tank capacitance (which consists for a large part of digitally switched varactors):

$$f = \frac{1}{2\pi\sqrt{LC}} \rightarrow \frac{\Delta f}{\Delta C} = \frac{-1}{4\pi C\sqrt{LC}} \tag{1.1}$$

If we go back to the example of the VCO with a frequency ratio of 2, we have seen that the tank capacitance has to change by a factor of 4. As can be seen from (1.1), the VCO frequency sensitivity will then change by a factor $4\sqrt{4} = 8$. In this example the nonlinearity of the CV curve of the varactor has been neglected, but typically the varactor is used in the middle of its tuning range, where this curve is rather linear. Such a large change in VCO gain presents serious problems for the design of the PLL in which it will be incorporated. It prevents keeping the PLL bandwidth constant and hence endangers the loop stability and an optimal phase noise performance.

The solution proposed here is to make the varactor size changeable. Instead of making one big analog varactor, a number of unit analog varactors are used. These

varactors can be controlled in two ways. Some units are used for analog continuous tuning, and their control node is connected to oscillator tuning voltage V_{tune}. The other units are used for fine-grain discrete tuning, and their control node is connected either to the power supply or to ground. At the lowest frequency the sensitivity is low, so a large analog varactor is needed. Most of the unit varactors will be connected to the analog control voltage. At high frequencies the sensitivity is relatively high and only a small analog varactor is needed. The other units can then be used as a discrete switched varactor, giving extra-fine discrete tuning curves.

1.3.1.3 Circuit Implementation Figure 1.3 shows a simplified view of the complete VCO architecture implemented. The inductor value was chosen small (0.75 nH) and is optimized for a wide tuning range. It has a symmetrical octagonal shape [8] and is implemented in the top metal layer, which has a thickness of 2 μm. The next metal level is used for the underpass connections only. The typical series resistance of the inductor is about 1 Ω.

The coarse frequency tuning is done with an array of 31 equal-sized varactors, controlled by the 5-bit control word D_{tune}. In combination with those varactors, active core units (control word D_{unit}) add the necessary negative resistance and also add some extra capacitance when the frequency is lowered. A total of 31 switched core units are employed, in parallel with a fixed negative resistance that has the size of about 10 units. This allows controlling the negative resistance generated by the VCO core over a factor of 4, as was required for a factor of 2 tuning of the oscillation frequency. Correspondingly, the total current of the active core will vary between 2.1 and 8.5 mA, whereas the bias circuit consumes 0.55 mA.

An analog varactor consists of 15 small units and is controlled by C_{kvco}. That digital code actually consists of two control words. Four bits are used to set the number of varactors that must be connected to the analog control voltage V_{tune}. Fifteen other bits are used to set the varactor control to power (1) or ground (0) in case it is not used for analog control. That creates a large set of extra-fine tuning curves that cover the range between two adjacent coarse-tuning settings.

As there are many control bits to set the proper frequency and gain of the VCO, and as the required settings of those bits are partially dependent on process, temperature, and voltage variations, a calibration sequence is needed to identify the correct setting for each desired center frequency. At power-up time, before actual operation, both the frequency and the frequency sensitivity of the VCO must be measured and stored in a look-up table for correct operation. The setting chosen must be such that the VCO gain is kept proportional to the frequency over the entire frequency range, as this keeps the PLL bandwidth constant [5].

Figure 1.4(a) shows a selected set of the measured frequency response of the VCO. Only some of the 32 coarse frequency steps are shown for most of the frequency range, showing a total tuning range from 3.14 to 5.2 GHz, equivalent to 49%. In the upper frequency range, a detail of the fine-tuning steps is also shown. At this high frequency, only 2 of the 15 small varactors are controlled by the analog tuning voltage. The others can be set digitally to 0 or 1, resulting in an extra set of 14 fine-tuning curves. The plot shows that there is sufficient overlap between consecutive curves. In

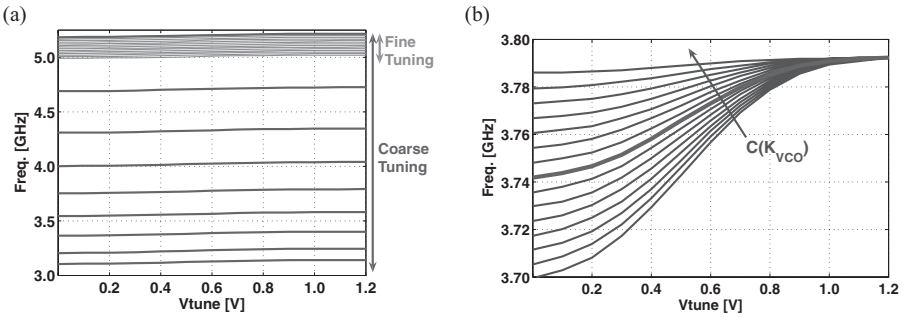

FIGURE 1.4 VCO measurement results: (a) set of coarse and fine tuning curves selected; (b) VCO gain settings.

the lower-frequency range (not shown), the coarse-tuning curves are closer together because the total capacitance in the tank is higher. But also, more analog varactor units are connected to the tuning voltage, leaving fewer analog units that are digitally controlled and hence fewer fine-tuning curves. Eventually, the entire frequency band can be covered continuously with the desired slope for the oscillator sensitivity.

The flexibility of the VCO gain is shown in Fig. 1.4(b). For a fixed coarse frequency setting, the number of analog varactor units is changed, giving different slopes of the frequency curves. The varactors that are not connected to the analog tuning voltage are biased at the power supply; hence all curves overlap at $V_{\text{tune}} = 1.2$ V. Clearly visible in the graph is the limited linear range of the MOS varactors used in the design. In the PLL the VCO settings are controlled so as only to be used in the most linear range of the tuning voltage, between 0.4 and 0.8 V, where K_{VCO} is almost constant.

Measured phase noise at an offset of 1 MHz ranges from -115 to -119 dBc/Hz for the upper and lower frequencies, respectively. This variation can be explained perfectly by the difference in $(f_0/\Delta f)^2$ [6], indicating indeed that the design is still limited by the limited Q of the inductor and that the use of the switched active core allows us to keep the current consumption optimal over the entire frequency range. The closed-loop integrated phase noise of the complete PLL is typically -36 dBc.

These measurements show that the VCO achieves continuous coverage over a very wide frequency range, with a fully controllable K_{VCO}, resulting in a stable and optimal PLL design for the entire tuning voltage range used.

1.3.2 0.1 to 6-GHz Quadrature Generation

As a result of the wideband VCO, the problem of LO carrier generation can be solved in a system that is not too complex and hence does not carry a large power penalty. The block diagram of the divide/multiply and quadrature (DMQ) is presented in Fig. 1.5. The DMQ contains several divide-by-2 blocks. They generate I and Q phases down to a division factor of 32. Each divider consists of two dynamic simplified flip-flops

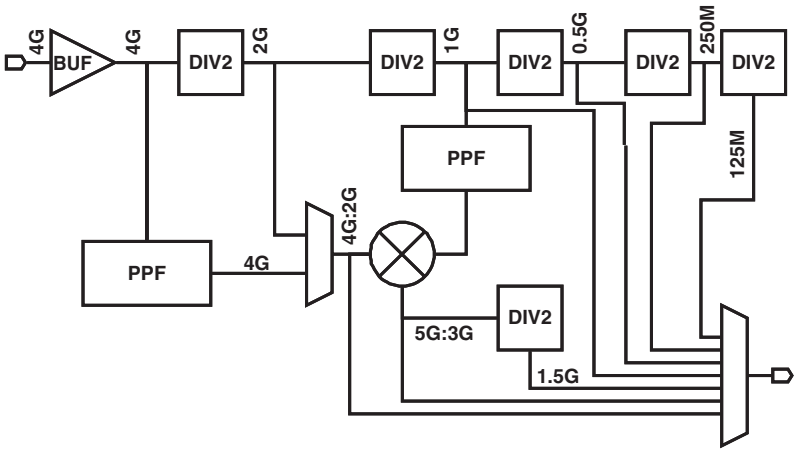

FIGURE 1.5 Block diagram of the DMQ circuit.

in feedback. The rail-to-rail operation of the latches ensures a minimal addition of phase noise, very important in cellular standards.

The DMQ further employs a single-side band (SSB) mixer. For a VCO at its 4-GHz center frequency, the mixer combines 4GHz or 2 GHz with 1 GHz to obtain 5 GHz or 3 GHz, respectively. The 2- and 1-GHz components are obtained by division of the VCO frequency. The 4-GHz quadrature phases needed for SSB mixer operation in the 5-GHz mode are generated through a polyphase filter (PPF1). This is implemented as a three-stage polyphase filter, with notches at negative frequencies of 3, 4, and 5 GHz. The circuit diagram of the SSB mixer is presented in Fig. 1.6. As both base frequencies used in the SSB mixing are square waves, they contain all odd harmonics. These are also combined in the mixer and will generate unwanted frequency components. To limit these, the 1-GHz (F1) component is first linearized by filtering out the third harmonic of the 1-GHz square wave. This negative frequency is attenuated by 40 dB with a two-stage polyphase filter with notches at -2 and -4 GHz. The output of this polyphase filter is a current whose four quadrature phases are injected directly into the SSB mixer. With the VCO at its center frequency, the mixer's switches are driven by either a 2- or 4-GHz rail-to-rail square wave. Cascode transistors below the switches provide a low impedant input for the linearized 1-GHz current. The bias current is provided by current sources at the bottom. Note that both current sources and cascode transistors are common to both the I and Q paths. The output of the mixer (F3) is amplified up to full rail swing with a differential pair followed by a string of inverters, of which the first is biased around its threshold voltage.

The SDR's LO frequency can be selected by a multiplexer integrated in the DMQ. This function is obtained by powering down the unused blocks and placing their outputs in a high-impedance state. In this way, no extra circuits are placed in the signal path.

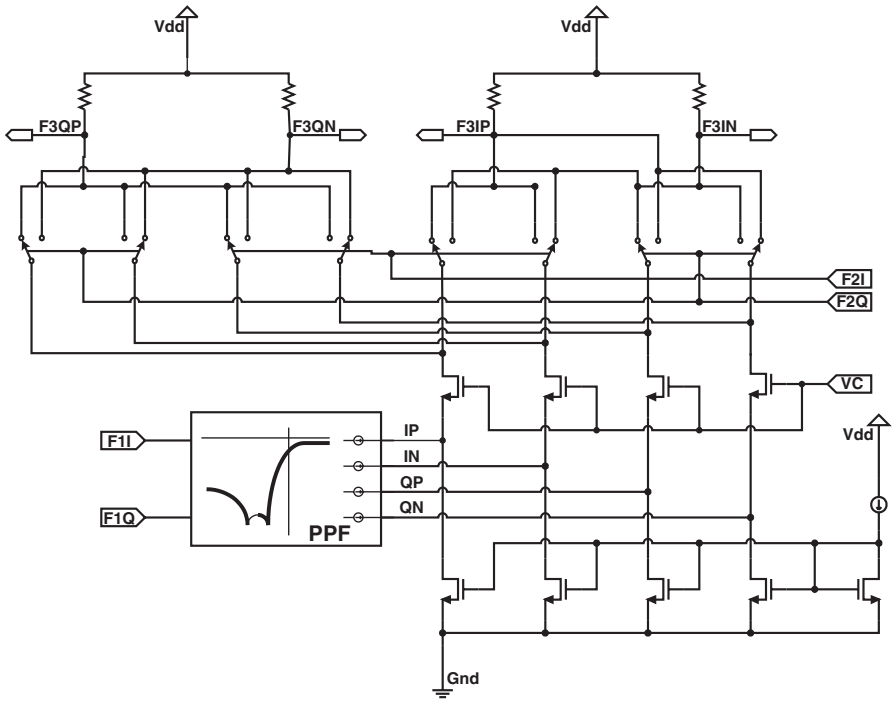

FIGURE 1.6 SSB mixer circuit diagram.

1.4 RECEIVER BUILDING BLOCKS

A key aspect for the receiver RF part is its interference robustness. The blocking requirements for simultaneous multi-mode operation imply the need for tunable narrowband circuits at the antenna interface. Either this function can be provided by a multi-band filtering block [9], in which case the receiver's input can be a wideband low-noise amplifier, or part of this burden can be taken up in the low-noise amplifier (LNA) design, as shown in the following section.

1.4.1 MEMS-Enabled Dual-Band Low-Noise Amplifier

In this first example the option of using microelectromechanical systems (MEMSs) switches to build a low-loss reconfigurable antenna filter section on a thin-film substrate is explored. This is especially relevant when simultaneously considering the design of the LNA, whose active CMOS part must be co-designed with the MEMS switch and the passive off-chip matching. The circuit schematic of Fig. 1.7 shows how multi-band operation is achieved independent of the inductive emitter degeneration [10]. A single-pole dual-throw (SPDT) MEMS switch is used to connect the LNA to either its 1.8-GHz matching circuit and antenna filter, or to its 5-GHz

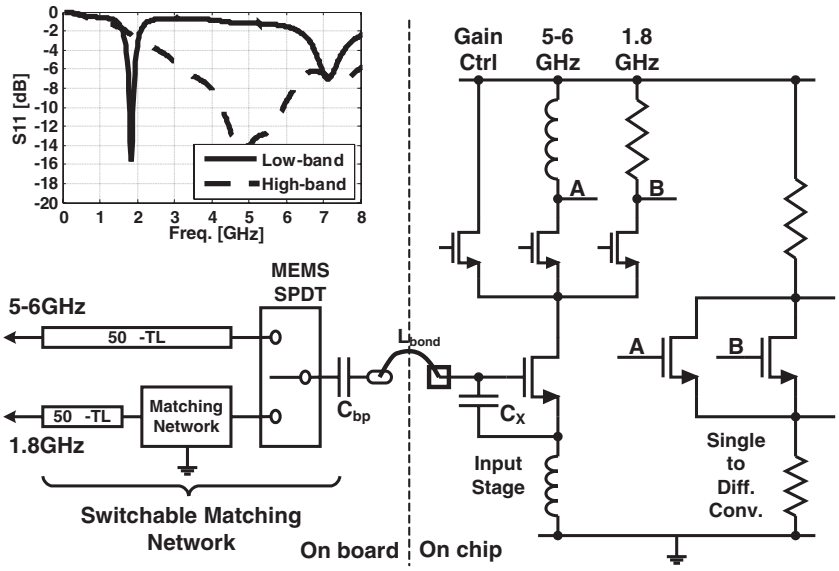

FIGURE 1.7 CMOS/MEMS co-designed dual-band LNA circuit schematic and input-matching measurement (inset).

section, which uses just the bonding wire for input matching. To prove this concept with a commercial component, a packaged MEMS switch [11] was mounted on a printed-circuit board (PCB) together with the CMOS die. Performance is affected only slightly, since the loss of the switch, including its package, was measured to be only 0.2 dB. This device has, of course, limited the validation of the proposed MEMS-enabled SDR receiver to only two bands. A mature technology that integrates MEMS switches in an MCM technology [12] will make it feasible to build more complex structures covering a broad range of frequency bands.

Optimal implementation of the switchable narrowband impedance matching at the LNA's input has been obtained by designing the on-chip part of the LNA such that no dedicated on-board matching components are needed in the 5 to 6-GHz band except for a simple series dc decoupling capacitor C_{bp}. The MEMS switch can be regarded approximately as a short 50-Ω transmission line, so putting the MEMS switch before the LNA chip will not change the chip's input return loss drastically. For the 1.8-GHz band, a simple matching network made up of one or two passive components can fulfill the matching requirement. The passband is quite narrow, making a simple passive matching network feasible. These components are less lossy at these low frequencies, and the receiver itself is also less noisy at lower frequencies, which assures a good overall noise figure (NF).

To maintain good matching conditions in the full implementation, the chip, including pad parasitics and ESD devices, has been designed in combination with the bondwire inductance, on-board components, and board parasitics. The RF bondpad is modeled by a 65-fF parasitic capacitance in series with a 50-Ω resistance. The ESD

diodes are sized $2 \times 0.6\,\mu m \times 12\,\mu m$ each and have in total a 60-fF parasitic capacitance. Each bondwire is modeled by a 1.3-nH inductance. PCB traces are modeled as small transmission lines when needed. Surface-mounted components on the board are characterized carefully with dedicated separate deembedding structures.

An extra on-chip capacitor C_X connecting the gate and source of the input transistor reduces the gate inductance needed for the input matching to a value of 1.3 nH at 5 GHz, which is conveniently implemented as a bonding wire. Otherwise, this inductance can be unrealistically large and the Q-factor of the input resonance network of the LNA would be too high to cover a 1-GHz band. Thanks to the 300-fF capacitance, the gate inductance can be implemented with a single bonding wire.

Internally, the LNA has two separate outputs to cover the required frequency range. A resistively loaded output is, of course, small in area and flexible in terms of wide bandwidth but can only provide enough gain at frequencies up to 2.5 GHz, due to the limitation of the 0.13-μm CMOS technology used. Therefore, a second output is added for the 5 to 6-GHz band with an LC-tuned load, and the selection of either one of those is done by proper biasing of the cascode transistors. A resistor in parallel with this inductor lowers its quality factor and hence increases its bandwidth in order to cover the 1-GHz bandwidth.

Gain switching is achieved by the current steering technique, when the third common-gate transistor is activated, which bypasses a certain fraction of the signal current to the power supply so as to reduce the gain. Finally, both outputs pass through a multiplexing single-to-differential converter.

The input stage is biased at 5.8 mA, and another 3.6 mA is used in the second stage, for an overall gain of 24 dB. As indicated in Fig. 1.7, S_{11} input matching better than -10 dB is achieved in both bands. The simulated LNA NF is around 2 dB, while the IIP3 value is -5 dBm in the low band and 3 dBm in the high band.

1.4.2 Wideband Low-Noise Amplifiers

Another option to demonstrate the SDR concept is to rely completely on the passives in the antenna interface for RF interference and blocking filtering. This makes the realization of the concept easier, as commercially available (multi-band) filtering blocks can be used in the implementation. Wideband low-noise amplifiers must now be used that cover an RF frequency range as large as possible for optimal flexibility, but on the other hand must still achieve state-of-the-art performance with respect to narrowband LNAs.

Covering the full 100 MHz to 6 GHz frequency range is challenging since achieving a low NF at hundreds of megahertz requires large transistors with low $1/f$ noise, while moving toward carrier frequencies of a few gigahertz requires fast transistors. Recently, several 90-nm wideband inductor-less feedback LNAs have been reported [13]. However, none of them achieved the performance targeted by a link budget analysis for an SDR-LNA below 500 MHz (for a 1.2-V power supply). LC-matched common-source (CS) LNAs typically cover a bandwidth from 3 to 10 GHz [14]. Extending the bandwidth down to 100 MHz would require prohibitively large inductors and thus chip area.

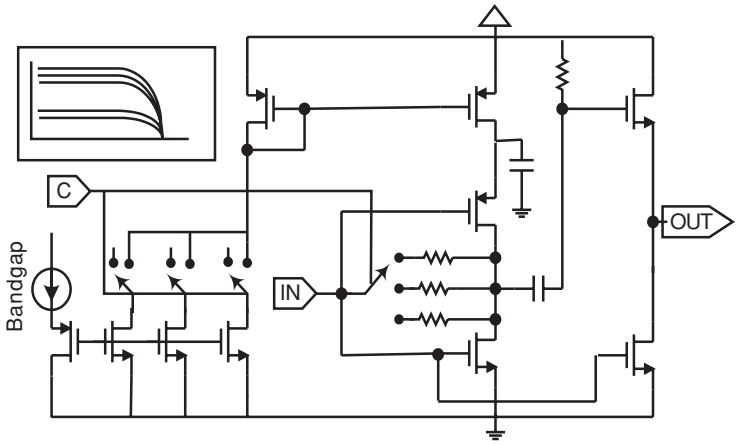

FIGURE 1.8 Schematic of a wideband resistive feedback LNA.

In the SDR front end presented, two LNAs are combined to cover the entire frequency range: An inductor-less feedback LNA with a small form factor (Fig. 1.8) covers frequencies from 100 MHz to 2.5 GHz, and a CS LC-matched LNA (Fig. 1.9) covers frequencies from 2.5 to 6 GHz. Only one LNA is powered at a time, to save power and provide filtering over half of the bandwidth.

The LNA with resistive feedback is based on the noise-canceling topology presented in [15]. It employs resistive feedback for wideband matching and noise canceling for low NF over a wide band. The resistive feedback design in general has lower gain and a higher noise figure than these of inductively matched narrowband designs, but it offers large savings in area. Moreover, the noise-canceling approach can reduce the NF, although a compromise must be taken with the linearity performance.

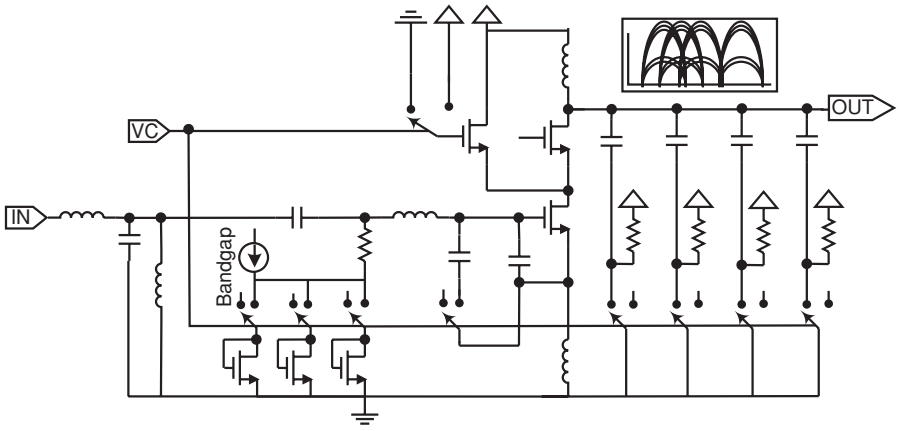

FIGURE 1.9 Schematic of a wideband LC-matched LNA.

The simplified circuit topology of the low-band LNA is shown in Fig. 1.8. For power efficiency, both NMOS and PMOS devices are used in the matching/amplifying stage. A digitally controlled bank of resistors allows us to switch from high- to intermediate- and low-gain modes. The biasing is done with a 3-bit programmable current source. This allows us to vary the gain in small steps around the different gain modes and to decreasing the power by half when switching from high- to low-gain mode. Taking into account the required gain and linearity, the maximum bandwidth achievable at the moment in the 0.13-μm CMOS technology is limited. More advanced technology nodes should overcome this issue in the future.

At a maximum gain of 22 dB (including the single-ended to differential conversion stage), typical simulation results achieve an NF of 2 dB and an IIP3 of −10 dBm at a power consumption of 12 mW. At reduced gain (10 dB), the linearity improves to +3 dBm while the power consumption decreases to 8 mW.

In a high-band LNA (see Fig. 1.9), broad input-matching bandwidth is achieved by taking up the input impedance of an inductively degenerated common-source stage into an LC bandpass filter [14]. Input matching from 6 GHz down to 2.4 GHz can be done with inductive elements of reasonable values, but extending that frequency band to lower values is practically not feasible. Also, the bonding wire and the ESD diodes are taken into account in the matching network.

At the output, a 4-bit programmable capacitor bank provides filtering. A pull-up resistor is added to obtain good linearity. Gain switching is achieved with a bypass cascode transistor that diverts a part of the signal current to the power supply for lower gain without influencing the input matching. Biasing is done with a 3-bit programmable on-chip voltage reference. Simulated values for NF and IIP3 are 2.4 dB and −10 dBm, respectively, at a maximum gain of 22 dB, with a power consumption of 12 mW.

1.4.3 Wideband Down-Conversion Mixer

As far as the dependence on RF carrier frequency is concerned, the down-conversion mixer poses no specific problems—as long as the correct LO frequency is applied, of course. Some standards do have specific blocking test cases that demand very good linearity, both second and third order, which are not yet targeted in this design.

The mixer, shown in Fig. 1.10, consists of a folded double-balanced Gilbert cell driving two current mirrors. The Gilbert cell is intrinsically wideband, as it has a capacitive input impedance and can be driven by a voltage source. Here it is used for wideband operation up to 6 GHz and is the core of the mixer. The current mirrors have a digitally programmable gain B.

An NMOS input pair is used as a transconductance, driving RF signal current into the core switch transistors that form the Gilbert cell. The input must be designed carefully, as it will determine both the noise and the linearity performance of the mixer. A rather large overdrive is needed for linearity, which in combination with the required low-noise performance results in a considerable biasing current of 5 mA.

Because of the low power supply voltage (1.2 V), a folded topology must be used, which does indeed have some drawbacks. The extra folding transistors will

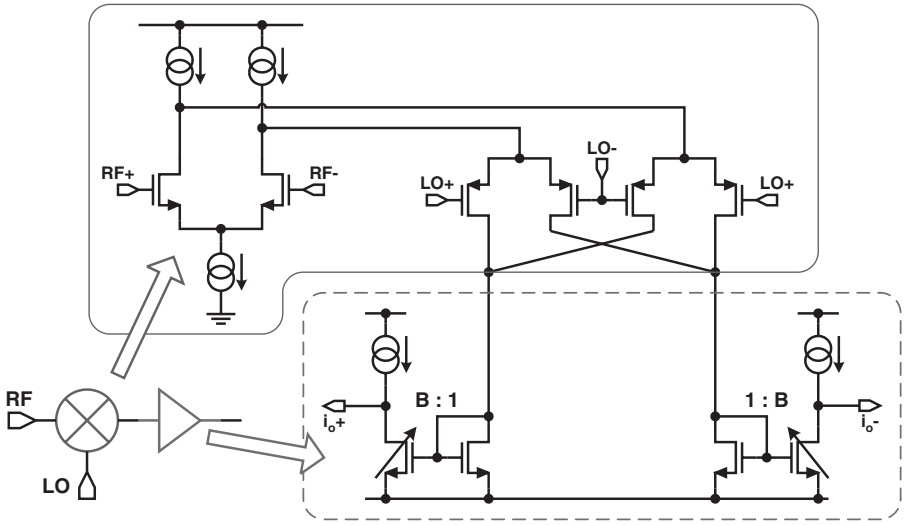

FIGURE 1.10 Schematic of a wideband down-converter.

contribute a certain amount of thermal noise, causing the overall receiver's NF to deteriorate. Some RF signal current will be lost in the parasitics at the folding node, lowering the mixer's conversion gain. The LO switching transistors are thus *folded* PMOS transistors, which also offers an advantage, as they can reduce the mixer's flicker noise. This can not be generalized,however; for examples, at even higher frequencies, NMOS transistors may behave better because of their higher switching speeds and smaller parasitics. The noise contributions in a switching mixer are not easy to understand or analyze [16] but can generally be kept within limits by using large LO signals and reduced dc current through the switching transistors.

The switchable gain of the mixer is achieved by two flexible current mirrors whose output transistors and current gain B are digitally programmable. Consequently, the mixer is actually a voltage-input and current-output building block. The current-mode output is indeed used to drive the subsequent lowpass channel filter. Too much voltage gain must still be avoided at this point to prevent clipping of the output due to strong interferers, and is only allowed after the first stage of the baseband lowpass filter.

1.4.4 Flexible Baseband Analog Circuits

To accommodate the channel bandwidths for a wide range of standards, the flexible baseband lowpass filter (LPF) should offer cutoff frequencies from a few hundreds of kilohertz to several megahertz. Selectivity scalability could allow us to change the filter order in case weaker interferers are detected or in case the standard selected requires less adjacent channel selectivity. The requirements in terms of noise and linearity for the analog baseband channel processing depend, of course, on the

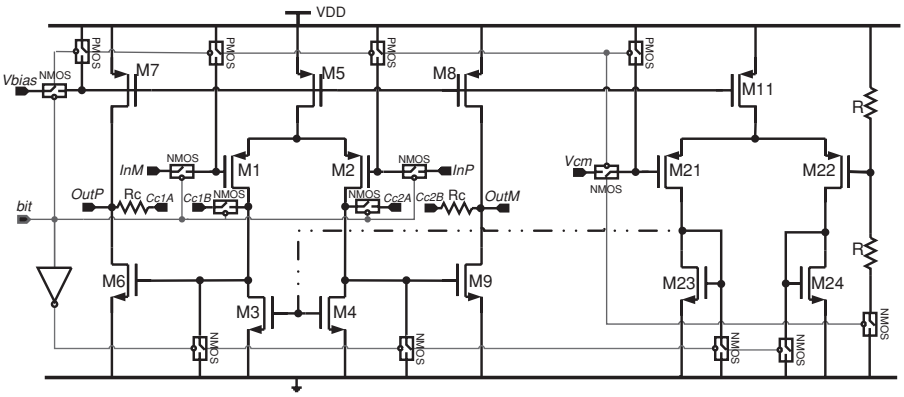

FIGURE 1.11 Switchable Miller op-amp schematic.

decisions taken in the link budget analysis for the entire receiver. In addition to that, an adaptive integrated noise level may lead to a further power saving in case a smaller signal-to-noise ratio (SNR) is required [17].

Adaptive control to provide flexible frequency discrimination and gain control may currently be met with better overall performance by employing proper tunable analog circuits. A possible solution implies the use of modular circuits made of the proper combination of basic units. This allows us to add the desired digital controls efficiently while minimizing layout issues. The concept of *component arrays* fits perfectly with these needs. An *array* is defined here as the parallel connection of dynamic analog blocks dimensioned in a binary-scaled fashion and activated whenever needed. A switchable op-amp is shown in Fig. 1.11. It is the basic unit of a flexible op-amp, which is made up of parallel connections of switchable op-amps in a binary-scaled array. Based on the standard Miller compensated architecture, this op-amp is switched on and off through a single bit. In particular, when the op-amp is off, all the PMOS gates are at V_{DD} and all the NMOS gates are grounded by means of MOS switches. Therefore, in the off-mode, the op-amp shows very high output impedance and zero power consumption. The switches that carry signal must be sized carefully, trading off their finite conductance for their nonlinear characteristic. The two Miller capacitors arrays Cc_{array} are connected at the nodes C_{1a}, C_{1b}, C_{2a}, and C_{2b}, while the node V_{cm} is connected to the output OutP and OutM by means of two resistors R_{cmfb}. In these conditions, the poles and the zero still maintain their original position as in a standard Miller op-amp.

This flexible op-amp is the basic active component of the biquadratic sections of the channel-select filter. It provides reconfigurable gain–bandwidth product (GBW), dynamic impedance scaling, and power scalability. The sensitivity of the dominant poles to supply voltage, process, and temperature variations can be minimized by employing a proper bias circuit [18].

Also all resistor and capacitors in the filter are implemented as arrays, such that their value can be controlled digitally over a very wide range. The capacitor arrays are

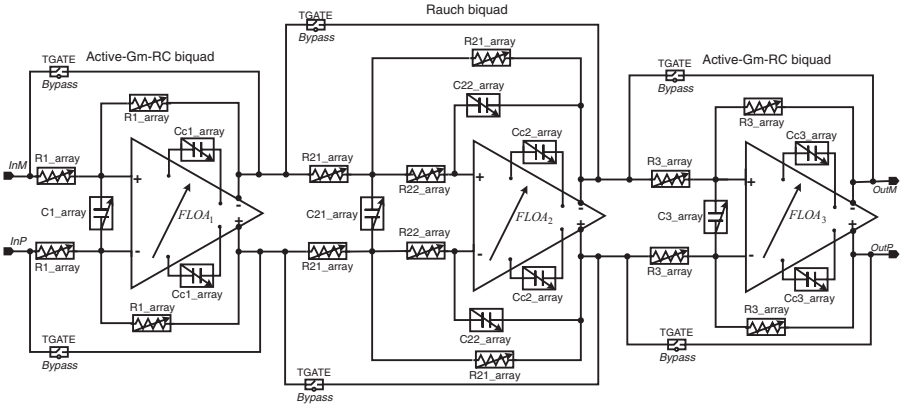

FIGURE 1.12 Flexible lowpass filter schematic.

the binary-weighted connection of basic units of metal interconnect capacitors. For better linearity performance, the NMOS control switches are always on the virtual ground side where the voltage swing is close to zero. The resistor arrays are built as a binary-weighted connection of polysilicon resistors. The control switches are implemented as straight NMOS-PMOS transmission gates; this solution assures a lower on-resistance and a better linear behavior.

Figure 1.12 is a schematic of the flexible baseband lowpass filter (LPF) based on an optimized cascade of active-G_m-RC [18] and Rauch biquadratic sections. The combination of this two biquadratic cells is based on power and linearity considerations. Active-G_m-RC cells guarantee a very good dynamic range with a limited cost in power. However, linearity in this cell is limited by the "weak" virtual ground of its op-amp. Therefore, the Rauch cell allows us to reach the required linearity for the overall filter. Both biquad topologies include the analog components arrays. This solution provides this baseband block with all the SDR programmability required.

The LPF provides the following features: coarse frequency tuning with adaptive power consumption by digitally controlling the resistor values and the number of switchable op-amp units. Fine frequency tuning for RC process deviation compensation is achieved by controlling the capacitor value. One of the power–performance trade-offs we implemented is to reduce the stopband attenuation performance when no large interferers or blockers signals are detected. This is accomplished by turning off and bypassing biquadratic sections. Low on-resistance bypass switches are implemented in each biquad. Furthermore, power consumption can be traded for increased kT/C noise by decreasing the capacitor sizes and the transconductance and increasing the resistor values at the same time. For proper interfacing with the down-conversion mixer, the biquads can also be placed in a current-input mode.

Figure 1.13 summarizes the measurements of a stand-alone flexible lowpass filter. For a sixth-order Butterworth selectivity and 85 μV_{rms} of input integrated in-band noise level, the cutoff frequency can be moved from 0.55 MHz to 17.6 MHz with a coarse step of 0.55 MHz, as shown in Fig. 1.13(a). The power consumption decreases

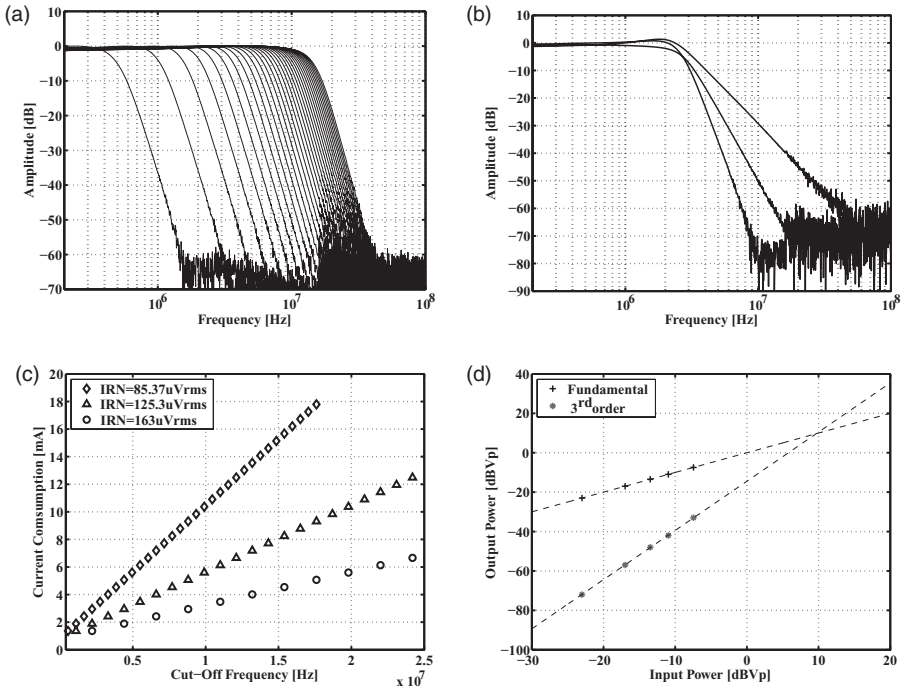

FIGURE 1.13 Flexible lowpass filter measurements: (a) cutoff frequency; (b) filter order; (c) noise–power trade-offs; (d) third-order intercept point.

linearly with reduced cutoff frequency; for example, the filter consumes 13.2 mW for WLAN 802.11a (11 MHz) and 3.6 mW for UMTS 3.86 (2.11 MHz), showing lower power consumption than that of comparable but less flexible designs [18]. In addition to this coarse frequency tuning, the flexible LPF provides fine frequency tuning by configuring its 7-bit capacitor arrays. By also taking this possibility into account, the effective frequency tuning range is included between 0.35 and 23.5MHz. As shown in Fig. 1.13(b), the transition band of the filter can be traded off for less power consumption in case no large interferers are detected or when less selectivity is required. Second-, fourth-, and sixth-order Butterworth-like selectivity are available by bypassing one or two biquadratic sections. Figure 1.13(c) shows how power consumption can be traded off for noise in case a different standard or relaxed sensitivity requires a lower SNR. By reducing the total capacitor size by a factor of 2 or 4 (and simultaneously increasing the resistor values to keep the filter bandwidth fixed), for examples, the power consumption decreases at the cost of a square-root increase in the total integrated noise level.

Figure 1.13(d) shows the IIP3 measurement at the maximum cutoff frequency. The input tones are at 8 and 9 MHz, so that the intermodulation products are well in-band. An IIP3 of 10 dBVp confirms the result expected. This linearity performance is nearly constant for all the cutoff frequency settings. This feature makes the proposed

design a very good alternative to the less-linear G_mC filters, even for flexible designs. I and Q mismatch measurements were also performed for every cutoff frequency; both amplitude and phase mismatch are well below 0.25 dB and 2.8°, respectively.

The variable-gain amplifier (VGA) that increases the filtered signal level to an amplitude fit for the dynamic range of the ADC is designed using the same philosophy. It is built from two cascaded inverting amplifiers using resistive feedback. Use of a flexible op-amp makes it possible to save power by adapting the op-amp bandwidth to the signal bandwidth expected. The gain of the VGA can be adapted in 3-dB steps from 0 to 39 dB. The gain switching time is constrained by the fast automatic gain control (AGC) operation and should be lower than 100 ns. In addition to that, the VGA provides different noise levels by changing the resistors arrays by fixed factors. A dc offset compensation loop is added that senses and removes the dc offset at power-up (in closed-loop operation) and holds the steady-state dc offset compensation value during the burst received (open-loop operation). This internal loop can also be deactivated, in which case a mixed-signal dc offset compensation loop is used that measures the dc level in the digital domain and uses a feedback DAC to compensate it at the input of the AGC.

1.4.5 Analog-to-Digital Conversion

The requirements posed on an ADC by the various standards differ widely, of course, depending on the signal bandwidth, the SNR needed, and even more on the amount of channel filtering and amplification of the preceding chain received. In the link budget used here, typically 8 to 10 bits are needed with sampling speeds up to 40 MS/s (two of them should be interleaved to achieve 80-MS/s operation for 802.11n systems). Those specs have long been the territory of pipeline architectures. But in scaled CMOS they become well in range of SAR ADCs. Most SAR ADCs use an operating principle similar to the charge redistribution architecture [19]. This requires fast-settling op-amps in both the input and reference voltages, able to settle their output voltage in a very short time while driving large capacitive loads. Also a high-speed clock for the controller that has to run at 10 times the sampling speed must be available.

A new SAR architecture is proposed that uses passive charge sharing (instead of active charge redistribution) to both sample the input signal and to perform binary-scaled feedback during the successive approximation [20]. The basic architecture depicted in Fig. 1.14(a) works completely in the charge domain. The input is sampled on a capacitor and during the SAR algorithm charge is added or subtracted until the result converges to zero. No active circuits are used to add and subtract these charges. Instead, simple passive switches do this. The only active element is the comparator itself, which is the basic principle that allows us to achieve the fundamental lowest limit on power consumption.

Operation of the ADC can be explained using the waveforms shown in Fig. 1.14(b). Starting from a reset state, at the start of conversion the input voltage is sampled by the charge on C_T, half of which afterward is transferred to C_S in a passive charge-sharing action. During conversion, the comparator determines the sign of the voltage,

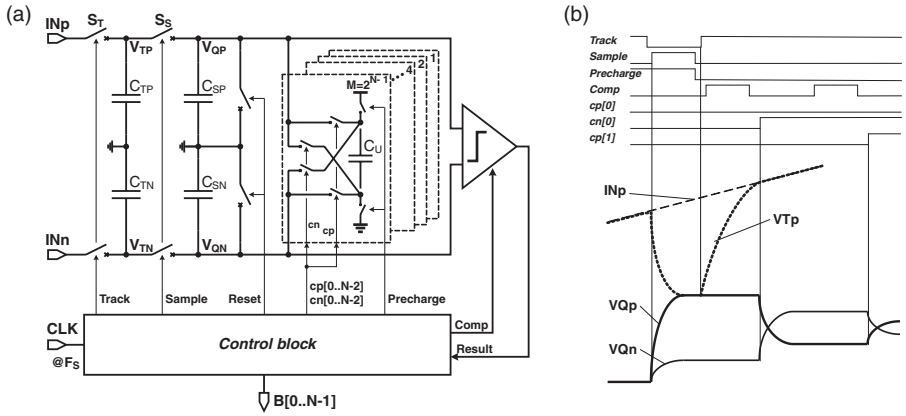

FIGURE 1.14 Charge-sharing SAR ADC: (a) basic architecture; (b) sample waveforms.

and the reference charges (which are stored on an array of binary-scaled precharged capacitors C_U) are added or subtracted passively. This charge sharing acts as the reference feedback DAC in the SAR algorithm, but requires no active power nor imposes constraints, on the reference voltage. The capacitor array size is determined by matching constraints, and the comparator offset needs to be calibrated, similar to the circuit proposed in [21]. To avoid the need of a high-speed clock (and its associated power consumption), an asynchronous controller was implemented. Timing of this controller uses a "valid" signal generated by the comparator when its output is ready.

The fundamental power limits of the original architecture have been removed by doing all the charge redistribution passively. The input is connected to the tracking capacitor for most of the conversion period, so a fast op-amp is not needed for input sampling. Second, settling problems in the reference voltage are avoided by precharging all capacitors to the same voltage before the conversion process, and this is signal independent. This way, the only remaining active elements, in the ADC are the comparator itself and the digital controller. Another advantage is the completely digital implementation, requiring only MOS switches and MOM capacitors, which makes it portable to new CMOS technologies.

The charge-sharing SAR ADC is implemented separately in a 90-nm 1P9M digital CMOS process. At a 50-MS/s conversion rate, the total current consumption is 0.7 mA from a 1-V power supply, divided over the various blocks as follows: digital 50%, comparator 35%, precharging 15%. Measured INL and DNL at 50 MS/s are below 0.6 LSB, as shown in Fig. 1.15(a). Despite this good linearity, the SNDR for low-frequency input signals is only 49 dB (ENOB = 7.8) because it is limited by underestimated comparator noise. At frequencies above 10 MHz, the nonlinearity of the input tracking switch becomes dominant because the low-V_T transistors intended to be used were not processed correctly. This causes the ENOB for a near-Nyquist input at 20 MS/s to deteriorate to 7.4, as shown in Fig. 1.15(b). With only 290 μW of power, the resulting FoM is only 65 fJ per conversion step. As none of the ADC

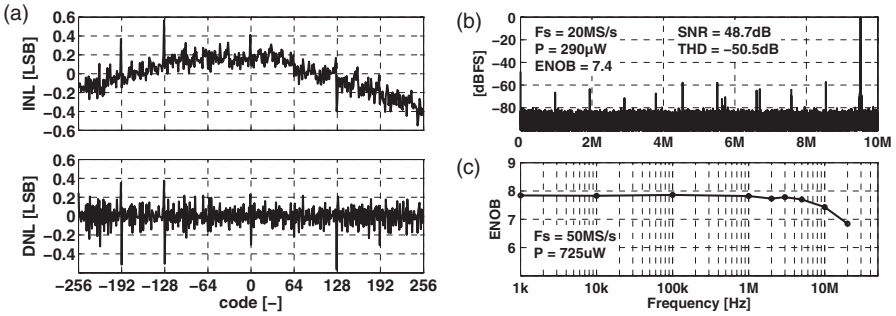

FIGURE 1.15 SAR ADC measurements: (a) INL/DNL; (b) near-Nyquist FFT at 20 MS/s; (c) ENOB vs. input frequency at 50 MS/s.

building blocks consume any static power, the FOM is maintained down to a very low conversion rate, allowing this ADC to be used in a wide variety of applications.

1.5 TRANSMITTER BUILDING BLOCKS

In the architecture presented, a direct-conversion transmitter architecture is used, as this offers the most potential for flexibility. The baseband section uses a programmable filter similar to the receiver filter discussed in Section 1.4.4 to remove the DAC aliases and a wideband direct up-conversion mixer. The pre-power amplifier (PPA) is the final block in the SDR transmit path and is discussed here in more detail.

The PPA includes extensive programmability of gain settings. The full circuit consists of four stages and is shown in Fig. 1.16(a). The output stage is an inductively loaded common-source amplifier with programmable bias current for optimal linearity vs. power trade-off. Some care must be taken in the reliability and lifetime of the output transistor, as it is possible that the output signal swings above the power supply and hence violates the reliability ratings of the technology. The use of a thick-oxide transistor, available for 3.3-V I/O compatibility, would solve this problem totally, but the large transistor and its associated parasitics would prevent the circuit from operating at high frequencies. However, in this PPA circuit these problems are much less severe than in, for example, a full-power CMOS PA. The effective output power is not as high, and for many applications the average swing is much lower than the peak. Furthermore, a nonnegligible voltage drop across the series resistance of the inductor sets the dc output voltage below the power supply.

The stage preceding the output stage provides gain programmability. Its circuit diagram is presented in Fig. 1.16(b). The core of the amplifier is a common-source stage with a PMOS resistive load. To control the gain, three additional PMOS transistors are placed in parallel with the main load. Their gates can be connected to V_{dd}, to turn them off and increase the gain, or to ground, to put them in the linear region and decrease the gain. However, changing the resistive load also has an impact on the dc voltage at the output of the amplifier and so, on its linearity. For a high

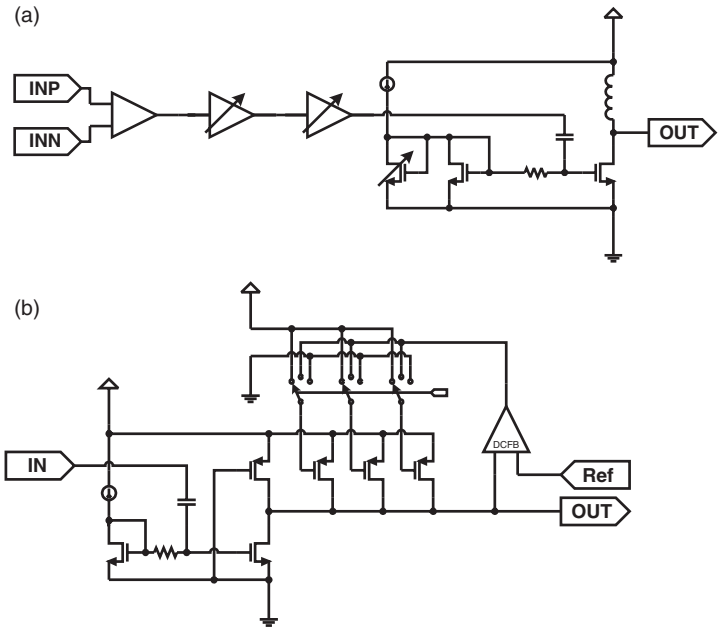

FIGURE 1.16 Pre-power amplifier: (a) block diagram; (b) variable-gain amplifier stage.

gain, the resistive load is large. The dc output voltage drops, compressing the drive transistor. This effect is reduced by keeping part of the bias current out of the PMOS load. In the stage presented, this is implemented by reusing the resistive load PMOS transistors that are turned off in high gain. Instead, the PMOS load transistors that do not contribute to the resistive load are used as current sources. Their gates are biased at a dc level, which is controlled by a dc level feedback circuit (DCFB). In this way, the high-gain linearity is enhanced without adding extra parasitic capacitances on the signal path. A slight disadvantage of using the PMOS transistors as a current source, rather than turning them off, is that the gain is reduced slightly, due to the limited output impedance of the current sources. Therefore, the possibility is left open to select whether the PMOS is turned off or used as a current source. Another point of attention is the reference level to which the dc output level is controlled. In high-bias-current conditions in combination with high gain, it is not possible to have a high-output dc voltage. The dc level feedback will clip to ground, and the bias PMOS transistors will act as resistors, reducing the gain rather than as current sources. To limit this effect, it is possible to program the target output dc level voltage to some extent. Finally, the total bias current through the amplifier can be controlled to optimize the power consumption for the linearity required.

The stage preceding the stage described previously has a similar structure but with extended gain programmability. A cascode transistor has been added to the main amplifier's path. This is possible, as the signal swing is still relatively small in this stage. In parallel to the main branch, several binary-weighted cascode transistors are

used to deviate part of the signal current to the supply, resulting in smaller gains. Although this approach may appear suboptimal from a power consumption point of view, the relative current that is lost for the complete PPA is minor, as the amplifier stages are scaled down from the output to the input. The performance of this circuit varies widely, of course, over carrier frequency, required output power, bias and gain settings, and so on. Simulation results indicate a total gain range of 50 dB and typical IM3 distortion levels of -35 dB at 0-dBm output power.

1.6 CALIBRATION TECHNIQUES

In a multi-mode zero-IF transceiver, design specifications normally span a broad range of present and future standards. Evidently, this overloads the design requirements and hampers effective design. An alternative is to design according to realistic design requirements and to calibrate the front end digitally for its eventual imperfections. The goal of calibration is to improve or optimize the performance of a full IC transceiver. In this sense, calibration relaxes the design requirements for multi-mode systems and enables the use of a low-cost architecture while being compliant with a broad range of standards. In the context of multi-mode systems, calibration should be dynamic, efficient, automatic, and implemented in the system. Calibration consists typically of two steps:

1. *Characterization.* The system imperfections are estimated, often at discrete time instances, namely at system startup, at mode handover, and/or at system-defined time instances. Most of today's calibration techniques require us to deactivate the system's normal operation, which is sometimes not allowed. Fortunately, some alternative techniques exist, allowing the system to remain operational during calibration.
2. *Compensation.* Once the imperfections of the receiver are characterized, the information obtained is used to optimize its performance. Typically, this is achieved by digital pre- and postcompensation of the baseband time-domain signal.

Calibration is one of the keys to enabling the realization of low-cost, high-performance SDR mobile terminals. However, the cost required must also be limited, and fully analog high-frequency calibration techniques are therefore rejected. The only relevant techniques in the SDR context are the digital calibration techniques, which will be discussed briefly.

1.6.1 Quadrature Imbalance

In homodyne receivers the RF signal is directly down-converted to a complex-valued baseband signal, which ideally must have equal amplitude and a phase difference of 90°. In a practical implementation, slight mismatches in amplitude and phase result

in quadrature imbalance for both transmitter and receiver chains. This imbalance is generally characterized by its amplitude mismatch ϵ and phase mismatch $\Delta\phi$, resulting in a negative frequency rejection (NFR) $= 10\log_{10}(\epsilon^2 + \tan^2 \Delta\phi)$.

Different quadrature imbalance characterization techniques exist at the receiver side, based mainly on adaptive filtering [22,23]. As these techniques exhibit slow convergence, they are applicable in streaming modes only. An alternative technique is presented in [24], where quadrature imbalance characterization is preformed based on one calibration measurement. This fast convergence builds on the realistic assumption of a smooth channel between the transmitter and receiver systems. The principal drawback of all the receiver characterization techniques mentioned is, however, the need for a quadrature imbalance-free generated transmitter signal and thus an ideal transmitter. A promising technique is presented in [25], where the quadrature imbalance of the transmitter and receiver systems is characterized separately based on a single calibration measurement. This technique might be perfectly suited in the multi-mode context. Realistic values of transmitter and receiver quadrature imbalance range up to 5% and 6° for ϵ and $\Delta\phi$, respectively. After calibration, the remaining transmitter and receiver quadrature imbalance should be lower than -35 dBc.

1.6.2 DC Offset

In direct-conversion receivers, parasitic coupling between the LO path and the RF path in both directions will cause self-mixing and creation of a dc offset at the baseband signal [26]. This dc offset decreases the effective dynamic signal swing, especially when the signal power received is low, and therefore reduces the gain and linearity performances of the receiver. Using a highpass filter to remove the dc offset from the baseband signal is not always an option; the targeted standards have a dense spectral occupation, and such filters cannot be implemented efficiently. Therefore, another calibration technique should be used. As the dc offset in the receiver path is generated primarily in the mixer, to limit its impact it should also be removed as close as possible to this mixer in the baseband path. Reference [27] suggests an architecture adding a digitally controlled complex compensation dc offset directly after the mixer. Building on a similar architectural approach, several authors [25,28] present characterization algorithms that find the optimal complex compensation dc value while keeping the system operational. Both techniques provide very fast convergence and are thus applicable in multi-mode transceivers.

1.6.3 Impact of LPF Spectral Behavior

Accurate control of the bandwidth of channel-select filters is required to guarantee, for example, sufficient suppression of adjacent channels. The spectral behavior calibration of the LPF can be performed completely at baseband; baseband switches are used to connect (part of) the analog baseband circuitry in between the transceiver's

baseband input and output. In a given configuration, the spectral behavior or transfer function can be characterized by comparing the digital signal before and after propagation through the analog circuit. When using a multi-tone signal with a relatively dense and sufficiently wide spectral content, the cutoff frequency, in-band ripple, and spectral phase relation can easily be characterized.

1.7 FULL SDR IMPLEMENTATION

The complete SDR front end (without ADC) has been implemented in a 1.2-V 0.13-μm CMOS technology [29]. In this section we describe the prototype with dual-wideband low-noise amplifiers (the one with a MEMS-enabled dual-band LNA is reported in [30]). Figure 1.17 shows a microphotograph of the SDR front end with highlights on the major circuits. The total die area is 3×3.8 mm^2, of which about 7.7 mm^2 is taken up by active circuits.

To verify receiver behavior over the complete input power range, a receiver budget measurement is shown in Fig. 1.18 for three different channel bandwidths. The SNDR is limited by thermal noise for low input powers. The $1/f$ noise corner is around 200 kHz, which explains the higher NF for the low-bandwidth mode. Realistic interferer and blocker levels are used, corresponding to a Bluetooth, UMTS, and WLAN scenario, respectively. To accommodate these interferers, the front end will reduce its LNA gain at higher input power levels, sometimes with a resulting small dip in the SNR. At high input powers, distortion is the limit.

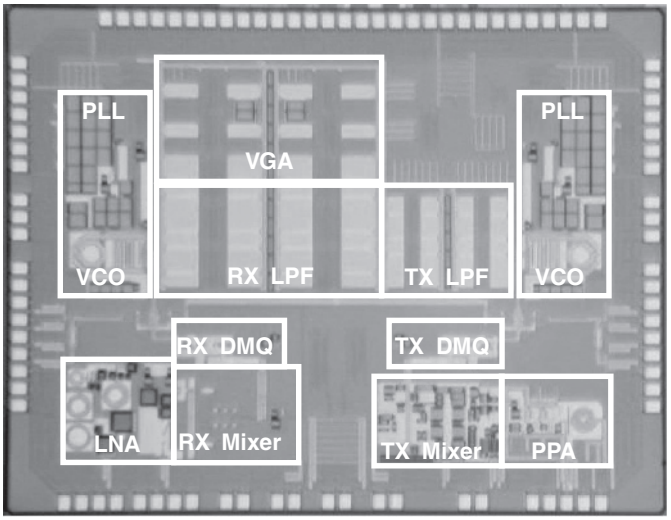

FIGURE 1.17 SDR chip microphotograph.

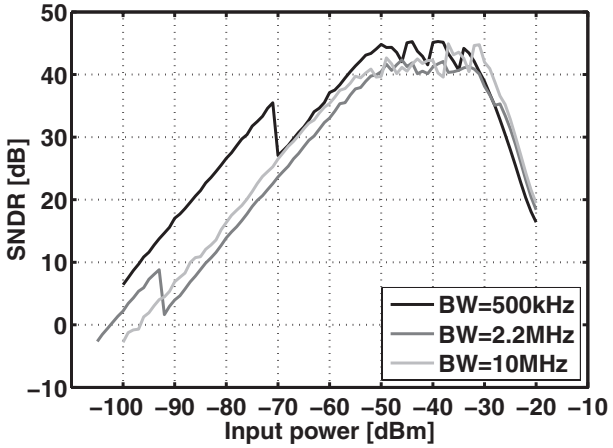

FIGURE 1.18 RX chain radio budget measurement.

The measured receiver noise figure, gain, and IIP3 as a function of the RF frequency are shown in Fig. 1.19 for a channel bandwidth of 20 MHz. A typical noise figure is around 5 dB up to a 5-GHz carrier frequency, above which the performance degrades, due to insufficient LO signal swing. The input IP3 varies from −8 to −4 dBm over the frequency range.

A measured 64QAM OFDM constellation and corresponding output spectrum are presented in Fig. 1.20. They correspond to an EVM of −29.5 dB for an output power of −0.5 dBm at 2.45 GHz. Performance varies, of course, over the carrier frequency and operation point chosen. Power can be saved at the expense of reduced linearity and degraded EVM. Table 1.1 presents a summary of the power consumption and performance of the various SDR circuits for various operating modes.

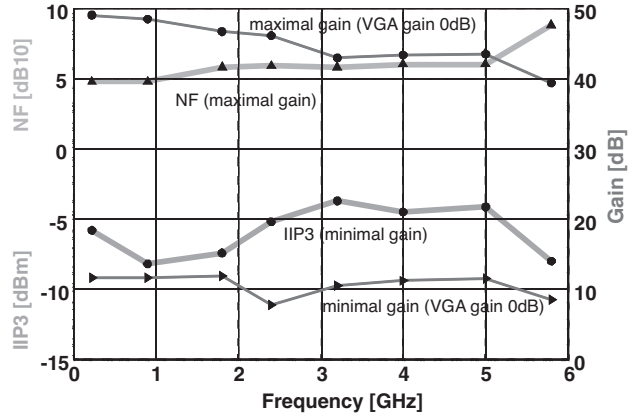

FIGURE 1.19 SDR receiver performance.

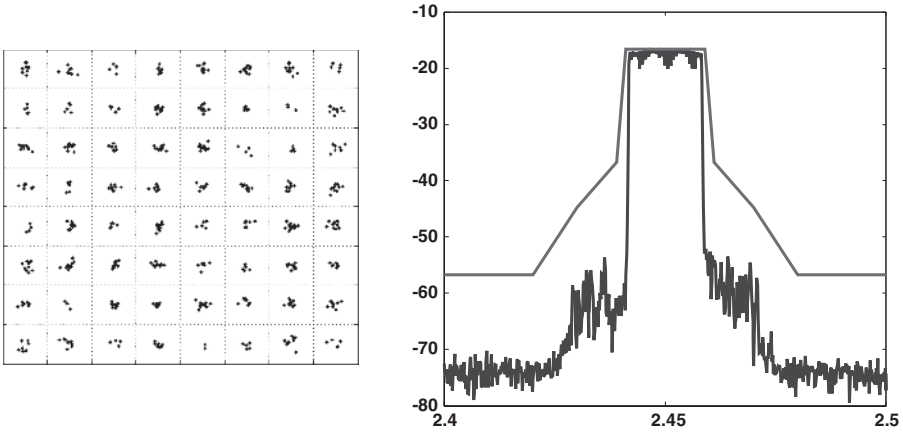

FIGURE 1.20 SDR TX constellation and output spectrum at 2.45 GHz ($P_{\text{out}} = -0.5$ dBm; EVM $= -29.5$ dB).

TABLE 1.1 SDR Performance Overview

Power supply = 1 .2 V

Current		(mA)	(min./typ./max.)		(mA)
Receiver	LNA	8/12	Transmitter	LPF	2×4
	Mixer	$2 \times 5/9/12$		Mixer	2.6 to 5.9
	LO buffer	$2 \times 3/4/7$		PPA	25 to 51
	LPF	$2 \times 0.3/10$			
	VGA	$2 \times 1/6$			
DMQ $\times$ 5, /4 @ 4.9 GHz		24	DMQ(cont.) /4 @ 1 GHz		4.8
$\times$ 3, /4 @ 3.0 GHz		16.5	/8 @ 500 MHz		4.5
$\times$ 3, /8 @ 1.5 GHz		16.5	/16 @ 250 MHz		4.4
IQ @ 4 GHz		14.3	/32 @ 125 MHz		4.4
/2 @ 2 GHz		6.8	Spur. tones $\times$ 3/4 mode		< -30 dBc

Receiver performance	
NF°(high gain)	4.8 to 8.5 dB
IIP3°(low gain)	-8.2 to -3 dBm
Gain	10 to 90 dB
Freq. range	0.1 to 6 GHz
Cutoff freq.	0.35 to 23 MHz
Current	27 to 82 mA

°f_{RF} 100 MHz to 6 GHz
Channel BW 20 MHz

Transmitter perf.	2.45 GHz	4.9 GHz
P1dB	5.8 dBm	1 dBm
OIP3	15.5 dBm	12 dBm

P_{out} (WLAN 64 QAM)	EVM	I_{vdd}
-0.5 dBm @ 2.45 GHz	-29.5 dB	51
-3.1 dBm @ 4.9 GHz	-29.2 dB	51
-6.2 dBm @ 5.24 GHz	-30.0 dB	51
-0.6 dBm @ 2.45 GHz	-26.5 dB	36
-2.3 dBm @ 2.45 GHz	-30.6 dB	36

Gain control @ 2.45 GHz	Range	Step
Mixer	20 dB	~ 5 dB
PPA	43 dB	< 2 dB

1.8 CONCLUSIONS

In this chapter, the basic architecture and implementation concepts of a software-defined radio analog front end have been presented. A direct-conversion transceiver with very flexible and reconfigurable building blocks has been analyzed at the system level in order to be able to cover the requirements imposed by a large set of communication standards in cellular, WLAN, WPAN, broadcasting, and positioning applications.

An important aspect of every SDR front end is the LO synthesis. Since many RF frequency bands need to be covered, ranging from 174 MHz up to 6 GHz, it is a very complex task to generate all local oscillator signals. A very wideband and flexible VCO has been demonstrated, which in combination with division and multiplication in quadrature makes this feasible. Several innovations in the various building blocks of the receiver and transmitter chain are also needed to achieve full SDR functionality. Some examples include the use of a MEMS switch in the input stage of an LNA, wideband LNAs, an ultraflexible baseband channel select filter, an innovative charge-sharing successive approximation ADC, and a programmable power amplifier driver. A versatile RX and TX path with programmability to address the various functional requirements of many different standards can offer an optimal power consumption simultaneously by trading in unnecessary performance at runtime. All these concepts are integrated in the world's first true SDR transceiver prototype, achieving good performance combined with extensive programmability. Although several other improvements will still be needed, the work presented has already taken an important step toward an energy-efficient multi-mode software-defined radio.

Acknowledgments

The author would like to thank all the members of the RF design team in the Wireless Research Group in IMEC who have made the SDR project possible: Boris Côme, Björn Debaillie, Vito Giannini, Michael Goffioul, Dries Hauspie, Mark Ingels, Michael Libois, Mingxu Liu, Pierluigi Nuzzo, Charlotte Soens, Geert Van der Plas, Gerd Vandersteen, Joris Van Driessche, and Piet Wambacq. The results on baseband analog circuits have been obtained in cooperation with Andrea Baschirotto, University of Salento, Italy. This research has been carried out in the context of IMEC's multi-mode multi-media program, which is partly sponsored by Samsung.

REFERENCES

1. M. Goffioul, G. Vandersteen, J. Van Driessche, B. Debaillie, and B. Côme, "Ensuring consistency during front-end design using an object-oriented interfacing tool called NETLISP," in *Design Automation Conference,* San Francisco, pp. 889–892, July 2006.
2. J. Van Driessche, J. Craninckx, and B. Côme, "Analysis and key specifications of a novel frequency synthesizer architecture for multi-standard transceivers," in *IEEE Radio and Wireless Symposium*, pp. 481–484, Jan. 2006.

3. K. Manetakis, D. Jessie, and C. Narathong, "A CMOS VCO with 48% tuning range for modern broadband systems," in *IEEE Custom Integrated Circuits Conference*, pp. 265–268, Oct. 2004.
4. M. Tiebout, "A CMOS fully integrated 1 GHz and 2 GHz dual band VCO with voltage controlled inductor," in *European Solid-State Circuits Conference*, pp. 799–802, Sept. 2002.
5. D. Hauspie, E.-C. Park, and J. Craninckx, "Wideband VCO with simultaneous switching of frequency band, active core and varactor size," *IEEE J. Solid-State Circuits*, vol. 42, pp. 1472–1480, July 2007.
6. J. Craninckx and M. Steyaert, *Wireless CMOS Frequency Synthesizer Design*. Dordrecht, The Netherlands: Kluwer Academic, 1998.
7. P. Andreani and A. Fard, "More on the $1/f^2$ phase noise performance of CMOS differential-pair LC-tank oscillators," *IEEE J. Solid-State Circuits*, vol. 41, pp. 2703–2712, Dec. 2006.
8. J. Craninckx and M. Steyaert, "A 1.8 GHz low-phase noise CMOS VCO using optimized hollow inductors," *IEEE J. Solid-State Circuits*, vol. 32, pp. 736–744, May 1997.
9. C. Nguyen, "Integrated micromechanical circuits for RF front ends," in *Proc. IEEE European Solid-State Circuits Conference*, Montreux, Switzerland, pp. 7–16, Sept. 2006.
10. M. Liu, M. Libois, M. Kuijk, A. Barel, J. Craninckx, and B. Côme, "MEMS-enabled dual-band 1.8 and $5-6$ GHz receiver RF front-end," in *IEEE Radio and Wireless Symposium*, Jan. 2007.
11. TeraVicta Technologies, *TT712-68CSP SPDT*7 GHz *RF MEMS Switch*. http://www.teravicta.com.
12. H. D. L. Santos, G. Fischer, H. Tilmans, and J. van Beek, "RF MEMS for ubiquitous wireless connectivity: II. Application," *IEEE Microwave Mag.*, vol. 5, pp. 50–65, Dec. 2004.
13. J.-H. Zhan and S. Taylor, "An inductor-less broadband LNA with gain step," in *Proc. IEEE European Solid-State Circuits Conference*, Montreux, Switzerland, pp. 344–347, Sept. 2006.
14. A. Bevilacqua and A. M. Niknejad, "An ultrawideband CMOS low-noise amplifier for 3.1–10.6 GHz wireless receivers," *IEEE J. Solid-State Circuits*, vol. 39, pp. 2259–2268, Dec. 2004.
15. F. Bruccoleri, E. Klumperink, and B. Nauta, "Wide-band CMOS low-noise amplifier exploiting thermal noise canceling," *IEEE J. Solid-State Circuits*, vol. 39, pp. 275–282, Feb. 2004.
16. H. Darabi and A. A. Abidi, "Noise in RF-CMOS mixers: a simple physical model," *IEEE J. Solid-State Circuits*, vol. 35, pp. 15–25, Jan. 2000.
17. V. Giannini, J. Craninckx, S. D'Amico, and A. Baschirotto, "Flexible baseband analog circuits for software-defined radio front-ends," *IEEE J. Solid-State Circuits*, vol. 42, pp. 1501–1512, July 2007.
18. S. D'Amico, V. Giannini, and A. Baschirotto, "A 4th-order active-G_m-RC reconfigurable (UMTS/WLAN) filter," *IEEE J. Solid-State Circuits*, vol. 41, pp. 1630–1637, July 2006.
19. J. McCreary and P. Gray, "All-MOS charge redistribution analog-to-digital conversion techniques: I," *IEEE J. Solid-State Circuits*, vol. 10, pp. 371–379, Dec. 1975.

20. J. Craninckx and G. Van der Plas, "A 65 fJ/conversion-step, 0-50 MS/s 0−0.7 mW 9bit charge sharing SAR ADC in 90nm digital CMOS," in *Technical Digest, IEEE International Solid-State Circuits Conference*, San Francisco, pp. 246–247, Feb. 2007.
21. G. Van der Plas, S. Decoutere, and S. Donnay, "A 0.16 pJ/conversion-step 2.5 mW 1.25 GS/s 4b ADC in a 90nm digital CMOS process" in *Technical Digest, IEEE International Solid-State Circuits Conference*, San Francisco, pp. 566–567, Feb. 2006.
22. A. Schuchert, R. Hasholzner, and P. Antoine, "A novel IQ imbalance compensation scheme for the reception of OFDM signals," *IEEE Trans. Consumer Electron.*, vol. 47, pp. 313–318, Aug. 2001.
23. M. Valkama, M. Renfors, and V. Koivunen, "Advanced methods for I/Q imbalance compensation in communication receivers," *IEEE Trans. Signal Process.*, vol. 49, pp. 2335–2344, Sept. 2001.
24. J. Tubbax, B. Côme, L. Van der Perre, S. Donnay, M. Engels, H. De Man, and M. Moonen, "Compensation of IQ imbalance and phase noise in OFDM systems," *IEEE Trans. Wireless Commun.*, vol. 4, pp. 872–877, May 2005.
25. B. Debaillie, P. Van Wesemael, and J. Craninckx, "Calibration method for RF imperfections enabling low-cost SDR," *IEEE International Conference on Communications*, Beijing, pp. 4899–4903, May 2008.
26. B. Razavi, *RF Microelectronics*. Upper Saddle River, NJ: Prentice Hall, 1998.
27. D. Haspeshgh, J. Ceuterick, L. Kiss, J. Wenin, A. Vanwelsenaers, and C. Enel-Rehel, "BBTRX: a baseband transceiver for a zero IF GSM hand portable station," in *IEEE Custom Integrated Circuits Conference*, May 1992.
28. B. Côme, D. Hauspie, G. Albasini, S. Brebels, W. De Raedt, W. Diels, W. Eberle, H. Minami, J. Ryckaert, J. Tubbax, and S. Donnay, "Single-package direct-conversion receiver for 802.11a wireless LAN enhanced with fast converging digital compensation techniques," in *IEEE MTT-S International Microwave Symposium Digest*, pp. 555–558, June 2004.
29. M. Ingels, C. Soens, J. Craninckx, V. Giannini, T. Kim, B. Debaillie, M. Libois, M. Goffioul, and J. Van Driessche, "A CMOS 100 MHz to 6 GHz software defined radio analog front-end with integrated pre-power amplifier," in *Proc. IEEE European Solid-State Circuits Conference*, Munich, Germany, pp. 436–436, Sept. 2007.
30. J. Craninckx, M. Liu, D. Hauspie, V. Giannini, T. Kim, J. Lee, M. Libois, B. Debaillie, C. Soens, M. Ingels, A. Baschirotto, J. Van Driessche, L. Van der Perre, and P. Vanbekbergen, "A fully reconfigurable software-defined radio transceiver in 0.13 μm CMOS" in *Technical Digest, IEEE International Solid-State Circuits Conference*, San Francisco, pp. 346–347, Feb. 2007.