1 Introduction to Highly Integrated and Tunable RF Receiver Front Ends

1.1 Introduction

With the ever-increasing demand for instant access to data over wideband communication channels, the quest for a universal mobile terminal capable of delivering the ultimate user experience has become imperative. Over the last decade, researchers were exploring the possibility of having a universal radio that can be programmed and reconfigured through software to operate on any bands, channel bandwidths, and modulations. Such a universal radio was named software-defined radio (SDR) [1–6]. The SDRs face unique challenges because their targeted applications are mostly in mobile handheld devices. They must be small and affordable, and must last longer between charges. The design of such a low-cost, low-power, and flexible radio that meets the tough requirements of individual standards is enormously challenging and was and still is a hot topic of research for circuit designers as well as system and hardware engineers. One common yet relatively simple example of an SDR is a 3G cell phone, which can support as many as 17 bands in three modes of operation, namely GSM, EDGE, and WCDMA/HSPA.

The most aggressive SDR architecture was proposed by Mitola in 1995 [1], and is shown in Fig. 1.1(a). The only analog blocks in the receiver and the transmitter are an ADC and a DAC, respectively. Such a transceiver provides maximum flexibility through the digital signal processor (DSP), and it is even capable of simultaneously detecting several standards. The receiver of such an ideal SDR, however, offers equal fidelity for the entire incident signal from the antenna, which is composed of the desired signal, possibly accompanied by some blockers. As shown in Fig. 1.2, in many wireless applications, the desired received signal can be very weak, whereas the blockers can be stronger by as much as 100 dB. These blockers can be created by nearby transmitters of the same communication standard, which in this case are called in-band blockers, or they can be out-of-band blockers generated by any of the other transmitters. The lack of any filtering in the SDR shown in Fig. 1.1(a) imposes an impractical dynamic range of about 100 dB for the ADC resolution. Based on the survey published in [7], as of today, such an ADC remains impractical, with an estimated power consumption of 2 kW, which is obviously not an acceptable power consumption level for mobile devices. Therefore, the SDR architecture perceived by Mitola will still remain a future dream despite its attractiveness as a true DSP-based solution.

To break this equal-fidelity reception that was detracting from Mitola's SDR receiver with the gigantic ADC power consumption, the strong received blockers must be

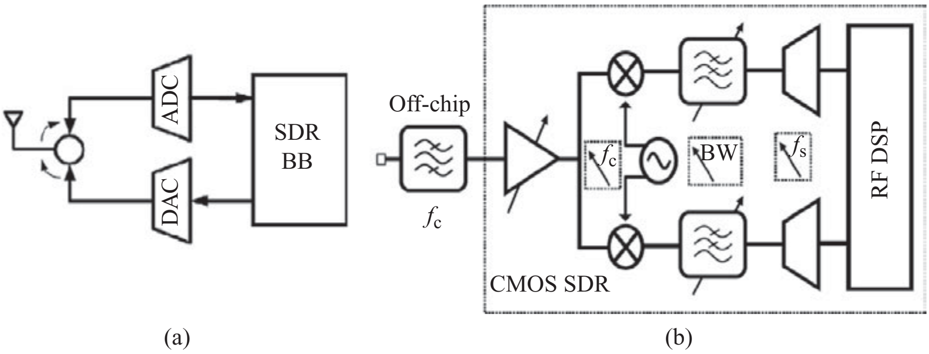

Figure 1.1 (a) Ideal SDR. (b) A more practical architecture. © 2010 IEEE. Reprinted, with permission, from [53].

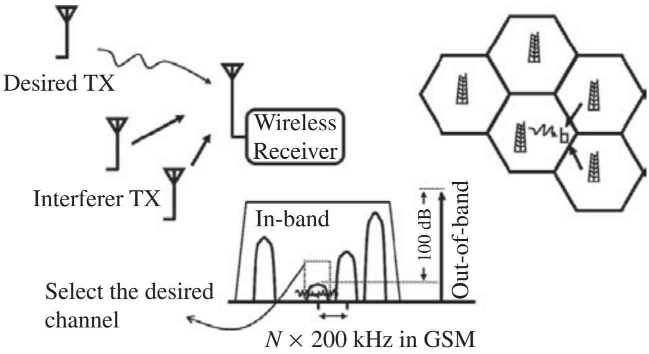

Figure 1.2 Blockers in wireless environments. © 2010 IEEE. Reprinted, with permission, from [53].

attenuated ahead of the ADC. This filtering could be balanced between the RF and analog baseband. Thus, a more practical receiver architecture shown in Fig. 1.1(b) is adopted in most SDRs today [5, 6], where, by means of downconversion,‘ a considerable portion of analog and digital signal processing is performed at a conveniently lower intermediate frequency (IF). What differentiates this receiver from other traditional radios is the added programmability in almost everything, including channel-select filter bandwidth or ADC sample rate to allow several modes of operation, as well as extended RF bandwidth of the transceiver front end and PLL range to support multiple bands.

The architecture of Fig. 1.1(b) solves the in-band blocker problem through programmable IF filtering, however, the out-of-band blockers remain a challenge. For example, in the case of GSM, this out-of-band blocker can be as strong as 0 dBm, which can compress the receiver front end excessively (explained more in the next section and Fig. 1.3), and thus desensitize it. Therefore, an additional front-end filter is needed to attenuate the blocker adequately [Fig. 1.1(b)] before it experiences the large gain of the low noise amplifier (LNA). The filter typically requires a narrow bandwidth

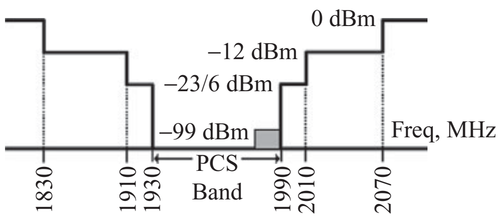

Figure 1.3 GSM out-of-band blocker profile. © 2010 IEEE. Reprinted, with permission, from [53].

set by a given application and a very sharp stopband. Consequently, due to its very high-quality nature, it is typically implemented externally, which adds considerably to the cost and size of the reference design. Additionally, the inevitable insertion loss of the filter increases the receiver noise figure directly.

Moreover, because the filter bandwidth and center frequency are inevitably not programmable, for every band or mode of operation, a dedicated input and a corresponding filter are needed. These items add further to the cost and, more importantly, oppose the promise of maximum hardware sharing offered by the SDR architecture shown in Fig. 1.1(b). This circumstance continues to be one of the greatest obstacles to realizing true software-defined radios.

In this chapter, we give a brief overview of several circuit design techniques proposed to address these great challenges, enabling highly programmable and tunable front-end filters integrated with the rest of the CMOS RF IC. We first start briefly to touch on the system-level requirements of the radio front end. The main focus is on cellular applications, which are the most challenging realization of SDR.

1.2 Front-end integration challenges and system requirements

Integration of external SAW filters involves unique circuit and system-level challenges. Because cellular is the most demanding standard in terms of blocking requirements, it is the main focus of this book. In this section, challenges and high-level requirements of the integration of the SAW filters for 2/3G transceivers are examined.

In the case of GSM/EDGE, the major challenge stems from out-of-band 0 dBm blockers that can be as close as 20 MHz or 80 MHz for low-band or high-band cases, respectively (Fig. 1.3). According to the 3GPP standard [8], while this out-of-band blocker hits the antenna, the desired signal can be as weak as −99 dBm, which is only 3 dB above the sensitivity level of −102 dBm. The standard identifies this 0 dBm blocker as a static sine wave. Therefore, the blocker imposes only compression issues and possibly puts some limitations on the local oscillator phase noise.

The RF filtering must provide at least 23 dB (approximately) of attenuation to the out-of-band blockers to reduce them to the level of the in-band blockers, which can be as large as −23 dBm at 3 MHz away (Fig. 1.3). These in-band blockers experience no RF filtering. Therefore, the receiver is expected to handle these −23 dBm in-band blockers.

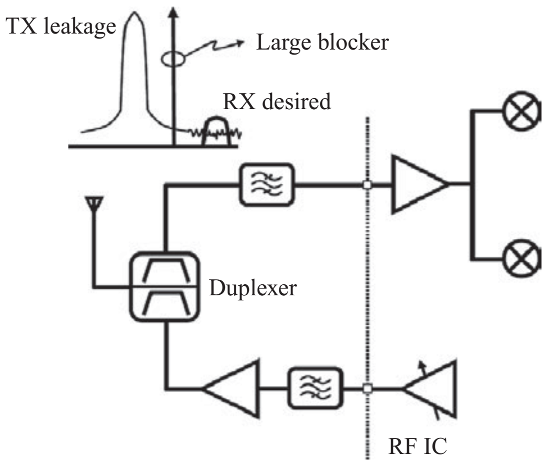

Figure 1.4 3G Full-duplex issue. © 2010 IEEE. Reprinted, with permission, from [53].

Of course, such a receiver can certainly handle the attenuated out-of-band blockers too. If realized by inductors and capacitors, the corresponding quality factor (Q) would be required to be greater than 100, making it impractical to be realized on-chip, especially in a regular bulk CMOS process.

Most advanced handsets today aim for −109 dBm sensitivity or better, even though a sensitivity of −102 dBm is specified in the standard. This translates to a total noise figure of 5 dB for the entire system (assuming 200 kHz bandwidth and 5 dB signal-to-noise ratio or SNR). Assuming a loss budget of 1.5 dB for the SAW filter and 1 dB for the antenna switch, the receiver noise figure must be about 2.5 dB. For a sensitivity of −109 dBm, removal of the SAW filter relaxes the receiver noise figure to 4 dB because the SAW filter introduces a loss of 1.5 dB.

In the case of 3G radios, the receiver and transmitter operate simultaneously (Fig. 1.4). The full-duplex issue raises some unique challenges. Ideally, an external duplexer realized by two highly selective filters separates the receive and transmit signals. In practice, due to the finite isolation of the duplexer, some of the strong TX signal leaks to the RX input, causing two issues. First, due to the front-end third-order nonlinearity, the leaked TX signal potentially can mix with a large out-of-band blocker (for example, the blocker at half-duplex frequency) and desensitize the RX. Second, the TX noise falling in the RX band effectively degrades the receive noise figure. To overcome these issues, two external SAW filters [9] are traditionally placed at the TX and RX ports to suppress the TX noise and leakage, respectively, thereby relaxing the phase noise and linearity requirements of the transceiver. In the case of the transmitter output, the SAW filter relaxes the noise requirement of the TX chain by providing some filtering. This RX band noise attenuation obviously comes with two drawbacks: an additional external component, and reduction of the transmitter efficiency due to passband loss of the SAW filter. On the other hand, the receiver SAW filter attenuates the TX residual

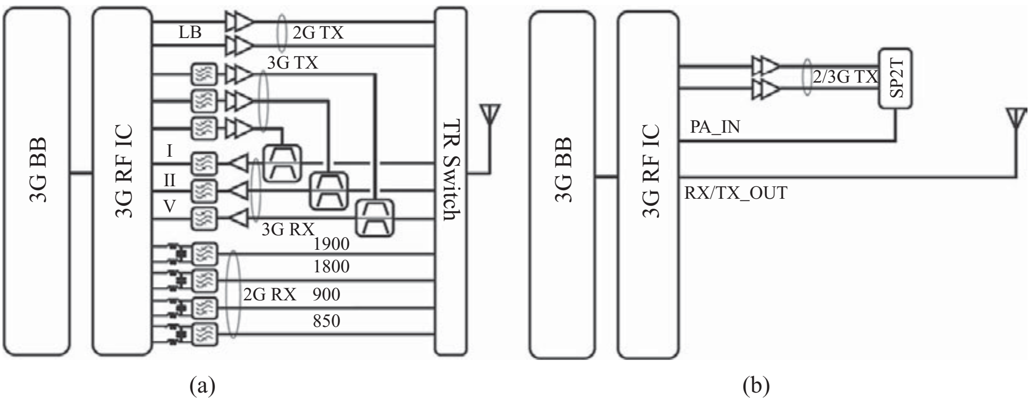

Figure 1.5 Examples of: (a) current, (b) future 3G radios. © 2010 IEEE. Reprinted, with permission, from [53].

leakage and any other blockers, thereby relaxing the linearity requirements of the RX chain.

Similar concerns are present in the case of the long-term evolution (LTE) standard. For the LTE standard, the out-of-band filtering requirements are the same as those required for 2G and 3G (−15 dBm, which is the worst case for LTE but is still the dominant requirement for GSM). However, the in-band blocking requirement for LTE is more stringent due to the wide channel bandwidth. The challenges of this stringent in-band blocking requirement mostly exist in the design of the integrated channel select filter, not the RF front end. Note that similar to 3G, LTE must support the FDD option as well. For other shorter-range standards (wireless personal area network [WPAN] or wireless local area network [WLAN]), the blocking requirements are far less stringent. For example, in Bluetooth applications, the out-of-band blockers are specified to be −10 dBm (from 0–2 GHz, and >3 GHz), as opposed to 0 dBm in the case of GSM. However, the recent demand for integrating WPAN and WLAN devices with cell phones, which is called coexistence, does impose more challenging blocking requirements for the aforementioned short-range standards. This is because the large TX signal of the cellular device is coupled with the other embedded devices (although the cellular device still has the most stringent in-band/out-of-band blocking requirements).

For these reasons, current cellular platforms use several external filters and duplexers to mitigate noise, compression, and linearity issues imposed by either the blockers or the TX leakage from the cellular radio itself. An example of a quad-band GSM/EDGE tri-band WCDMA radio is shown in Fig. 1.5(a), which uses as many as 10 SAW filters, three duplexers, and several matching components.

Besides the obvious size and substantial cost implications, the presence of these external components is contrary to the hardware-sharing concept provided in SDRs. In this chapter we aim briefly to discuss techniques to eliminate these external passive components, and, ultimately, introduce a single-input 3G SDR with all external components integrated [Fig 1.5(b)]. These techniques will be described thoroughly in subsequent chapters.

1.3 2G receiver SAW elimination

In a high-band GSM receiver, the RF filter must attenuate the out-of-band blocker at 80 MHz away by about 23 dB or more (see Fig. 1.3). Note that passband of the PCS band is 60 MHz wide. The low band has a bandwidth of 35 MHz but a more stringent stopband of 20 MHz away. Realization of such a filter with LC structure demands a very high quality factor (Q) for the filter components. This inevitably high Q requirement calls for Q enhancement techniques, but they have proven to be insufficient to meet the stringent GSM noise and linearity requirements [10–12]. Moreover, if implemented on-chip, any small variations in the values of the capacitors and inductors can cause significant changes in the frequency response due to the high-Q nature of the filter. In this section, we discuss several other ideas proposed recently to eliminate the front-end SAW filters. Of the few techniques described, the M-phase filtering followed by linear LNA receivers seem to be the most promising in the realization of true SAW-less GSM receivers, whereas the active blocker-cancellation technique and having the mixer first may not be as attractive for cellular applications.

1.3.1 Mixer-first receivers

One simple and basic way of enhancing the receiver linearity is to reduce the gain in front of the downconversion mixers that are typically the bottleneck. As shown in Fig. 1.6, a very aggressive way of reducing this gain is to remove the low noise amplifier. Mixer-first receivers, also known as LNA-less receivers, proposed in [13, 14], have demonstrated promising linearity due to the lack of the high gain of the LNA upfront. However, as expected, these receivers suffer from less than good noise figures (NF). Note that because a SAW filter has a typical insertion loss of about 1.5 dB, a SAW-less GSM receiver can enjoy a more relaxed NF of about 4 dB, as derived in Section 1.2. Still, removing the LNA will inevitably degrade the noise figure, unless noise contributions of the mixer buffer and the LO chain are sufficiently reduced by carrying much higher currents in these blocks. In addition, to improve the receiver noise figure, mixer switches that are connected to the receiver input must be sufficiently large. The large switch sizes could potentially exacerbate the LO-to-RF feedthrough to an unacceptable level.

Another remaining challenge is the harmonic mixing, as the LNA-less receiver is quite wideband and provides little filtering at the LO harmonics. Due to harmonic mixing, blockers located at the harmonics of the LO can be down-converted to the baseband and aliased on top of the desired signal. The 8-phase mixing scheme proposed in [14] helps substantially at the expense of increasing the LO chain power consumption. The 8-phase design removes the $3f_{\mathrm{LO}}$ and $5f_{\mathrm{LO}}$ blockers inherently and shifts the closest folding harmonic to $7f_{\mathrm{LO}}$, although in practice the harmonics rejection at the $3f_{\mathrm{LO}}$ and $5f_{\mathrm{LO}}$ blockers is finite, limited by the mismatches in the LO phases.

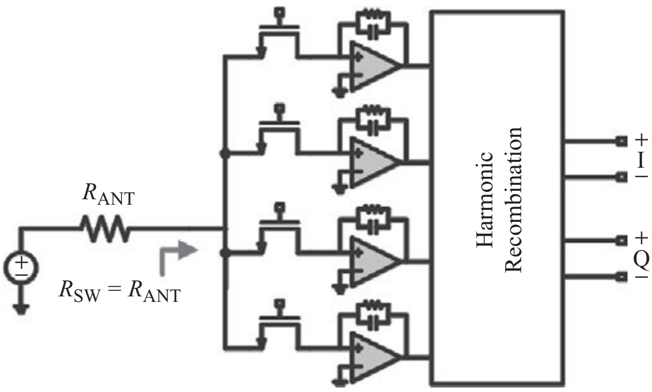

Figure 1.6 Mixer-first receiver front end. © 2010 IEEE. Reprinted, with permission, from [53].

1.3.2 Active blocker cancellation

On-chip active blocker cancellation potentially can be an alternative to replacing external SAW filters. Active blocker cancellation can be in two major forms: feedforward-based cancellation and feedback-based cancellation. The receivers [15, 16] use feedforward cancellation to achieve on-chip high-Q bandpass filters, whereas feedback cancellation is used in the receiver reported in [17–19]. Active blocker-cancellation techniques will be explained in detail in the next chapter. In this chapter, we explain briefly the feedforward blocker-cancellation technique utilized in [15]. In this approach, the LNA is kept in the receiver front end so as not to compromise the NF at the sensitivity level. However, before they reach the LNA output and cause compression, the out-of-band blockers are canceled by exploiting the feedforward cancellation technique, as illustrated in Fig. 1.7.

The feedforward path must suppress the desired signal and should allow only blockers to pass through. To remove the blocker from the desired signal, a sharp notch filter in the feedforward path is required. Because such a high-Q stopband filtering is not realizable at the RF, the notch is realized by a frequency translational loop where the low-Q baseband bandpass filter (BPF) response appears at RF through subsequent downconversion and upconversion by the same receive LO clocks. This configuration effectively leads to a very sharp RF filter in which the center frequency is controlled precisely by the LO and the bandwidth is controlled conveniently and set by the lower −3 dB cutoff frequency of the low-Q, low-frequency BPF (Fig. 1.7).

Because quadrature LO signals are available in the receive chain anyway, no additional LO phases are needed. The actual filter realization is shown in Fig. 1.8. The feedforward path resembles a linear time-invariant system in which the impulse response is: $h_{RF}(t) = h(t) * \cos(\omega_{LO}t)$, where $h(t)$ is the original BPF impulse response. In other words, the baseband frequency high-pass response is translated to $\pm f_{LO}$ to create a stopband or notch response.

The feedthrough of the LO signals to the receiver input is a potential concern similar to the mixer-first approach. In this case, however, due to much better isolation between

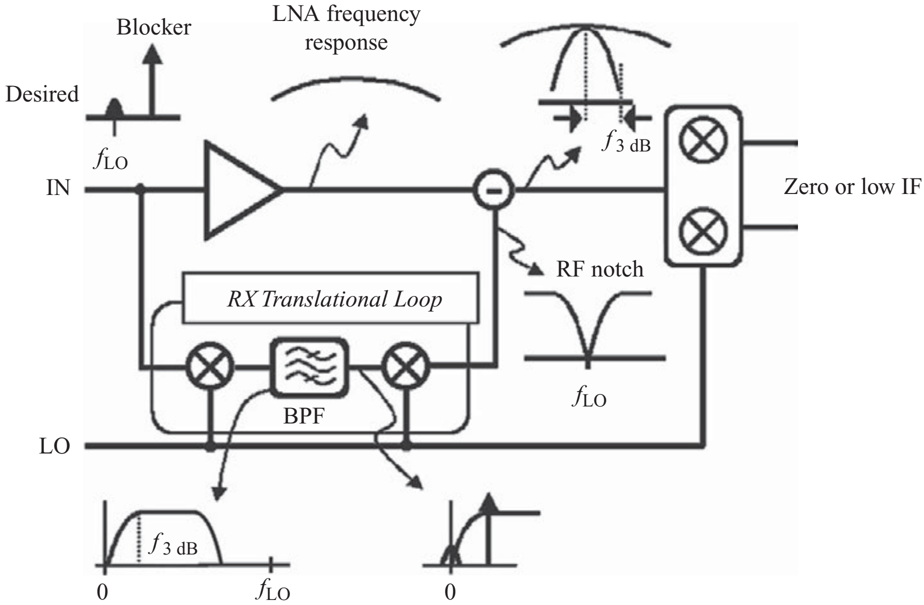

Figure 1.7 Feedforward blocker-cancellation concept. © 2010 IEEE. Reprinted, with permission, from [53].

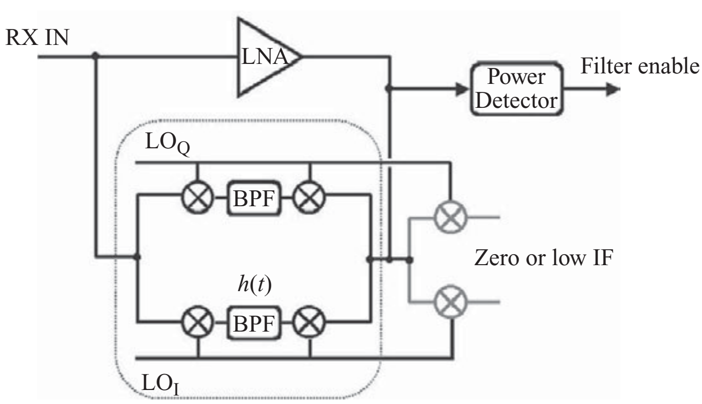

Figure 1.8 Actual realization of feedforward filter. © 2010 IEEE. Reprinted, with permission, from [53].

the LO ports and the input, the LO feedthrough is inherently low and can be lowered adequately by performing careful and symmetric layout techniques.

Despite a better NF of 4.2 dB achieved here [15] compared to the LNA-less approach, this scheme suffers from two issues:

- The lack of any input filtering imposes relatively challenging linearity on the LNA input devices, thus compromising the NF for the linearity.

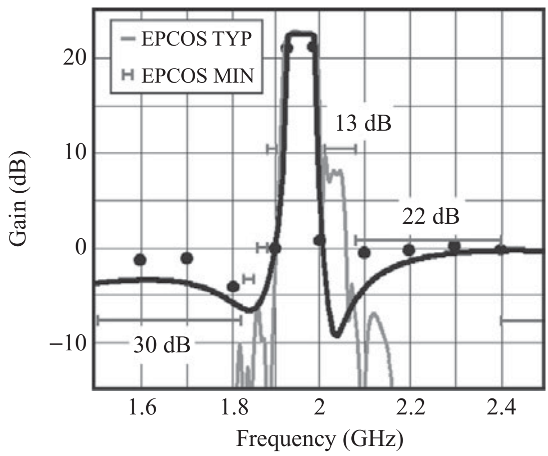

Figure 1.9 Feedforward filter response compared to EPCOS SAW. © 2010 IEEE. Reprinted, with permission, from [53].

- The inevitable phase and gain mismatches between the main and feedforward paths limit the amount of filtering to a low of about 20 dB, as shown in Fig. 1.9. Although this value is marginally adequate and comparable to the worst case for commercial external SAW filters, a higher rejection would be helpful through tighter control of gain and phase matchings between the two paths. The latter issue can be alleviated by using *adaptation techniques* to improve matching at the expense of some complexity.

1.3.3 *N*-phase filtering

The N-path filtering concept was introduced as early as 1960 in [20] and was used in switched capacitor (SC) filters in the early 1980s [21]. Consider the low-pass SC filter of Fig. 1.10 in which a sampling frequency of f_C is applied. As imposed by the Nyquist limit (NY), the maximum allowable frequency is $f_C/2$, leaving only one replica of the low-pass response. Now assume that there are N replicas of the same filter, each turning on at $1/N$th of the clock cycle in a periodic manner. It can be shown [21] that now the Nyquist limit is extended by N times, allowing other replicas of the filter response, particularly the one at f_C, to be extracted. Inspired by the N-path filtering concept, a class of very useful and intriguing filters called N-phase filters can be synthesized. These filters are also called M-phase filters, and in this book, we adopt this terminology. These filters and all variations thereof are the subject of most chapters of this book.

An example of a four-phase realization is shown in Fig. 1.11, which is composed of four baseband impedances and four switches driven by nonoverlapped clock phases with 25% duty-cycle clocks. It can be shown [22] that the input impedance seen from the RF side is roughly equal to the baseband impedance Z_{BB} frequency-translated to $\pm f_{LO}$ in series with the switch resistance, R_{SW}, where f_{LO} is the four-phase clock frequency

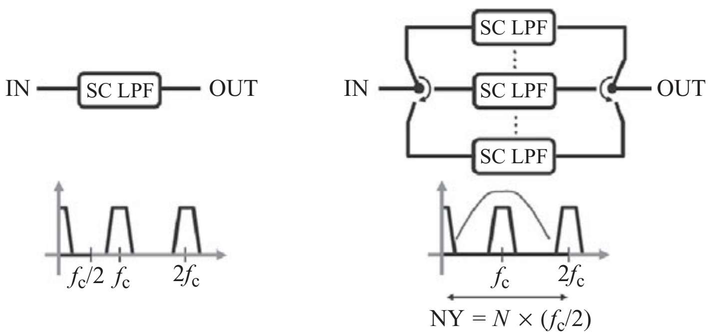

Figure 1.10 N-path filtering concept. © 2010 IEEE. Reprinted, with permission, from [53].

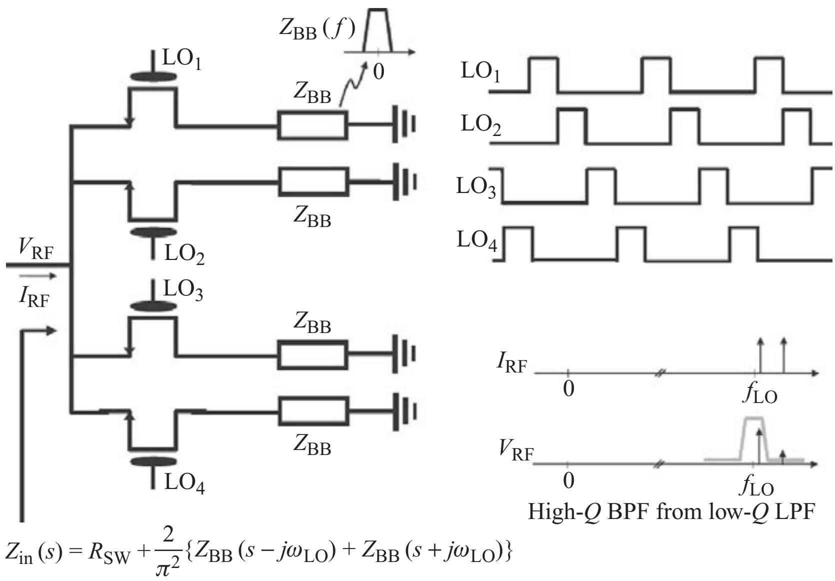

Figure 1.11 Implementation of four-phase filter. © 2010 IEEE. Reprinted, with permission, from [23].

applied to the switches:

$$Z_{in}(s) \cong R_{SW} + \frac{2}{\pi^2}\{Z_{BB}(s - j\omega_{LO}) + Z_{BB}(s + j\omega_{LO})\} \tag{1.1}$$

To extend the far-out flat region and improve the stopband rejection, it is critical to minimize the switch resistance with respect to the total RF impedance attached to the input of the filter. For typical values of switch sizes and RF impedance, a stopband rejection of up to 20 dB is feasible in one stage of filtering.

Shown in Fig. 1.12 is the actual circuit implementation of the differential four-phase filter, in which the baseband impedance Z_{BB} is simply reduced to a capacitor C_{BB}.