A "Divide and Conquer" Technique for Implementing Wide Dynamic Range Continuous-Time Filters

Yorgos Palaskas, *Member, IEEE*, Yannis Tsividis, *Fellow, IEEE*, Vladimir Prodanov, and Vito Boccuzzi

Abstract—This paper presents a technique for implementing analog filters with wide dynamic range and low power dissipation and chip area. The desired dynamic range of the filter is divided into subranges, each covered by a different filtering path optimized specifically for this subrange. This results in small admittance levels for the individual filtering paths and correspondingly small power dissipation and chip area. The system provides undisturbed output during range switching, contrary to conventional automatic gain control (AGC)/filter arrangements that generate disturbances every time the gain of the AGC changes. We also report on a low-noise highly linear CMOS transconductor useful for high-frequency applications. A chip implementing the ideas of this paper was fabricated in a 0.25-μm digital CMOS process. The intended application of the filter is channel selection in an 802.11a/Hiperlan2 Wireless Ethernet receiver. The chip dissipates 9 mA, occupies an area of 0.7 mm^2, and maintains a signal/(noise + IM3 distortion) ratio of at least 33 dB over a 48-dB signal range, with good blocker immunity. This performance represents at least an order of magnitude improvement over existing channel selection filters, even those that do not achieve disturbance-free operation.

Index Terms—Analog filters, automatic gain control (AGC), channel selection filters, companding.

I. Introduction

ANALOG filters are often used in wireless receivers for channel selection, that is, to reject strong interfering channels (blockers) before analog-to-digital (A/D) conversion. This relaxes the resolution requirements of the A/D converter, which would otherwise have to digitize both the desired channel and the strong blockers, and results in reduced power dissipation and chip area for the overall receiver. Analog filters, however, require large power dissipation and chip area on their own, especially when they exhibit high selectivity, high frequency of operation, and/or large dynamic range [1], [2]. The designer typically has limited control over the choice of selectivity and operating frequency, since these are dictated by system-level decisions. The chip area and power dissipation can be reduced by relaxing the dynamic range requirements of the filter. The scheme of Fig. 1 is sometimes used to this end. Here, an automatic gain control (AGC) circuit is included before the filter to ensure that the strength of the signal presented to the latter is op-
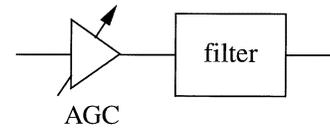
Fig. 1. Conventional AGC-filter arrangement.

timal at all times, i.e., as high above noise as possible, but still below the overload level. This allows the filter to have a relatively high noise level, resulting in small power dissipation and chip area.

Due to the dynamical nature of filtering, whenever the AGC gain changes disturbances are generated at the output of the AGC-filter combination of Fig. 1. These disturbances follow the natural response of the filter, and thus have frequency components within the passband, which is exactly where the desired output signal also resides. These in-band disturbances might affect the decodability of the received signal since they cannot be removed, e.g., with subsequent filtering. The duration of the disturbance, that is, the time it takes for the filter to settle to the new gain setting, is directly proportional to the quality factors of the filter. Therefore, filters with high selectivity, like those typically used for channel selection in wireless receivers, would be more prone to this problem. The settling problem is also expected to be more pronounced if the AGC has discrete gain steps, which is usually the case in CMOS implementations.

The disturbances of the previous paragraph would be innocuous if the gain adjustment were done during the "preamble" of the data frame, if one is available (the preamble is a time slot where no useful information is transmitted and is used for clock synchronization, AGC settling, etc.). However, certain wireless standards do not allow for a preamble. Also, there are cases where the AGC-filter settling problem discussed above would be an issue even if a preamble is used. For example, when a blocker suddenly appears at the input of the system, a new gain adjustment is required mid-frame, in which case the data packet might have to be dropped and retransmitted [3]. Another practical case where a preamble might not be adequate is when the strength of the desired channel itself varies significantly within a data frame. Finally, the duration of the preamble might not be enough for settling of the AGC-filter combination.

Most wireless receivers in the literature are evaluated under "static" conditions, that is, with fixed levels of input signals. This is the reason why the effects described in the previous two paragraphs have received limited attention. In this paper, we propose a technique that allows for uninterrupted disturbance-free operation of the filter over the entire time, while still maintaining the power and chip area advantages of the AGC-filter scheme of Fig. 1.
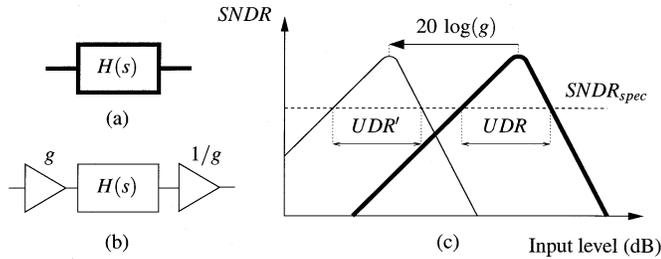
Fig. 2. Inserting an arbitrary filter between two amplifiers with gains g and $1/g$, and its effect on SNDR and UDR.

Sections II and III describe the proposed technique and how this can be used in channel selection filters. Section IV describes the implementation of a prototype filter used to demonstrate the technique. This section also presents a highly linear high-frequency transconductor and a low-power common-mode stabilization scheme. Finally, Section V reports measured results from the fabricated system.

II. "DIVIDE AND CONQUER" TECHNIQUE

In the "divide and conquer" technique [4], gain switching transients are avoided by using multiple filtering paths that operate all the time. Each of these filtering paths covers a fraction of the desired dynamic range and, therefore, requires much less power dissipation and chip area than what a single filter covering the whole dynamic range would need.

In order to understand the technique, we first define the signal-to-noise-plus-distortion ratio (SNDR) of the filter as the desired output signal power divided by the sum of the corresponding noise and distortion powers. Consider first a single filter path as in Fig. 2(a). The thick line in Fig. 2(c) shows how the SNDR varies with the input level. For weak inputs, the SNDR increases with a slope of $+1$ dB/dB, since the desired output power increases with $+1$ dB/dB, while the noise is constant. For strong inputs, the SNDR decreases with a slope of -2 dB/dB, since the desired output power increases with $+1$ dB/dB, while the distortion power (assumed to be dominated by third-order intermodulation, IM3) increases with $+3$ dB/dB. Typical filtering applications require that the SNDR at the output of the filter be higher than a minimum specified value $\text{SNDR}_{\text{spec}}$. The range of signals over which the SNDR requirement is met is defined as the usable dynamic range (UDR) of the filter. The definitions of $\text{SNDR}_{\text{spec}}$ and UDR are shown in Fig. 2(c).

Assume now that the filter of Fig. 2(a) is inserted between two amplifiers with gains g and $1/g$ as shown in Fig. 2(b). Note that this arrangement preserves the transfer function of the system. The output $1/g$ amplifier might not be needed in certain applications, in which case it can be removed without altering the results that will be presented below (see also the discussion near the end of Section III). As was mentioned in Section I, filters with high selectivity tend to generate considerable noise and distortion. We can then neglect the noise and distortion of the input and output amplifiers for simplicity, in which case the SNDR plot of the modified system of Fig. 2(b) will simply be a shifted version of the original SNDR plot, as shown by the thin line in
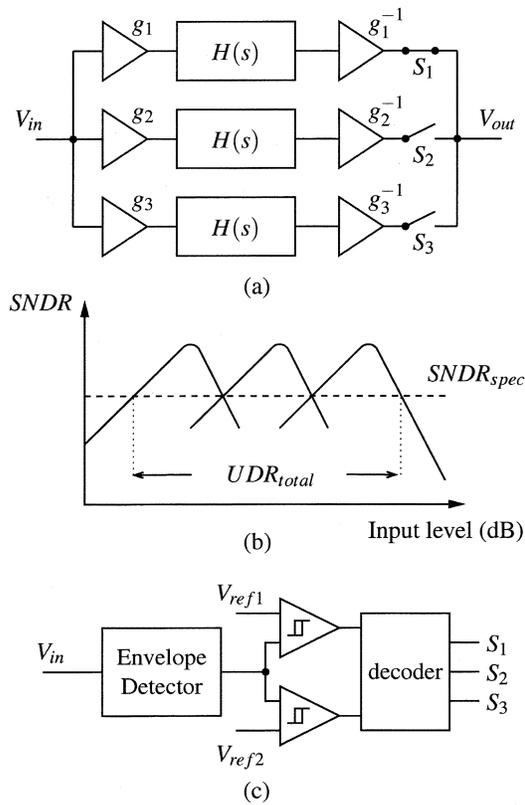

Fig. 3. The "divide and conquer" technique schematically.

Fig. 2(c). We observe that, under our assumption that the amplifiers are ideal, the UDR of the modified system will be the same as that of the original one.

Consider now the configuration of Fig. 3(a). Here, three identical filters are inserted between amplifiers with appropriate gains. The input is applied to all three filtering paths, and all filters operate continuously. By suitably selecting the gains we can have the SNDR plots shown in Fig. 3(b). Observe that the required UDR is split between the three filters. For any given input level, one of the filters is operating at an appropriate SNDR level, and this filter is selected by closing the corresponding switch at the output. For example, if the input signal is large then a path with small input gain is chosen in order to avoid overloading the core filter. On the other hand, if the input is small, then a path with a large input gain is chosen, for good SNR. The decision as to which filter is the optimal one can be made using the circuit of Fig. 3(c) which detects the strength of the input signal and compares it with the boundaries of the three ranges.

If the input level changes, then a different filter takes over by closing the appropriate switch. This, however, will not introduce undesired transients at the output, since the gains of the individual filtering paths are fixed and, therefore, no settling of the filters is required. The overall output simply switches between the outputs of two of the filtering paths, which in fact are equal since all filtering paths have the same transfer function.

We next consider the power dissipation and total capacitance of the system in Fig. 3(a). Assume, as an example, that the desired UDR of the system is 45 dB. Each of the three filters in Fig. 3(a) would then need a UDR of $45/3 = 15$ dB. Let the power dissipation of each of these filters be P, and the
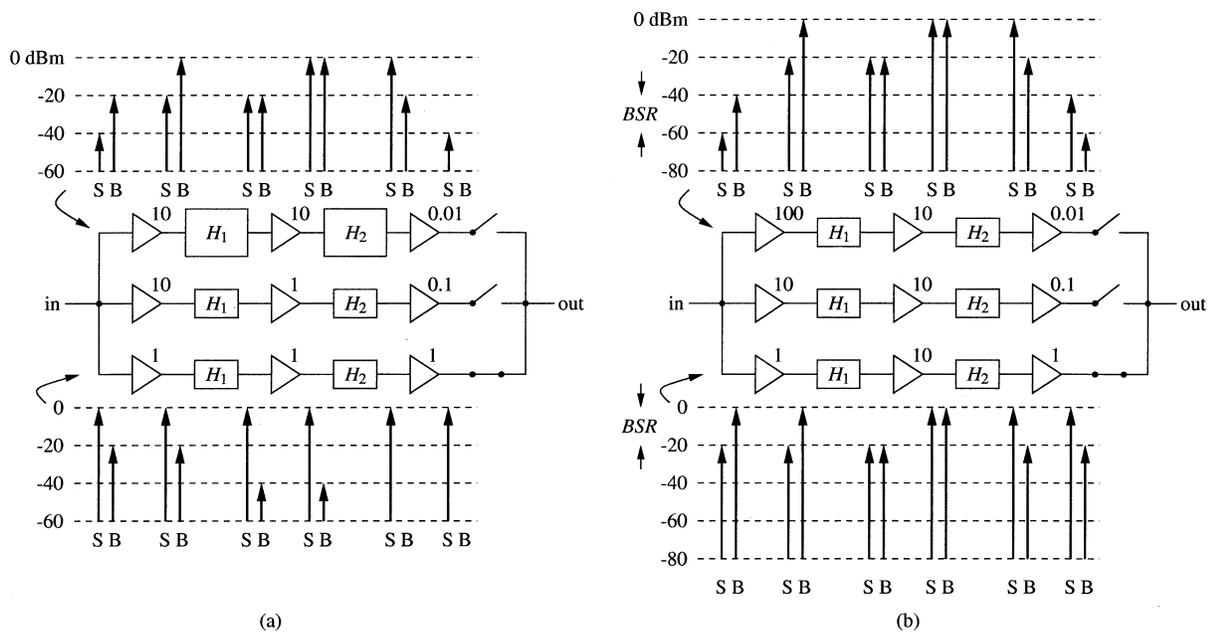
Fig. 4. (a) The scheme of Fig. 3(a) modified for the *fixed blocker* case. (b) The same scheme for the *fixed BSR* case. S stands for signal and B stands for blocker.

corresponding capacitance C. If a conventional filter were used to cover the desired UDR of 45 dB, this filter would need to have 30 dB larger UDR than each of the filters in Fig. 3(a). This would require 1000 times smaller noise power, which could be achieved at the expense of 1000 times more capacitance and power dissipation [1], [2]. Therefore, a conventional filter would have a power dissipation of $1000P$ and a total capacitance of $1000C$, as opposed to $3P$ and $3C$, respectively, using the scheme of Fig. 3(a). This amounts to about 333 times improvement in power dissipation and capacitance (and thus chip area, assuming the latter is dominated by capacitance). Note that the power dissipation of the input/output amplifiers and control circuit [Fig. 3(c)] has been neglected in the above calculation. Therefore, a smaller savings ratio is expected in a real implementation; nevertheless, saving ratios of one or two orders of magnitude are possible.

Instead of three, more filters can be used in the scheme of Fig. 3(a), which would result in larger power and chip area savings. Reference [4], however, shows that increasing the number of filters results in "diminishing returns," while the more frequent switching between the filtering paths would eventually lead to worse performance due to nonidealities such as switching transients, imperfect matching between the filters, etc. A small number of filters, typically three, is usually the best choice. The technique of Fig. 3 is in a sense similar to the "divide and conquer" algorithms used in computer programming, e.g., the fast Fourier transform. In both cases, a difficult problem is broken into many smaller ones that are much easier to solve, thus saving power/chip area in one case, and computation time in the other, although we use many more filters, or solve many more problems.

In practice, one should make the UDRs of the individual filtering paths overlap somewhat. This overlap is needed to introduce a hysteresis effect in the operation of the system, in order to prevent it from "hopping" between two filtering paths because of the ripple of the envelope detector, or when the input amplitude is near the transition region between two subranges.

When the input is large, paths with large input gains will be overloaded. This is not a problem since these paths will only be used when the input becomes small again. As a precaution one could use limiting devices across the nodes of these filters to ensure that they will never severely overload, in which case it could take some time for them to recover when needed.

III. "DIVIDE AND CONQUER" IN CHANNEL SELECTION FILTERS

Some modifications of the basic scheme of Fig. 3 are required when implementing channel selection filters in which large blockers might be present at the input of the system. Before presenting the corresponding design considerations, some clarifications are in order. Typical wireless standards, e.g., 802.11a, specify the blocker levels when the desired channel is at (or near) the reference sensitivity power. It is not specified, however, what the blocker levels should be when the power of the desired channel is above reference sensitivity. Two possibilities have been considered: (1) the blocker power is assumed to be fixed at the level specified at reference sensitivity and (2) the blocker power is allowed to follow the desired signal power with a fixed blocker to signal ratio (BSR). We will be referring to these two cases as *fixed blocker* and *fixed BSR* case, respectively.

Fig. 4(a) shows how the system of Fig. 3(a) should be adapted for the *fixed blocker* case. As seen in Fig. 4(a), each filtering path is implemented as filtering sections interleaved with amplification. Specific gain values, power levels, etc., have been used in Fig. 4 to keep the discussion as simple and intuitive as possible; the actual design of each filtering path, e.g., gain allocation, can be performed using the analytical results of [5]. As seen in Fig. 4(a), the three filtering paths have different amplifications *within* the filter. For example, when the desired signal is weak [top part of Fig. 4(a)], the input is amplified before being applied to the first filtering section $H_1(s)$. The blocker at the output of $H_1(s)$ is shown attenuated because of the filtering action of $H_1(s)$. It is then allowable to amplify before filtering section $H_2(s)$, and so on. The bottom part of Fig. 4(a) shows

the corresponding situation for a strong desired signal, accompanied by the same blocker as before. In this case, it is not allowable to amplify anywhere within the filter since this would lead to overloading due to the desired signal. A careful analysis of the system of Fig. 4(a) [6] reveals that the upper filtering path (the one that is optimized for small desired signals) has much tighter noise requirements than the lower paths, since the signal level is limited well below the blocker, and the noise must be kept well below the signal level. The noise of this filter can be made small by scaling all its transconductances and capacitances by the same factor. This is shown schematically in Fig. 4(a) by depicting the filtering sections of the upper path larger in size.

Fig. 4(b) shows the situation for the *fixed BSR* case. Observe that now the blocker at the input is always BSR dB stronger than the desired signal (compare the left and right parts of the figure). Because the BSR is the same for all power levels, the amplification within the filter is the same for all filtering paths, and the only difference between them is in the input and output amplifiers. In this case the noise requirements of the three filtering paths are identical, since the signal presented at the input of the $H_1(s)$ filtering section is the same for all of them. Therefore, identical filters can be used in all paths. In the rest of this paper we have adopted the worst case scenario of a fixed BSR in order to extend the applicability of the results as much as possible, and also for simplicity in the implementation.

As seen in Fig. 4(b) all filtering paths include amplification within the filter to take advantage of the attenuation of the blocker along the filter. If the strength of the input is detected and used to determine the optimal path, this amplification might lead to overloading due to the desired signal when the blocker is small or not present at all (our assumption that the blocker is BSR dB stronger than the desired channel refers to the worst case). This, for example, will be the case in Fig. 4(b) if the desired signal is at -20 dBm and the blocker is absent; the system would then select the middle path which would result in overloading of H_2. This problem can be resolved by detecting the strength at the *output* of the system: the input can be large due to an in-band signal or a blocker; at the output the blocker has been rejected eliminating this ambiguity. Note that the system should be designed with enough headroom to accommodate the blocker, when the latter is present.

When the scheme of Fig. 4(a) or (b) is used to implement a channel selection filter, the output amplifiers might not be needed and they can be removed. This reduces the resolution requirements of the A/D converter that typically follows the filter in a wireless receiver. The schemes of Fig. 4 then perform a coarse AGC/filtering operation. A fine AGC operation can still be performed before the A/D converter if desired. The coarse AGC action of the schemes of Fig. 4 then reduces the required gain variation of the fine AGC circuit, which results in smaller complexity and power dissipation for the latter. If the output amplifiers are removed from Fig. 4 the resulting systems are no longer linear time invariant (LTI). Whether this is a problem or not depends on the details of the demodulation algorithm used in the receiver. If desired, the LTI behavior of the system can be restored by implementing the output amplification in the digital domain, that is, after the A/D converter, an arrangement that does not compromise the desirable coarse AGC action of the system. Note that in this case the output of the filter would not be continuous. This does not increase the speed requirements of
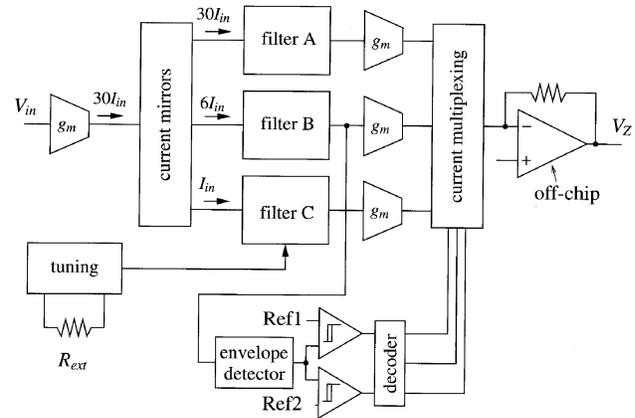

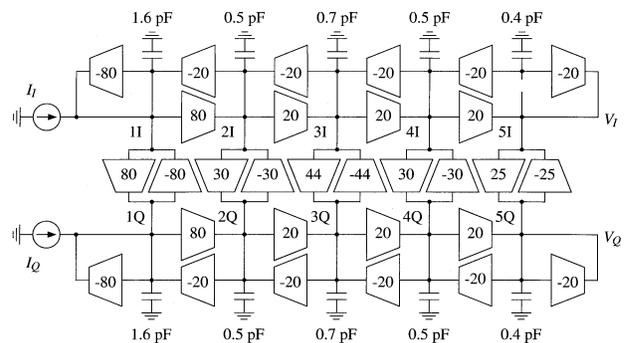Fig. 5.   Simplified block diagram of the fabricated system.

Fig. 6.   Complex filter implementation. All transconductances are in μA/V.

the A/D since the discontinuity does not carry useful information—only the values before and after the discontinuity matter as long as the A/D can properly sample those values.

IV. PROTOTYPE FILTER DESIGN

In this section, we discuss the implementation of a system demonstrating some of the ideas presented above. The system has been designed to perform partial channel selection in a 802.11a/Hiperlan2 low-IF [7] receiver. Due to space limitations, we restrict our attention to the most demanding data rate of 54 Mb/s. Based on the overall receiver planning (e.g., noise figure and gain allocation), it was found that the filter should maintain an SNDR$_{spec}$ of 33 dB over a UDR of at least 45 dB. The blockers were specified to be $+4$ dBc and $+23$ dBc for the adjacent and alternate channels, respectively, with respect to the desired channel power.

Fig. 5 shows a simplified block diagram of the fabricated system. Weak signals are processed by filter A, which has the largest input gain and, therefore, the best noise performance, medium strength signals are processed by filter B, and strong signals are processed by filter C. Each of the core filters was designed to implement a complex (or image-reject or polyphase) [7] bandpass transfer function with center frequency $+10$ MHz and bandwidth 18 MHz. Fig. 6 shows the G_m-C implementation of one of the core filters [8].

This design assumes that the blocker power is allowed to increase proportionally with the power of the desired channel as the latter is raised from reference sensitivity (fixed BSR scenario). As explained in Fig. 4(b), the optimal performance in
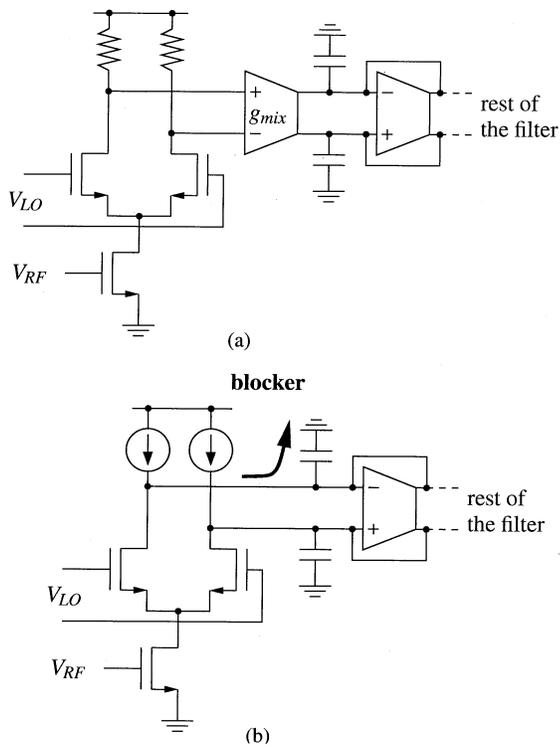
Fig. 7. (a) Conventional arrangement used to interface the mixer with the channel selection filter in a receiver. (b) Arrangement used in this design: the mixer current is directly fed to a current-input filter. In (b), the blocker current is shunted to ground through the input capacitors before it gets a chance to develop a voltage and generate distortion.

this case is obtained by using the same amplification within the core filters. The allocation of gain within the core filters was performed using the analytical optimization algorithm of [5]. Note that no explicit amplifiers are needed since the same function can be performed by progressively increasing the impedance level within the filter (smaller capacitors near the output). The core filter that resulted from the optimization was scaled to achieve the desired SNDR$_{\text{spec}}$ of 33 dB over at least $45/3 = 15$ dB of UDR. An additional 7-dB margin was added to the UDR of each subfilter to allow incorporating overlap/hysteresis between the individual subranges, and to account for offsets, inaccurate MOS models, etc. The small capacitance values of the filters (see Fig. 6) are made possible thanks to the use of the divide-and-conquer technique, which allows the filter noise floor to be rather high. In order to avoid overloading the filter when a blocker is *not* present at the input (see Section III), the decision as to which is the optimal path is taken by detecting the envelope at the *output* of filter B.

A simple envelope detector and comparators with hysteresis were included in the fabricated prototype. The envelope detector is the simple version of the circuit presented in [9]. The comparator consists of two competing current sources, one corresponding to the envelope and the other to a reference current. Hysteresis was introduced by making the reference current depend on the present state of the comparator.

The three core filters of Fig. 5 are current driven to take advantage of the attenuation of the blocker at the input node [10]. This is explained in Fig. 7(a) and (b), which depict a simple single-balanced Gilbert mixer driving a voltage-input and a
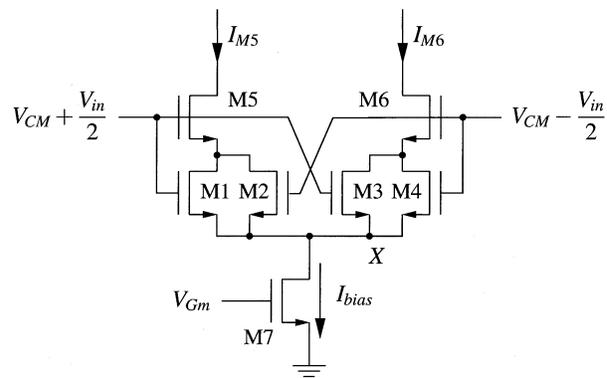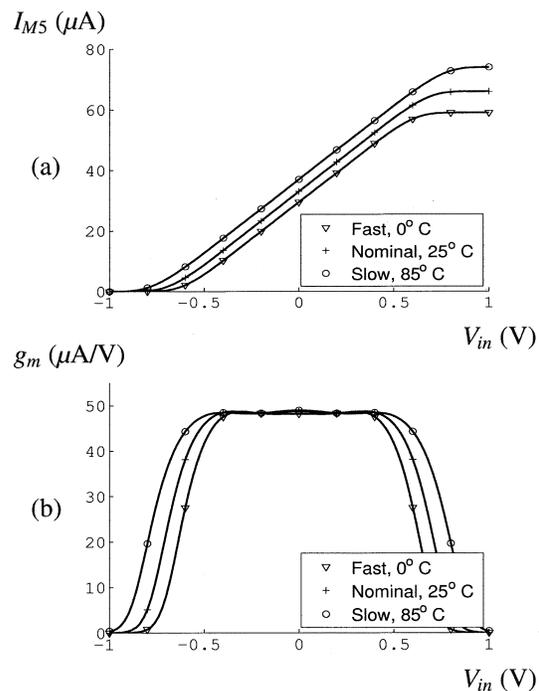

Fig. 8. Transconductor used in the filter [10].

Fig. 9. (a) Simulated single-ended output current of the transconductor as a function of the differential input voltage V_{in}. (b) Simulated g_m of the transconductor versus V_{in}.

current-input filter, respectively. In the conventional arrangement of Fig. 7(a) a large blocker will develop a large voltage at the input of transconductor g_{mix} resulting in distortion. In the current-input case of Fig. 7(b), on the other hand, the blocker current is shunted to ground through the input capacitor. The blocker will, therefore, be attenuated somewhat before it gets a chance to develop a voltage and generate distortion.

The input currents for the three filtering paths, with the appropriate scaling factors, are derived from a current mirror as shown in Fig. 5. In an actual receiver, the input of the current mirror would be provided by a Gilbert-type mixer. Since high-frequency mixers typically suffer from significant nonidealities due to the small sizes of the transistors, we chose to use a simple G_m stage to provide the input current. In this way, it is possible to evaluate the performance of the filter alone, without the mixer obscuring the results. The outputs of the three filters are multiplexed in the current domain and driven to an off-chip transimpedance amplifier.
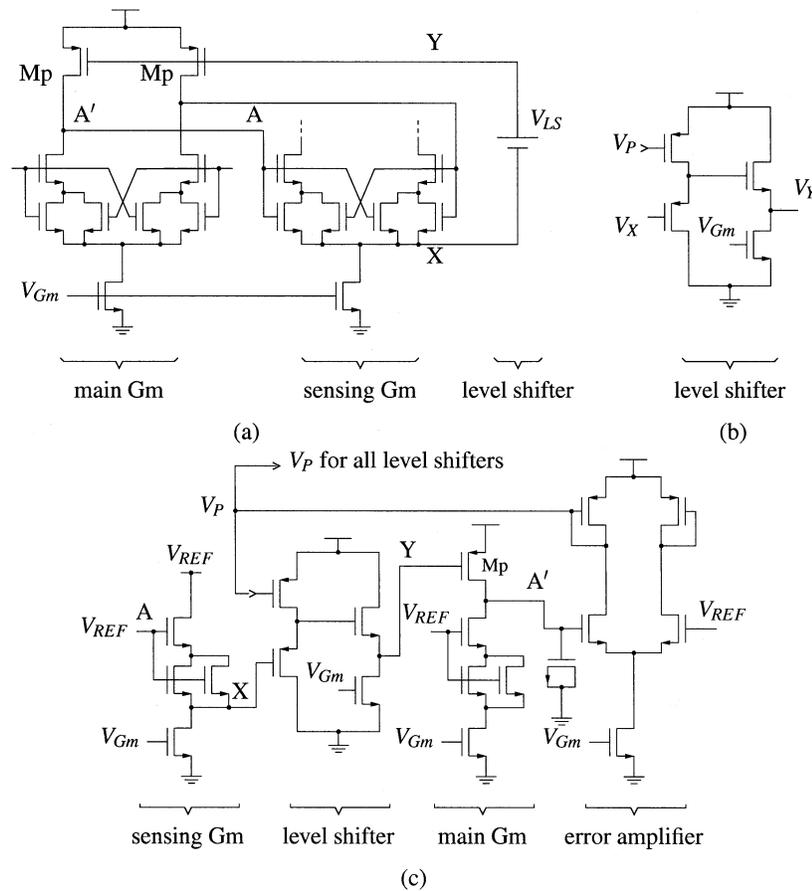
Fig. 10. (a) CMFB arrangement used in the design. (b) Level shifter. (c) Replica biasing circuit.

The integrating capacitors (nMOS transistors in inversion) were very stable in the process used (3σ variation of $\pm 5\%$ over process). This allowed us to perform automatic tuning of the center frequency by simply "locking" a replica G_m to a stable external reference resistance [11] and using the resulting bias voltage to bias the transconductors of the filters. No Q tuning was necessary because of the large dc gain and good high-frequency behavior of the transconductors used in the design (see the next section).

A. Transconductor

The transconductor used in the fabricated filter is shown in Fig. 8 [10]. The linearization mechanism in Fig. 8 is related to that in [12]. Transistors M5 and M6 operate as source followers, forcing the drain–source voltages of M1–M4 to approximately follow the input voltage. When V_{in} is small, transistors M1–M4 are in triode and behave as linear resistors, thus linearizing the $I_{\text{out}} - V_{\text{in}}$ characteristic through source degeneration. As V_{in} is increased from zero, the effective g_m of the transconductor initially starts rolling off as would be the case with any source degenerated differential pair. As V_{in} is increased further, transistor M2 is driven into saturation. With appropriate sizing and biasing of the transistors this effect can give a "bump" in the g_m curve resulting in an almost equiripple $g_m - V_{\text{in}}$ characteristic, as shown in Fig. 9.

The advantage of the G_m stage of Fig. 8 is that, contrary to [12], the noise of the tail current source is common-mode and cancels in the differential output current. This point is very

important for circuits operating with low supply voltages. If the noise of M7 in Fig. 8 did not cancel, one would have to bias M7 with a small g_m to achieve low noise in the first place. But a small g_m requires a significant gate overdrive, thus reducing the available headroom. One could argue that the V_{DS} voltages of M1–M4 also reduce the available headroom. This, however, is not a significant problem since a moderate V_{DS} of about 200 mV can result in the very good linearity performance of Fig. 9. Due to the cancellation of the noise from M7, the transconductor of Fig. 8 achieves a low excess noise factor of 2.7, including the contribution from a properly designed pMOS load.

Due to its simplicity, the transconductor of Fig. 8 exhibits very good high-frequency performance. The input impedance is mainly capacitive, and can easily be incorporated in the integrating capacitors of the filter. Large devices can then be used for good matching and large dc gain. The $I_{\text{out}} - V_{\text{in}}$ characteristic has a parasitic pole/zero pair due to the resistive degeneration. It can be shown that the pole and the zero frequencies track (to a first-order approximation) the desired time constants of the filter, resulting in a very robust frequency response for the overall filter.

B. Common-Mode Feedback

The common-mode feedback (CMFB) circuit used in the fabricated chip is shown in Fig. 10(a). The common-mode (CM) voltage at the output of the main G_m is first sensed at the common-source node X of an identical sensing G_m. Transconductors that are already part of the filter are used for
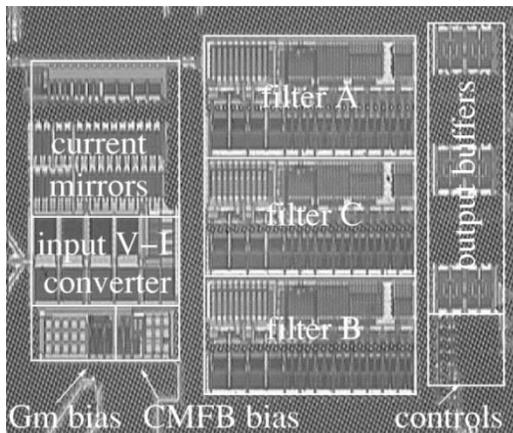
Fig. 11.   Microphotograph of the fabricated chip.

Fig. 12.   Setup used to test the fabricated chip.

CM sensing, as opposed to using dedicated G_m's just for this purpose [13], [14]. The detected CM voltage V_X is fed back to the gates of transistors Mp in Fig. 10(a) through a unity-gain buffer/level shifter V_{LS}. The level shifter was implemented as a cascade of a pMOS and an nMOS source followers as shown in Fig. 10(b). This was done because the desired level shift V_{LS} could not be achieved with a single source follower, either an nMOS or a pMOS.

The CMFB loop of Fig. 10(a) has a dominant pole corresponding to node A′, which is the only high impedance point in the loop and, furthermore, the grounded integrating capacitors are connected there. The CMFB loop is, therefore, inherently fast and stable, and thus requires small power dissipation (about 10 μA per node in the fabricated system). This is in contrast to conventional CMFB circuits where the loop is closed through a high gain amplifier that compares the output CM voltage with a desired CM reference voltage [15]. In the latter case, the extra pole introduced by the CMFB amplifier significantly complicates the stabilization of the loop and typically results in significant power dissipation. The dc loop gain of the CMFB loop of Fig. 10(a) is $g_{Mp}r_o$, where g_{Mp} is the g_m of transistor Mp and r_o is the output resistance at node A′. In our case this loop gain was in the order of 40 dB. This amount of CMFB loop gain might be somewhat smaller than what could be obtained using a conventional CMFB scheme, but it is more than enough as far as CM stabilization is concerned. Other CMFB circuits that use unity-gain blocks instead of high gain amplifiers in the feedback path have been presented in [16]–[18].

The CMFB circuit of Fig. 10(a) sets the CM level to a stable value, which would, however, drift due to process and temperature variations if a fixed V_{LS} had been used. Fig. 10(c) shows a scheme that adjusts V_{LS} so that the output CM voltage equals a certain desired value V_{REF} in spite of process and temperature variations. This scheme is an adaptation of the one used in [10] and is conceptually similar to the replica biasing CMFB circuits presented in [17] and [18].

To understand the operation of the circuit of Fig. 10(c), let us first consider the CMFB loop of Fig. 10(a). Assume that the loop is broken between the output of the main and the input of the sensing G_m [nodes A′ and A, respectively, in Fig. 10(a)], and that the desired CM voltage V_{REF} is applied at the inputs of the sensing G_m. If the level shift V_{LS} has the right value then the return voltage at the output of the main G_m should
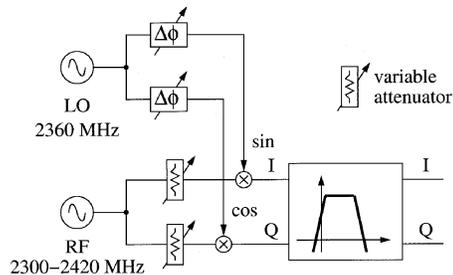
also be V_{REF}. The biasing circuit of Fig. 10(c) contains a CM replica of the circuit of Fig. 10(a). The return voltage at node A′ in Fig. 10(c) is compared with the desired reference V_{REF} and the error is amplified and fed back to adjust the level shift V_{LS} until the return voltage becomes V_{REF}. The adjustment of V_{LS} is achieved by changing the bias voltage V_P of the level shifter. The resulting bias voltage V_P is used to bias *all* level shifters in the filter [Fig. 10(b)].

V. TESTING PROCEDURES AND MEASURED RESULTS

In this section, we present frequency and time domain measurements from the fabricated prototype. A microphotograph of the chip is shown in Fig. 11, where the reader can identify the three filters, A, B, and C, along with the other circuits composing the system (see Fig. 5).

The setup used to test the fabricated system is shown in Fig. 12. This is essentially the front-end of a full low-IF receiver. The RF signal, situated around 2.36 GHz, is downconverted in quadrature with two mixers. The local oscillator (LO) inputs of the mixers are derived by phase shifting a 2.36-GHz LO. As explained in [7], mismatches between the I and Q paths of a complex filter can result in the image signal leaking into the passband. In order to accurately assess the image leak of the actual filter, the I and Q inputs provided by the test setup should be matched to a very high degree of accuracy; otherwise, mismatches in the test setup can cancel/mask mismatches in the actual filter and make the latter appear much better/worse in terms of image leak than it really is. Using variable attenuators and phase shifters, we were able to achieve an I/Q imbalance at the input of the filter of less than ± 0.02 dB/± 0.1 degrees over the image frequency band -20 to 0 MHz. This level of accuracy is adequate if the image leak of the filter is in the order of -50 dB (see below).

Frequency Response: Fig. 13(a) shows the amplitude response of the three filtering paths, both I and Q. The passband gains of the individual paths are seen to differ by about 15 dB. Apart from this intended difference, all six curves of Fig. 13(a) are seen to be very well matched with each other. This is more evident in Fig. 13(b) where we plot the error between the amplitudes responses with filter B as the reference (Q output only). The amplitude error is seen to be smaller than ± 0.2 dB across the entire passband (1–19 MHz).

Fig. 14(a) shows the phases of the I and Q outputs of the three filtering paths. The phases of the I and Q outputs differ by 90° as expected for a complex filter. It is also seen in Fig. 14(a) that the phases of all I outputs match very well with each other, and similarly for the Q outputs. Fig. 14(b) shows the phase error between the Q outputs of the three filters. The maximum phase
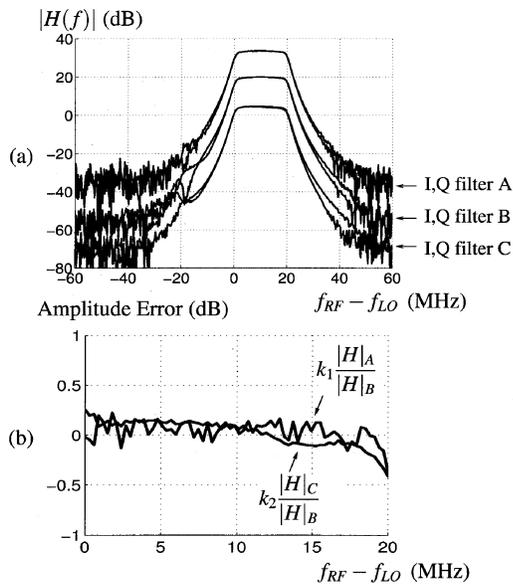
Fig. 13. (a) Amplitude response and (b) amplitude errors between the three filtering paths.
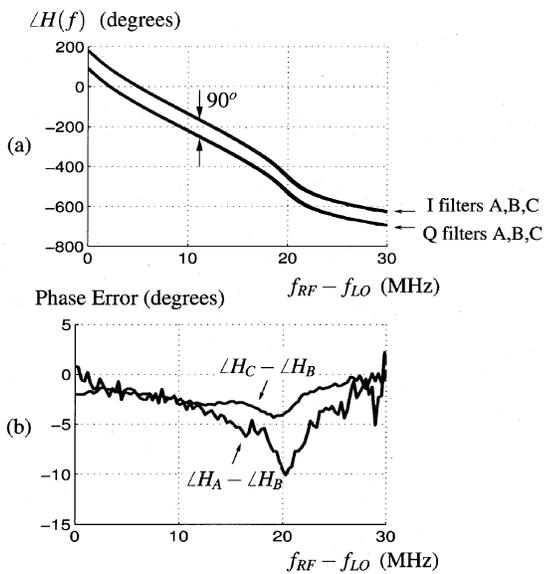

Fig. 14. (a) Phase response and (b) phase errors between the three filtering paths.

error in the passband is about $6°$, which corresponds to about $\pm 0.5\%$ group delay error.

The amplitude and phase errors of Figs. 13(b) and 14(b), respectively, represent undesired deviations from the ideal behavior. Whether the measured errors are acceptable for a particular application or not can only be determined at the system level by evaluating the resulting performance degradation, e.g., of the overall receiver bit-error rate (BER). The mismatch between the three filters could have been improved by splitting each filter in two subfilters and laying them out in a common centroid fashion ABCCBA, where, e.g., A denotes the two half-filters composing filter A.

From Fig. 13(a), the image leak of the filters can be estimated to be better than -54 dB [6]. This level of image leak was achieved through very careful layout of the core filter.
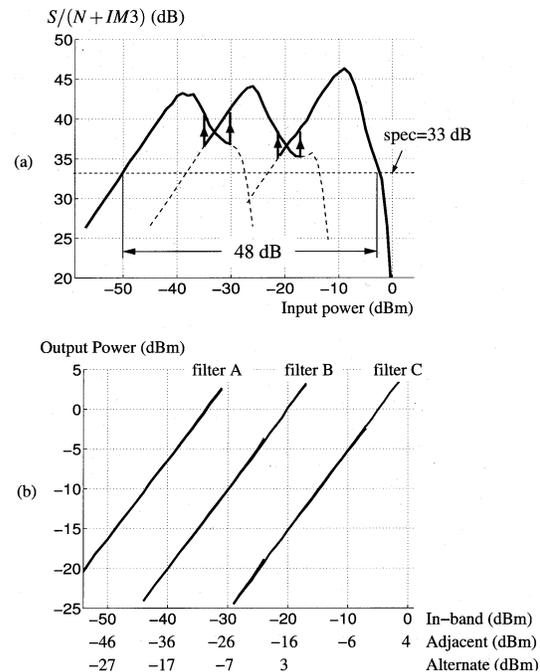

Fig. 15. (a) Measured SNDR of the fabricated system. (b) Output of the three filters with and without blockers.

Dynamic Range: Fig. 15(a) shows the SNDR plots for the three filtering paths (dashed lines), and for the overall system of Fig. 5 (solid line). The distortion is IM3, with the two input tones at $+10$ and $+11$ MHz. Fig. 15(a) clearly shows the hysteresis effect implemented in the system. The SNDR is higher than the minimum required 33 dB over a 48-dB range of input signals.

Fig. 15(b) shows the output power when the input consists of the desired signal ($+10$ MHz) alone, the desired signal together with the adjacent channel blocker (-10, $+30$ MHz), and the desired signal together with the alternate channel blocker (-30, $+50$ MHz)—a total of five superimposed curves per filter. The power of the alternate channel blocker was set 23 dB higher than the power of the desired signal, and 4 dB higher for the adjacent channel blocker. No measurements are reported when the blocker power exceeds the maximum specified in-band power of -1 dBm. This is a realistic assumption since the maximum in-band power corresponds to the linearity limit of the circuits that precede the filter in the receiver, e.g., the downconversion mixer. Fig. 15(b) only reports measurements that fall within the range where each filter is used. It is seen that the presence of the blockers causes negligible compression to the desired signal, which is exactly the behavior the system was designed for.

Transient Effects: Fig. 16(a) shows the output V_Z of the system (see Fig. 5) when a 10-MHz sinusoid with a slowly varying envelope is applied at the input. The coarse AGC action of the system can clearly be seen. For example, at $t \simeq 3.6$ μs, the input signal (not shown) starts becoming too small, which would lead to insufficient SNR. A different filtering path with a higher input gain and, thus, a higher SNR, is automatically selected at $t \simeq 3.6$ μs, and the output envelope in Fig. 16(a) is seen to increase accordingly. Similarly, at $t \simeq 15$ μs the input starts becoming too large and a path with a smaller input gain is selected. Note that the amplitude where the switching occurs is different in the two cases due to hysteresis effect incorporated
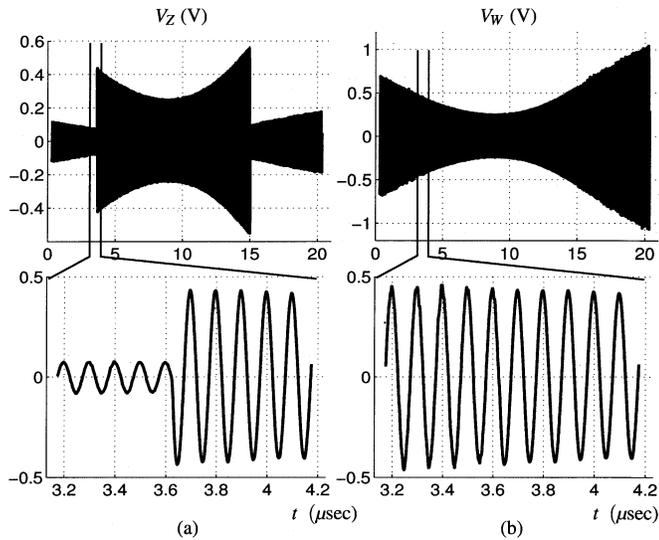
Fig. 16. (a) Output of the fabricated system V_Z (see Fig. 5) in response to a 10-MHz sinusoid with slowly varying envelope. (b) Signal of (a) after sampling and amplification by the inverse of the input gain. The input applied to the system (not shown) was similar to the waveform in (b).
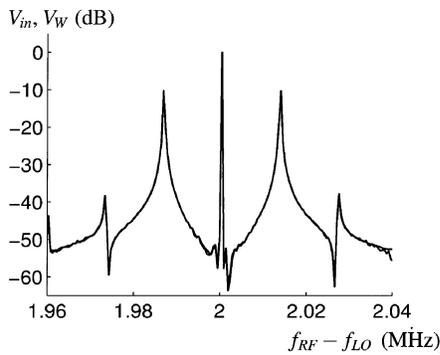

Fig. 17. Spectra: V_{in} is a 2-MHz input with a slowly varying envelope and V_W is the resulting output after amplification by the inverse of the input gain. The two spectra practically coincide for at least 45 dB relative to the peak.

TABLE I
PERFORMANCE SUMMARY FOR THE FABRICATED IMAGE REJECTION
FILTER/COARSE AGC PROTOTYPE

Technology	0.25 μm Digital CMOS
Chip Area (excl. in/out buffers)	0.7 mm^2
Supply Voltage	2.5 V $\pm$ 10 %
Filter Type (I+Q)	10th order Complex
	0.1 dB Chebychev
Center Frequency	+10 MHz
Bandwidth	18 MHz
Adjacent Channel Attenuation	38 dB
Passband Transresistance	25 / 5 / 0.8 kΩ
Power Dissipation	23 mW
$SNDR_{spec}$	33 dB
Usable Dynamic Range	48 dB
Adjacent Channel Blocker	+4 dBc
Alternate Channel Blocker	+23 dBc
Image Leakage	-54 dB

TABLE II
COMPARISON WITH OTHER RECENTLY PUBLISHED FILTERS

Reference	Power (mW)	Ord.	Cutoff (MHz)	*DR* (dB)	FoM (aJ)	Norm. FoM
Yoshizawa JSSC 02	11.6	7	1.92	74.5	31	65
Itakura JSSC 99	31	4	0.2	88.5	55	116
Andreani JSSC 02	7.4	14	3.5	67.5	27	57
Andreani JSSC 02	6	18	3.5	62.5	54	114
Khorramabadi ISSCC 96	10.5	7	0.62	80.5	22	46
Jussila ESSCIRC 00	11.6	5	2.1	82.8 [a]	6	12
Behbahani JSSC 00	119 [b]	8	15	82.8 [a,b]	5.2	11
This work	22.5	10	19	84	0.47	1

[a] Noise at max gain; distortion at min gain. [b] Two filters being turned on/off depending on blocker – assume blocker is present 50% of time.

in the design. The waveform of Fig. 16(a) has been sampled and amplified by the inverse of the input gain, after subtracting the offsets which are mostly due to the output buffers. The result is shown in Fig. 16(b). No disturbance is observed in the resulting waveform, as expected from the discussion in association with Fig. 3.

The experiment of Fig. 16 was repeated to obtain the corresponding spectra. Due to limitations in our test setup (insufficient resolution at high sampling rates/large sampling intervals) we had to use a smaller input frequency in the frequency-domain experiment of Fig. 17 than in the time-domain experiment of Fig. 16 (2 and 10 MHz, respectively). Fig. 17 shows the spectra of the 2-MHz input V_{in}, and the signal after the $1/g$ amplification V_W. These spectra were obtained with postprocessing in Matlab. The two spectra in Fig. 17 practically coincide for at least 45 dB relative to the peak, demonstrating that the fabricated system together with digital domain amplification (see the last paragraph of Section III) indeed behaves like an LTI system achieving disturbance-free operation.

Summary and Comparison: Table I summarizes the measured performance of the fabricated coarse AGC/complex filter

prototype. We next compare the performance of the fabricated prototype with that of other recently published channel selection filters (Table II). The figure of merit (FoM) used as the basis for the comparison is defined as [19]

$$\text{FoM} = \frac{\text{Power Dissipation}}{(\text{No. of poles}) \times (\text{Upper 3-dB Frequency}) \times (\text{DR})}. \quad (1)$$

The DR is the ratio of the maximum signal (defined as causing 33 dB of signal/IM3-distortion) to noise. The DR is defined with respect to in-band signals only. This FoM does some injustice to filters with high quality factors [1], to bandpass/complex filters, and to filters with good blocking performance; our design has all these properties, so the comparison is conservative. Also, some of the filters reported in Table II include variable gain amplifiers at the input or within the filter. The FoM reported in Table II refers to the best case noise (gain at maximum) *and* the best case distortion (gain at minimum) of these filters, so again the comparison is conservative. It is seen that still the fabricated

prototype is at least one order of magnitude more power efficient than all filters in Table II, even those that do not provide disturbance-free operation.

VI. CONCLUSION

We have presented a simple CMOS-compatible technique for implementing analog filters with wide dynamic range and low power dissipation and chip area. Multiple filtering paths with small power and chip area are used, each optimized for portion of the desired total dynamic range. Depending on the strength of the signals processed by the system, one of these paths is automatically selected and connected to the output. Since the subfilters operate continuously, no settling transients are observed when switching to a different filtering path, contrary to the case of conventional AGC-filter combinations. We have discussed the applicability of this scheme in the case of channel selection filters. We have presented a highly linear low-noise high-frequency transconductor circuit, as well as a power-efficient common-mode stabilization scheme that utilizes unity-gain buffers. A 9-mA 0.7-mm^2 1–19-MHz complex filter has been implemented in a 0.25-μm digital CMOS process to demonstrate the ideas in this work. The fabricated prototype maintains a signal/(noise + IM3 distortion) of at least 33 dB over a 48-dB range of signals with good blocker immunity.

REFERENCES

[1] Y. Palaskas and Y. Tsividis, "Power-area-DR-frequency-selectivity tradeoffs in weakly nonlinear active filters," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 1, Bangkok, Thailand, 2003, pp. I453–I456.

[2] E. Vittoz, "Future of analog in the VLSI environment," in *Proc. IEEE Int. Symp. Circuits and Systems*, New Orleans, LA, 1990, pp. 1372–1375.

[3] F. Behbahani, A. Karimi-Sanjaani, W. Tan, A. Roithmeier, J. Leete, K. Hoshino, and A. Abidi, "Adaptive analog IF signal processor for a wide-band CMOS wireless receiver," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1205–1217, Aug. 2001.

[4] G. Palaskas and Y. Tsividis, "A "divide and conquer" technique for the design of wide dynamic range continuous-time filters," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 1, Sydney, Australia, May 2001, pp. 252–255.

[5] Y. Palaskas and Y. Tsividis, "Dynamic range optimization of weakly nonlinear, fully balanced, Gm-C filters with power dissipation constraints," *IEEE Trans. Circuits Syst. II*, vol. 50, pp. 714–727, Oct. 2003.

[6] G. Palaskas, "Syllabic companding filters with emphasis on wireless applications," Ph.D. dissertation, Columbia Univ., New York, NY, 2003.

[7] J. Crols and M. Steyaert, "Low-IF topologies for high performance analog front ends of fully integrated receivers," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 269–282, Mar. 1998.

[8] P. Andreani, S. Mattisson, and B. Essink, "A CMOS gm-C polyphase filter with high image band rejection," in *Proc. Eur. Solid-State Circuits Conf.*, Stockholm, Sweden, Sept. 2000, pp. 244–247.

[9] K. Koli and K. Halonen, "Low voltage MOS-transistor-only precision current peak detector with signal independent discharge time constant," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 3, Hong Kong, June 1997, pp. 1992–1995.

[10] V. Prodanov, G. Palaskas, J. Glas, and V. Boccuzzi, "A CMOS AGC-less IF strip for Bluetooth," in *Proc. Eur. Solid-State Circuits Conf.*, Villach, Austria, Sept. 2001, pp. 488–491.

[11] C. Laber and P. Gray, "A 20-MHz sixth-order BiCMOS parasitic-insensitive continuous-time filter and second-order equalizer optimized for disk-drive read channels," *IEEE J. Solid-State Circuits*, vol. 28, pp. 462–470, Apr. 1993.

[12] F. Krummenacher and N. Joel, "A 4 MHz CMOS continuous time filter with on-chip automatic tuning," *IEEE J. Solid-State Circuits*, vol. 23, pp. 750–757, June 1988.

[13] J. Silva-Martinez, M. Steyaert, and W. Sansen, "A 10.7-MHz 68-dB SNR CMOS continuous-time filter with on-chip automatic tuning," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1843–1853, Dec. 1992.

[14] T. Itakura, T. Ueno, H. Tanimoto, A. Yasuda, R. Fujimoto, T. Arai, and H. Kokatsu, "A 2.7-V, 200-kHz, 49-dBm, stopband-IIP3, low-noise, fully balanced gm-C filter 1C," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1155–1159, Aug. 1999.

[15] B. Razavi, *Design of Analog CMOS Integrated Circuits*. New York: McGraw-Hill, 2000.

[16] J. O. Voorman, A. van Bezooijen, and N. Ramahlo, "On balanced integrator filters," in *Integrated Continuous-Time Filters*, Y. Tsividis and J. O. Voorman, Eds. New York, NY: IEEE Press, 1993.

[17] B. S. Song and P. Gray, "Switched-capacitor high-Q bandpass filters for IF applications," *IEEE J. Solid-State Circuits*, vol. 21, pp. 924–933, Dec. 1986.

[18] H. Khorramabadi and P. Gray, "High-frequency CMOS continuous time filters," *IEEE J. Solid-State Circuits*, vol. 19, pp. 939–948, Dec. 1984.

[19] A. Yoshizawa and Y. Tsividis, "Anti-blocker design techniques for MOSFET-C filters for direct conversion receivers," *IEEE J. Solid-State Circuits*, vol. 37, pp. 357–364, Mar. 2002.

Yorgos Palaskas (S'98–M'02) was born in Greece in 1973. He received the Diploma in electrical and computer engineering from the National Technical University of Athens, Greece, in 1996, and the M.S. and Ph.D. degrees, both in electrical engineering, from Columbia University, New York, NY, in 1999 and 2002, respectively. His Ph.D. work was in the area of syllabic companding filters for wireless applications.

During the summer of 1999, he was with Texas Instruments, Inc., where he worked on disk drive electronics. During the summers of 2000 and 2001, he was with Agere Systems (formerly Bell Labs), doing research on integrated IF filters. Since January 2003, he has been with the Communications and Interconnect Technology Laboratory, Intel Corporation, Hillsboro, OR, where he is working on wireless transceiver architectures.
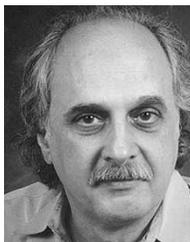
Yannis Tsividis (S'71–M'74–SM'75–F'86) received the B.S. degree from the University of Minnesota, Minneapolis, in 1972 and the M.S. and Ph.D. degrees from the University of California at Berkeley in 1973 and 1976, respectively.

He is the Charles Batchelor Professor of Electrical Engineering at Columbia University, New York, NY. He has worked at Motorola Semiconductor and AT&T Bell Laboratories, and has taught at the University of California at Berkeley, the Massachusetts Institute of Technology, Cambridge, and the National Technical University of Athens, Greece. His latest book is *Operation and Modeling of the MOS Transistor*, 2nd ed. (Oxford, U.K.: Oxford Univ. Press, 2003).

Dr. Tsividis is the recipient of the 1984 IEEE Baker Best Paper Award, the 1986 European Solid-State Circuits Conference Best Paper Award, and the 1998 IEEE Circuits and Systems Society Guilllemin-Cauer Best Paper Award. He is a corecipient of the 1987 IEEE Circuits and Systems Society Darlington Best Paper Award, and received the Presidential Teaching Award from Columbia University in 2003.

Vladimir Prodanov received the B.S. degree from the Technical University, Sofia, Bulgaria, and the M.S. and Ph.D. degrees from the State University of New York at Stony Brook in 1991, 1995, and 1997, respectively.

He was with Bell Labs—Research, Lucent Technology, from 1997 to 2001. He is currently with Agere Systems (formerly Lucent Microelectronics), Allentown, PA. He is also an Adjunct Professor at Columbia University, New York, NY, where he teaches analog and RF circuit design.

Vito Boccuzzi was born in 1954. He received the B.S. degree in electrical engineering from the New Jersey Institute of Technology, Newark, in 1980 and the M.S. degree from Adelphi University, Garden City, NY, in 1990.

From 1980 to 1996, he was with AIL Systems (Airborne Instrumentation Laboratory) working in the instrumentation division as an Analog Designer for synthesized signal generators and spectrum analyzers and then in the radar division where he was involved in the design of receivers for airborne doppler radars. In 1996, he joined Bell Laboratories, Lucent Technologies (now Agere Systems), Allentown, PA, as a Member of Technical Staff. He is currently involved with the Communication Circuit Research laboratory in the design and test of prototype analog ICs for wireless.