




THOMAS H. LEE

The Design of CMOS Radio-Frequency Integrated Circuits

Second Edition

CAMBRIDGE

THE DESIGN OF CMOS RADIO-FREQUENCY INTEGRATED CIRCUITS

SECOND EDITION

THOMAS H. LEE

Stanford University



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 1998
© Thomas H. Lee 2004

This book is in copyright. Subject to statutory exception and
to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 1998
Second edition first published 2004

Printed in the United States of America

Typeface Times 10.75/13.5 and Futura System AMS-T_EX [FH]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication data
Lee, Thomas H., 1959–

The design of CMOS radio-frequency integrated circuits / Thomas H. Lee. – 2nd ed.
p. cm.

ISBN 0-521-83539-9

1. Metal oxide semiconductors, Complementary – Design and construction. 2. Radio
frequency integrated circuits – Design and construction. 3. Radio – Transmitter-receivers.
I. Title.

TK7871.99.M44L44 2004
621.39'732 – dc22

2003061322

ISBN 0 521 83539 9 hardback




FIGURE 19.4. Image-reject mixer.

in the IF chain. Furthermore, the receiver is now susceptible to interference from sources within the relatively large bandwidth of the front end, so linearity requirements become more severe (perhaps much more so). Choosing to use upconversion therefore depends on whether the increased susceptibility to interferers is acceptable in exchange for improved image rejection and simplified LO design.

19.2.5 DUAL CONVERSION

We've noted that a low intermediate frequency is demanding of preselector and LO performance, but places no great demands on IF filtering. We've also noted that a high IF increases the difficulty of realizing the channel-selective IF filters, but mitigates the image rejection problem and therefore relaxes requirements on the preselector. A high IF also happens to simplify LO design owing to the reduction in required oscillator tuning range.

A *dual-conversion* receiver uses two IFs in an effort to combine the benefits of both downconversion and upconversion architectures by using both techniques together. In a dual-conversion superheterodyne receiver, a first mixer produces a high IF to take care of the image rejection issue, while a second mixer and low IF ease the channel selection problem. Some receivers employ a third IF to provide even greater flexibility in the tradeoff. Most modern high-performance receivers implemented in discrete technologies use a dual-conversion architecture, with the frequency plan determined in large part by what high-quality filters are available at reasonable cost.

19.2.6 THE IMAGE-REJECT RECEIVER

The image-reject receiver employs a *complex mixer* that exploits the relationship between the desired signal and the undesired image to cancel the image during the mixing process. Thus, no preselector filter is needed (at least in principle), and design of the overall architecture can proceed with diminished concern for the image problem. In particular, a relatively low IF may be chosen to ease the requirements on IF filtering, analog-to-digital conversion, and subsequent baseband processing. See Figure 19.4.




FIGURE 19.5. Weaver architecture.

To perform this miracle, the RF signal feeds two mixers. Two LO signals, in quadrature, drive the other mixer ports. The desired difference frequency component is the same in both mixer outputs, but the image terms are in quadrature. In the classic implementation shown, a constant-gain, frequency-independent 90° phase shifter allows simple addition of the two signals to cancel the undesired image terms.

Because it is somewhat difficult to build broadband quadrature phase shifters, the alternative architecture (due to Weaver²) shown in Figure 19.5 is often attractive. In the Weaver architecture, a pair of quadrature mixing operations eliminates the need for the phase shifters. As long as the two pairs of LO signals are truly in quadrature and the gains in the two paths are perfectly matched, there will be perfect cancellation of the unwanted image.

Effect of Gain and Phase Errors

Complete image rejection in the preceding architectures depends on perfect phase quadrature and perfect gain matching. Of course, it is more realistic to assume that neither condition is satisfied in practice. The ultimate rejection of the image then depends on the levels of matching achieved. To quantify the matching requirement, first consider explicitly the effect of gain errors in the block diagram of Figure 19.4. Specifically, but without loss of generality, let the I and Q signals model all the error as follows:

$$I = B \cos(\omega_{LO}t), \quad (10)$$

$$Q = A \sin(\omega_{LO}t). \quad (11)$$

If the RF input signal is $\cos(\omega_{RF}t)$, then the output of the I mixer is

$$\frac{B}{2} [\cos(\omega_{RF}t + \omega_{LO}t) + \cos(\omega_{RF}t - \omega_{LO}t)] \quad (12)$$

² D. K. Weaver, "A Third Method of Generation and Detection of Single-Sideband Signals," *Proc. IRE*, December 1956, pp. 1703–5.

and the output of the Q mixer is

$$\frac{A}{2}[\sin(\omega_{\text{RF}}t + \omega_{\text{LO}}t) - \sin(\omega_{\text{RF}}t - \omega_{\text{LO}}t)]. \quad (13)$$

The low-pass filters remove the sum frequency component, while the 90° phase shifter converts (minus) sine to cosine.

The sensitivity of the difference frequency component in Eqn. 13 to the sign of the argument is ultimately the source of this mixer's image rejection property, since an RF signal above the LO frequency by an amount equal to the IF produces an IF signal with a sign opposite to that produced by mixing down the image. Phase shifting, and then summing with a cosine-mixed version of the LO, results in a reinforcement of the desired component and a simultaneous suppression of the image.

For an RF signal above the LO, the overall output may be expressed as

$$\frac{1}{2}[A \cos(\omega_{\text{IF}}t) + B \cos(\omega_{\text{IF}}t)]; \quad (14)$$

the image RF signal produces an output given by

$$\frac{1}{2}[-A \cos(\omega_{\text{IF}}t) + B \cos(\omega_{\text{IF}}t)]. \quad (15)$$

Let us define the *image rejection ratio* (IRR) as the power ratio of the desired output to the undesired output. Then

$$\text{IRR}_{\text{gain}} = \left[\frac{A+B}{A-B} \right]^2 = \left[\frac{1+B/A}{1-B/A} \right]^2 = \left[\frac{1+(1+\epsilon)}{1-(1+\epsilon)} \right]^2 \approx \frac{4}{\epsilon^2}; \quad (16)$$

in the last approximation, it is assumed that the gain error ϵ is much smaller than unity. As we shall see shortly, the gain error of any practical image-reject mixer *must* satisfy this inequality in order to meet realistic performance specifications. Hence, the approximation is guaranteed to be valid in all cases of practical interest.

One may carry out an analogous derivation assuming perfect gain matching but imperfect 90° phase shift. In that case, the IRR due to a departure $\Delta\phi$ (in radians) from quadrature may be expressed as

$$\text{IRR}_{\text{phase}} = 1 + 4(\cot \Delta\phi)^2 \approx \frac{4}{(\Delta\phi)^2}. \quad (17)$$

Again assuming small gain and phase errors, the net IRR may be computed using the sum of the errors given by Eqn. 16 and Eqn. 17:

$$\text{IRR}_{\text{tot}} \approx \frac{4}{(\Delta\phi)^2 + \epsilon^2}. \quad (18)$$

It is quite difficult to achieve much better than about a 0.1% gain error and a 1° phase error, particularly at high frequencies, without some form of calibration. These numbers correspond to an IRR of about 41 dB. More typically, image rejection is rarely significantly better than about 35 dB. Because most receivers require much larger (e.g. 80-dB) image rejection, the image-reject architecture alone rarely provides sufficient rejection. Autocalibration can help close the gap, but it is unlikely to




FIGURE 19.6. Passive analog quadrature generator.

provide an additional 40 dB of image rejection. Hence, additional filtering is usually required.

Aside from that limitation, an additional drawback of the image-reject mixer is that mixers are noisy and can be power-hungry, so adding another one can be unattractive from the perspective of power consumption and dynamic range.

Quadrature Generators

The image-reject mixer requires the ability to generate quadrature phase shifts. Although it isn't difficult to provide a 90° shift over a broad frequency range, it is much harder to maintain a constant amplitude response at the same time. In fact, no finite network can provide both a constant phase shift and a constant gain magnitude over an infinite frequency range, so one must settle for approximations that work well only over some limited range.

The RC – CR network shown in Figure 19.6 is a popular method for generating quadrature signals over a narrow frequency band. The phase shift is 90° at all frequencies, but the magnitude response is not constant.

The phase shift of the I branch is zero at DC, and asymptotically heads down toward -90° as the frequency approaches infinity. The phase shift of the Q branch starts off at $+90^\circ$ at DC, and heads downward toward zero with the same functional shape as the I branch. Hence, even though the phase shift of *each* RC branch is not constant, the 90° *difference* between them is.

Unfortunately, the amplitude of each output changes dramatically with frequency, since the I branch is low-pass in nature whereas the Q branch is a high-pass filter. The amplitudes of the two outputs are equal only at the pole frequency:

$$\omega = \frac{1}{RC}. \quad (19)$$

There is, of course, a 3-dB attenuation at this frequency, so one must accept this significant loss when using this network. Furthermore, this network is thermally noisy, and its contribution to overall noise figure must be kept under control.

The design of this quadrature generator is straightforward in principle: the resistor and capacitor values are chosen to set the pole frequency equal to the operating frequency. In so doing, the effect of parasitic loading must be taken into account. Even so, an amplitude mismatch is all but inevitable because of component tolerances. If necessary, one may restore amplitude equality with a variable gain amplifier in an




FIGURE 19.7. Digital quadrature generators.

AGC loop or with a limiter. Clearly, this type of compensation is effective only in generating one frequency at time, such as in producing a quadrature LO.

In such single-frequency applications, it may be appropriate to consider alternative quadrature generation methods that automatically provide equal-amplitude outputs. Quadrature oscillators, consisting of two integrators in a feedback loop, can inherently produce equal-amplitude sine waves in quadrature (see Chapter 17).

In those instances where square-wave signals are acceptable, one may use digital circuits to eliminate the amplitude matching problem. Ring oscillators, in which differential gain stages are used (in a two- or four-stage ring, for example), can also produce excellent quadrature outputs. Such an oscillator must be embedded within a PLL in order to set the frequency.

Other digital options that do not require a PLL involve clocked circuits (see Figure 19.7). The first circuit produces quadrature outputs whose frequency is one fourth that of the clock, so the clock frequency may have to be exceedingly high. One attribute of the circuit, though, is that it is insensitive to the duty cycle of the incoming clock, since state changes occur only on the rising edge.

The output of the second circuit is at one half the clock frequency, so this circuit demands less of the circuitry. However, since state changes occur on both rising and falling edges of the clock, the quality of the quadrature outputs is now sensitive to the duty cycle of the clock. If the clock and its inverse are generated with the crude method shown, further degradation results from the propagation delay of the inverter.

Various other phase-shifting methods have evolved to handle those cases where one must provide a quadrature relationship to a relatively broadband signal. A common example is the passive *RC* network shown in Figure 19.8.

The geometric mean of the *RC* frequencies is chosen equal to the desired center frequency. The detailed selection of the individual stage *RC* time constants to satisfy an arbitrary gain constancy constraint is quite involved, but an extremely crude rule of thumb often suffices for a first-cut design. Each stage provides reasonably constant




FIGURE 19.8. Two-stage broadband quadrature generator.

gain (within about 0.2 dB) over roughly a 10% bandwidth, so the required number of stages is determined by the bandwidth over which a constant gain is required. The two-stage design shown can provide relatively constant gain over approximately a $\pm 20\%$ bandwidth about the center frequency by staggering the two time constants. For example, to design for a 1-GHz center frequency, one might choose R_1C_1 to correspond to about 900 MHz, while R_2C_2 might be selected to correspond to about 1.1 GHz.³ A significant disadvantage of this network, however, is its attenuation and high noise.

This type of network provides multiple output phases and thus is known as a *polyphase filter*. It is exceptionally good at providing quadrature outputs yet has a versatility that extends well beyond that function, as we'll see shortly.

To understand first why the polyphase network can generate quadrature so well, consider redrawing the network as shown in Figure 19.9 (only one stage is shown for clarity). Suppose that the input signals are not quite in quadrature. As a specific example, suppose that Q_{in} is closer to I_{in} than it should be. If we imagine this imperfection as a slight clockwise rotation of Q_{in} and Q_{in-b} in the octagonal layout on the right, then we see that Q_{out} moves a little clockwise from where it should be, but so does I_{out} . This circuit thus improves quadrature as signals pass through it.⁴ Cascading such stages provides additional improvement. The ultimate limitations

³ This network is evidently due to M. J. Gingell, "Single Sideband Modulation Using Sequence Asymmetric Polyphase Networks," *Electrical Communication*, v. 48, 1973, pp. 21–5. Practical examples that function over a decade frequency range for SSB generation using Hartley's phasing method may also be found in various editions of the *ARRL Handbook for Radio Amateurs*, from about 1981 to 1992. See also British Patent #1,174,710 (filed 7 June 1968, granted 17 December 1969) as well as its U.S. counterpart, #3,559,042 (filed 19 May 1969, granted 26 January 1971).

⁴ As a consequence, the requirements for component matching are looser in the early stages than in the later ones.




FIGURE 19.9. Two equivalent ways to draw a single-stage polyphase filter.

are imposed by mismatch in the components and in the loads on the various output nodes. These loads include, of course, ever-present parasitics. As long as one accommodates all of these effects in the design of the networks, excellent quadrature is assured.⁵ For this reason, the polyphase network is extremely popular among hams for SSB generation using the phasing method of Hartley.⁶ There, a multistage cascade of filters provides good quadrature over the baseband bandwidth, and a separate polyphase filter is often used to generate good quadrature for the mixers.

The same properties that make it attractive for quadrature and SSB generation make the polyphase filter useful for image rejection. Just as SSB generation requires (in effect) the ability to distinguish between negative and positive frequency baseband components, image rejection requires the ability to distinguish between signals that are negatively and positively offset from a carrier. However, an ordinary filter by itself cannot make such a distinction; the magnitude response is a function of frequency alone. An ordinary image-reject filter may thus have to exhibit extraordinary steepness in order to function (as in Carson's SSB generation method).

Another way of using ordinary filters to provide image rejection is to decompose a signal into its separate orthogonal components through a quadrature downconversion, leaving us with two real signals. A pair of real filters may then process those in-phase and quadrature signals separately from that point on. This method is the

⁵ Additionally, one must always bear in mind the lossiness (and attendant noise) of these networks.

⁶ In discrete implementations, the Weaver method has not found wide use for SSB generation. Recall that, unlike the methods of Carson and Hartley, mixer offsets in a Weaver modulator produce an in-band tone. This artifact is generally much more objectionable than carrier leakage, and the cancelling of multiple offsets that is needed to eliminate it is more bother than most wish to bear. Consequently, the phasing method has shown enduring popularity among amateur radio enthusiasts. The polyphase filter's invention has largely taken care of what was once the most difficult part: designing the phase shifter.




FIGURE 19.10. Idealized computer mouse (only one of two wheel assemblies shown).

essence of the Weaver architecture, and fundamentally it seeks to implement a *complex filter* with a pair of real filters.

Another option is simply to implement a complex filter directly. Such a filter operates on two (or more) inputs simultaneously and generally produces two (or more) outputs. The most natural and common choice is to drive a complex filter with in-phase and quadrature components that represent the real and imaginary parts of a complex signal. Such a filter's magnitude response is a function not only of frequency but also of the phase relationships between the two inputs. This dual sensitivity means that such a filter potentially responds differently to positive and negative frequency components, allowing us to discriminate against (say) an image signal.

Conventional demonstrations of complex filters usually begin with the theory of complex signals – invoking Hilbert transforms, if not a host of trigonometric identities. The reader is encouraged to investigate that rigorous approach independently; it's good for you, like broccoli. Here, however, we pursue a somewhat different approach and study how a conventional computer mouse works.

A mouse?

Yes, and here's why. A mouse needs to respond not only to speed (frequency), but also to direction (the sign of that frequency); a mouse that cannot distinguish between upward and downward motions is not a very useful device. These are the same requirements that we must satisfy to generate SSB as well as to reject images. So, if we understand how a mouse works, we'll have developed the key insights we need to understand complex filters.

Construction details vary among manufacturers, but Figure 19.10 captures the essential features of a basic mechanical mouse (real mice have more teeth). In this imagined implementation of a mouse, optically opaque teeth are disposed around a wheel that are set up to interrupt intermittently the light paths of a pair of photoemitters and photodetectors as the wheel rotates. Thus, when one emitter–detector path




FIGURE 19.11. Polyphase image-reject filter.

is fully obscured (corresponding to a logic zero, say), the other is on the verge of unblocking. Depending on the direction of rotation, the I signal will lead or lag the Q signal. For example, if the wheel is rotating counterclockwise, the rising edge of I precedes that of Q by 90° . If the rotation is clockwise, then the rising edge of I lags that of Q by 90° . This arrangement thus produces a pair of output signals that tells us both the frequency and direction of rotation. These attributes are precisely those used by a polyphase filter.⁷

Observe that if we have access to only one of the outputs then the only information conveyed is the magnitude of the frequency (speed of mouse motion). Any subsequent processing of the mouse signal is therefore limited to operations that take frequency – and only frequency – as an input. However, with access to *both* I and Q we could build a filter that would, for instance, make the mouse respond only to motion in one direction. Disregarding the undesirability of such a mouse, we recognize the deeper principle: Knowledge of both I and Q enables us to perform the desired discrimination against, say, a negative-frequency “image.”

As a specific electronic example, consider the application of a polyphase filter displayed as Figure 19.11. This system first uses a pair of quadrature mixers to resolve an input signal into its in-phase and quadrature components. The polyphase filter then accepts *both* components as filter inputs and processes them as an ensemble.

⁷ You may see reference to terms such as “sequence asymmetry” and the like. The term *sequence* may be taken as synonymous with *frequency* for the most part. Strictly speaking, what is meant is the sequence of vectors encountered as the phasor rotates. Hence, a positive sequence might be defined as I , Q , I_b , and Q_b (corresponding to a counterclockwise rotation in our mouse diagram). *Asymmetry* refers to the possibility that the frequency response magnitude may be different for positive and negative frequencies (sequences). Finally, the term *polyphase* derives from the multiplicity of input or output phases that such a filter processes.




FIGURE 19.12. Idealized frequency response of polyphase image-reject filter.

As with the mouse example, the filter has the ability to treat negative and positive frequency components differently. In this particular case, the frequency response is indeed different for positive and negative frequencies, as seen in Figure 19.12.

We see that the polyphase filter used this way is actually a bandstop filter. By choosing the zero location(s) properly, we may suppress the negative-frequency components and thereby reject the undesired image. This property has been exploited to enable good performance from a low-IF (also known as a near-zero IF, or NZIF) architecture.⁸ As noted earlier, one may cascade several stages to obtain broader band rejection.

One impediment to doing so is that deriving the transfer functions by hand for polyphase filters of varying order is an unpleasant exercise. It's bad enough for the second order, but it's completely hopeless for the third order and above. It's so unpleasant, in fact, that we won't even leave it to the reader as a homework problem, tempting as that may be. Many published examples seem to have been designed by trial and error (at least, these publications are conspicuously silent about the specific methods used). Instead of making you (or us) derive them, we will simply provide explicit formulas that have been kindly derived elsewhere.⁹ In all that follows, the transfer functions are defined in terms of the complex representations of signals:

$$V_{in}(t) = I_{in}(t) + jQ_{in}(t), \quad (20)$$

$$V_{out}(t) = I_{out}(t) + jQ_{out}(t), \quad (21)$$

For the first-order filter, the transfer function is

$$G_1(j\omega) \equiv \frac{V_{out}(j\omega)}{V_{in}(j\omega)} = \frac{1 + \omega RC}{1 + j\omega RC}. \quad (22)$$

⁸ J. Crols and M. Steyaert, *CMOS Wireless Transceiver Design*, Kluwer, Dordrecht, 1998.

⁹ H. Kobayashi et al. cite the use of *Mathematica* to perform the actual derivation in their paper, "Explicit Transfer Function of RC Polyphase Filter for Wireless Transceiver Analog Front-End," *Proceedings of APASIC*, 2002.

The second-order filter's transfer function is a little more involved, but still tractable:

$$G_2(j\omega) = \frac{(1 + \omega R_1 C_1)(1 + \omega R_2 C_2)}{1 - \omega^2 R_1 C_1 R_2 C_2 + j\omega(R_1 C_1 + R_2 C_2 + 2R_1 C_2)}. \quad (23)$$

The transfer function for the third-order filter, on the other hand, is big enough that we can't even write it down in one go. We need to break it into pieces:

$$G_3(j\omega) = \frac{N(j\omega)}{D(j\omega)}, \quad (24)$$

where

$$N(j\omega) = (1 + \omega R_1 C_1)(1 + \omega R_2 C_2)(1 + \omega R_3 C_3). \quad (25)$$

The denominator is

$$D(j\omega) = D_R(j\omega) + jD_I(j\omega), \quad (26)$$

with

$$D_R(j\omega) = 1 - \omega^2[R_1 C_1 R_2 C_2 + R_2 C_2 R_3 C_3 + R_1 C_1 R_3 C_3 + 2R_1 C_3(R_2 C_1 + R_2 C_2 + R_3 C_2)] \quad (27)$$

and

$$D_I(j\omega) = \omega[R_1 C_1 + R_2 C_2 + R_3 C_3 + 2(R_1 C_2 + R_2 C_3 + R_1 C_3)] - \omega^3 R_1 C_1 R_2 C_2 R_3 C_3. \quad (28)$$

Note that all of these filters exhibit notches at negative frequencies, as in the first example. Observe further that the notch locations are simple RC products, where each RC is that of an individual section. These observations help to converge on a final design much more quickly by guiding the choice of an initial set of component values. It is common practice (but by no means essential or necessarily optimum) to choose all the capacitors equal in value and then to vary only the resistance values. As a further guide to designing broadband polyphase filters, choose the geometric mean of the zero frequencies equal to the center frequency of the filter. Then, vary the ratio of the zero frequencies and investigate the worst-case stopband rejection. If insufficient, reduce the ratio and add more cascaded sections. In general, you will need to supplement this work with an accommodation of component tolerances and mismatch. Monte Carlo simulation techniques work well for these relatively simple structures.

19.2.7 DIRECT CONVERSION

The upconversion and dual-conversion receiver architectures seek to solve the image rejection problem by using a high IF to displace the image enough to allow a simple filter to provide the necessary rejection. An alternative choice is to use a downconversion architecture in which the intermediate frequency is zero. Since a signal and its image are separated by twice the IF, a zero IF implies that the desired signal is its own image. Therefore, goes the argument, there is no image to reject and hence no




FIGURE 19.13. Direct-conversion receiver.

need for a front-end image-reject filter; front-end filtering requirements thus become especially easy to satisfy.¹⁰ Furthermore, with a zero IF, all subsequent baseband processing can take place at the lowest possible frequency. For example, relatively low-frequency analog-to-digital converters (ADCs) and digital-signal processing engines can implement the channel filter, as well as perform demodulation and ancillary housekeeping functions. Hence, no external filters are required (in principle, anyway). Furthermore, the flexibility inherent in digital approaches opens the possibility for a “universal” receiver, one that can accommodate many different standards with one piece of hardware; see Figure 19.13.

The most general direct-conversion receiver requires *two* mixers and LOs. The reason is that the phase of the LO with respect to the incoming RF signal is important. If the phases are coincident (or anticomcident), the demodulated signal is of maximum strength. If the phase relationship happens to be a quadrature one, the demodulated signal is zero.

In order to accommodate arbitrary phase relationships between the RF and LO signals, then, a mixer must be augmented with another one driven by an LO in quadrature with the first one. By combining the outputs of these two mixers, correct demodulation is possible with arbitrary input phase. Mismatch in the *I* and *Q* paths is not nearly as serious here as in an image-reject architecture because – rather than attempting to reject the image – this architecture actually *exploits* the “image.” The only effect of mismatches is some distortion.

With all of these attributes, it seems as if the direct-conversion receiver has no peer, especially for amenability to IC implementation. Indeed, the ubiquitous pager is often a direct-conversion receiver. The simple architecture (and simple signaling scheme) thus appears to permit a highly integrated, low-cost realization.

¹⁰ It is nonetheless generally good practice to use some type of filter in any case – to avoid, for example, out-of-band interferers overloading the front end.