Naval Research Laboratory
Washington, DC 20375-5320

NRL/5320/MR—2024/4

Wideband IQ Downconversion: Part I—Sample Rates, IF Frequencies, and Basic Architectures

J. O. COLEMAN (RETIRED)

Advanced Radar Systems Branch
Radar Division

May 30, 2024

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
30-05-2024	NRL Memorandum Report	October 1, 2012	September 30, 2014

4. TITLE AND SUBTITLE

Wideband IQ Downconversion: Part I—Sample Rates, IF Frequencies, and Basic Architectures

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
		1E88

6. AUTHOR(S)

J. O. Coleman*

7. PERFORMING ORGANIZATION / AFFILIATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Ave SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/5320/MR—2024/4

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR / MONITOR'S ACRONYM(S) NUMBER	11. SPONSOR / MONITOR'S REPORT NUMBER(S)
Office of Naval Research One Liberty Center 875 N. Randolph Street, Suite 1425 Arlington, VA 22203-1995	ONR	

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTAL NOTES

*Retired

14. ABSTRACT

This tutorial explores high-level digital-signal-processing (DSP) architectures for wideband IQ demodulators through ten detailed design examples. A fair bit of basic multirate DSP is reviewed to keep the discussion self contained. While this report presents itself as Part I, no Part II was ever written, so comments on Part II can be viewed as suggestions for further work.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	c. THIS PAGE		
U	U	U	SAR	42

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Greg Tavik	(202) 404-1945

This page intentionally left blank.

CONTENTS

This page intentionally left blank.

WIDEBAND IQ DOWNCONVERSION: PART I—SAMPLE RATES, IF FREQUENCIES, AND BASIC ARCHITECTURES

1. INTRODUCTION

This Part I document, together with the hoped-for eventual Part II, is structured as a tutorial on the design of high-performance IQ demodulators with focus on the DSP. Tutorials require examples, and the examples here, in addition to supporting the tutorial, are intended to provide a basic collection of immediately usable demodulator designs, a collection of designs representing a meaningful portion of the parameter space. In the course of developing the concepts behind these sorts of designs, a fair bit of basic multirate digital signal processing (DSP) is reviewed in the interests of keeping the whole thing self contained, though that review was never a specific goal of the effort.

This document's intended audience is nonspecialists who are willing to work to change that status. It's aimed in particular at (1) RF- and microwave-system designers who are familiar with DSP basics but who need to push a bit deeper in this one area in order to understand what sorts of IQ demodulators are feasible and (2) DSP or other signal-processing engineers who find themselves supporting a detailed IQ-demodulator design effort even though they have had little or no prior opportunity to work with bandpass sampling, multirate processing, or IQ demodulation.

Each of this document's example designs outputs a complex baseband signal—an IQ signal with its real and imaginary parts termed I (in-phase) and Q (quadrature) components—in discrete time containing the signal content of an entire input IF or RF band of interest. Consistent with this "entire band" notion, these designs are nontunable. A subsequent system aimed at extracting one or more subbands of particular interest may well involve tuning, but these designs do not. Each of the designs here is wideband at its input in that the input signal band's width is a substantial fraction of its center frequency. Each of the designs is wideband at its output in that the output sample rate exceeds that signal-band width relatively modestly: from less than 5% to fully 25% with an average of just over 11%. All of the designs are based on bandpass sampling.

Every IQ demodulator discussed here formally includes analog antialiasing filtering, an A/D converter, and DSP. The focus here is on the DSP. The choice of the A/D converter, though crucially important in any actual system, is unambiguously outside the scope of this document. The analog filtering is crucial, but the reader is assumed to either have basic analog (probably RF) filter expertise or to have ready access to it.

In spite of its antialiasing function, the analog filtering is everywhere here termed the "IF filter" in respect of the traditional role of this type of demodulator. Further, "IF filter" here actually means the entirety of the pre-A/D filtering, all referred to the A/D input. In most cases its net frequency response will be dominated by an IF filter actually located at the A/D input, but in some cases significant contributions to the net response will also be made by an RF filter or, in the case of a multiple-conversion receiver, an earlier IF filter.

Manuscript approved May 22, 2024.

For each of the ten designs presented explicitly, several more are implicit, because each of the explicit designs can be trivially modified to operate in a different Nyquist zone. In other words, the input signal band can be moved by any integral multiple of half the A/D sampling rate. The only system changes required are a different IF filter and, in some cases, a change to single digital frequency shift, a change that amounts to nothing more than inverting the polarity of every other A/D output. Rather than offer a proof, I encourage the reader to walk through the logic of any of the design examples here with such a shifted signal band in mind.

In this document frequencies are never given units explicitly, but the implicit unit everywhere is 1/400 of the desired width of the signal band, a choice that makes all signal bands in this document appear to be 400 units wide. This is an arbitrary choice aimed at convenience in diagram labeling, convenience that arises in two ways. First, the number 400 is sufficiently composite to permit many examples to be constructed in which all frequencies have integer values. Of course integer values are completely unnecessary, but they are a practical convenience when annotating diagrams in which space is at a premium. Second, the number 400 is small enough to keep most of those integers in a reasonable range for display.

This document is Part I and focuses on high-level architecture, the subject of the next section. Structuring and optimizing the component digital filters is deferred to Part II, as is discussion of DSP implementation. This document is somewhat informal and has no traditional "summary" or "conclusions" section. One can reread this section and the introduction to the next in lieu of a summary. One conclusion is or will soon be obvious—IQ demodulation can be accomplished using many combinations of IF frequency and sampling rate—and further conclusions on these matters will only be formed, by any of us, in the fullness of time.

Throughout this document, links are this color.

2. HIGH-LEVEL ARCHITECTURE

High-level architecture here is a view of the system as a single signal path comprising processing steps like filtering, sampling, modulation by a complex exponential, decimation, and occasionally zero interpolation. The purpose of this high-level view is to get the sequence of steps right and to illustrate and justify required relationships among the various filter cutoff frequencies, modulation frequencies, sample rates, decimation and interpolation ratios, etc. Details like depth of aliasing suppression sought, passband ripple allowed, etc. are not a part of this discussion and are dealt with in Part II, where actual filter design is discussed.

There is some subtlety to what is covered here versus there, however. Many of the digital-filter designs of that Part II discussion will be based on so-called frequency-response masking (FRM) approaches that have their own interesting signal-flow architectures, architectures that use parallel signal-flow paths or processing "arms." We defer discussion of FRM architectures for now, however, because the intent here is instead to do as much as possible with a particular simple graphical approach to signal-flow architecture, an approach that is somewhat awkward to apply to systems with parallel arms. This section is about understanding as much as possible using simple tools.

That said, the single-path restriction of this section actually restricts less than one might expect. For example, some classic discussions of IQ demodulators refer to two parallel signal paths containing real signals that here are instead be taken to be a single path containing a complex signal. Further, in these

systems the combination of a digital filter immediately followed by decimation by some integer M is quite common, and such a combination is naturally realized using what to the nonspecialist would certainly look like a multi-arm system. So the single-arm restriction should not be dwelled upon, as it's really just a formal way to limit the high-level architecture discussion to what the tools used here can conveniently handle.

This section's core is the collection of ten example high-level architectures, all but one of which are presented in Section 2.3. These design examples will be presented roughly in order of increasing complexity and with each design adding one or two major new design concepts or tools to our repertoire.

Before the example designs are presented, Section 2.1 outlines the high-level functioning of nontunable IQ demodulators in general, using spectral sketches to represent the processing steps as well as signals. Section 2.2 then reviews the important notion of an Nth-band filter. The spectral-sketch approach is further elaborated in the detailed discussions of the Section 2.3 example designs, particularly the first, Design A in Section 2.3.1. The many example designs of Section 2.3 are followed in Section 2.4 by a general derivation of the key frequency relationships illustrated by the examples taken as a group. That section concludes with this document's final design example.

Figures 1 and 2 present a sort of global view of the ten example architectures. First, to aid in seeing the frequency relationships of the design examples as specific instances of the general relationships derived in Section 2.4, Fig. 1 presents those general relationships here, well in advance of their derivation. Second, Fig. 2 lists and plots key parameters of each example design so that they can be easily compared to each other. Figure 2 also serves as a *de facto* index to the parts of this document that cover those designs.

Readers eager to compare the merits of the example designs will naturally want to see estimates of the computational complexity of the DSP needed to realize each of those architectures. That, however, is not actually knowable at this first level of the design process, as it depends on the detailed filter designs used in realization, the topic of Part II. To get there, we must do this first. For what it's worth, experience suggests that each IQ demodulator presented here can likely be implemented, with aliasing in internal downsampling suppressed by well over 100 dB, with perhaps 100 coefficient multiplies per complex output sample, though in some cases sophisticated filter architectures will be needed to achieve that level of computational efficiency.

2.1 Spectral Sketches

The IQ demodulators of interest in this document are those equivalent in net effect to the sequence of frequency-domain steps sketched schematically in Fig. 3 from top to bottom. All horizontal axes there are identical and represent frequency with the origin a bit right of center, but to reduce clutter, only the first axis is actually labeled. The first line and lines marked on the left with $=$ represent signal spectra, and lines marked instead with $\star$ or $\times$ represent operations on signals. These representations are schematic, not literal, and they are intended only to show locations of passbands, stopbands, etc. and general features of signal spectra.

The first line of Fig. 3 represents the IF input comprising a desired signal and representative out-of-band spurious signals. The next line is identified by the $\times$ on the left as a filter frequency response to be multiplied by the signal spectrum just above it. This particular frequency response is unusual in that no symmetry about the origin is indicated, so as shown it is not that of an ordinary analog RF filter because its impulse response is complex. If we had such an implausible frequency response available to us, here it would

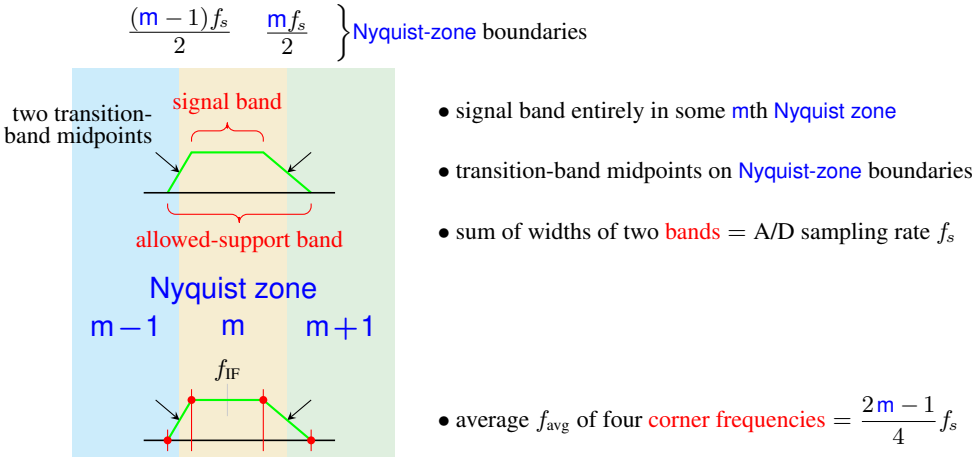

Fig. 1 — General relationship of the A/D sampling rate to the IF-filter specification. In the common special case in which transition bands have equal widths, average-of-corners frequency f_{avg} is coincident with the IF frequency, the center of the signal band, so the last rule becomes $f_{IF} = (2m-1)f_s/4$.

Design	*Section*	f_s	f_{IF}	f_{avg}	f_{out}
A (Fig. 4, p. 6)	2.3.1 (p. 8)	1000	750		500
B (Fig. 5, p. 14)	2.3.2 (p. 13)	1260	1575		420
C (Fig. 6, p. 16)	2.3.3 (p. 15)	1760	1320		440
D (Fig. 7, p. 18)	2.3.4 (p. 17)	2640	3300		440
E (Fig. 8, p. 19)	2.3.4 (p. 17)	2640	3300		440
F (Fig. 9, p. 21)	2.3.5 (p. 20)	2574	3300	$3217\frac{1}{2}$	429
G (Fig. 11, p. 26)	2.3.6 (p. 22)	1475	3300	$3318\frac{3}{4}$	$491\frac{2}{3}$
H (Fig. 12, p. 29)	2.3.7 (p. 28)	1200	3300		450
I (Fig. 13, p. 30)	2.3.7 (p. 28)	$1885\frac{5}{7}$	3300		$419\frac{1}{21}$
J (Fig. 15, p. 37)	2.5 (p. 36)	1400	1000	1050	420

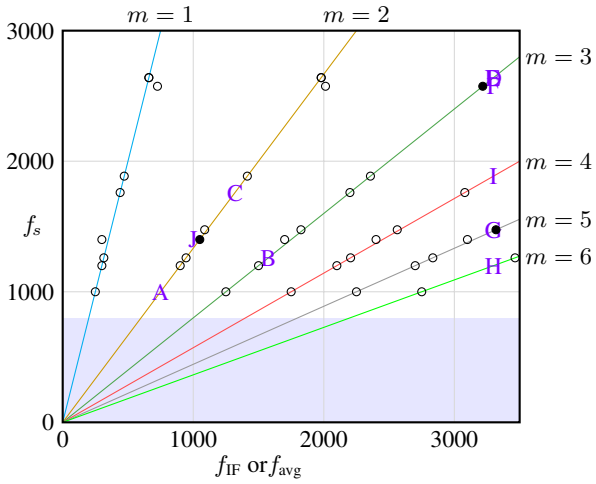

Fig. 2 — Parameters for the ten example designs: sample rate f_s, (complex) output sample rate f_{out}, IF frequency f_{IF}, and where different from the latter, average-of-corners frequency f_{avg} of Fig. 1. In the plot, coordinate (f_{IF}, f_s) is marked with the design name, and where different (f_{avg}, f_s) is marked with •. Any design can be modified to pair the original f_s with an f_{IF} in a new Nyquist zone. Such modified (f_{IF}, f_s) pairs are marked with ○. Assuming a signal-band width of 400, the part of f_s above the plot's lower band is the combined width of the lower and upper IF-filter transition bands.

block the negative-frequency signal component while passing the corresponding positive-frequency signal component, technically making it a sort of analytic bandpass filter. Further, some parts of the interference would be eliminated, and other parts, those falling into the filter's transition bands, would be attenuated. In this particular design, the positive-frequency transition bands have counterparts at negative frequencies that we can think of as transition bands as well. The latter are not something we particularly desire, but they are not harmful, and we'll see soon enough that there's most often no way not to end up with them, at least in part, so we may as well sketch them in now.

The third line is marked on the left with convolution symbol $\star$ to identify it as a spectral convolution, in this case with many spectral impulses. The "$\cdots$" ellipses on the right indicate that the sketch is to be considered periodic, so that it continues in the same way in both directions. The filter-output spectrum, discussed above but not actually shown, is convolved spectrally with the impulses to copy it many times over, yielding the output signal spectrum sketched on the fourth line. That spectrum is readily interpreted as a demodulated and sampled version of both the desired signal and partially attenuated interfering signals that have been aliased into out-of-band portions of the output spectrum.

While it is hoped that this spectral-sketch notation for signal-processing flow will be reasonably intuitive for most readers, that notation, as elaborated on considerably in Section 2.3.1 below, is developed in detail in [1], with many examples. There the subtleties of amplitude scaling are dealt with carefully as well, but here we will not discuss scaling at all and simply leave such matters to implementors.

Key to the Fig. 3 spectral-sketch notation is that in it we can describe what we want to do while worrying little about the difficulty of realizing some of the operations used, in particular the net filtering shown. In the first example design below, we'll carefully transform this system into one containing only realizable operations and in doing so will obtain a version of the classic IQ demodulator. As in most engineering discussion, we'll be somewhat informal about language and for brevity will sometimes say "signal" rather than "signal spectrum" and "filter" rather than "filter frequency response." The context should make clear which is intended.

2.2 Nth-band Filters

Various design examples in Section 2.3 below contain digital filters that can optionally be made halfband, thirdband, fourthband, etc., so let us quickly review the basic idea and purpose of an Nth-band digital filter.

A digital filter is Nth-band if and only if its *Nyquist sum*, the sum of N copies of its frequency response, copies shifted in frequency by each of $0, \frac{1}{N}, \frac{2}{N}, \ldots, \frac{N-1}{N}$ times the period of the response, has the frequency response of a pure delay, perhaps with gain. (An example will soon make this clearer.) At that Nyquist-sum delay, the filter impulse response has a sample equal $1/N$ of the Nyquist-sum gain. More importantly, the definition of the Nyquist sum above has a mathematical counterpart in the time domain: starting at that $1/N$ sample and counting outward in either direction, every Nth sample of the impulse response is zero.

When an Nth-band filter is chosen for an application, the idea is generally that it's worth tolerating its frequency-response quirkiness—more on that shortly—in order to obtain those zero samples in its impulse response, which save computation in the implementation, since one hardly needs a hardware multiplier to multiply by zero. This is particularly valuable for halfband filters, because nearly half their impulse-response samples are zero.

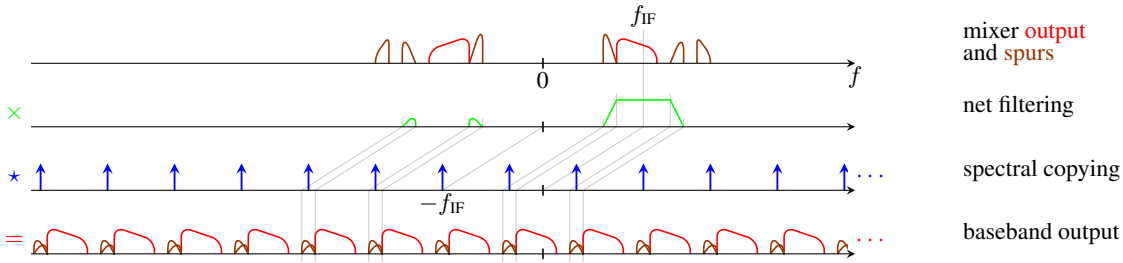

Fig. 3 — Classic IQ demodulation from an IF at an odd multiple (here 3) of a quarter of the IF-sampling rate to a baseband output at half that rate. Only functionality is shown here. The actual demodulator structure is derived in Fig. 4.

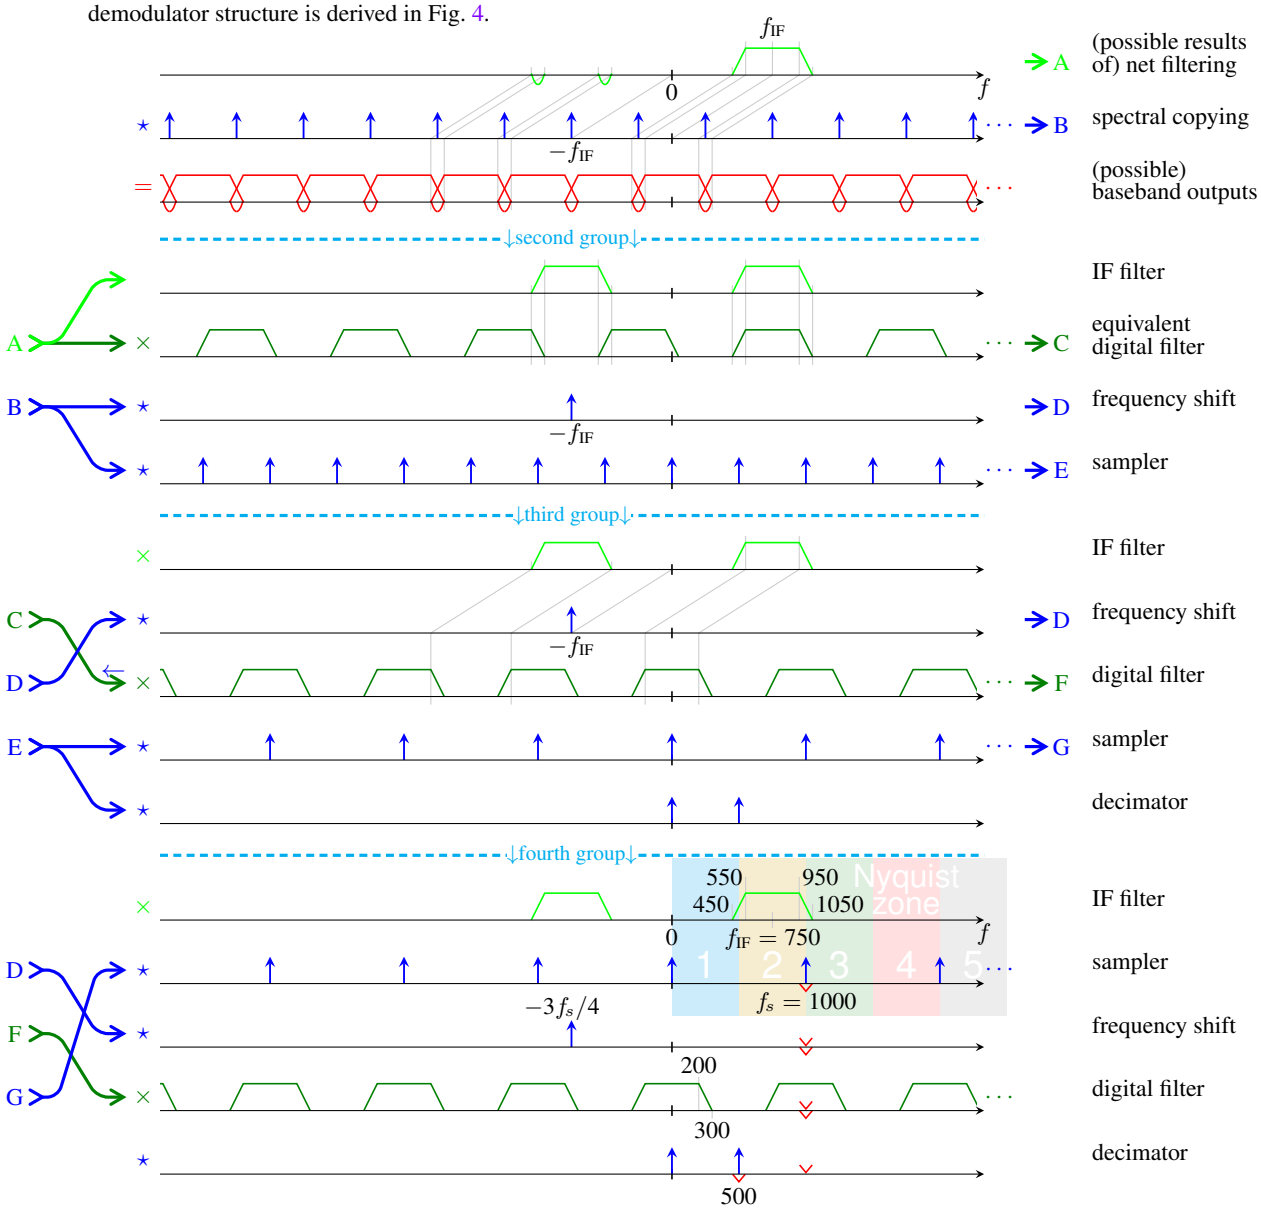

Fig. 4 — Design A. Transformation of the Fig. 3 functionality (first group) through intermediate equivalent systems (second and third groups) into a realizable system (fourth group). Frequencies are relative to an arbitrary reference frequency and so are dimensionless.

The issue with the quirky frequency response is that one can no longer specify passband and stopband characteristics independently. For example, consider this specification for a filter frequency response.

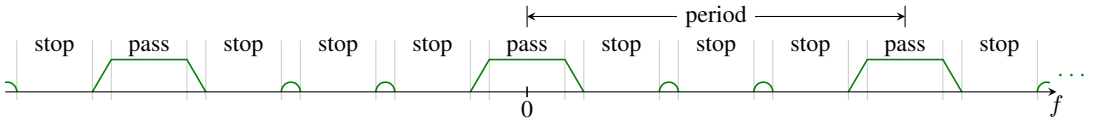

Each period contains one passband and three stopbands, and these four specified bands are identical in width. The passbands are separated from adjacent stopbands by transition bands, and adjacent stopbands are separated by don't-care bands sketched here as little hills. All these transition and don't-care bands have the same widths, so the passband and stopband centers are separated by integral multiples of one fourth of a period. Filter specifications structured much like this, though with varying numbers of stopbands per period, appear throughout this document.

The catch is that in most applications far greater levels of passband ripple than stopband ripple can actually be tolerated. We might, for example, want stopband frequency responses confined between $\pm 10^{-5}$, a -100 dB specification, while tolerating any passband response confined between 1 ± 10^{-2}, a peak-to-leak passband ripple level of some 0.17 dB.

If we include a fourthband requirement in the filter optimization, however, the implied Nyquist sum combines four frequency-response terms the behavior of which we can easily sketch:

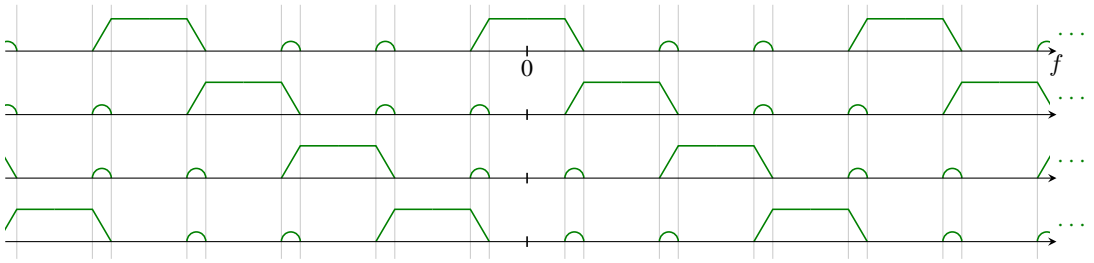

In designing this to be a fourthband filter, we can simply set this sum to unity everywhere and then put the desired net gain and delay into the filters after the fact. To reach unity in each of the original passbands and stopbands, the sum combines what were originally the three distinct stopbands, now aligned vertically, and one passband. If these four terms are constrained in filter optimization to sum to one—a fourthband constraint—it is enough to specify three of the terms explicitly. This amounts to specifying any three of a four-band set comprising three stopbands and a passband. Conventionally, one specifies the stopbands and lets the appropriate passband follow automatically.

But what does this do to ripple levels? If we use the same $\pm 10^{-5}$ stopband bounds as before, we are now guaranteed a passband confined between $1 \pm 3 \times 10^{-5}$. This is a far, far flatter passband than we actually need! Should we celebrate? No, because a shorter impulse response results when the passband specification

is loose than when it is tight. So our fourthband constraint, while putting some nice zeros into the impulse response, has also made that impulse response a bit longer. The zeros reduce the number of coefficient multiplies required per output sample, but the greater impulse-response length increases it. Which is the dominant effect?

The fourthband filter also allows the filter optimization fewer degrees of freedom with which to determine the transition and don't-care bands, because in the Nyquist sum above, the response in those bands sums to unity as well. If the small-hill transition bands peak at tiny levels, as is often the case, this sum is more or less just a sum of the two transition bands, and the fourthband requirement will result in each transition-band response crossing $1/2$ (and therefore each other) right in the middle of the band. This is not likely where it would have ended up anyway, so the fourthband requirement has imposed a hidden requirement in the transition bands as well as in the passbands and stopbands. And additional requirements always come with a cost, here either poorer stopband suppression or a longer impulse response. We don't want to give up suppression, so again we have the possibility of a slightly longer impulse response.

Does this tell us that an Nth-band requirement makes a filter less computationally efficient? No, because having every Nth impulse-response sample be zero automatically is quite the benefit, particularly when N is small. One can never say for certain whether an Nth-band filter is a win without simply designing the filter with and without the Nth-band feature and comparing the computational resources required, but a halfband filter is nearly always a win, and Nth-band filters for $N > 2$ certainly can be.

In discussing the example high-level architectures of Section 2.3 below, we will take care to point out when filters can be Nth-band. In the example filter designs of Part II we will design using Nth-band filters where possible (in part to avoid having to decide on passband ripple specifications). We must always remember though that in a real design context it may be possible to do better by removing the Nth-band requirement, particularly for larger N.

2.3 Architectures for IQ Demodulation

The first of the ten design examples is discussed in far greater detail than the others. This is primarily to elaborate on the spectral-sketch notation introduced above, but it will also review some key concepts in digital signal processing.

2.3.1 Design A—Classic Decimation by Two

We begin by walking through the derivation of the the most widely used of the nontunable systems for direct IQ demodulation of an IF signal. Sometimes the correspondence between systems in the literature and this system is not at all obvious, but still this approach is mathematically equivalent to many of them. And the approach to its derivation taken here will serve as a foundation for discussion of the more interesting systems to follow. Further, this derivation introduces the spectral-sketch approach that we'll use throughout the Section 2 discussion on high-level architecture.

Given a desired output baseband sample rate, the classic approach uses the lowest A/D sampling rate possible, so it is not surprising that it evolved when A/D performance presented the primary constraint on obtainable performance. As we'll see below, the price paid for this minimization of pressure on A/D performance is the imposition of rather severe requirements on the IF filter. The possibility of sampling with a faster A/D converter as a way to ease filter requirements—later design examples—is a key motivation

behind this document. Faster A/D conversion also benefits spurious response by spreading A/D-conversion harmonics over a wider frequency range and then filtering more of them out digitally [2].

Still, it is important to walk through the classic design so that the origin of its inadequacies can be seen. This will also introduce key design principles and spectral-sketch representations of useful signal-processing operations not yet encountered.

The design process is laid out in Fig. 4 (page 6), which presents a sequence of systems representing an evolution from the functional demodulator system of Fig. 3, which is not directly realizable, to a functionally equivalent but realizable system by working through two intermediate systems. The four systems are shown in sequence from top to bottom in Fig. 4 as four groups of "spectral lines" with ----- separators between groups. The operation of this particular design's final form, at the bottom of Fig. 4, could certainly have been examined directly. A careful derivation from a simple functional description is presented instead, however, to illustrate a design process that will, later when we look at more involved demodulator designs, be the simplest way to gain solid understanding of a system and, in particular, to understand its key frequency relationships.

The First Group of Spectral Lines

These lines, at the top, duplicate the system of Fig. 3 with two small notational changes. The first change is that the first two lines of Fig. 3 become a single line, the first, in Fig. 4. Basically, we sketch the filter frequency response but interpret it as also representing the full range of possible spectral components of its own input signal. In other words, we assume that any nonstopband portion of the filter's input spectrum might be occupied by signal components, whether desired components or spurious interference components. This convention applies only to the first spectral line in a group, and it is easy to spot notationally: not only is this line first, the usual $\times$ marking to the left of a filter frequency response is missing.

The second change is that some small filter transition bands are sketched below the axis, in this case in the first line of the first group. These are bands we don't actually need for functionality but that we realize through experience are either unavoidable or both harmless and of some benefit, generally to filter optimization. However, the apparent negative amplitude is not intended to indicate a negative, or 180° phase, frequency response. In fact at this point we don't know the phase of any portion of any of the frequency responses, as these high-level schematic sketches make no attempt to represent phase. We associate negative spectral amplitudes with those small bands in the diagram only for the convenience of making it easier, elsewhere in the diagram, to spot signals that result from aliasing of those bands' contents. In the first group of Fig. 4, for example, those aliased components are visible in the third line, in the out-of-band portion of that output spectrum. We do sketch these below-the-axis bands in a simple schematic way, using a simple "hump" of a fixed size even though in implementation the frequency responses in the different areas so flagged might be quite varied.

We will not have occasion in this document to sketch frequency responses any part of which are actually intended to be negative, so we can reserve negative spectral amplitudes for representing these harmlessly aliasable spectral regions.

The Second Group of Spectral Lines

On the right the first two lines of the first group are designated A and B. In the second group each of those lines is split, as shown on the left, into two lines.

The second group's analog IF filter and its equivalent digital filter C[1] in cascade have the frequency response of the first group's net filtering A. A digital filter always has a periodic frequency response and so is shown with a "$\cdots$" marking on the right. An analog filter's response never has a "$\cdots$" periodicity marking. Any analog filter that is intended in implementation to be a single, ordinary filter with a real impulse response will always have a conjugate-symmetric frequency response, which property is represented schematically here with even symmetry about the origin.

In this particular design, the splitting of net filtering A into an IF filter and digital filter C is specifically intended to align the negative- and positive-frequency passbands of the IF filter with a stopband and a passband of digital filter C. The idea is that because filter A has a complex impulse response, a purely analog realization would be awkward[2] and prone to precision issues. So instead we use complex digital filter C to obtain a frequency response not symmetric about the origin, and we use the stopbands of a real IF filter to suppress the unavoidable extra passbands in digital filter C that arise through periodicity. The two filters can do together what neither can do alone, have a nonperiodic frequency response and a complex impulse response.

The particular classic design in Fig. 4 aims at the lowest reasonable sample rate—to be discussed shortly—and therefore, at this stage, the smallest possible period for digital filter C.

Similarly, spectral-copying impulses B in the first group are just the convolution $\star$ in the second group of the one impulse of frequency shift D and the infinite number of impulses of sampler E. Splitting the one line into two in this way centers one impulse of the periodic set on the origin to make this operation into ordinary sampling. The amount of the frequency shift required to do this is not unique, of course, even though the time-domain operation represented ultimately will be. Here the amount of the shift is just the IF frequency for conceptual simplicity.

The Third Group of Spectral Lines

In the third group, digital filter C and frequency shift D are interchanged to put the frequency shift first, so that the frequency shift translates the center of the IF signal spectrum (not shown) from the IF frequency to the origin. The passband of digital filter C of the most interest, in order to have the same effect on signals, must likewise have its center be translated from the IF frequency to the origin in creating digital filter F. The imposition of this downward frequency shift in the course of reversing C and D on the left is flagged with a "$\leftarrow$" over the right end of the line there for digital filter C.

The purpose of reversing C and D here is to give the resulting filter a real impulse response. This is not strictly necessary at this point, because the ultimate computational complexity of filtering the newly

[1] When an actual, realizable system is analyzed, that system might contain a "filter" that is related to an "equivalent filter" that appears in some later analysis. Here, however, the usual analysis steps are reversed into a synthesis procedure, making the equivalent come first. The digital filter to which equivalent digital filter C is equivalent will appear in the third group.

[2] It requires using two wires to carry the signal, one each for the real and imaginary parts. In one particular case, it leads to a classic "phasing method" of demodulating a single-sideband (SSB) signal.

complex signal with the newly real filter is actually identical to the computational complexity of filtering the original real signal with the original complex filter. Basically, the frequency shift can be realized in any of several equivalent places in this system and in others presented later. In all of these designs, however, we do the shift early so as to work with real digital filters exclusively. We'll see in Part II that this in no way prohibits moving the shifting computations to more convenient places at later stages of the design process, so doing it this way now is simply a convenient custom.

In a very similar way, the infinitely many impulses of sampler E are split in the third group into the infinitely many impulses of sampler G and the two impulses of a decimator. Of course sampling at some rate f_s and immediately decimating by two is equivalent to just sampling at $f_s/2$. For samplers the impulse spacing is the sample rate, so representing a low-rate sampler by one M times as fast followed by decimation by M will always require "factoring out" (in a convolution sense) a set of M impulses to represent the decimation.

The high-speed sampler that results is uniquely determined by decimation ratio M, but there are actually infinitely many ways a set of M impulses can be chosen to represent the decimator. Here, for example, those impulses could be moved from the frequencies shown to the negatives of those frequencies with no harm done. In the mathematics of all this, the periodicity of the complex exponential turns out to ensure that all these spectral representations of decimation correspond to identical time-domain operations in the implementation, so the choice in the spectral sketches is purely one of convenience. In this document, decimating by M from a sample rate f_s is always represented by impulses at frequencies 0, f_s/M, $2f_s/M$,..., $(M-1)f_s/M$.

The actual reason for splitting sampler E into sampler G and decimation is to put digital filter F adjacent to sampler G in the sequence of processing steps. Once they are adjacent, we observe something important: the frequency response of digital filter F would be unchanged if translated by the frequency of any spectral impulse of sampler G. Whenever such a property holds, we can interchange a spectral multiplication by a frequency response and spectral convolution with impulses without changing the frequency response in the process, a very convenient fact termed the *most noble identity* in [1]. When the impulses represent a decimator in particular, it is conventionally termed the *noble identity for decimation*. The simplest case in which the most noble identity applies is the one pictured here, where the impulses and the frequency response have the same spectral period. Of course we desire this exchange because ultimately we need sampler G to precede digital filter F for realizability. Digital filters need digital inputs.

The Fourth Group of Spectral Lines

In the fourth group, the use of the most noble identity just discussed appears on the left as as line G crossing line F. Line G also crosses line D, an interchange of operation order allowed by the commutativity of convolution.

We adopt the convention that the final group, here the fourth, always represents an implementable system and so is shown with explicit frequency markings to guide implementation. As a part of that, input and output sample rates for the various discrete-time operations are marked with $\overset{\vee}{-}$ and $\underset{\vee}{-}$ respectively. Visually the caret looks vaguely like an arrow in the direction of signal flow, and that arrow appears on the input or output side respectively of the frequency axis interpreted as a system or a block in a sort of top-to-bottom block diagram.

Here the fourth group contains five signal-processing steps.

1. An analog IF filter operates on an analog input signal and produces an analog output.
2. A sampler samples that output at $f_s = 1000$. Spectral convolution with unit impulses at integer multiples of spectral interval f_s corresponds to time-domain multiplication with impulses of area $1/f_s$ at integer multiples of time interval $1/f_s$. The time-domain signal created by this multiplication is an impulse stream with signal sample values encoded into the impulse areas. This is a convenient mathematical representation of a digital signal, but of course physically those sample values are obtained by an A/D converter.
3. The new discrete-time signal is shifted down in frequency by $3f_s/4$ by multiplying by $\mathrm{e}^{-j2\pi(3f_s/4)t}$. This operation's input and ouput are impulse trains that are nonzero only at times $t = n/f_s$, where integer n is the sample index, so the operation effectively just multiples sample n by $\mathrm{e}^{-j2\pi 3n/4} = j^n$. This will turn out in the Part II discussion of computational structures to be extraordinarily simple to implement.
4. The frequency-shifted signal is digitally filtered. In most designs of this type this filter is made halfband.
5. The filtered signal is decimated by two.

We'll see in the Part II discussion of computational structures that filtering followed by decimation, as in the last two steps here, will be always implemented together for efficiency. The step 4 frequency shift can be integrated into the filter-decimation computation as well.

In Figure 4 then, the system represented by the first group represents the overall goal of the design, and the system represented by the equivalent last group can be mapped relatively straightforwardly—the details will be covered in the Part II discussion of filter design and computational structures—into a realization. The middle groups are the derivation that ties these systems together. We will derive all of our IQ demodulators in this fashion, though with far less explanation. The point of the spectral-sketch approach, after all, is to let the sketches do the talking.

Variations

The bandpass sampling in this system as shown uses the second Nyquist zone, but in principle one could use any other Nyquist zone by simply changing the IF filter and the amount of the frequency shift. All available choices, however, imply an IF frequency at an odd multiple of one fourth of the A/D rate, the rate at which sampler G of the third and fourth groups operate. This is just the second frequency relationship of Fig. 1 for the special case in which the transition bands of the IF filter have equal widths.

In addition to those IF-frequency options, there is one very useful specialization. Making the digital filter of Fig. 4 system a halfband filter, one of the Nth-band filter types discussed above in Sections 2.2, almost certainly greatly benefits computational efficiently. In addition, a little z-transform algebra suffices to show that with a halfband filter that system is equivalent to the common IQ demodulator system in which the real and imaginary signal paths contain a delay and a Hilbert-transform filter respectively. The Hilbert-transform representation offers no real advantage over the halfband-filter representation, however. In fact the halfband-filter version is much simpler to understand.

This Makes a Poor Wideband Design

The first two lines of the second group in Fig. 4 reveal that the width of the signal band, here 400, plus the width of the transition band, here 100, is necessarily equal to half the A/D sample rate. This is just the first frequency relationship of Fig. 1. For this demodulator structure it means that wideband versions must have very narrow transition bands, both for the digital filter, which is seldom a problem, and for the analog IF filter, where the implied cost/performance tradeoff is often less favorable.

2.3.2 Design B—Decimation by Three

The structure of Design A of Fig. 4, developed in Section 2.3.1, makes perfect sense when obtaining a faster A/D and faster DSP hardware is more expensive than obtaining a high performance IF filter. This is less often the case today than it was in times past, so here we begin considering alternative structures, ones that ease the requirements on the IF filter by using faster sampling together with decimation ratios higher than two. We leave composite ratios to subsequent examples and here consider only prime ratios. It will turn out that composite decimation ratios make for more economical implementations than do all but quite small primes, so in this design we assume a decimation ratio of three.

In this example we also explore obtaining extreme output bandwidth by choosing a sample rate barely over the signal bandwidth times the decimation ratio so that the width of the signal band is barely below the output sample rate. This choice will lead to an an interesting example filter design in Part II.

The high-level architecture appears in Fig. 5 (page 14),which parallels Fig. 4 in structure. The way in which we arrive at this design is best seen in the third group, starting at the output.

1. Given a desired output bandwidth, here 400, choose an output sample rate just a bit larger, here 420. The latter is just 5% greater than the former.[3] The last line of the first group shows that the passband and stopband corners of the digital filter, here 200 and 220, are determined respectively by the output bandwidth and output sample rate.

2. Choose the decimation ratio to be a small prime, here 3. Now input sample rate $f_s = \langle\text{output sample rate}\rangle \times \langle\text{decimation ratio}\rangle = 420 \times 3 = 1260$ is determined. In the third group we see that this is also the period of the digital filter response.

3. Let's assume the IF filter has symmetric transition bands. Then, consistent with the Fig. 1 frequency relationships, the third group in Fig. 5 shows that twice the IF frequency should be chosen to be an odd multiple of half the digital filter's frequency-response period of $f_s = 3 \times 420 = 1260$. Here we set $f_{\text{IF}} = 5f_s/4 = 1575$, but in principle we could decrease this f_{IF} by $f_s/2$ or increase it by any of $f_s/2, f_s, 3f_s/2, 2f_s, 5f_s/2o, \ldots$ without disrupting the key frequency relationships of the third group. Of course in half those cases the digital filter's frequency response would need shifting to keep passbands aligned correctly.

[3] Of course the conventional view is that a larger "excess bandwidth" would lead to a simpler digital filter, because its transition bandwidth would be a larger fraction of the period of its frequency response. But as we will see in Part II, recent advances in digital-filter structures allow a digital filter to be designed to meet even extreme specifications while remaining computationally efficient.

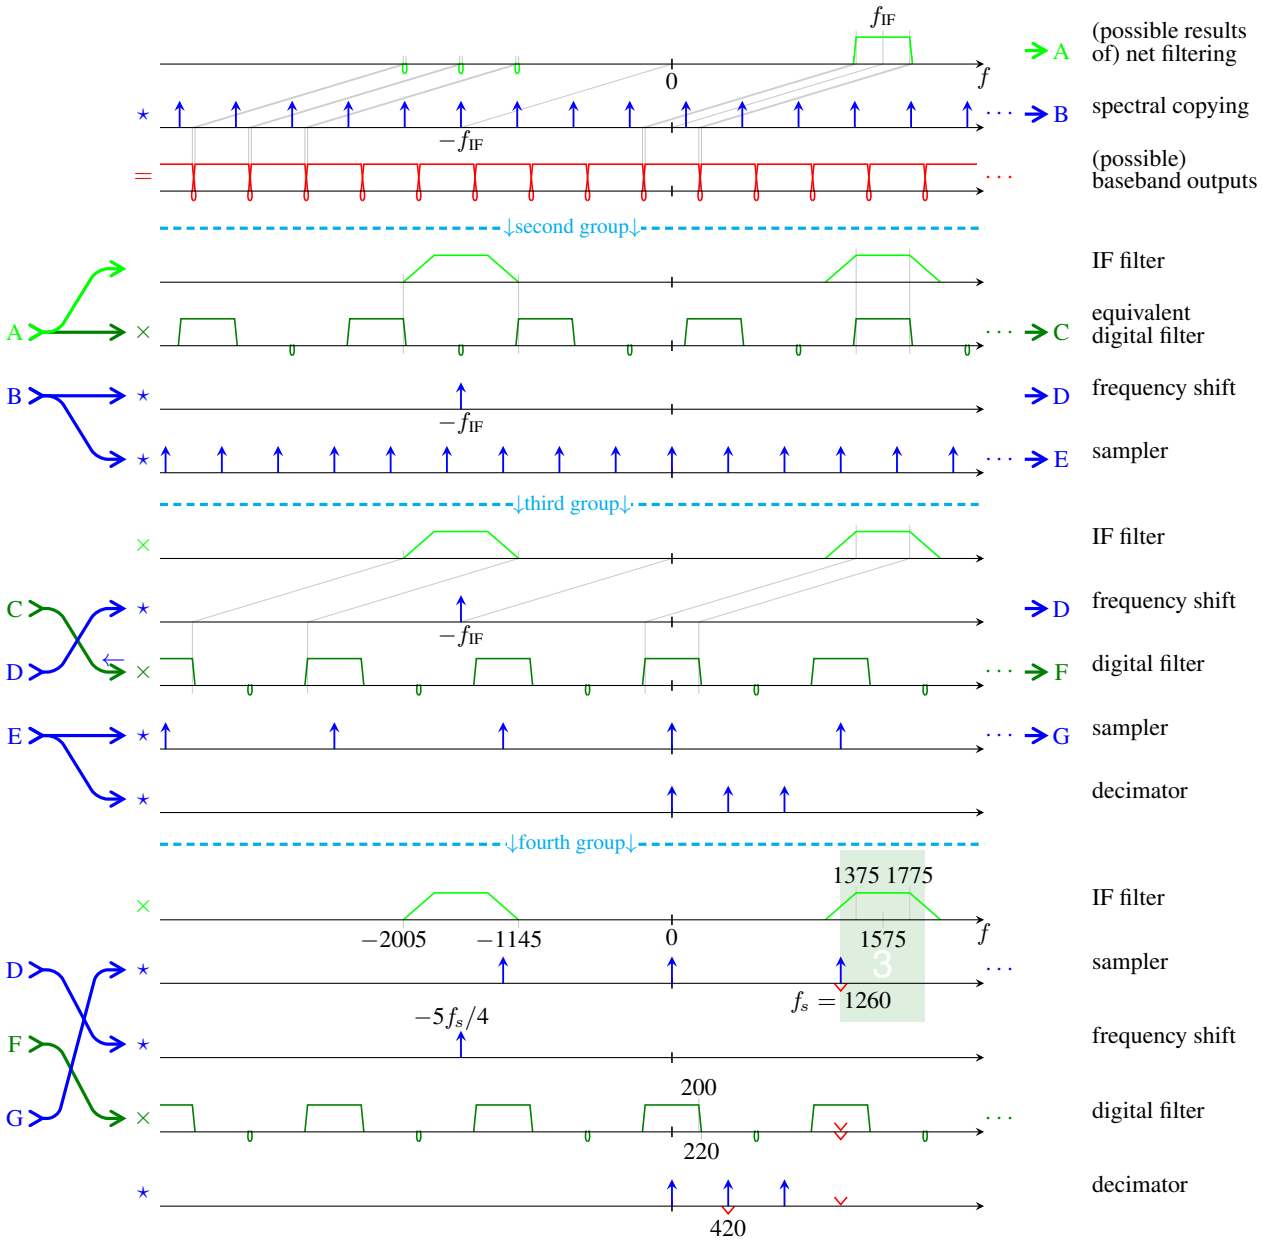

Fig. 5 — Design B. Relative to the conventional system of Design A in Fig. 4, here a (2.1×) higher IF allows the RF filter a wider (2.5×) transition band, and (26%) faster sampling allows the IF filter to have (3.875×) wider transition bands. A higher decimation ratio actually reduces the output sample rate (by 16%) used for DSP computation. The minimum front-end A/D bandwidth is $71/38\times \approx 1.87\times$ higher. As shown this system uses bandpass sampling in the third Nyquist zone, but straightforward changes to the IF filter and frequency shift would allow any other Nyquist zone to be used instead.

In the third group we see what is perhaps the key frequency relationship here: the width of the signal band plus the width of the support—the set of frequencies where it is nonzero—of the (one-sided) IF-filter frequency response is equal to the input sample rate. This is, of course, just the first rule in Fig. 1.

Now the purpose of this design is apparent: the signal bandwidth is the same as in Design A of Fig. 4, but now, because of the higher decimation ratio, the sample rate has increased substantially. This gives the digital filter's frequency response a larger period, which in turn allows wider IF-filter transition bands.

The digital filter specifications here are compatible with using a thirdband filter, one of the Nth-band filter types discussed above in Sections 2.2).

2.3.3 Design C—Decimation by Four in Two Steps

The design of Fig. 6 (page 16) increases the decimation ratio yet again, to four this time, to allow an even gentler transition-band rolloff in the IF filter. Two new features of this design stand out.

First is a minor aspect of the logic of the design. In earlier designs, factoring a frequency shift out of spectral-copying impulses B in the second group had two purposes. It created slow sampler F, which can be implemented in various ways, and it created frequency shift D that could be interchanged with equivalent digital filter C in the third group to transform the latter into net digital filter E, which has frequency-response symmetry indicating a real impulse response. Here, however, only the latter applies, as the spectral-copying impulses B are already those of a slow sampler, even before the split.

Second and more important is that because the overall decimation ratio is a product of very small integers, the digital filtering can be implemented extremely efficiently. The most efficient case is when the overall decimation ratio is a power of two, as it is here. The key is that decimation is split into stages. Instead of one decimation step using four spectral impulses, we have first decimation J and a second decimation, each using two spectral impulses and so each representing decimation by two. First decimation J reduces the sample rate from 1760 to 880, and the second decimation reduces it again from 880 to 440.

The benefit of splitting the decimation is indirect: it leads to splitting the third group's net digital filter E into the fourth group's combination of shaping filter G with a smaller frequency-response period and masking filter H with a larger frequency-response period. Roughly speaking, shaping filter G gives the passband and transition bands of net digital filter E their shape, hence the name. The stopbands of masking filter H then mask out alternating passbands of shaping filter G so that the product of those two responses has the form of net digital filter E in the third group. In the last group the filtering steps and decimation steps alternate because in arriving at that final design the most noble identity is used three times: when line J crosses line G, when line I crosses line G, and when line I crosses line H. Of course invocations of the most noble identity in these derivations correspond exactly to the crossings of lines representing spectral-impulse convolutions with lines representing digital-filtering operations. In these diagrams such a crossing is easily identified by color.

For ease of comparison with later designs, let us look more carefully at the filter factorization. A period of the frequency response of net digital filter E has width f_s and contains four uniformly sized and spaced bands with design specifications, one passband and three stopbands. These specified bands have widths equal to that of the signal band and are separated by small transition or don't-care bands of identical width. The latter are also don't-care bands for the two component filters, shaping filter G and masking filter H.

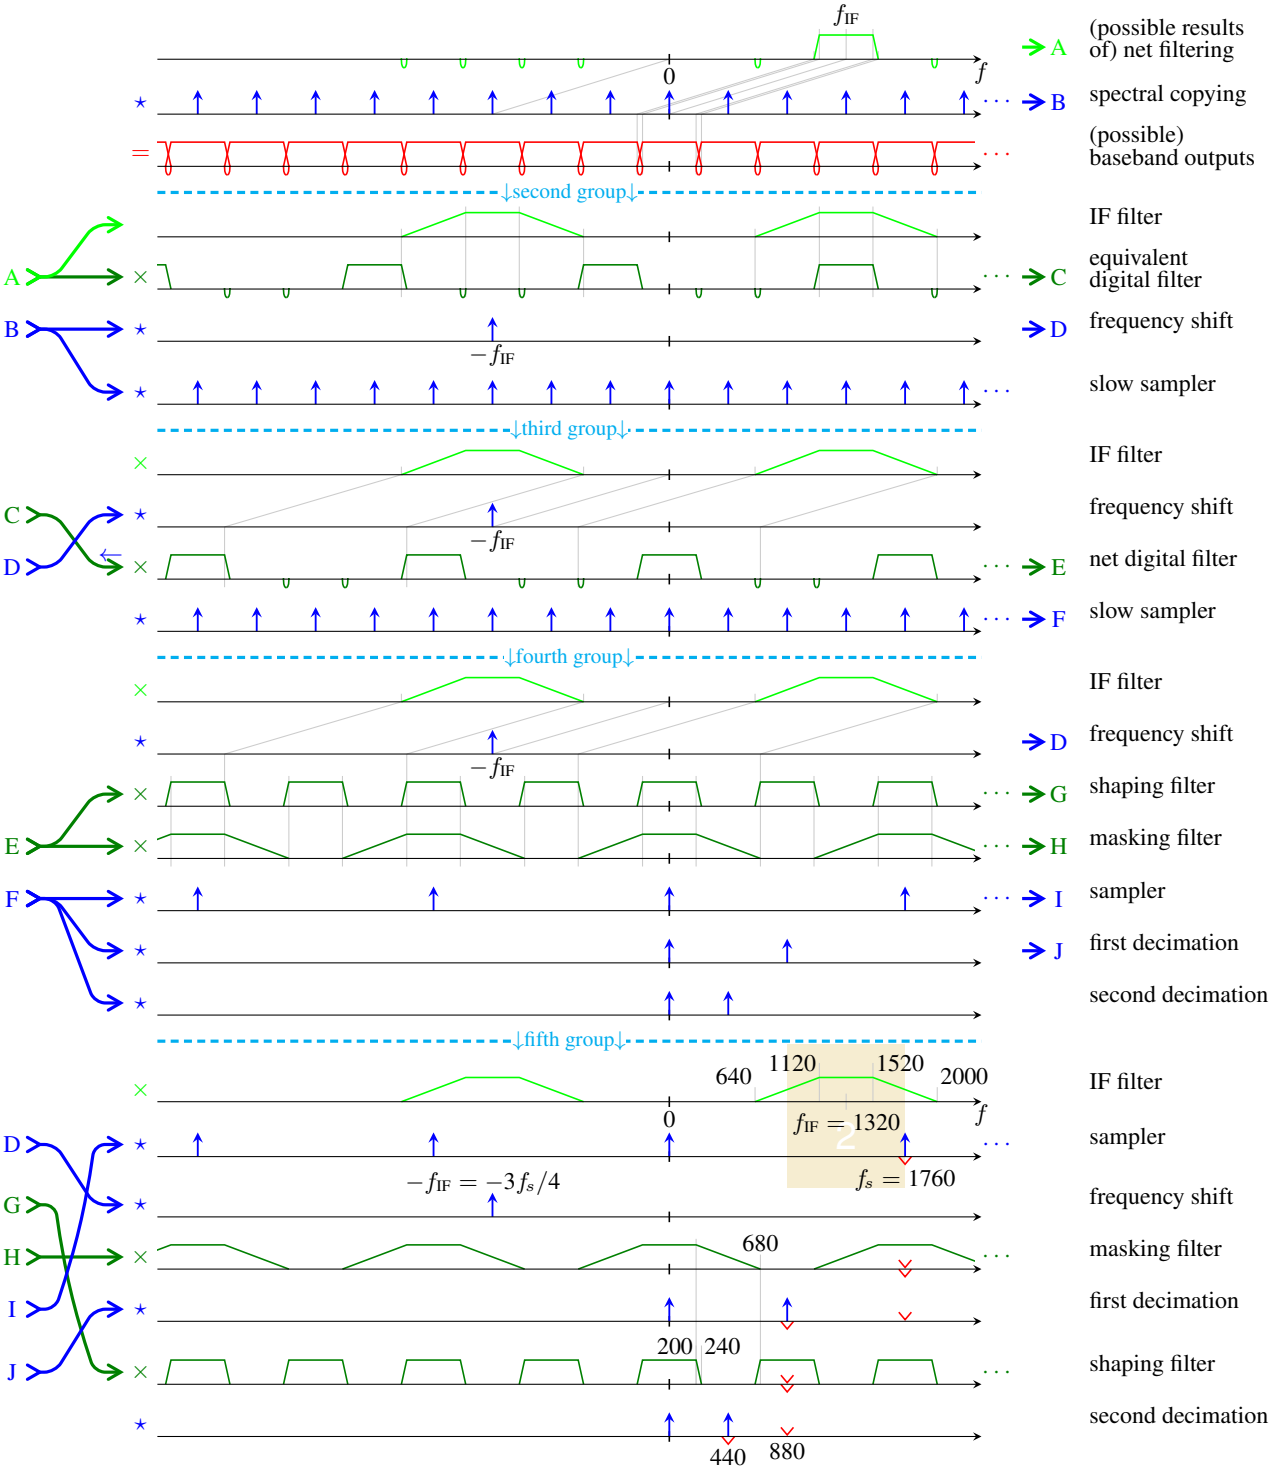

Fig. 6 — Design C. Even wider IF filter transition bands than in Design B of Fig. 5 are permitted here because of a sampling rate that is higher yet, high enough to make decimation by four natural. The digital filter and subsequent decimation by four are split for computational efficiency into two stages of filtering and decimation by two using filters that can be made halfband. Bandpass sampling here uses the second Nyquist zone, a mathematically arbitrary choice governed in practice by hardware considerations.

Table 1 — Band structure of the Design C digital filters in Fig. 6.

filter	band specifications				period	Nth-band option
net digital E	*pass*	*stop*	*stop*	*stop*	four bands	fourthband
shaping G	*pass*	*stop*	*pass*	*stop*	two bands	halfband
masking H	*pass*	—	*stop*	—	four bands	halfband

Equivalent digital filter C of the second group and net digital filter E of the third group have transition or don't-care bands of identical width spaced uniformly across the frequency response, each separated from its neighbors by a fully specified "do care" band, either a passband or stopband, of width equal to the signal band. The contents of all these don't-care bands are ultimately aliased together but end up outside the signal band in the output. That's what makes them don't-care bands.

We can represent the specified bands using the first line of Table 1. The rest of the table shows how the corresponding bands of component filters G and H are specified in order to obtain the desired behavior in net digital filter E. Certainly both component filters must have passbands where net digital filter E does. Further, each stopband of net digital filter E must also be a stopband for at least one of component filters G and H. Given those basic constraints, the goal is to make the implementations of the component filters computationally efficient by giving their impulse responses relatively few nonzero samples. This is done for shaping filter G by giving its frequency response a small period to make its impulse-response samples widely spaced. This is done for masking filter H by keeping the latter's impulse response quite short using huge frequency-response transition regions. Those transition regions are largely the result of giving masking filter H a don't-care region, marked "—" in the table, wherever shaping filter G has a stopband.

Shaping filter G and masking filter H here can each be made halfband. If both are made halfband, net digital filter E will be fourthband automatically.

The substantial benefits of this multirate approach to filtering are in the computational efficiency of the DSP required to implement it. Seeing this with greater clarity properly requires a specific design context, so we defer that discussion to the Part II material on filter design. Similar multirate strategies and the halfband filters so useful in them are discussed in [1, Part IV].

As before, with this approach we are free to change Nyquist zones by changing the IF frequency by any integral multiple of half the sampler's rate and modifying the IF filter and frequency shift accordingly. Further, new and extraordinarily efficient halfband-filter architectures make it quite practical to use a much smaller excess output bandwidth, here arbitrarily set to 10%, though there is no obvious reason to do so. Changing the excess output bandwidth would change f_s and f_{IF} as well.

2.3.4 Designs D and E—Decimation by Six in Two Steps

Designs D and E, respectively in Figs. 7 and 8 (pages 18 and 19), increase the decimation ratio yet again, now to six, using slightly different approaches. This is the highest decimation ratio yet, so the transition bands of the IF filter are the widest here as well. IF frequency f_{IF} was put in the third Nyquist zone, which given this relatively high sampling rate suggests RF sampling, but the IF could easily be lowered

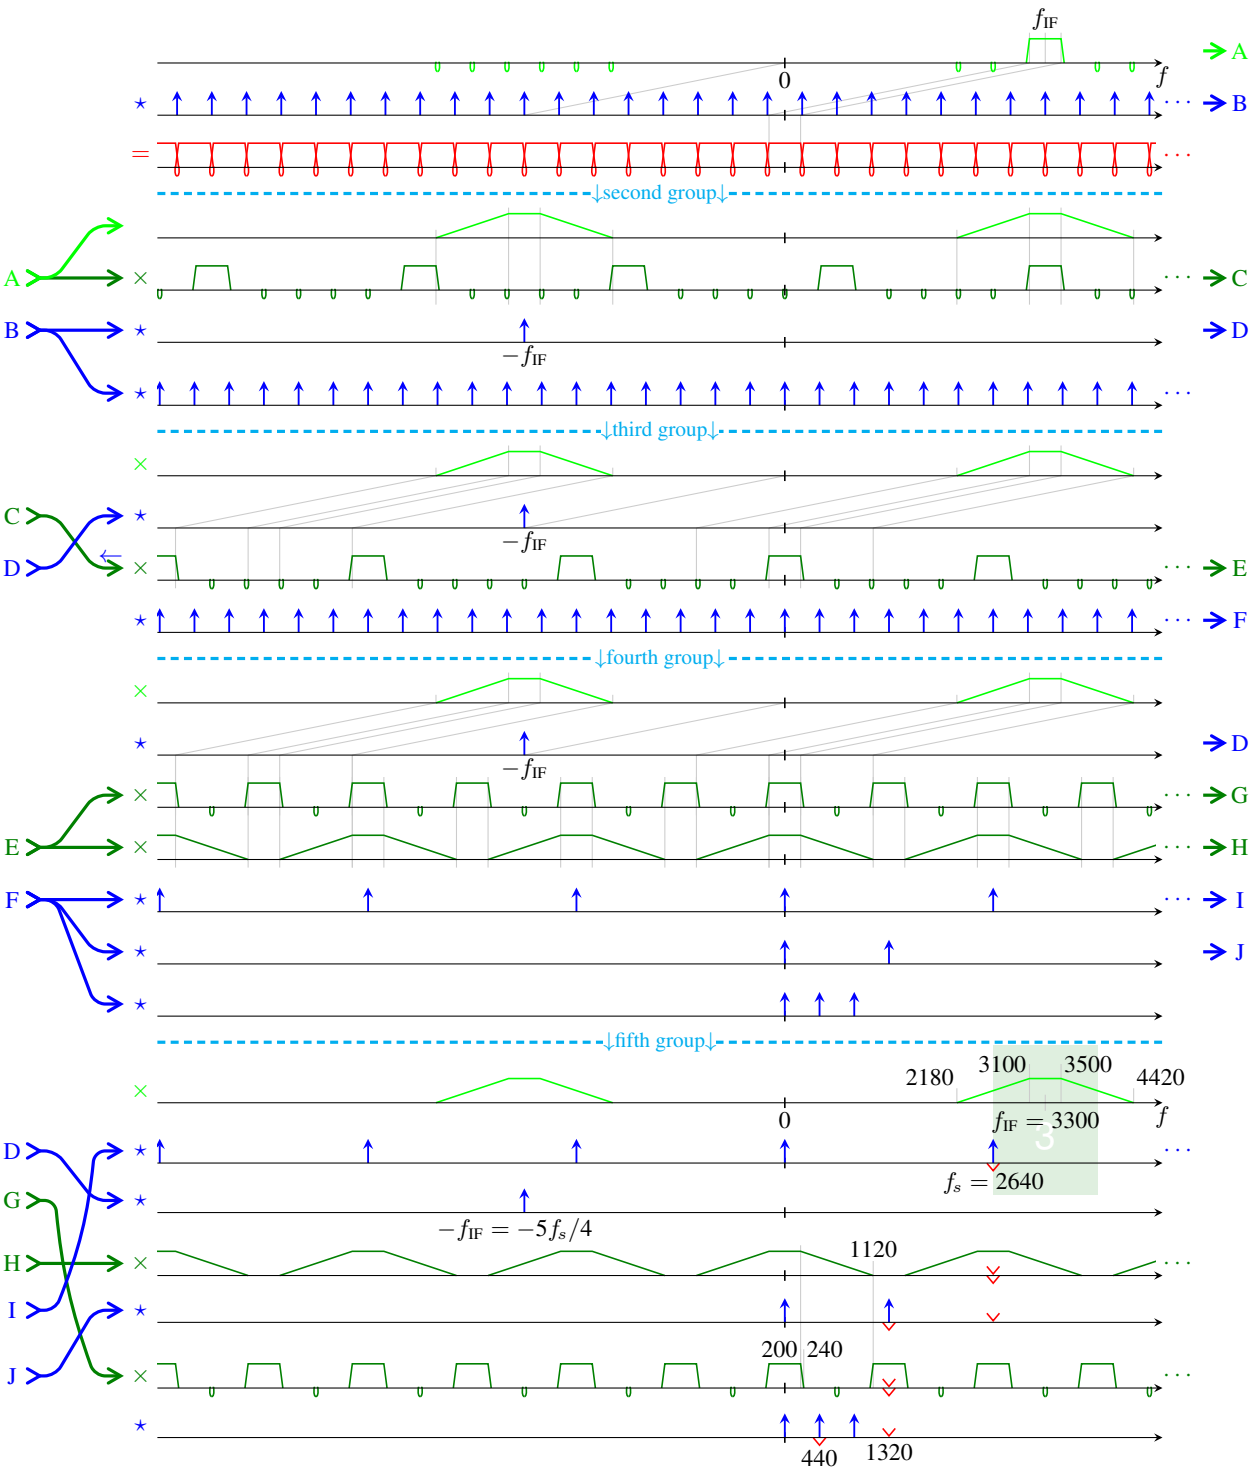

Fig. 7 — Design D. Decimation by three and then by two or by six overall, to allow very wide IF-filter transitions. Shaping filter G and masking filter H can be thirdband and halfband, respectively, for computational efficiency. Use of the third Nyquist zone is illustrated but not required. See Figs. 8 and 9 for variants.

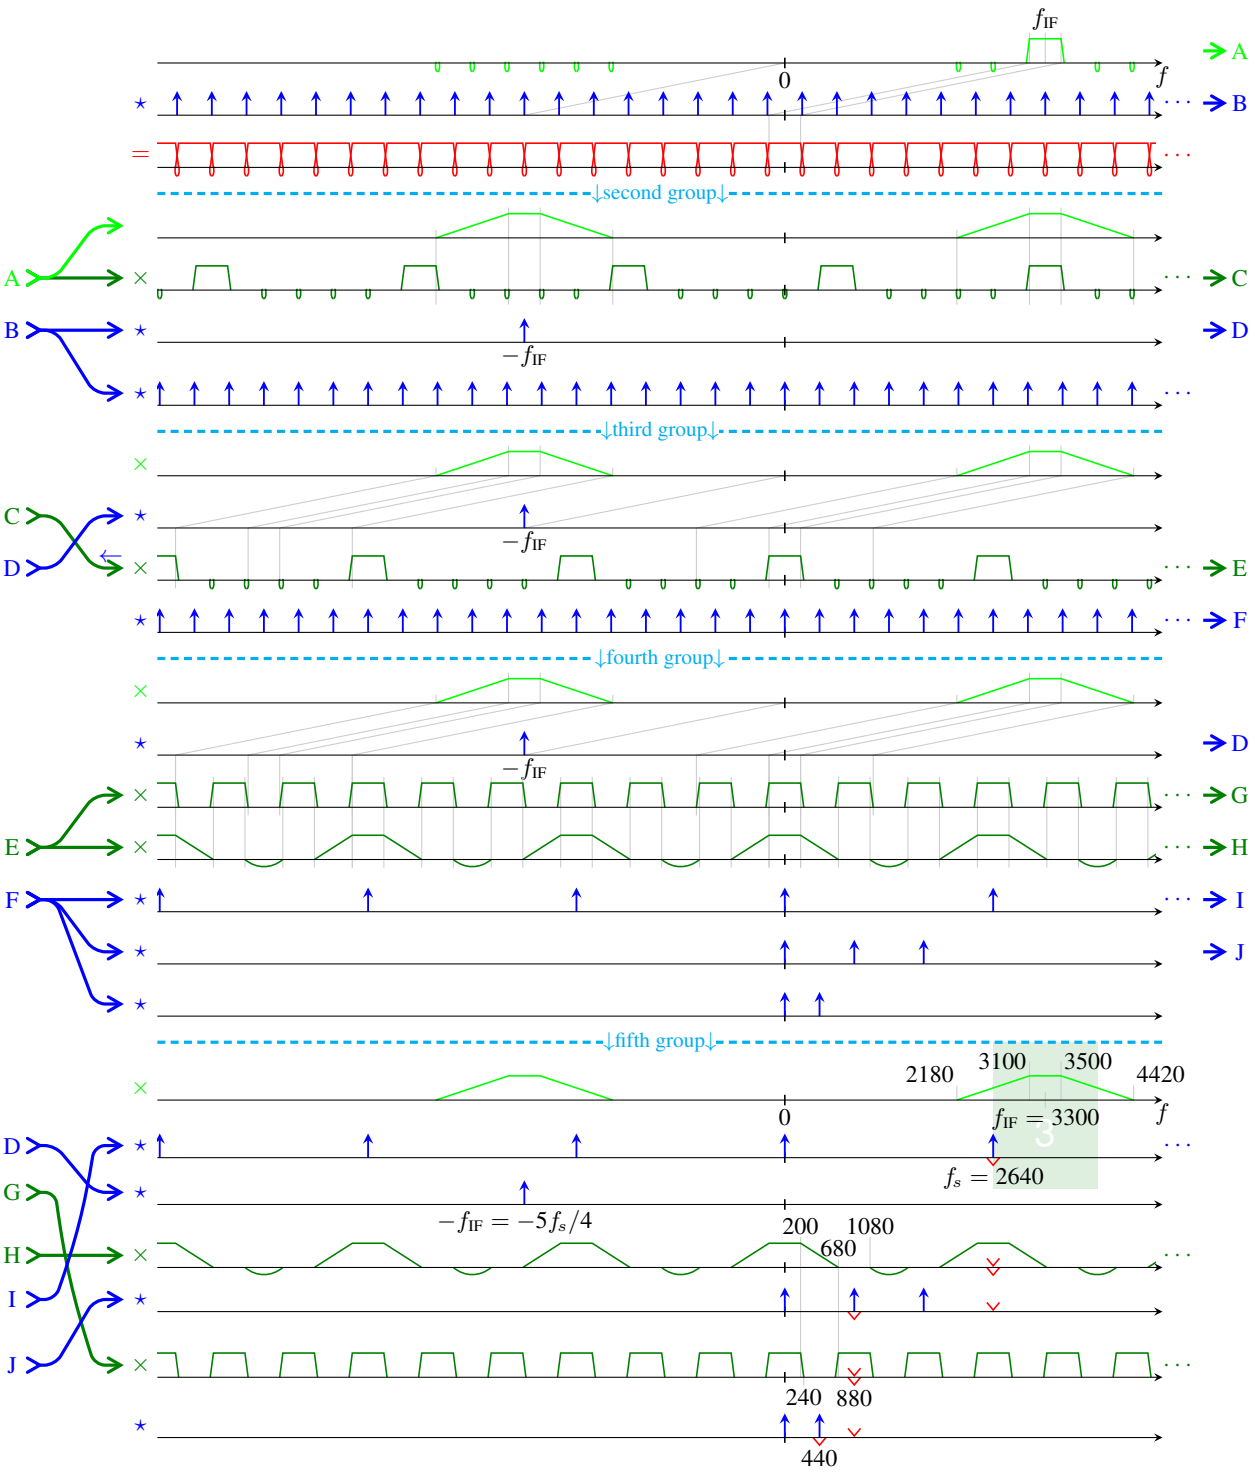

Fig. 8 — Design E. This variation on Design D of Figs. 8 differs in the fourth and fifth groups because it decimates by two and then by three, or by six overall, so that here shaping filter G and masking filter H can be halfband and thirdband, respectively. See Figs. 7 and 9 for variants.

Table 2 — Digital-filter band structure for Designs D and E.

Design D of Fig. 7

filter	band specifications						period	Nth-band option
net digital E	*pass*	*stop*	*stop*	*stop*	*stop*	*stop*	six bands	sixthband
shaping G	*pass*	*stop*	*stop*	*pass*	*stop*	*stop*	three bands	thirdband
masking H	*pass*	—	—	*stop*	—	—	six bands	halfband

Design E of Fig. 8

filter	band specifications						period	Nth-band option
net digital E	*pass*	*stop*	*stop*	*stop*	*stop*	*stop*	six bands	sixthband
shaping G	*pass*	*stop*	*pass*	*stop*	*pass*	*stop*	two bands	halfband
masking H	*pass*	—	*stop*	—	*stop*	—	six bands	thirdband

by $f_s/2 = 1320$ to 1980 to create what would presumably then be an IF-sampling system operating in the second Nyquist zone.

The differences between Designs D and E are in their fourth and fifth line groups, where the third group's net digital filter E and slow sampler F are split into components differently in the fourth group. The particulars of factoring slow sampling into fast sampling and decimation here allow noble-identity compatibility with the chosen filter factorizations, which are key.

Table 2 outlines the arrangement of the specified bands in the two designs. Having each of shaping filter G and masking filter H be Nth-band as indicated on the right in Table 2 will make net digital filter E a sixthband filter automatically.

There is no way to know which of Designs D and E is more computationally efficient short of simply designing both systems and comparing.

2.3.5 Design F—A Nonideal Ratio of IF to Sampling Rate

Each of Designs A through E has $f_{\text{IF}} = (2m - 1)/4$ for some positive integer m, a choice that gives the IF filter's transition bands equal widths above and below the passband. Here, however, we examine what must change when that relationship does not hold. In particular, we modify Design D of Fig. 7 to pair a slightly lower sample rate f_s with the same f_{IF} as before. The result appears in Fig. 9 (page 21).

The first thing that must change is the extent of the transition bands of the IF filter. This arises in the second group, where the positive-frequency passband of the IF filter is aligned with a passband of equivalent digital filter C, exactly as before. The negative-frequency content of the frequency response of IF filter A, excluding only that small part of the content that will alias out of the signal band in the final output, must fall in the interval extending between adjacent passbands of equivalent digital filter C as before. However, the negative-frequency passband of A is no longer centered in that interval, because of the new relationship between f_s and f_{IF}. This is how the initial sampling rate—the frequency-response period of equivalent digital filter C—drives the requirement that the two transition bands of IF filter A have unequal widths.

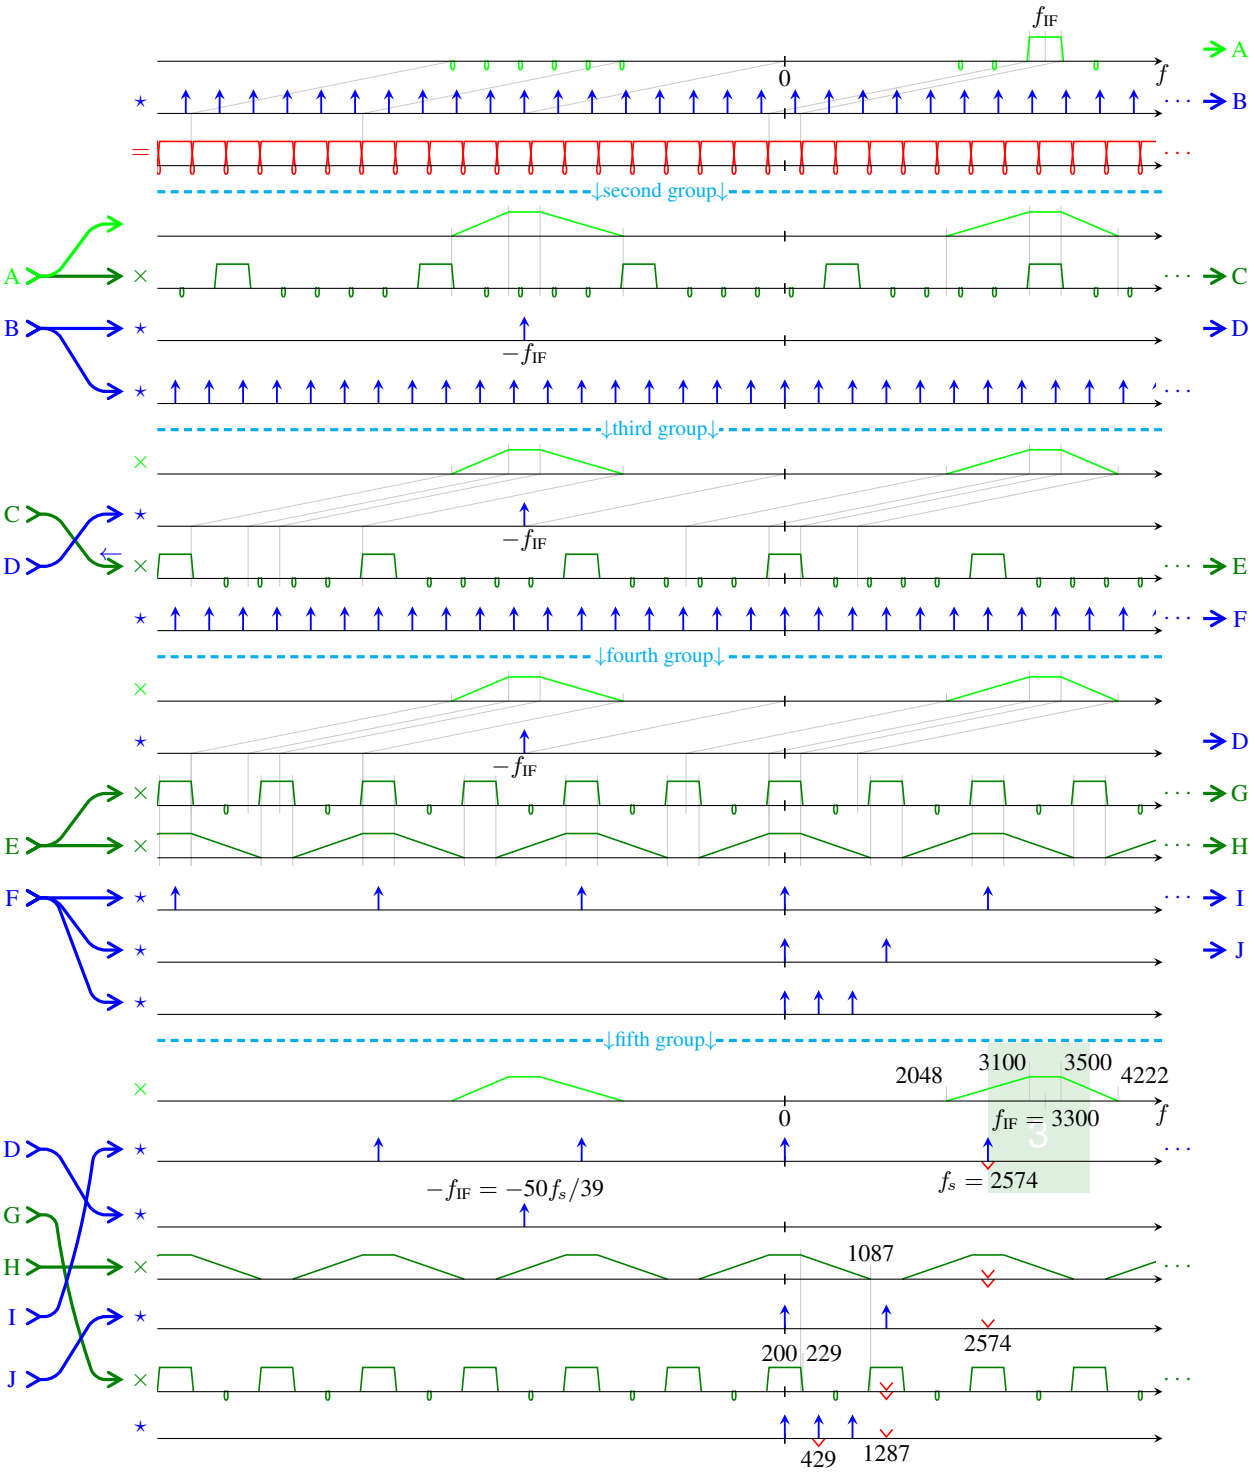

Fig. 9 — Design F. Like Design D of Fig. 7 but with the input sample rate lowered slightly, perhaps to respect some limit on A/D speed. This yields an asymmetric IF-filter transition-band specification and a less-convenient frequency shift. Placing the IF here at $f_{IF} = 50f_s/39$ gives the shift's complex exponential 39 distinct values.

The second thing that must change as consequence of breaking the "standard" relationship between f_{IF} and f_s is the amount of the frequency shift. Here it no longer corresponds to a complex exponential with an nth sample equal to j^n or j^{-n} as before, so the trivial implementations associated with the latter are no longer available. If $f_{\text{IF}} = Jf_s/K$ with positive integers J and K relatively prime, frequency-shift complex exponential $\mathrm{e}^{-j2\pi f_{\text{IF}}t}$ at time $t = n/f_s$ becomes $\mathrm{e}^{-j2\pi Jn/K}$, which can take on just K distinct values. At first glance, implementation would therefore require a short table—K complex entries—and a complex multiply per sample (at the relatively high input sample rate). However, in many cases—the details are a topic for the Part II discussion of computational structures—simple structural transformations can refer the shift to a lower-rate part of the system, for the sake of computational efficiency.

Can the Nyquist zone be changed as with earlier designs? Can the IF frequency be changed in this scheme with minimal change elsewhere? Yes, it can. The second group of Fig. 9 shows that the negative-frequency portion of IF filter A's frequency response falls into equivalent digital filter C's third intra-passband interval to the left of the signal band. It could just as well be placed in the second or fourth such interval, with the horizontal positioning of the first two lines of the group then adjusted as a unit to center the frequency response of IF filter A on the origin. This change in f_{IF} would of course also change the amount of frequency shift D, but nothing else would change. The shaping and masking filters and the decimations would be the same. Given the particular f_s value pictured then, $f_{\text{IF}} = 3300$ could be replaced with $f_{\text{IF}} = 2013$ or $f_{\text{IF}} = 4587$.

2.3.6 Design G—IF-filter Transition Bands of Equal Ratiowise Widths

IF filters, particularly in certain frequency ranges, are sometimes designed to have rational transfer functions arrived at through lowpass-to-bandpass transformation. The usual transformation gives the two transition bands identical ratios of upper and lower edge frequencies, making those transition bands equal in width ratiowise but not differencewise. Their widths appear equal on a log frequency scale but not on a linear frequency scale. Figure 10 shows two ways to parameterize corner frequencies for such an IF filter. In the first corner-frequency parameterization shown there, for example, each transition-band has an upper-to-lower edge-frequency ratio of β/α. In the second parameterization, that ratio is f_1/f_0 for each transition band.

Moving back and forth between the two parameterizations will prove useful below. It uses three relationships that each equate the same quantity in the two systems. For the first, second, and third relationship, those quantities are the ratio of the upper stopband corner to the lower passband corner, the upper passband corner frequency, and the ratio of the upper stopband and passband corner frequencies:

$$\alpha\beta = \frac{m}{m-1}, \qquad (1) \qquad\qquad \alpha f_c = m f_0, \qquad (2) \qquad\qquad \frac{\beta}{\alpha} = \frac{f_1}{f_0}. \qquad (3)$$

To obtain formulae for obtaining the second parameterization from the first, solve (1) for m, solve (2) for f_0, and solve (3) for f_1. To obtain formulae for obtaining the first parameterization from the second, divide (1) by (3) and solve for α, multiply (1) by (3) and solve for β, and solve (2) for f_c.

It is no surprise that (1) requires $m > 1$. Figure 10 shows that $m > 1$ is also required for the lower stopband corner to be a positive frequency and for the lower transition band to have a nonzero width.

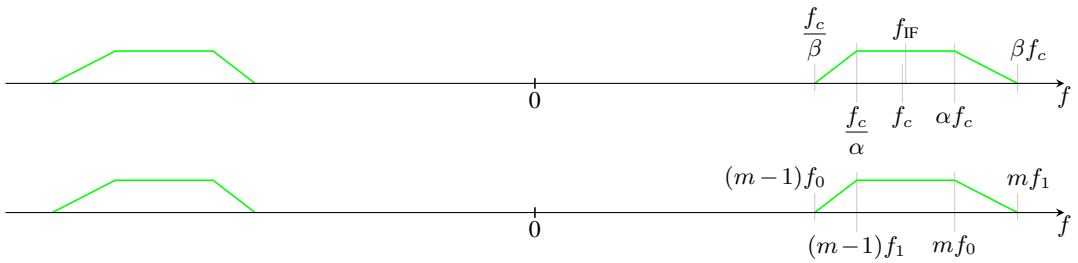

Fig. 10 — Two parameterizations of the IF-filter corner frequencies produced by a standard lowpass-to-bandpass transformation. In general parameter m is real with $m > 1$, but the approach of Design G of Fig. 11 further restricts it to integer values. As a practical matter, that restriction quantizes β given an arbitrary α.

Ratiowise equality of transition-band widths is useful well beyond systems based strictly on lowpass-to-bandpass transformation, because it's a reasonably natural way to balance stopband-suppression resources between upper and lower transition bands in the design of the IF filter. Here, however, let us take the lowpass-to-bandpass transformation as motivating and begin by examining that transformation. Unlike the usual presentations [3], this one focuses on verifying the key ratiowise-equality property of the transition-band widths. Readers in no need of convincing that systems designed with such a transformation possess this property may wish to skip this next small section.

The Lowpass-to-Bandpass Transformation

This derivation uses the first Fig. 10 parameterization, though it is moved to radian frequency with $\omega = 2\pi f$ for convenience.

The design process begins with some prototype lowpass filter's rational transfer function $L(s)$. Given a positive constant α and positive radian frequencies ω_p and ω_c, we construct the new transfer function as

$$H(s) = L\left(\frac{-\omega_p\,\alpha\,(\omega_c^2 + s^2)}{\omega_c\,s\,(\alpha^2 - 1\,)}\right). \tag{4}$$

Suppose ξ is radian frequency normalized by ω_c so that $s = j\xi\omega_c$. At that frequency our new transfer function has value

$$H(j\xi\omega_c) = L\left(j\omega_p\,\frac{\alpha\,(\,1 - \xi^2)}{\xi\,(\alpha^2 - 1\,)}\right). \tag{5}$$

Especially interesting are the values $\xi = \pm 1/\alpha$ and $\xi = \pm\alpha$, because they result in $H(\pm j\omega_c/\alpha) = L(\pm j\omega_p) = H(\mp j\omega_c\alpha)$. The sign flip in the second equality is a tad awkward, but a prototype filter with a real impulse response gives the numerator and denominator polynomials of $L(s)$ real coefficients so that conjugate $L^*(j\omega) = L(-j\omega)$, and the definition (4) of $H(s)$ then similarly implies $H^*(j\omega) = H(-j\omega)$. Using a "$-$" in the argument then as a way to conjugate these functions, we we can certainly write

$$|H(j\omega_c/\alpha)|^2 = |L(j\omega_p)|^2 = |H(j\omega_c\alpha)|^2. \tag{6}$$

Since $H(s)$ has the same magnitude response at $\omega_c\alpha$ and ω_c/α that lowpass prototype $L(s)$ has at ω_p, it is convenient to let ω_p be some important reference frequency for the prototype. Let us take it to be the passband corner in particular so that ω_c/α and $\omega_c\alpha$ become the passband corners of bandpass response $H(s)$.

To get the stopband corners then, write the argument to $L(\ )$ in (4) as a product

$$-\frac{\omega_c^2+s^2}{\omega_c s}\times\frac{\omega_p}{\alpha-\frac{1}{\alpha}}$$

of two factors, with the first factor containing all dependency on s and with the second, a constant, containing all dependency on ω_p and α. We can rewrite that constant to replace passband cutoff ω_p with the lowpass prototype's stopband cutoff ω_b and to replace α with a constant β, greater than unity, chosen to obtain the equality

$$\frac{\omega_p}{\alpha-\frac{1}{\alpha}}=\frac{\omega_b}{\beta-\frac{1}{\beta}}. \tag{7}$$

A suitable β clearly exists, because $\beta-1/\beta$ on the right can be made to equal any positive constant whatever by appropriate choice of β greater than unity. Given these new constants ω_b and β then, we can rewrite (4) as

$$H(s)=L\left(\frac{-\omega_b\beta(\omega_c^2+s^2)}{\omega_c\, s\,(\beta^2-\ 1\ )}\right).$$

so that (4) through (6) apply exactly, by all the same arguments, if we simply replace α with β and ω_p with ω_b. In particular,

$$|H(j\omega_c/\beta)|^2=|H(j\omega_c\beta)|^2=|L(j\omega_b)|^2.$$

Now we can interpret ω_c/β and $\omega_c\beta$ as the stopband corners of bandpass response $H(s)$.

Rearranging (7) yields a small result that will prove useful later: the ratio of the difference between stopband corners to the difference between passband corners is just

$$\frac{\beta\omega_c-\frac{\omega_c}{\beta}}{\alpha\omega_c-\frac{\omega_c}{\alpha}}=\frac{\omega_b}{\omega_p}. \tag{8}$$

Finally, substituting $\xi=\pm1$ into (5) yields $H(\pm j\omega_c)=L(0)$, the DC value of $L(s)$, revealing that the lowpass prototype's DC magnitude response is mapped to frequency ω_c in the bandpass filter. The geometric mean of the passband cutoff frequencies is ω_c, as is the geometric mean of the stopband cutoff frequencies.

The IQ-Demodulator Architecture

Design G in Fig. 11 illustrates the building of a demodulator architecture around an IF filter characterized by the second parameterization of Fig. 10 specialized to integer m, which yields a demodulator operating in the mth Nyquist zone. The Fig. 1 frequency relationships hold exactly, with the integer m being the same in Figs. 1 and 11. As for all designs thus far, the output sample rate is $1/M$ of the input sampling rate f_s for some decimation ratio M.

The number of design degrees of freedom in this system are surprisingly few. Assuming the desired width—call it f_{wid}—of the signal band and the Nyquist-zone number m are known (more on the latter shortly), two cases are described below, cases corresponding to one of f_{IF} and f_s being given with the other then determined by m. Begin with key definition of $f_{\text{wid}} = mf_0 - (m-1)f_1$ as the difference of the passband corner frequencies in Fig. 11.

The Two Cases

Then if f_{IF} is given, take its definition as the average of the passband corners in Fig. 11, so that $2f_{\text{IF}} = mf_0 + (m-1)f_1$, as key as well. Those two key definitions can be almost trivially solved as simultaneous equations for

$$f_0 = \frac{f_{\text{IF}} + f_{\text{wid}}/2}{m} \tag{9}$$

$$f_1 = \frac{f_{\text{IF}} - f_{\text{wid}}/2}{m-1}. \tag{10}$$

Corner-frequency ordering in Fig. 11 requires $f_0 < f_1$, which by (9) and (10) is equivalent to

$$m < \frac{f_{\text{IF}}}{f_{\text{wid}}} + \frac{1}{2}. \tag{11}$$

As per Fig. 11, sample rate f_s is computed as

$$f_s = f_0 + f_1. \tag{12}$$

Then if f_s is given, take its definition $f_s = f_0 + f_1$ from Fig. 11 as key as well. Solve those two key definitions as simultaneous equations for

$$f_0 = \frac{(m-1)f_s + f_{\text{wid}}}{2m-1} \tag{13}$$

$$f_1 = \frac{mf_s - f_{\text{wid}}}{2m-1}. \tag{14}$$

IF frequency f_{IF} is computed as the average of the passband corners in Fig. 11:

$$f_{\text{IF}} = \frac{mf_0 + (m-1)f_1}{2}.$$

In principle Nyquist-zone number m has no upper limit and therefore neither does f_{IF} here. In practice one simply imposes a practicality limit somewhere.

The four IF-filter corner frequencies in Fig. 11 (or in the IF-filter parameterizations of Fig. 10) can be computed from m, f_0, and f_1. Net decimation factor $1/M$ can be chosen in the usual way. (Section 2.4 ultimately extends the latter choice to ratio L/M.) In either case above, (13) and (14) must hold. Substituting them into $f_1 > f_0$, which is just the requirement that the stopband corners must be outside the passband corners, shows that inequality to be equivalent to $f_s > 2f_{\text{wid}}$, the expected Nyquist sampling limit.

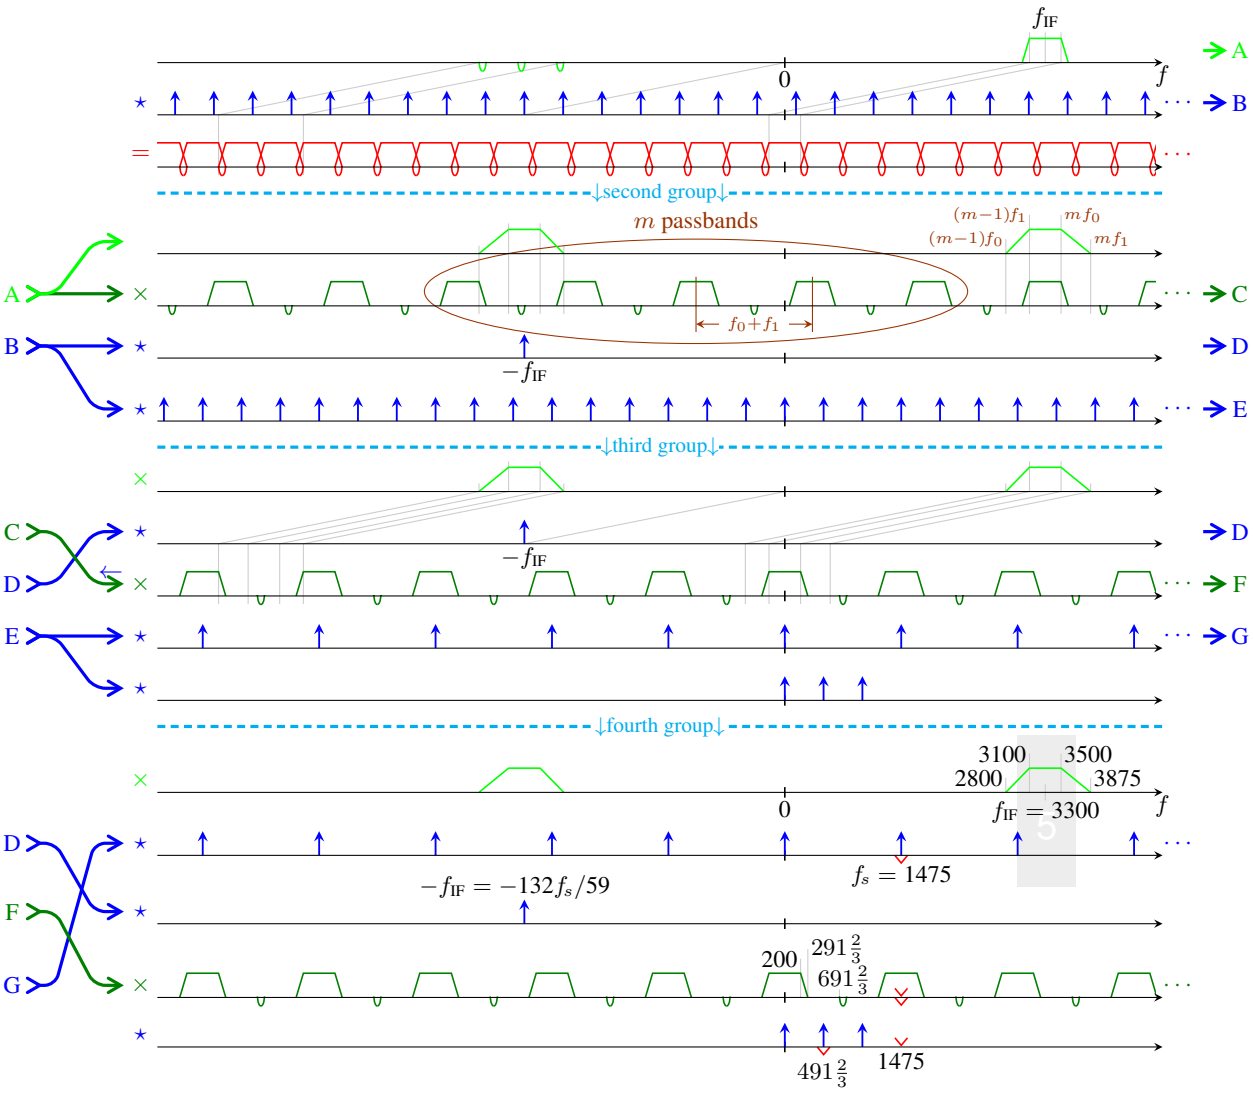

Fig. 11 — Design G. Like Design F of Fig. 9, here IF-filter transition bands have unequal widths computed as differences. Here, however, they do have equal widths computed as ratios, which is convenient when the transfer function of the IF filter is obtained from a standard lowpass-to-bandpass transformation. Second-group relationships show how to construct a compatible decimation-based demodulator for the mth Nyquist zone using frequencies f_0 and f_1 such that IF-filter bandedges bandedges $(m-1)f_0$, $(m-1)f_1$, mf_0, and mf_1 of the IF filter form a positive and strictly-increasing sequence. The m passbands of the second group are counted from just left of the IF-filter frequency response's positive-frequency support to just left of its negative-frequency support.

Table 3 — Parameter choices for demodulating a 3100–3500 signal band using an IF filter with transition bands of equal ratiowise widths.

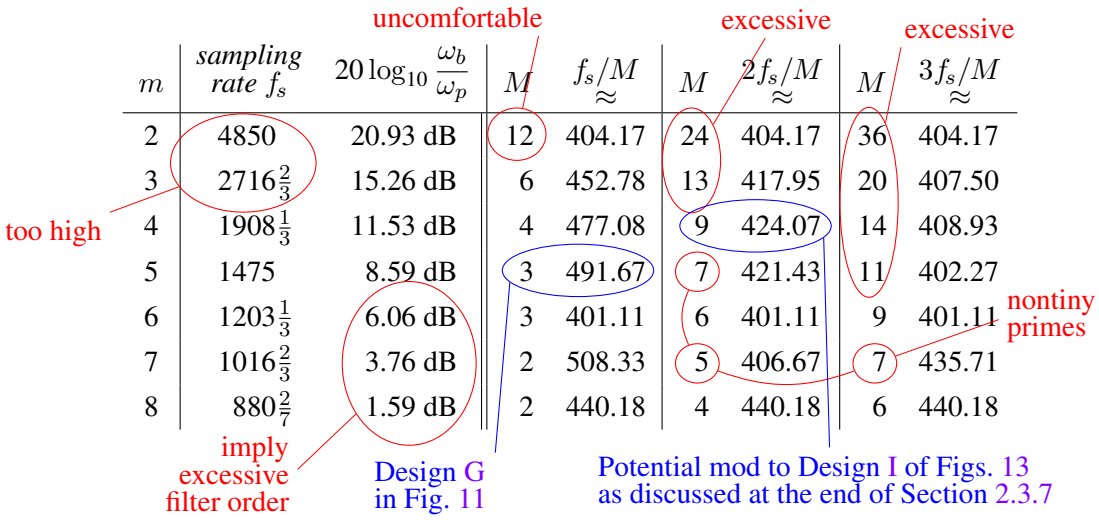

m	*sampling rate* f_s	$20\log_{10}\frac{\omega_b}{\omega_p}$	M	$f_s/M \approx$	M	$2f_s/M \approx$	M	$3f_s/M \approx$
2	4850	20.93 dB	12	404.17	24	404.17	36	404.17
3	$2716\frac{2}{3}$	15.26 dB	6	452.78	13	417.95	20	407.50
4	$1908\frac{1}{3}$	11.53 dB	4	477.08	9	424.07	14	408.93
5	1475	8.59 dB	3	491.67	7	421.43	11	402.27
6	$1203\frac{1}{3}$	6.06 dB	3	401.11	6	401.11	9	401.11
7	$1016\frac{2}{3}$	3.76 dB	2	508.33	5	406.67	7	435.71
8	$880\frac{2}{7}$	1.59 dB	2	440.18	4	440.18	6	440.18

We can estimate the stopband suppression obtained in the IF filter using the order of the lowpass prototype filter and, from (8),

$$\frac{\omega_b}{\omega_p} = \frac{\text{difference between stopband corners}}{\text{difference between passband corners}} = \frac{f_s - f_{\text{wid}}}{f_{\text{wid}}} = \frac{f_s}{f_{\text{wid}}} - 1 \tag{15}$$

Nyquist limit $f_s > 2f_{\text{wid}}$ ensures that this exceeds unity. All-pole lowpass transfer functions, including classics like the Butterworth, Type-I Chebyshev, etc., have asymptotic stopband suppression of 20 dB per frequency decade per pole. A reasonable estimate then of stopband suppression when such prototypes are used is the prototype filter order times $20\log_{10}(\omega_b/\omega_p)$. For Type-II Chebyshev or Elliptic or other pole-zero prototype filters, the relation between stopband suppression and ω_b/ω_p is not so simple, but even then stopband suppression remains a function of filter order and ratio ω_b/ω_p or, equivalently, $20\log_{10}(\omega_b/\omega_p)$.

We can arrive at Design G of Fig. 11 in particular by first assuming $f_{\text{wid}} = 400$ and $f_{\text{IF}} = 3300$ and then choosing Nyquist zone m and decimation ratio M. Let's walk through the process. First, the upper bound of (11) here becomes $m < 8.75$, so it is enough to consider $m = 2, \ldots, 8$. The part of Table 3 to the left of the $\|$ double line tabulates these m values and the values that then result for sampling rate f_s using (9), (10), and (12) and the estimated IF-filter stopband suppression per net prototype pole using (15). (More on the other columns shortly.) We might eliminate the two smallest m values as candidates, supposing each to require an excessive sampling rate f_s relative to A/D technology we're willing to consider. We might then eliminate the three highest m values as each requiring more than six prototype poles to get to a hypothetical minimum stopband suppression of 50 dB. These choices would leave only $m = 4$ and $m = 5$ as possibilities.

The architectures presented so far, including Design G of Fig. 11, use a combination of filtering and decimation by M to transform a signal at the input sampling rate f_s into one at an output sample rate of f_s/M. That output rate must exceed signal-band width f_{wid} to avoid signal aliasing at the system output,

and that constrains the decimation ratio:

$$M < \left\lfloor \frac{f_s}{f_{\text{wid}}} \right\rfloor .$$

The first of three double columns to the right of the $\|$ double line in Table 3 shows that bounding value of M and the associated output rate f_s/M. (The other columns are discussed below, at the end of Section 2.3.7.) Both quantities are reasonable for both $m = 4$ and $m = 5$. For this design, the latter was chosen to hold decimation to a single stage for simplicity, since decimation is not the topic of most interest in this discussion.

While Design G of Fig. 11 is relatively simple, the output sample rate is the highest since Design A of Fig. 4. The next section shows how to obtain more choices of output sample rate.

2.3.7 Designs H and I—IQ Demodulators Embedding Rational Sample-Rate Conversion

In Designs H and I in Figs. 12 and 13, the ratio of the input sample rate f_s to the output sample rate is respectively $8/3$ and $9/2$. In previous designs this was a decimation ratio and therefore an integer, and for practicality that integer needed to be highly composite. This is the first time a ratio of such integers is permitted instead.

Consider Design H of Fig. 12 in particular. The major change from previous designs comes in the second group, where a new operation has been introduced on the second line. Because it is spectral convolution with impulses at three uniformly spaced points beginning with the origin, it certainly looks like a decimation. But our decimations have always arisen from being factored out of lower-speed sampling operations to create higher-speed sampling operations. Here, however, the opposite happens: in the fourth group this operation is combined with high-speed sampling to create lower-speed sampling. Further, in the second and third groups, where this new operation stands alone, it does not operate on a discrete-time input. Its input is analog. Expressing system operation in terms of such a continuous-time operation is certainly not standard-issue DSP theory, and for that reason the operation doesn't have a standard name. It will act a little like sampling but clearly isn't sampling, so both here and where it is introduced in [1, Part VI] it is termed *faux sampling*.

Let's walk through the operation of the second group in Fig. 12. As before, the output of the IF filter falls into two spectral regions, one at positive frequencies and one at negative frequencies. We are accustomed to using an equivalent digital filter like C to pass the signal band in one of those components, by convention the one at positive frequencies, while rejecting everything else, or at least everything else that might end up in the signal band as a result of slow-sampling operation E.

What is new here is that faux sampling turns each spectral component from the IF filter into into three spectral copies that are identical except for frequency offsets. We still need equivalent digital filter C to pass one signal band on the right, but we now have three such bands to choose from, as each of the three spectral copies contains identical information. Here we take the signal band from the leftmost of the three, the one at the lowest frequency, but we could just as well have picked one of the others (and below in Design J of Fig. 15 we actually will).

Just as in earlier designs, once we have arranged for equivalent digital filter C to pass a signal band, we need to arrange for that filter to also suppress everything else, at least everything else that slow-sampling

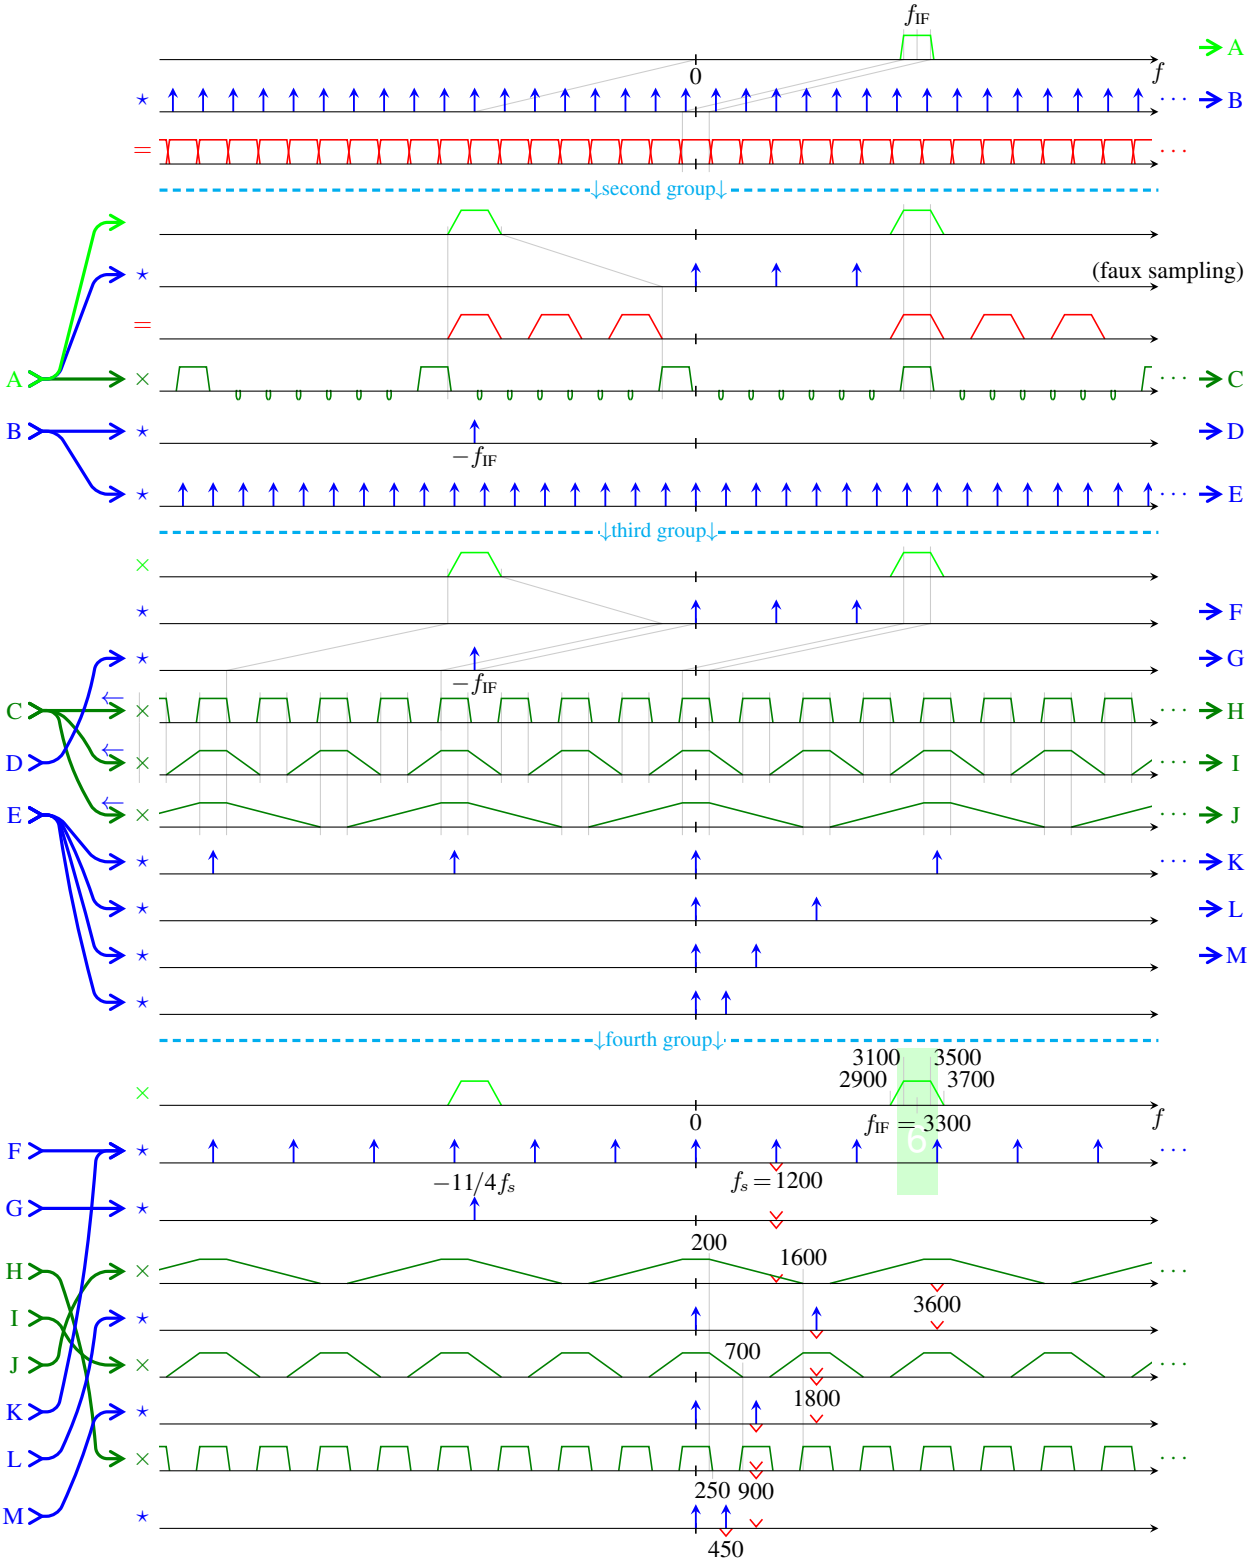

Fig. 12 — Design H. It has three decimations and two masking filters. Input and output sample rates have the ratio 8/3, which can be nonintegral because of the faux sampling—three impulses only—in the second group that looks like but never functions as a decimation. In the fourth group it combines with fast sampling K to become sampling at lower rate f_s. It also results in zero interpolation at the first masking-filter input. Each digital filter can be halfband.

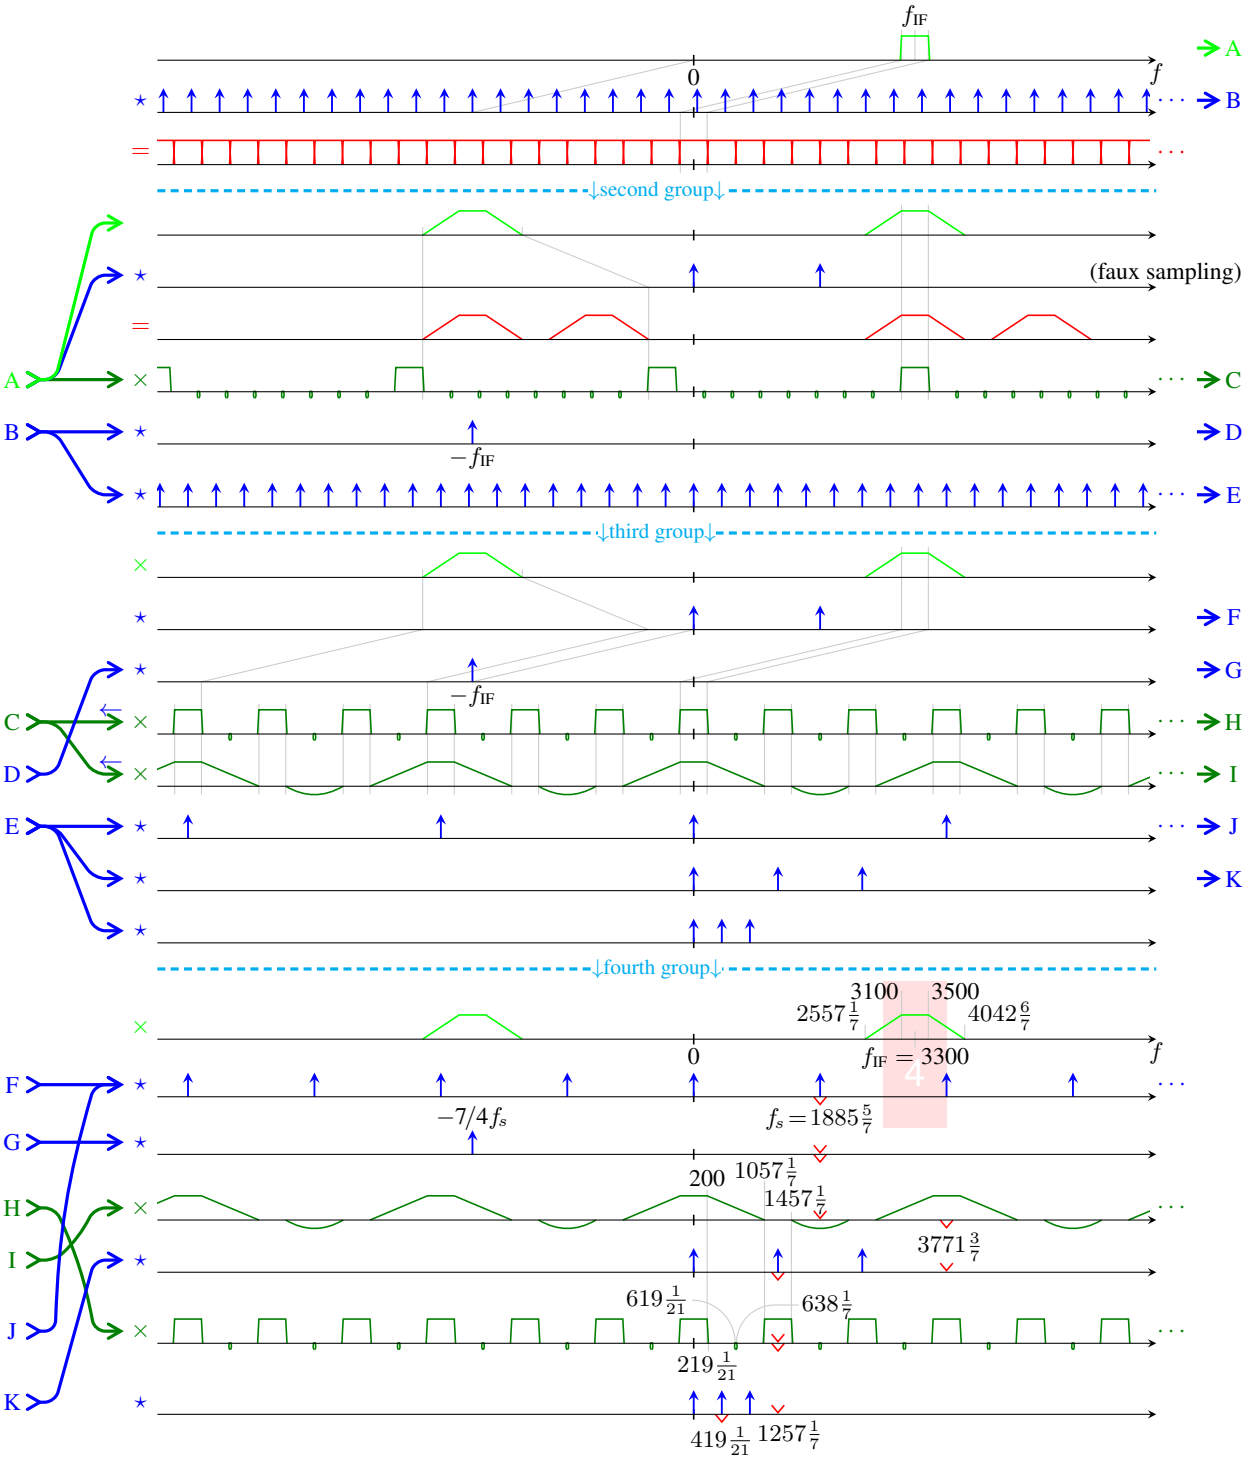

Fig. 13 — Design I. Like Design H this design uses faux sampling in the second group and therefore zero interpolation at the masking-filter input in the fourth group. Here, however, faux sampling uses just two spectral impulses, so zero interpolation is by a factor of two. Both the shaping and masking filters can be thirdband.

operation E would map into the output signal band. What is different in this design is that there is much more spectral material to be rejected. At the faux-sampling output, we need to reject all three negative-frequency spectral components, we need to reject the two positive-frequency spectral components that we are not using at all, and, as usual, we need to steepen the transition bands around the signal band by rejecting most of those transition bands' spectral content.

Once equivalent digital filter C has passed what we want and rejected what we don't in the second group, we proceed much as in earlier designs, except for a minor notational change. In those designs, reversing the order of frequency shift D and equivalent digital filter C shifted the latter's frequency response to create a "net digital filter" that was subsequently split into a shaping and masking filters. Here, however, the order reversal, the frequency-response shifting, and the splitting of that net digital filter into components all happens at once, in the third group.

Also, for the first time here the filter splits three ways, into a shaping filter and two masking filters, with the latter having frequency responses with different periods. Splitting three ways instead of two is not related to faux sampling. It's just that here the frequency-response period of equivalent digital filter E is large, eight times the output sampling rate, so expressing it as a product of small factors naturally requires more factors: $8 = 2 \times 2 \times 2$. Of course this also means that slow-sampling operation E must split—this happens in the third group—into a fast-sampling operation K and three decimations by two.

In the fourth group faux sampling does again introduce something new. There fast-sampling operation K moves up quite far and in the process crosses another spectral convolution and all three filtering operations. The latter crossings of course rely on the most noble identity and are why that sampling operation had to be so fast, as fast as the largest period of the frequency responses crossed. Once these crossings have made K and faux sampling adjacent steps, they are combined to produce lower-speed sampling, at speed $f_s = 1200$.

The rest of the system is much as in earlier designs except that now the frequency response period of the first masking filter is three times the input sampling rate f_s. At first glance the discrete-time input signal to the filter doesn't seem to have time-domain impulses at a high enough rate, but we can simply suppose those impulses to be present but to have zero areas. After all, 0 and $0 \times \delta(t - \tau)$ are certainly indistinguishable mathematically. We implement this scheme using digital filtering of the impulse areas, as usual, but in digital filtering it is customary to acknowledge these assumed-in zero-valued input samples by saying we are "zero interpolating" the input signal to bring it up to the required rate. In implementation those zeros are neither actually created nor multiplied by coefficients. The idea of zero interpolation is a convenience to clarify analysis and to guide the correct structuring of the implementation.

In Design H of Figs. 12 each of the three filters can be made halfband.

This section's other design example, Design I of Fig. 13, is very similar, except that its faux sampling operation uses just two spectral impulses and—these two facts are unrelated—the system uses just two decimations. As a result of the latter, there is only one masking filter. Each decimation is by three, and each of the two filters can be made thirdband.

Modifying Designs H and I to Accomodate IF-Filter Transition Bands of Equal Widths Ratiowise

The examples of this section can be easily modified to accomodate IF-filter transition bands of equal ratiowise widths as discussed above in Section 2.3.6 and as seen in Design G of Fig. 11. These modified

designs are not shown in this document, but it is easy enough to work out roughly how they would look. In each of these designs, using L faux-sampling impulses gives the frequency response of equivalent digital filter E in the second group a period of Lf_s, and this also becomes the zero-interpolated sample rate in the fourth group (where it is flagged by the ▽ tick mark on the frequency response of the first masking filter, in the fourth line). Net decimation by some overall ratio M then yields an output sample rate of Lf_s/M. Modifying Designs H and I to incorporate IF-filter corner-frequency structure of the form used in Design G, the second parameterization of Fig. 10, preserves these basic ratios, but with f_s changed just a little, as the latter will be determined by (9), (10), and (12) using Nyquist-zone number m and the passband corner frequencies.

Consider such a modification of Design I of Figs. 13 in particular. It operates in Nyquist zone $m = 4$, so the $m = 4$ row of Table 3 (p. 27) applies. The rightmost columns of that table, not discussed earlier, apply when $L = 2$ and $L = 3$. Each shows the decimation factor M and the resulting output sample rate when interpolated rate Lf_s is decimated as much as possible consistent with an output sample rate above the formally required minimum (for the example of that table) of 400. The $m = 4$, $L = 2$, and $M = 9$ combination flagged in Table 3 can be used here. Using the second parameterization of Figure 10 and Nyquist-zone number $m = 4$, the stopband corner frequencies move from $2557\frac{1}{7}$ and $4042\frac{6}{7}$ to $3500(m-1)/m = 2625\frac{1}{7}$ and $3100m/(m-1) = 4133\frac{1}{3}$. Table 3 shows that sampling rate f_s would be increased from the $1885\frac{5}{7}$ of Figs. 13 to $1908\frac{1}{3}$, and the output sample rate would increase from $419\frac{1}{21}$ to $424\frac{2}{27}$, so the changes, other than to the IF filter, are quite modest.

A similar modification of Design H of Figs. 12 is equally straightforward but not quite so easily tied to Table 3. Design H operates in Nyquist zone $m = 6$ with $L = 3$. The rightmost column of the $m = 6$ row of Table 3 notes that $M = 9$ would yield an output sample rate of roughly 401.11, which is clearly not enough over the 400 minimum for practicality. Decimation by $M = 8$ instead, as in Design H, would be fine, however. These modifications to Design H would raise the initial sample rate f_s from 1200 to the $1203\frac{1}{3}$ shown in the table and would increase output sample rate Lf_s/M from 450 to $451\frac{1}{4}$. The IF-filter stopband corners move from 2900 and 3700 to $3500(m-1)/m = 2916\frac{2}{3}$ and $3100m/(m-1) = 3720$.

The change in IF-filter corner frequencies was greater in the the Design H modification than in the earlier Design I modification because the difference between arithmetic and ratiowise transition-band symmetry grows smaller as the Nyquist zone number m increases. While there is other material to cover before examining Design J of Fig. 15 (p. 37) below in detail, we can look ahead and see that symmetry difference most clearly there, because that design operates in the the second Nyquist zone.

2.4 General Design Relationships

Section 2.3 above presents many design examples and wraps up with two designs in Section 2.3.7 that use faux sampling, Designs H and I in Figs. 12 and 13. In the second group of each of those designs, there are two spectral relationships of particular interest.

Relationship 1: The signal band passed by equivalent digital filter C is from the lowest-frequency spectral component that faux sampling creates from the positive-frequency component of the IF output spectrum.

Relationship 2: All components that faux sampling creates from the negative-frequency portion of the IF output spectrum fit neatly between two of the periodic passbands of equivalent digital filter C.

It turns out that neither of those two relationships is actually required. Figure 14 illustrates the general case, not with a complete design example—that comes below in Design J of Fig. 15—but with key features from the first four lines of what in each of Designs H and I was the second group. The spectral relationships generalize to these key features sketched in Figure 14.

- Faux sampling features L impulses, spaced by convention (alternative conventions are possible, but all ultimately yield the same systems) uniformly from the origin rightward at spacing equal to the final system's input sample rate f_s and where any integer $L \geq 1$ is permitted. The case of no faux sampling at all corresponds to $L = 1$. In the final system L takes the role of the interpolation ratio at the input to the first digital filter.

- The frequency response of equivalent digital filter C has period Lf_s.

- That digital filter obtains its signal band from the "$(\rho + 1)$th" spectral copy of the positive-frequency IF-filter output—the part on the "ρight"—where $0 \leq \rho < L$. (Designs H and I in Figs. 12 and 13 both used $\rho = 0$, as per Relationship 1 above.)

- The Pth periodic filter passband left of the signal band, where $P \geq 1$, falls just to the left of the "$(\ell + 1)$th" spectral copy of the negative-frequency IF-filter output—the part on the "ℓeft"—where $0 \leq \ell < L$. In particular, the center of that Pth filter passband is exactly $f_s/2$ from the midpoint of the support of the nearest spectral copy. (Designs H and I in Figs. 12 and 13 both used $\ell = 0$, as per Relationship 2 above.)

While this is far from obvious, these relationships actually hold even in the extreme case of $P = 1$, $\ell = 0$, and $\rho = L - 1$. That parameter combination will turn out to put the signal band in the first Nyquist zone—established in Fig. 17 below—so that the IF filter's lower transition band actually straddles the frequency origin.

Our first, preliminary look at these relationships, in Fig. 1, characterized the positive-frequency portion of the IF-filter frequency response by a signal band and an allowed-support band. Here in Fig. 14, the centers of those bands are further labeled f_{IF} and f_{sup}, respectively, with average-of-corners frequency f_{avg} then just the midpoint between f_{IF} and f_{sup}. Arithmetically symmetric transition bands result in $f_{\text{avg}} = f_{\text{sup}} = f_{\text{IF}}$, but more generally, including when transition bands are ratiowise symmetric as discussed in Section 2.3.7, these three frequencies are distinct.

2.4.1 The Fig. 1 Frequency Relationships

Here the general structure of Figure 14 is shown to satisfy the four frequency relationships of Fig. 1. The importance of this is that by establishing that a design with the Figure 14 structure necessarily has parameters for which the Fig. 1 rules all hold, we automatically establish the equivalent contrapositive statement: if parameters violate any Fig. 1 rule, no design of the Figure 14 structure can attain those parameters. This effectively makes the Fig. 1 relationships into design rules.

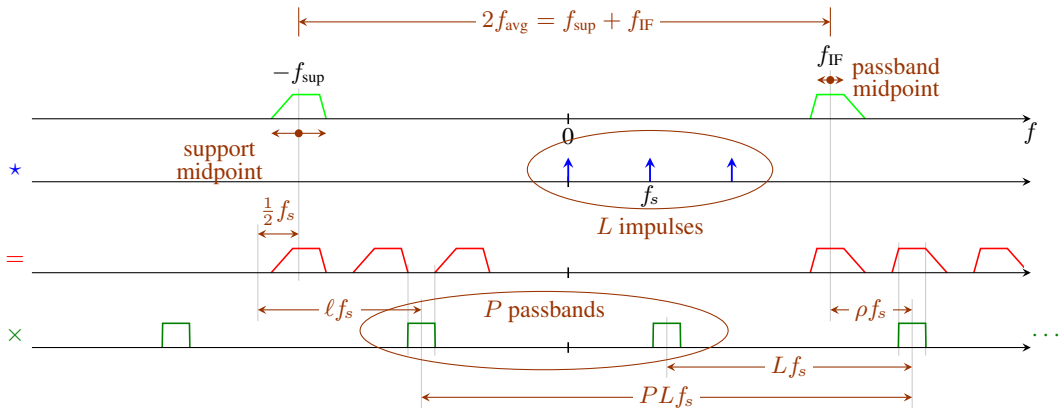

Fig. 14 — Derivation of allowed combinations of f_s and f_{avg}. Here f_{avg} (top) should be thought of as a dependent variable determined by integer design variables according to annotated relationship $f_{\text{avg}} = f_s(PL + \ell - \rho - \frac{1}{2})$, where $P \geq 1$, $L \geq 1$, $0 \leq \ell < L$, and $0 \leq \rho < L$.

The four relationships are addressed in reverse order.

The average-of-corners rule

The purpose of Figure 14 is to simplify the derivation of the relationship between average-of-corners frequency f_{avg}, sampling rate f_s, and integer design parameters L, ρ, P, and ℓ. Begin with the large interval of width PLf_s marked at the bottom of the figure. Adjoining that interval on the left is an interval of width ℓf_s. The width $\ell f_s + PLf_s$ of the combined interval can also be expressed, as shown in the remainder of the figure, as $f_s/2 + f_{\text{sup}} + f_{\text{IF}} + \rho f_s$. Equating the two width expressions and solving for f_{avg},

$$f_{\text{avg}} = \frac{f_{\text{sup}} + f_{\text{IF}}}{2} = \big(2(PL + \ell - \rho) - 1\big)\frac{f_s}{4}. \tag{16}$$

This is just $f_{\text{avg}} = (2m - 1)f_s/4$ with

$$m = PL + \ell - \rho \tag{17}$$

so f_{avg} lies exactly at the center of the mth Nyquist zone for sampling at rate f_s, the system's initial sampling rate. Allowable values for the parameters on the right side of (17) permit m to take any positive integer value.

Signal-Band and Support-Band Widths Sum to the Sample Rate

In the third line of Fig. 14, copies of the spectral shape of the frequency response of the IF-filter—consider either positive or negative frequencies but not both—repeats at intervals of f_s. Starting at any stopband corner, one can move to another one copy away, a distance of f_s, by first traversing the adjacent support band and then the smaller interval between copies of the basic spectral shape. By construction that smaller interval has width equal to that of the signal band. This can be seen in how the fourth-line passband on the left aligns with such an interval on the third line.

All this really proves is that the IF-filter transition bands have been drawn so as to make this rule true. However, it is clear graphically that we could not make either transition band wider without intruding on

either the Pth digital-filter passband or, if $\ell = 0$, the one to its right. In effect the rule then is ultimately enforced by the demands of anti-aliasing filtering.

Transition-Band Midpoints lie on Nyquist-Zone Boundaries

Label the top and bottom corner frequencies of the positive-frequency signal band or passband $\bar{f}_t$ and $\bar{f}_b$ respectively. Likewise, label the top and bottom corner frequencies of the positive-frequency support band $\underline{\mathrm{f}}_t$ and $\underline{\mathrm{f}}_b$ respectively.

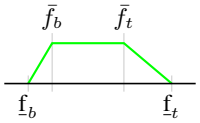

The rule just proved and the definition of average-of-corners frequency f_{avg} can be written in these terms as

$$f_s = \overbrace{\bar{f}_t - \bar{f}_b}^{\text{signal-band width}} + \overbrace{\underline{\mathrm{f}}_t - \underline{\mathrm{f}}_b}^{\text{support-band width}} \tag{18}$$

$$f_{\text{avg}} = \frac{1}{2}\Big(\underbrace{\frac{\bar{f}_t + \bar{f}_b}{2}}_{\text{signal-band midpoint}} + \underbrace{\frac{\underline{\mathrm{f}}_t + \underline{\mathrm{f}}_b}{2}}_{\text{support-band midpoint}}\Big). \tag{19}$$

The upper and lower edges of the Nyquist zone are at $f_{\text{avg}} \pm f_s/4$ or, using (18) and (19),

$$\text{lower edge} = \frac{\bar{f}_t + \bar{f}_b + \underline{\mathrm{f}}_t + \underline{\mathrm{f}}_b}{4} - \frac{\bar{f}_t - \bar{f}_b + \underline{\mathrm{f}}_t - \underline{\mathrm{f}}_b}{4} = \frac{\bar{f}_b + \underline{\mathrm{f}}_b}{2} = \text{midpoint of lower transition band}$$

$$\text{upper edge} = \frac{\bar{f}_t + \bar{f}_b + \underline{\mathrm{f}}_t + \underline{\mathrm{f}}_b}{4} + \frac{\bar{f}_t - \bar{f}_b + \underline{\mathrm{f}}_t - \underline{\mathrm{f}}_b}{4} = \frac{\bar{f}_t + \underline{\mathrm{f}}_t}{2} = \text{midpoint of upper transition band.}$$

The Signal Band is Contained in a Single Nyquist Zone

The expressions just derived for Nyquist zone boundaries allow the required nesting of intervals to be expressed as

$$\big[\bar{f}_b, \bar{f}_t\big] \subset \left[\frac{\underline{\mathrm{f}}_b + \bar{f}_b}{2}, \frac{\bar{f}_t + \underline{\mathrm{f}}_t}{2}\right].$$

This is satisfied automatically because $\underline{\mathrm{f}}_b < \bar{f}_b$ and $\bar{f}_t < \underline{\mathrm{f}}_t$ are implicit in the definitions of the four quantities involved.

2.4.2 *Choosing Parameters for a Particular Nyquist Zone*

The general faux-sampling scheme of Fig. 14 is easily tailored so as to put the signal band in the mth Nyquist zone as long as $m \geq L$, because for any positive integer L and any integer m, we can always decompose m as $m = \lfloor m/L \rfloor L + (m \bmod L)$. Comparing with (17), we see that

$$\text{if } m \geq L, \text{ we can set} \qquad P = \left\lfloor \frac{m}{L} \right\rfloor, \qquad \ell = m \bmod L, \qquad \rho = 0$$

to put each of P, ℓ, and ρ in the required range.

We always have another option. If we clearly understand the modulo-L remaindering operation to produce a result in $0, \ldots, L-1$ even for negative arguments, in other words if we use the standard mathematical definition of "modulo" rather than the incorrect one favored by designers of programming languages, we can decompose much as before but now using $-m$ to obtain $-m = \lfloor -m/L \rfloor L + (-m \bmod L) = -\lceil m/L \rceil L + (-m \bmod L)$. Taking negatives then yields $m = \lceil m/L \rceil L - (-m \bmod L)$, so that

$$\text{for any } m, \text{ we can set} \qquad P = \left\lceil \frac{m}{L} \right\rceil, \qquad \ell = 0, \qquad \rho = (-m) \bmod L. \tag{20}$$

When L divides m exactly, the two options yield the same result: $P = m/L$ with $\ell = \rho = 0$, examples of which we saw in Designs H and I in Figs. 12 and 13 discussed in Section 2.3.7. Formally speaking, in most cases we are free to add the same integer to ℓ and ρ as long as we respect that they each must be kept in the range $0, \ldots, L-1$, but as a practical matter, there does not appear to be any reason to do this. Requiring at least one of ℓ and ρ to be zero harmlessly removes the pointless ambiguity.

The phrase "choose any L" above might appear daunting, but in fact we generally have an L value in mind from the outset, because L becomes an interpolation ratio that ultimately gives equivalent digital filter C a frequency response of period Lf_s. The usual design process will then ultimately also involve decimation by some net ratio M to obtain output sample rate Lf_s/M. The mathematics do not require it, but as a practical matter $L = 1, \ldots, 3$ are typically good choices, as are those M values comprising products of perhaps one, two, three, or maybe four factors of 2 or 3 or, as one of multiple factors, even a factor of 5.

2.5 Design J—Rational Rate Conversion and IF-filter Transition Bands of Equal Ratiowise Widths

Design J in Fig. 15 (page 37) features IF-filter transition bands of equal ratiowise widths as in Design G but, unlike that earlier design, also embeds rational sample-rate conversion in the spirit of Designs H and I. The general scheme is exactly that of Fig. 14 as just discussed in Section 2.4. The embedding of rational sample-rate conversion here is a bit more interesting than in Designs H and I, because those earlier designs both—in Fig. 14 terms—had $\ell = 0$ and $\rho = 0$. Here a disciplined design process leads us to use $\rho = 1$.

That process centers on the second group in Fig. 15. We take IF frequency $f_{\text{IF}} = 1000$ as given and arbitrarily set $m = 2$ to use the second Nyquist zone. The chosen frequency f_{IF} falls in the second Nyquist zone as long as $1000 < f_s < 2000$, and in fact we can expect our design process to put f_s somewhere near the center of that interval and thereby give us adequate widths for IF-filter transition bands.

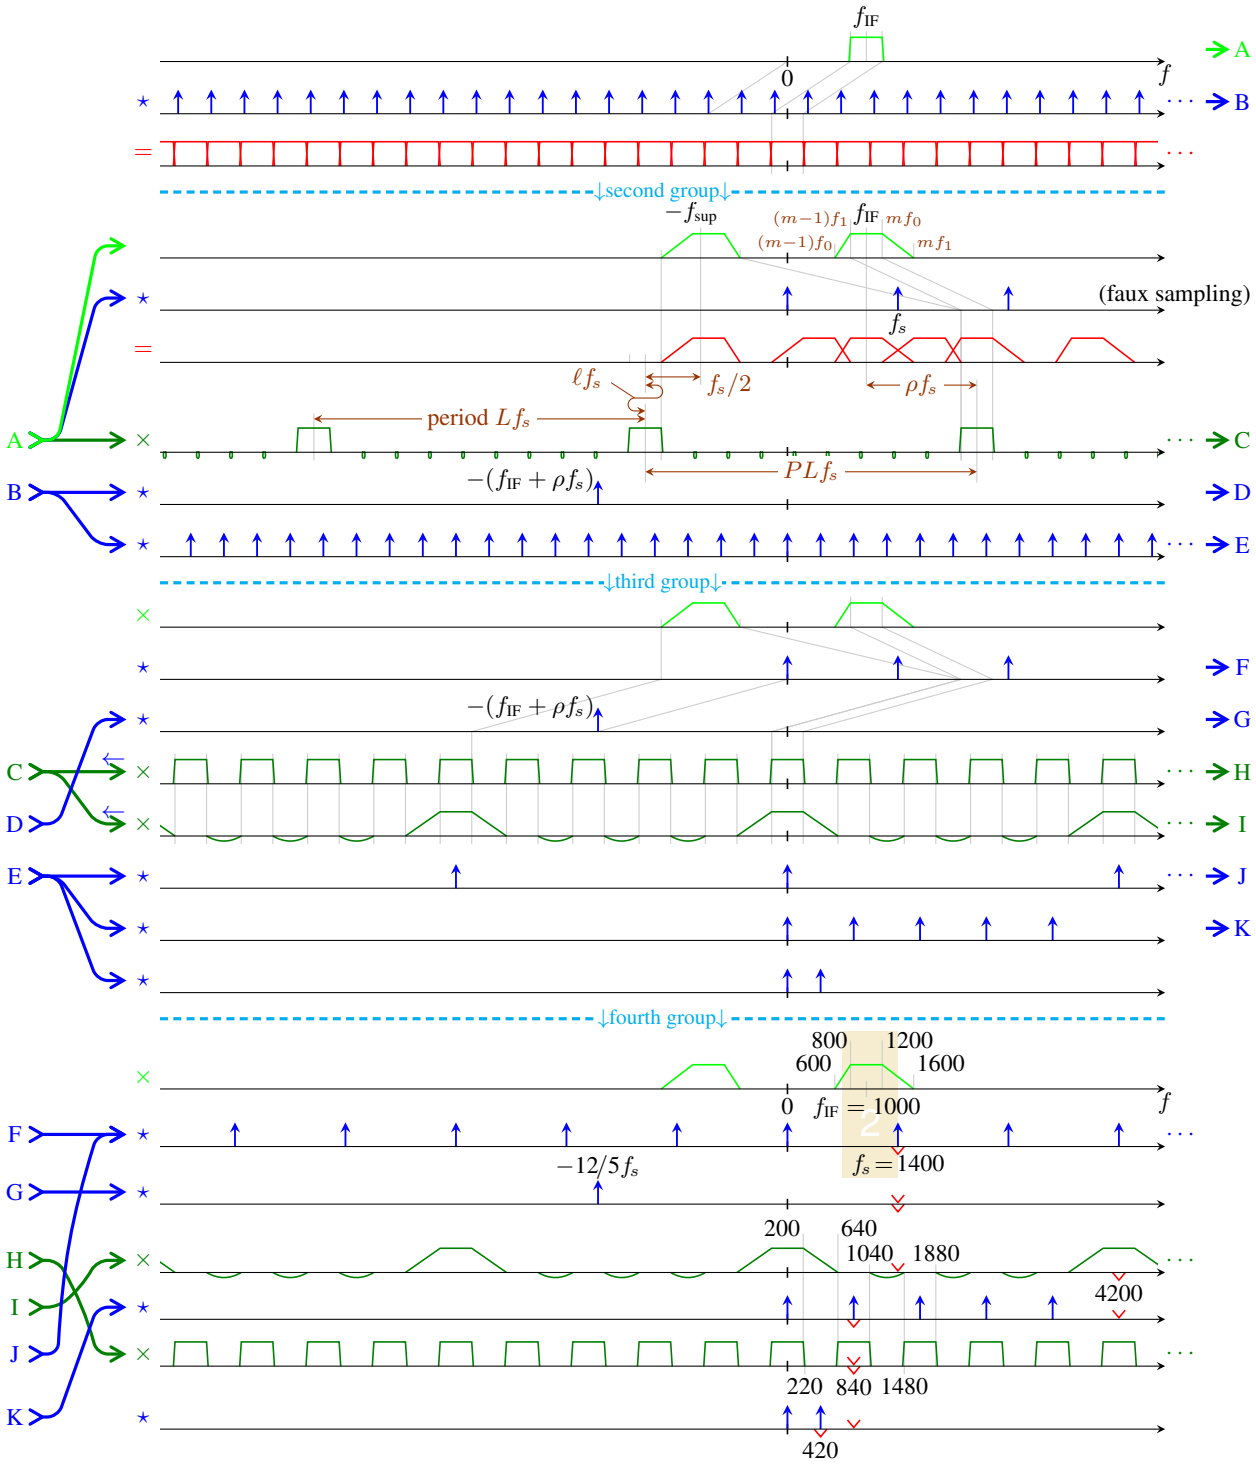

Fig. 15 — Design J. Second-group annotation here details the general faux-sampling scheme of Fig. 14 but here with $L = 3$, $P = 1$, $\ell = 0$, and $\rho = 1$ and in combination with ratiowise equal-width transition bands in the spirit of Fig. 11 to effect sampling in Nyquist zone $m = PL + \ell - \rho = 2$. Zero interpolation by $L = 3$ and decimations by 5 and 2 yield a factor of $3/10$ net sample-rate change. Here masking and shaping filters can be fifthband and halfband respectively.

More specifically, (9) and (10) with $f_{\text{wid}} = 400$ and $m = 2$ yield $f_0 = 600$ and $f_1 = 800$, from which (12) gives $f_s = 1400$. From (15) we then have $\omega_b/\omega_p = 2.5$, so an IF filter designed using lowpass-to-bandpass transformation should be capable, depending on filter type, of at least $20\log_{10}(\omega_b/\omega_p) \approx 8$ dB of stopband suppression per prototype pole. A prototype of perhaps 6th to 8th order should be adequate, depending on filter type.

We now have all four IF-filter corner frequencies, 600, 800, 1200, and 1600, so we can draw the first line of the second group in Fig. 15. As expected, the transition-band midpoints of 700 and 1400 are the boundaries of the second Nyquist zone, as per the Fig. 1 rules. The allowed-support band then extends from 600 to 1600 with a midpoint $f_{\text{sup}} = 1100$. The average-of-corners frequency f_{avg} of Fig. 1 is 1050 and is, as expected, exactly at the center of the second Nyquist zone.

Next we choose the number L of faux-sampling impulses, which ultimately becomes an interpolation rate, along with net decimation rate M. We can tabulate potential output sample rates $f_{\text{out}} = Lf_s/M$ as a function of L and M, keeping anything above 400 but below 550 (an arbitrary cutoff) as candidate values. Surprisingly few values survive this screen.

M	$L = 1$	$L = 2$	$L = 3$
3	$466\frac{2}{3}$		
4			
5			
6		$466\frac{2}{3}$	
7			
8			525
9			$466\frac{2}{3}$
10			420

We could choose $L/M = 1/3$ or $L/M = 3/8$, but the interesting option is $L/M = 3/10$ which yields the relatively attractive $f_{\text{out}} = 420$. We therefore proceed with $L = 3$ and $M = 10$, even though the factor of 5 in M is something new for us.

Now that we have $m = 2$ and $L = 3$, the remaining faux-sampling parameters of Fig. 14 are chosen as per the discussion of Section 2.4.2. Because $m < L$, we use (20) and so set $P = 1$ with $\ell = 0$ and $\rho = 1$. In the second group of Fig. 15, we can now sketch the $L = 3$ faux-sampling impulses and the potential faux-sampling output on the second and third lines respectively. The template in Fig. 14 shows us where to sketch the passbands of digital filter C in the fourth line. We will ultimately decimate by $M = 10$, so we give the frequency response of C eight tiny don't-care bands below the axis per period. Frequency shift D on the fifth line is of course chosen to move passband center at $f_{\text{IF}} + f_s$ down to the spectral origin.

From here on things are relatively routine. There is only one real choice: since we are decimating by $M = 10$, we will either decimate by 2 and then 5 or by 5 and then 2. This design takes the latter approach, as experience suggests that the filtering will be more computationally efficient. "Suggests" is the key word here, however, and to be sure one must carry out the design both ways.

ACKNOWLEDGMENTS

This work was supported by the FlexDAR component of the InTop program of the Office of Naval Research.

REFERENCES

1. J. O. Coleman, "Signals and systems II," *IEEE Potentials*, **29**(1, 2, 3, 4, 5, and 6), 2010, in each of the year's six issues.

2. ——, "Front-end IMD in receive arrays: How beamforming sums combine ADC spurs," Naval Research Laboratory, Washington DC, USA, Memo Report MR/5320--13-9420, estimated mid. 2013.

3. H. Y.-F. Lam, *Analog and Digital Filters: Design and Realization* (Prentice-Hall, 1979).