

Wiley Trading Series

+ DOWNLOADABLE SOFTWARE



CYCLE ANALYTICS FOR TRADERS

ADVANCED TECHNICAL
TRADING CONCEPTS

JOHN F. EHLERS

WILEY

CYCLE ANALYTICS FOR TRADERS



CYCLE ANALYTICS FOR TRADERS

Advanced Technical Trading
Concepts

John F. Ehlers

WILEY

Cover image: Chart © iStockphoto.com/Andrey Prokhorov
Cover design: Wiley

Copyright © 2013 by John F. Ehlers. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Ehlers, John F., 1933-

Cycle analytics for traders : advanced technical trading concepts / John F. Ehlers.

pages cm

ISBN 978-1-118-72851-2 (cloth) — ISBN 978-1-118-72841-3 (ebk) — ISBN 978-1-118-72860-4 (ebk) 1. Technical analysis (Investment analysis) 2. Investment analysis. I. Title.

HG4529.E388 2014

332.63'2042 — dc23

2013034306

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

Preface	ix
About the Author	xiii
CHAPTER 1 Unified Filter Theory	1
Transfer Response	1
Nonrecursive Filters	3
Recursive Filters	8
Generalized Filters	10
Programming the Filters	11
Wave Amplitude, Power, and Decibels (dB)	13
Key Points to Remember	13
CHAPTER 2 SMAs, EMAs, or Other?	15
Simple Moving Averages (SMAs)	15
Exponential Moving Averages (EMAs)	18
Weighted Moving Averages (WMAs)	21
Median Filter	22
Key Points to Remember	23
CHAPTER 3 Smoothing Filters on Steroids	25
Nonrecursive Filters	25
Modified Simple Moving Averages	29
Modified Least-Squares Quadratics	30
SuperSmoother	31

	SuperSmoother Filter Applications	34
	Key Points to Remember	36
CHAPTER 4	Decyclers	39
	Decycler Construction	39
	Decycler Application	41
	Decycler Oscillator	43
	Key Points to Remember	45
CHAPTER 5	Band-Pass Filters	47
	Band-Pass Filter	47
	Band-Pass Filter Q	51
	Automatic Gain Control (AGC)	54
	Spectral Dilation Removal	56
	Band-Pass Filter	56
	Measuring the Cycle Period	58
	Key Points to Remember	61
CHAPTER 6	Market Structure and the Hurst Coefficient	63
	Fractal Dimension	65
	Computing the Hurst Coefficient	67
	The Hurst Coefficient in Action	68
	Drunkard's Walk Hypothesis for Market Structure	70
	Key Points to Remember	74
CHAPTER 7	Spectral Dilation	77
	Frequency Content of Indicator Outputs	77
	Roofing Filter as an Indicator	80
	Impact of Spectral Dilation on	
	Conventional Indicators	83
	Key Points to Remember	88
CHAPTER 8	Autocorrelation	91
	Background	91
	Autocorrelation	93
	Autocorrelation Periodogram	102
	Autocorrelation Reversals	110
	Key Points to Remember	113

CHAPTER 9	Fourier Transforms	115
	Spectral Dilation	116
	Discrete Fourier Transform (DFT)	117
	Key Points to Remember	124
CHAPTER 10	Comb Filter Spectral Estimates	125
	Spectral Dilation	125
	Computing a Comb Filter Spectral Estimate	126
	Key Points to Remember	133
CHAPTER 11	Adaptive Filters	135
	Adaptive Relative Strength Index (RSI)	135
	Adaptive Stochastic Indicator	142
	Adaptive CCI (Commodity Channel Index)	147
	Adaptive Band-Pass Filter	152
	Adaptive Indicator Comparison	157
	Key Points to Remember	158
CHAPTER 12	The Even Better Sinewave Indicator	159
	Even Better Sinewave Approach	160
	Even Better Sinewave Description	160
	Using the Even Better Sinewave Indicator	162
	Key Points to Remember	164
CHAPTER 13	Convolution	165
	Theoretical Foundation	165
	Heat Map Display	168
	Computing Convolution	169
	Key Points to Remember	174
CHAPTER 14	The Hilbert Transformer	175
	Analytic Signals	176
	Hilbert Transformer Mathematics	177
	Computing the Hilbert Transformer	181
	The Hilbert Transformer Indicator	183
	Using the Hilbert Transformer to Compute the Dominant Cycle	186

Dual Differentiator	187
Phase Accumulation	189
Homodyne	192
Key Points to Remember	194

CHAPTER 15 Indicator Transforms 195

Fisher Transform	195
Inverse Fisher Transform	198
Cube Transform	200
Key Points to Remember	202

CHAPTER 16 SwamiCharts 203

SwamiCharts Overview	204
SwamiCharts RSI	205
SwamiCharts Stochastic	210
Roll Your Own SwamiCharts	216
Key Points to Remember	216

CHAPTER 17 Swing-Trading Strategies 217

Conventional Wisdom	219
Anticipating the Turning Point	220
Sine Wave Uniqueness	221
Safety Valve	224
Exiting a Trade	225
Stop Loss	225
Evaluating a Trading Strategy	226
Monte Carlo Evaluation	227
Stockspotter.com	228
Key Points to Remember	229

About the Website	231
Index	233

P R E F A C E

It has been over 10 years since *Rocket Science for Traders* was published. In those days, technical analysis was primarily the province of futures traders, while portfolio theory and fundamental analysis comprised the conventional wisdom for equity traders. However, there have been profound changes in the marketplace since that time. Futures trading has lost popularity because of scandals involving segregated customer accounts, the stock market's relentless trend upward has been broken with the result that buy-and-hold is no longer a valid investment strategy, and new trading vehicles such as exchange-traded funds (ETFs) have evolved. In addition, commission rates have decreased, and the Internet has made electronic trading available to everyone. All this has caused investors to be more involved in the trading process and interested in self-directed trading. Major brokerage houses have responded by including technical analysis tools in their trading platforms.

This is a technical resource book written for self-directed traders who want to understand the scientific underpinnings of the filters and indicators they use in their trading decisions rather than to use the trading tools on blind faith. There is plenty of theory and years of research behind the unique solutions provided in this book, but the emphasis is on simplicity rather than mathematical purity. In particular, the solutions use a pragmatic approach to attain effective trading results. The concepts are presented so they can be understood with only a background in algebra. The writing style in the book is intentionally terse so the reader doesn't need to wade through a mountain of words to find the ideas being presented. EasyLanguage computer code is used to calculate and display the indicators. From my viewpoint,

EasyLanguage is just a dialect of Pascal with key words for trading. Therefore, the code should be nearly as readable as English.

Cycles are unique because they are one of the few characteristics of market data that can be scientifically measured. However, cycle measurement is extremely complex. In the most general sense, there is a triple infinity of parameters—period, phase, and amplitude—that must be identified simultaneously to completely describe the cycles. Additionally, market cycles are ephemeral and are often buried in pure noise. So the compromises begin. One of the first realizations that a trader must make is that cycles cannot be the basis of trades all the time. Sometimes the cycle swings are swamped by trends, and it is folly to try to fight the trend. However, the cyclic swings can be helpful to know when to buy on a dip in the direction of the trend. Traditional indicators such as Stochastics, relative strength index (RSI), moving average convergence/divergence (MACD), and commodity channel index (CCI) are subject to the same constraints, and therefore this book will lead to a greater understanding of all technical indicators.

Most important, *Cycle Analytics for Traders* will allow traders to think of their indicators and trading strategies in the frequency domain as well as their motions in the time domain. This new viewpoint will enable them to select the most efficient filter lengths for the job at hand. The descriptions are written for understanding at several different levels. Traders with little mathematical background will be able to assess general market conditions to their advantage. More technically advanced traders will be able to create indicators and strategies that automatically adapt to measured market conditions by using combinations of computer code that are described.

So what should a trader take away from this book? These are a few of the new concepts that I have ranked in priority:

- An awareness of Spectral Dilation, and how to eliminate it or to use it to your advantage.
- How to use automatic gain control (AGC) to normalize indicator amplitude swings.
- Thinking of prices in the frequency domain as well as in the time domain.
- An awareness that all indicators are statistical rather than absolute, as implied by their single-line displays.
- Several advanced cookbook filters. These include the SuperSmoother, roofing filter, even better sinewave, decycler, and Hilbert Transform Indicator.

- Several different methods of estimating market spectra and sifting out the dominant cycle, with the autocorrelation periodogram being the preferred method.
- How to use transforms to improve the display and interpretation of indicators.

The concepts I have developed and derived from scientific principles are new and useful aids to short-term trading. Ultimately, trading comes down to buying and selling decisions. These decisions are never easy, and in the final chapter I unite the concepts with a few tips and tricks that I have acquired in my years of trading. Above all, trading should be approached as a statistical process. Even with a good performance of 60 percent winning trades, 60 percent is a lot closer to 50–50 than it is to 100 percent regarding a single event. Therefore, the performance judged by a few trades is invalid, and I would encourage readers to stick with a trading strategy they have developed with a profitable history, albeit hypothetical, and let the statistics be the light to success in the long run.

As evidence of my warped sense of humor, each chapter starts with a “Tom Swifty” pun that encapsulates the entire content of the chapter and I hope serves as an anchor for the reader’s memory. I think the computer code is often the most succinct and efficient method of describing a concept. Accordingly, my style is to be brief, with plenty of poetic license with mathematical notation in an effort to convey the concepts to most traders. Each chapter concludes with the significant points to remember from that chapter.

I wish you all good trading.

John F. Ehlers
August 2013

ABOUT THE AUTHOR

John Ehlers is an electrical engineer, receiving his BSEE and MSEE from the University of Missouri. He did his doctoral work at The George Washington University, specializing in Fields & Waves and Information Theory. He has retired as a senior engineering fellow from Raytheon. He has been a private trader since 1976.

John is a pioneer in introducing the MESA cycles-measuring algorithm and the use of digital signal processing in technical analysis. He developed maximum entropy spectrum analysis (MESA) over three decades ago. The program has evolved with the increased capacity of modern computers.

John has written extensively about quantitative algorithmic trading using advanced DSP (digital signal processing) and has spoken internationally on the subject. His books include *MESA and Trading Market Cycles*, *Rocket Science for Traders*, and *Cybernetic Analysis for Stocks and Futures*. His approach is unique. Any technique must first work on theoretical waveforms before testing against real-world data is attempted.

John is a cofounder of StockSpotter.com.

Unified Filter Theory

“It is too complex,” said Tom simply.

Simplicity is at the heart of the concept of linear systems. Input data are supplied to the system, and the system provides the resultant as an output. There is only one input and only one output. However, the system between the input and output can be as complex as desired. The output divided by the input is the transfer response of the system. It is this transfer response that describes the action of the system.

In this chapter you will find the difference between nonrecursive filters and recursive filters, and combinations of the two, enabling you to select the best filter for each application. In addition, you will find that the responses in the time domain and in the frequency domain are intimately connected. When designing filters for trading, it is beneficial to consider the response in both of these domains. It is important to remember that no filter is predictive—filter responses are computed on the basis of historical data samples.

By thinking in terms of the transfer responses, you will easily make the transition between filter theory and programming the filters in your trading platform.

■ Transfer Response

Consider a four-bar simple moving average. The input data are the last four samples of price. The filter output is one-fourth of the most recent price data plus one-fourth of the data sample delayed by one bar plus one-fourth

of the data sample delayed by two bars plus one-fourth of the data sample delayed by three bars. If we allow the symbol Z^{-1} to represent one unit of delay, then we can write an equation for the transfer response of a simple moving average (SMA) as:

$$(\text{Output} / \text{Input Data}) = 1/4 + Z^{-1}/4 + Z^{-2}/4 + Z^{-3}/4 \quad (1-1)$$

The values of $1/4$ are called the coefficients of the filter. In general, the filter coefficients sum to 1, so the ratio of the input to the output is 1 if the input is a constant. If we choose to generalize the filter to be other than an SMA, the values of the coefficients can be arbitrarily assigned. Further, we can extend the filter to have any arbitrary length. In this case, the filter transfer response can be written as:

$$H(z) = b_0 + b_1 * Z^{-1} + b_2 * Z^{-2} + b_3 * Z^{-3} + b_4 * Z^{-4} + \dots + b_N * Z^{-N} \quad (1-2)$$

The interesting thing about this equation is that we have now written the transfer response as a generalized algebraic polynomial. The polynomial can have as high an order as desired.

The filter generality can be extended by writing the transfer response as the ratio of two polynomials as:

$$H(z) = \frac{b_0 + b_1 * Z^{-1} + b_2 * Z^{-2} + b_3 * Z^{-3} + b_4 * Z^{-4} + \dots + b_N * Z^{-N}}{a_0 + a_1 * Z^{-1} + a_2 * Z^{-2} + a_3 * Z^{-3} + a_4 * Z^{-4} + \dots + a_N * Z^{-N}} \quad (1-3)$$

This equation completely describes the transfer response of any filter. The only thing that differentiates one filter from another is the selection of the coefficients of the polynomials. It is immediately apparent that the more fancy and complex the filter becomes, the more input data is required. This is really bad for filters used in trading because using more data means the filter necessarily has more lag. Minimizing lag in trading filters is almost more important than the smoothing that is realized by using the filter. Therefore, filters used for trading best use a relatively small amount of input data and should be not be complex.

Although mathematicians will cringe at the notation, filter computations can perhaps be better understood by simplifying Equation 1-3 as:

$$\frac{\text{Output}}{\text{Input}} = \frac{\text{Numerator}}{a_0 + \text{Denominator}}$$

Clearing fractions by cross multiplying, we get an equation useful for programming:

$$a_0 * Output + Denominator * Output = Numerator * Input$$

$$a_0 * Output = Numerator * Input - Denominator * Output$$

$$Output = (Numerator * Input - Denominator * Output) / a_0 \quad (1-4)$$

Equation 1-4 says that the filter output is comprised of two parts. The first part, the numerator term, uses only input data values. If that is the only term used in the filter, the filter is said to be nonrecursive. The second part, the denominator term, consists of previously computed values of the output. Filters using any previously computed values of the output are said to be recursive. The distinction is important because it is difficult to create recursive filters in some computer languages used for trading. Parenthetically, the coefficient a_0 is usually unity to keep things simple.

■ Nonrecursive Filters

A nonrecursive filter is one where the output response depends only on the input data and does not use a previous calculation of the output to partially determine the current value of the output. Nonrecursive filters have wide applications and therefore have acquired many different names. Among the aliases are:

- Moving average filters
- Finite impulse response (FIR) filters
- Tapped delay line filters
- Transversal filters

SMA filters are a special case of moving average filters where all the filter coefficients have the same value.

One of the most important filter characteristics to a trader is how much lag the filter introduces at the output relative to the input. A nonrecursive filter whose coefficients are symmetrical about the center of the filter always has a lag equal to the degree of the filter divided by two. For example, a nonrecursive filter of degree six will have a three-bar delay. This delay is constant for all frequency components. Since lag is very important, and

since lag is directly related to filter degree, filters used for trading most generally are simple and are of low degree.

If the a_0 coefficient equals one and all the other “ a ” coefficients are zero, the most general transfer response is just the simple polynomial in the numerator. From the fundamental theorem of algebra, the polynomial can be factored into as many complex roots as it has degrees. In other words, the transfer response can be written as:

$$H(z) = (c_0 + Z^{-1}) * (c_1 + Z^{-1}) * (c_2 + Z^{-1}) * (c_3 + Z^{-1}) * \dots * (c_N + Z^{-1}) \quad (1-5)$$

The coefficients may be complex numbers rather than real numbers. In this case, the roots of the polynomial are called the zeros of the transfer response. For example, the four-bar SMA transfer response is a polynomial of degree three and therefore has three roots factored as:

$$H(z) = (1/4) * (Z^{-1} + 1) * (Z^{-1} + \sqrt{-1}) * (Z^{-1} - \sqrt{-1}) \quad (1-6)$$

This transfer response has one real root and two complementary imaginary roots. If we substitute an exponential as $\exp(-j2\pi f) = Z^{-1}$, in the real root portion of Equation 1-5, we get using DeMoivre’s theorem:

$$\begin{aligned} Z^{-1} + 1 &= 0 \\ e^{-j2\pi f} + 1 &= 0 \\ \frac{e^{-j2\pi f/2} + e^{j2\pi f/2}}{2} &= 0 \\ \text{Cos}(\pi f) &= 0 \end{aligned} \quad (1-7)$$

This equation can be true only when the frequency is half the sampling frequency. Half the sampling frequency is the highest frequency that is allowable in sampled data systems without aliasing, and is called the Nyquist frequency. In our case, the sampling is done uniformly at once per bar, so the highest possible frequency we can filter is 0.5 cycles per bar, or a period of two bars. Equation 1-7 shows that the zero in the transfer response occurs exactly at the Nyquist frequency. We have succeeded in completely canceling out the highest possible frequency in the four-bar SMA.

We can see the frequency characteristic of the transfer response by starting with a five-element SMA and then generalizing.

$$H(z) = \frac{1 + Z^{-1} + Z^{-2} + Z^{-3} + Z^{-4}}{5}$$

Multiplying both sides of this equation by Z^{-1} and subtracting that multiplicand from both sides of the equation, we obtain

$$H(z)(1 - Z^{-1}) = \frac{1 - Z^{-5}}{5}$$

$$H(z) = \frac{1 - Z^{-5}}{5(1 - Z^{-1})}$$

$$H(z) = \frac{Z^{5/2} - Z^{-5/2}}{5 * (Z^{1/2} - Z^{-1/2})}$$

We get the frequency response of this five-element SMA by making the substitution $Z^{-1} = \exp(-j2\pi f)$, where f is the sampling frequency. Then,

$$H(2\pi f) = \frac{e^{j5\pi f/2} - e^{-j5\pi f/2}}{5 * (e^{j2\pi f/2} - e^{-j2\pi f/2})}$$

$$H(2\pi f) = \frac{\text{Sin}(5\pi f / 2)}{5\text{Sin}(\pi f)}$$

The equality of the exponential expressions and the sine equivalent will be recognized by readers familiar with complex variables as DeMoivre's theorem. For readers without this math background, please accept it on faith.

Generalizing this result for an N -length SMA, we have the transfer response of an SMA in the frequency domain expressed as:

$$H(2\pi f) = \frac{\text{Sin}(N\pi f / 2)}{N\text{Sin}(2\pi f / 2)}$$

But since the Nyquist frequency is half the sampling frequency, the transfer response in the frequency domain is

$$H(2\pi f) = \frac{\text{Sin}(N\pi f)}{N\text{Sin}(\pi f)} \tag{1-8}$$

where f = frequency relative to the sampling frequency

The important conclusion from this discussion is that we can think of the transfer response with equal validity in the time domain or in the frequency domain.

When we plot the response of the four-element SMA as a function of frequency in Figure 1.1, we see that we not only have a zero at the Nyquist frequency, but also at a frequency of 0.25.

The horizontal axis is plotted in terms of frequency rather than the cycle period that is most familiar to traders. Frequency and period have a reciprocal relationship, so a frequency of 0.25 cycles per bar corresponds to a four-bar period. The vertical axis is the amplitude of the output relative to the amplitude of the input data in decibels. A decibel (dB) is a logarithmic measure of the power in the output. Figure 1.1 shows that there are zeros in the filter transfer response in the frequency domain as well as in the time domain.

The concept of thinking of how a filter works in the frequency domain as well as how it works in the time domain is central to the understanding of the indicators that will be developed. Low frequencies near zero are passed from input to output with little or no attenuation. Since higher frequencies are blocked from being passed to the output, the SMA is a type of low-pass




FIGURE 1.1 Frequency Response of a Four-Bar Simple Moving Average

filter—passing low frequencies and blocking higher frequencies. Low-pass filters are data smoothers that remove the higher-frequency jitter in the input data that often makes the data hard to interpret. The penalty traders pay for this smoothing is the lag introduced in the transfer response.

Low-pass filters are not the only filters that can be generated with the generalized transfer response of Equation 1-3. Suppose we arrange to have the coefficients to be as:

$$b_0 = 0.5$$

$$b_1 = -0.5$$

$$a_0 = 1$$

All other coefficients are equal to zero.

Then the frequency response of the filter is shown in Figure 1.2.

In this case, the higher frequencies are passed, and the lower frequencies are severely attenuated by the filter. This is an example of a high-pass filter. Since trends can be viewed as pieces of a very long cycle, a high-pass filter is basically a detrender because the low-trend frequencies are rejected in its transfer response.




FIGURE 1.2 Frequency Response of a Two-Bar Difference Filter

Since the coefficients of a simple high-pass filter are equivalent to just taking the difference of two consecutive samples of input data, the difference operation can be viewed as analogous to a derivative function in the calculus. This concept enables a high-pass filter to be used in several different ways in trading to attempt to create a predictive waveform. If the input data are assumed to be in a trend, then the difference between any two data samples is constant. In this case, adding the difference to the current bar data predicts the value of the input data for the next sample. Alternatively, if the input data are assumed to be a quiescent sine wave, a trader can use the relationship from calculus as:

$$\frac{d(\sin(2\pi ft))}{dt} = \frac{1}{2\pi f} \cos(2\pi ft) \quad (1-9)$$

If the frequency of the sine wave is known, the high-pass filter not only provides a waveform that leads the input data waveform by 90 degrees, but also provides the means to normalize the output amplitude to the amplitude of the input data.

Returning to Equation 1-2 for a generalized nonrecursive filter, and factoring out a $Z^{-N/2}$ term, we obtain:

$$H(z) = (b_0 Z^{N/2} + b_1 * Z^{N-1/2} + \dots + 1 + \dots + b_{(N-1)} Z^{-(N-1)/2} + b_N * Z^{-N/2}) * Z^{-N/2} \quad (1-10)$$

Since $Z^{-N/2}$ is a pure delay term, and since $\exp(-j2\pi f)$ can be substituted for Z^{-1} , Equation 1-10 is proof that nonrecursive filters having coefficients symmetrical about the center of the filter will have a constant delay at all frequencies. Further, that delay will be exactly half the degree of the transfer response polynomial.

■ Recursive Filters

A recursive filter is one where the output response depends not only on the input data but also on previous values of the output. Strictly recursive filters are characterized by using only a constant in the numerator and multiple terms in the denominator of Equation 1-3. Recursive filters also have wide applications and therefore have acquired many different names. Among the aliases are:

- Exponential moving average filters
- Infinite impulse response (IIR) filters

- Ladder filters
- Autoregressive filters

If the b_0 coefficient is a constant and all the other “ b ” coefficients are zero, the most general transfer response is just the simple polynomial in the denominator. This polynomial can be factored into as many complex roots as it has degrees. In other words, the transfer response can be written as:

$$H(z) = b_0 / ((c_0 + Z^{-1}) * (c_1 + Z^{-1}) * (c_2 + Z^{-1}) * \dots * (c_N + Z^{-1})) \quad (1-11)$$

The coefficients may be complex, rather than real, numbers. In this case, the roots of the polynomial are called the poles of the transfer response because a zero in the denominator of the transfer response causes the transfer response to go to infinity at that point. One can visualize the transfer response as the canvas of a circus tent in the context of complex numbers, and the poles in the transfer response are analogous to the tent poles. While it is possible to choose coefficients that cause the transfer function to “blow up,” frequencies are constrained to be real numbers, and therefore it is relatively easy to avoid the complex pole locations.

Consider the special case of a recursive filter where

$$\begin{aligned} b_0 &= \alpha \\ a_0 &= 1 \\ a_1 &= -(1 - \alpha) \end{aligned}$$

Then, Equation 1-3 becomes

$$\begin{aligned} \frac{Output}{Input} &= \frac{\alpha}{1 - (1 - \alpha) * Z^{-1}} \\ Output * (1 - (1 - \alpha) * Z^{-1}) &= \alpha * Input \\ Output &= \alpha * Input + (1 - \alpha) * Z^{-1} * Output \end{aligned}$$

Then, using the conventional notation that $Output[1]$ equals the output one bar ago:

$$Output = \alpha * Input + (1 - \alpha) * Output[1] \quad (1-12)$$

Equation 1-12 is exactly the equation for an exponential moving average (EMA). Note that the sum of all of the coefficients on the right-hand side of Equation 1-11 sum to 1 so that the filter has no noise gain.




FIGURE 1.3 Exponential Moving Average Frequency Response for $\alpha = 0.2$

Figure 1.3 shows the frequency response of the EMA when $\alpha = 0.2$. The EMA is a type of low-pass filter, passing the lower-frequency components of the input data and attenuating its higher-frequency components. If α is made to be smaller, fewer of the lower-frequency components are allowed to pass, and the high-frequency components are attenuated to a greater degree. Conversely, if α is made to be larger, there is less smoothing, and therefore more higher-frequency components of the input are allowed to pass to the output. There are no zeros (or poles) in the transfer response.

The lag of EMA filters will be derived in Chapter 2.

■ Generalized Filters

A generalized filter uses both the numerator and denominator of Equation 1-3 to achieve a wider range of responses other than low-pass filtering and high-pass filters. Some familiar aliases for these generalized filters are:

- Autoregressive moving average (ARMA) filters
- Autoregressive integrated moving average (ARIMA) filters