Preface

The great advancements in the design of microchips, digital systems, and computer hardware over the past 40 years have given birth to *digital signal processing* (DSP) which has grown over the years into a ubiquitous, multi-faceted, and indispensable subject of study. As such DSP has been applied in most disciplines ranging from engineering to economics and from astronomy to molecular biology. Consequently, it would take a multi-volume encyclopedia to cover all the aspects and ramifications of DSP, and such a treatise would require many authors. This textbook focuses instead on the fundamentals of DSP, namely, on the representation of signals by mathematical models and on the processing of signals by discrete-time systems. Various types of processing are possible for signals but the processing of interest in this volume is almost always linear and it typically involves reshaping, transforming, or manipulating the frequency spectrum of the signal of interest. Discrete-time systems that can reshape, transform, or manipulate the spectrum of a signal are known as *digital filters*, and these systems are the primary focus of this edition as they were in the author's previous textbooks on the subject, namely, *Digital Filters: Analysis and Design*, McGraw-Hill, 1979, *Digital Filters: Analysis, Design, and Applications*, McGraw-Hill, 1993, and *Digital Signal Processing: Signals, Systems, and Filters*, McGraw-Hill, 2006. These volumes were actually different editions of the same textbook and could easily be titled Digital Filters, First, Second, and Third Edition, respectively, with the present volume being the Fourth Edition. The subtitle *Analysis, Design, and Signal Processing Applications* of this edition is included just to emphasize the strong focus of the textbook on the analysis, design, and applications of digital filters. An important difference between the present and previous editions is that the present edition pays much more attention to the practical side of the subject matter, e.g., design procedures and algorithms, relegating to references the more theoretical aspects of the subject matter such as theorem proofs and derivations.

This author considers the processing of continuous- and discrete-time signals to be different facets of one and the same subject of study without a clear demarkation where the processing of continuous-time signals by analog systems ends and the processing of discrete-time signals by digital systems begins. Discrete-time signals sometimes exist as distinct entities that are not derived from or related to corresponding continuous-time signals. The processing of such signals would result in transformed discrete-time signals that would, presumably, be enhanced or more desirable versions of the original signals. Obviously, reference to an underlying continuous-time signal would be irrelevant in such a case. However, more often than not discrete-time signals are derived from corresponding continuous-time signals and, as a result, they inherit the spectral characteristics of the latter. Discrete-time signals of this type are often processed by digital systems and after that they are converted back to continuous-time signals. A case in point can be found in the recording industry where music is first sampled to generate a discrete-time signal which is then encoded and stored in a computer file or recorded on a CD as a long series of numbers. When the music is played back, the discrete-time signal is decoded and then converted back into a continuous-time signal. In order to preserve the spectrum of the underlying continuous-time signal, e.g., that delightful piece of music, through this series of signal manipulations, special attention must be paid to the spectral relationships that exist between continuous- and discrete-time signals. These relationships are examined in great detail in Chaps. 3, 4, and 6. In the application just described, part of the processing must be performed by analog filters. As will be shown in Chap. 6, there is often a need to use a bandlimiting lowpass analog filter before sampling and, on the other hand, the continuous-time signal we hear through our stereo systems is produced by yet another lowpass analog filter. Therefore, knowledge of analog filters is prerequisite if we are called upon to design DSP systems that involve continuous-time signals in some way. Knowledge of analog filters is crucial in another respect: some of the better recursive digital filters can be designed only by converting analog into digital filters, as will be shown in Chaps. 11, 12, 13, and 18.

The prerequisite knowledge for the book is a typical undergraduate mathematics background of calculus, complex analysis, and simple differential equations. At certain universities, complex analysis may not be included in the curriculum. To overcome this difficulty, the basics of complex analysis are summarized in Appendix A which can also serve as a quick reference or refresher. The elliptic approximation in Chap. 11 is based on the theory of elliptic functions which is not normally included in undergraduate curricula. However, as shown in Chap. 11, transfer functions that would satisfy arbitrary prescribed specifications can easily be obtained without a background knowledge of elliptic functions; this can be done by using a step-by-step design procedure that involves rapidly converging series which represent the elliptic functions involved. Readers interested in the basics of elliptic functions are referred to the previous editions of this book. Chapter 14 requires a basic understanding of computer arithmetic and random processes. These subjects are treated in detail by specialized courses on these subjects. However, courses on these subjects may not be part of the program of every student who wants to take a course on DSP or digital filters. To circumvent this difficulty, the basics of computer arithmetic and random processes are reviewed immediately before they are needed in Chap. 14.

Chapter 1 provides an intuitive overview of filter-based DSP. It starts with a classification of the types of signals encountered in DSP. It then introduces in a heuristic way the characterization of signals in terms of frequency spectrums. The filtering process as a means of transforming or altering the spectrum of a signal is then described. The second half of the chapter provides a historical perspective of the evolution of analog and digital filters and their applications. The chapter concludes with three specific applications, namely, (1) processing of electrocardiograms, (2) processing of stock-exchange data, and (3) processing of DNA and protein sequences, which illustrate the scope, diversity, and usefulness of digital filters in DSP.

Chapter 2 deals with the fundamentals of discrete-time systems. Topics considered include basic system properties such as linearity, time invariance, causality, and stability; characterization of discrete-time systems by difference equations; representation by networks and signal flow graphs and analysis by node-elimination techniques. Time-domain analysis is introduced at an elementary level. The analysis is accomplished by solving the difference equation of the system by using induction. Although induction is not known for its efficiency, it is an intuitive technique that provides the newcomer with a clear understanding of the basics of discrete-time systems and how they operate, e.g., what are initial conditions, what is a transient or steady-state response, what is an impulse response, etc. The chapter continues with the representation of discrete-time systems by convolution summations on the one hand and by state-space characterizations on the other. The exposition is done without the use of transform theory which often clouds processes that are in reality very easy to grasp even for the newcomer.

Chapters 3 and 4 deal with the elements of transform theory that are prerequisite for the analysis and design of digital filters. Chapter 3 describes the Fourier series and transform as the principal mathematical entities for the spectral characterization of continuous-time signals. The Fourier transform is deduced from the Fourier series through a limiting process whereby the period of a periodic signal is stretched to infinity. This chapter also describes a relatively recent way of looking at impulse functions, which provides practical solutions to some classical DSP problems without compromising the mathematical principles associated with convergence. Chapter 4 deals with the Z transform which is easily the most important mathematical tool for the representation of discrete-time signals. The Z transform is viewed as a Laurent series and that immediately causes the Z transform to inherit the mathematical properties of the Laurent series. By this means, the convergence properties of the Z transform are more easily understood and, furthermore, a host of algebraic techniques immediately become applicable for the inversion of the Z transform. The chapter also deals with the use of the Z transform as a tool for the spectral representation of discrete-time signals.

The application of transform theory to discrete-time systems is covered in Chap. 5. By applying the Z transform to the convolution summation, a discrete-time system can be represented by a transfer function that encapsulates all the properties of a linear system, e.g., time-domain response, stability, steady-state sinusoidal response, and frequency response. The chapter also includes stability criteria and algorithms that can be used to decide with minimal computational effort whether a discrete-time system is stable or not. The concepts of amplitude and phase responses and their physical significance are illustrated by examples as well as by 2- and 3-dimensional MATLAB plots that show clearly the true nature of zeros and poles. Chapter 5 also delineates the standard first- and second-order transfer functions that can be used to design lowpass, highpass, bandpass, bandstop, and allpass digital filters. The chapter concludes with a discussion on the causes and elimination of signal distortion in discrete-time systems such as amplitude distortion and delay distortion.

Chapter 6 introduces the class of impulse-modulated signals which have the extraordinary property that they are both sampled and continuous in time. As such, they share characteristics with both continuous- as well as discrete-time signals. Therefore, these signals provide a bridge between the analog and digital worlds and thereby facilitate the DSP practitioner to interrelate the spectral characteristics of discrete-time signals to those of the continuous-time signals from which they were derived. The chapter also deals with the sampling process in detail, the use of digital filters for the processing of continuous-time signals, and the characterization and imperfections of analog-to-digital and digital-to-analog converters.

Chapter 7 presents the discrete Fourier transform (DFT) and the associated fast Fourier-transform method as mathematical tools for the analysis of signals on the one hand and for the software implementation of digital filters on the other. The chapter starts with the definition and properties of the DFT and continues with the interrelations that exist between the DFT and (1) the Z transform, (2) the continuous Fourier transform, and (3) the Fourier series. These interrelations must be thoroughly understood, otherwise the user of the fast Fourier-transform method is likely to end up with inaccurate spectral representations for the signals of interest. A unique feature of the treatment of the subject matter of this chapter is that the DFT is applied to a periodic version of a given discrete-time signal and not to the discrete-time signal itself as is often done in other textbooks. That way consistency is maintained between the discrete-time signal to be transformed and the inverse DFT which is periodic by definition.

Chapter 8 introduces the window technique as a tool that can be used along with the DFT for the processing of continuous-time signals of long or infinity duration, which are only approximately bandlimited. Through the use of a class of functions known as *window* functions, often referred to as *windows* for short, a signal of infinite duration can be forced to become bandlimited in such a way as to preserve the frequency spectrum of the continuous-time signal as far as possible. The chapter deals at length with the characterization, properties, and spectral characteristics of state-of-the-art windows such as the Kaiser, Dolph-Chebyshev, and ultraspherical windows. The chapter also describes a technique that can be used to determine the required minimum sampling frequency to achieve a prescribed precision in the frequency spectrum of the discrete-time signal relative to that of the underlying continuous-time signal.

Chapters 1 to 8 deal essentially with the characterization and properties of continuous- and discrete-time, periodic and nonperiodic signals and with the properties of discrete-time systems in general. Chapters 9 to 19, on the other hand, are concerned with the design of various types of digital filters. The design process is deemed to comprise four steps, namely, approximation, realization, implementation, and study of system imperfections brought about by the use of finite arithmetic. Approximation is the process of generating a transfer function that would satisfy the required digital-filter specifications. Realization is the process of converting the transfer function or some other characterization of the digital filter into a digital network or structure. Implementation can take two forms, namely, software and hardware. In a software implementation, a difference equation or state-space representation of the digital filter is converted into a computer program that simulates the performance of the digital filter whereas in a hardware implementation, a digital network is converted into a piece of dedicated hardware that can perform the required processing. System imperfections are almost always related to the use of finite-precision arithmetic and manifest themselves as numerical errors in filter parameters or the values of the signals being processed.

Although the design process always starts with the solution of the approximation problem, the realization process is much easier to deal with and for this reason it is treated first in Chap. 9. As will be shown, several realization methods are available that lead to a great variety of digital-filter structures. Chapter 9 also deals with the special class of systolic structures which happen to have some special properties that make them amenable to integrated-circuit implementation.

Chapter 10 is concerned with closed-form methods that can be used to design nonrecursive filters. The chapter starts by showing that constant-delay (or linear-phase) nonrecursive filters can be easily designed by forcing certain symmetries in the impulse response. The design of such filters through the use of the Fourier series in conjunction with the window method is then described. The chapter focuses on the design of nonrecursive filters that would satisfy prescribed specifications using either the Kaiser or the ultraspherical windows. The design method based on the ultraspherical window is enhanced by incorporating an unconstrained optimization technique and by this means more economical nonrecursive filters can be designed. The chapter concludes with a method based on the use of classical numerical analysis formulas that can be used to design specialized nonrecursive filters that can perform interpolation, differentiation, and integration.

The approximation problem for recursive filters can be solved by using direct or indirect methods. In direct methods, the discrete-time transfer function is obtained directly in the z domain usually through iterative optimization methods. In indirect methods, on the other hand, the approximation problem is solved in the continuous-time domain. First, a continuous-time transfer function is obtained which is then converted into a discrete-time transfer function through a series of transformations. The basic concepts pertaining to the characterization of analog filters and the standard approximation methods used to design analog lowpass filters, i.e., the Butterworth, Chebyshev, inverse-Chebyshev, elliptic, and Bessel-Thomson methods, are described in detail in Chap. 11. The chapter concludes with certain classical transformations that can be used to convert a given lowpass approximation into a corresponding highpass, bandpass, or bandstop approximation. The chapter is focused on the application of the various approximation methods rather than their theoretical aspects and derivations.

Chapter 12 deals with the approximation problem for recursive digital filters. Indirect methods are described by which a given continuous-time transfer function can be transformed into a corresponding discrete-time transfer function, e.g., the invariant impulse-response, matched-z-transformation, and bilinear-transformation methods. The chapter concludes with certain transformations that can be used to convert a given lowpass digital filter into a corresponding highpass, bandpass, or bandstop digital filter.

A detailed procedure, with design examples, that can be used to design Butterworth, Chebyshev, inverse-Chebyshev, and elliptic digital filters that would satisfy prescribed specifications is found in Chap. 13. The chapter also introduces group-delay equalization as a means of designing digital filters that have constant group-delay or, equivalently, linear phase response. Such filters are preferred in applications where group-delay signal distortion is undesirable.

A digital filter implemented in hardware form is essentially a specialized processor and as for any digital computer, its precision depends on the length of the registers used to store numbers, commonly referred to as the *word length*. Economics dictate that the word length for hardware implementations of digital filters be as short as possible. However, as the word length is reduced, numerical errors tend to increase and, in effect, there is a trade-off between the cost of hardware and the amount of precision that can be achieved. The effects of using a finite word length in digital filters along with relevant up-to-date methods of analysis are discussed in Chap. 14. The topics considered include coefficient quantization and methods to reduce its effects, signal scaling, product quantization and methods to reduce its effects, parasitic and overflow limit-cycle oscillations and methods to eliminate them. It should be mentioned that finite word-length effects are also present in software implementations of digital filters but since even low-cost laptop computers have word lengths that provide precision equivalent to about 15 decimal places, finite word-length effects are rarely significant enough to worry about. Chapter 14 also includes a family of digital-filter structures that offer reduced sensitivity to coefficient-quantization effects.

Chapters 15 to 17 deal with the solution of the approximation problem for digital filters by using iterative optimization methods. Chapter 15 describes a number of efficient algorithms based on the Remez exchange algorithm that can be used to design nonrecursive filters of the standard types, e.g., lowpass, highpass, bandpass, and bandstop filters, and also specialized filters, e.g., filters with arbitrary amplitude responses, multiband filters, and digital differentiators. Chapter 16, on the other hand, considers the design of recursive digital filters by using unconstrained optimization methods. To render this material accessible to readers who have not had the opportunity to study optimization before, a series of progressively improved but related algorithms is presented starting with the classical Newton algorithm for convex unconstrained problems and culminating in a fairly sophisticated, practical, and efficient quasi-Newton algorithm that can be used to design digital filters with arbitrary frequency responses. Chapter 16 also provides an introduction to the design of equalized recursive filters which are often used in applications where a constant group delay is required. Applications of such filters can be found in the field of audio signal processing where it is important to minimize the relative group delay among frequency components to prevent perceptible audio distortion, and in digital communications where variations in the group delay can result in signal spreading thereby causing inter-symbol interference between time-concentrated information symbols. Perfectly linear-phase filters which have a constant group delay at all frequencies can be easily designed as nonrecursive filters by using the methods described in Chaps. 10 and 15. However, in many applications a perfectly linear-phase response is not required and filters that have approximately linear-phase response or, equivalently, approximately constant group delay, are quite acceptable. In such applications, recursive filters are more attractive than their nonrecursive counterparts for two main reasons. First, they can satisfy the required filter specifications with a much lower filter order thereby reducing the computational requirement and/or the complexity of hardware and, second, they have a much smaller group delay. Chapter 17 deals with the design of equalized recursive digital filters using constrained optimization. The constraints imposed ensure that the resulting equalized recursive filter is stable and constraints can also be imposed to limit the maximum pole radius to control the sensitivity of the filter to coefficient quantization or to prevent transition band amplitude-response anomalies which are sometimes undesirable.

Chapter 18 is in effect a continuation of Chap. 9 and it deals with the realization of digital filters in the form of wave digital filters. These structures are derived from classical analog filters and, in consequence, they have certain attractive features, such as low sensitivity to numerical errors, which make them quite attractive for certain applications. The chapter includes step-by-step procedures by which wave digital filters satisfying prescribed specifications can be designed either in ladder or lattice form. The chapter concludes with a list of guidelines that can be used to choose a digital-filter structure from the numerous possibilities described in Chaps. 9, 14, and 18.

Chapter 19 deals with some of the numerous applications of digital filters to signal processing. The applications considered include downsampling and upsampling using decimators and interpolators, the design of quadrature-mirror-image filters and their application in time-division to frequency-division multiplex translation, Hilbert transformers and their application in single-sideband modulation, and two-dimensional digital filters.

The purpose of Appendix A is two-fold. First, it can be regarded as a brief review of complex analysis for readers who have not had the opportunity to take a course on this important subject. Second, it can serve as a reference monograph that brings together those principles of complex analysis that are required for the analysis and design of digital filters.

The book can serve as a text for undergraduate or graduate courses and various scenarios are possible depending on the background preparation of the class and the curriculum of the institution. Some possibilities are as follows:

- **A Series of Two Undergraduate Courses:** First-level course: Chaps. 1–7, second-level course: Chaps. 8–14. Possible titles for the two courses could be Digital Signal Processing I and II.
- **A Series of Two Graduate Courses:** First-level course: Chaps. 5–12, second-level course: Chaps. 13–18. Possible titles for the two courses could be Digital Signal Processing I and II or Digital Signal Processing for the first-level course and Digital Filters for the second-level course.
- **One Undergraduate/Graduate Course:** Assuming that the students have already taken relevant courses on signal analysis and system theory, a one-semester course could be offered comprising Chaps. 5–12 and parts of Chap. 14. The title for the course could be Digital Signal Processing or Digital Filters.

The book is supported by a collection of online materials posted at **http://www.ece.uvic.ca/~dsp/** which include:

- PDF slide presentations for classroom use,
- the author's DSP software package *D-Filter*[1] which implements most of the analysis and design methods and techniques found in the book,
- MATLAB programs developed by the author that can be used to analyze and design digital filters, and
- end-of-chapter solutions under password that can be downloaded by instructors at recognized academic institutions.[2]

I would like to thank Drs. Stuart Bergen, Rajeev Nongpiur, and Parameswaran Ramachandran for providing support materials and computer programs on the design of ultraspherical window functions, on the design of equalized recursive filters, and on genomic signal processing, respectively, to Drs. Esam Abdel-Raheem, Majid Ahmadi, Stuart Bergen, Sunder Kidambi, Parameswaran Ramachandran, and Wu-Sheng Lu for reviewing parts of the manuscript and providing helpful comments, to Lynne Barrett for her considerable help with the proofreading, Neha Bhargava and Anju Joshi, Project Managers at Cenveo Publisher Services, for their contributions to the production of the book and last, but not least, to Michael McCabe, Senior Editor, Professional and International Group, McGraw-Hill Education, for seeing the project through to a successful conclusion. Thanks are also due to numerous others who have contributed to the evolution of this book as mentioned in the acknowledgments of the previous editions.

Andreas Antoniou

---

[1] This is currently being ported into MATLAB.

[2] The password can be obtained by completing the Instructor's Form and sending it to the author at: **aantoniou@shaw.ca.**