Preface

For many years, rumors have been circulating in the realm of digital signal processing about quantization noise:

(a) the noise is additive and white and uncorrelated with the signal being quantized, and

(b) the noise is uniformly distributed between plus and minus half a quanta, giving it zero mean and a mean square of one-twelfth the square of a quanta.

Many successful systems incorporating uniform quantization have been built and placed into service worldwide whose designs are based on these rumors, thereby reinforcing their veracity. Yet simple reasoning leads one to conclude that:

(a) quantization noise is deterministically related to the signal being quantized and is certainly not independent of it,

(b) the probability density of the noise certainly depends on the probability density of the signal being quantized, and

(c) if the signal being quantized is correlated over time, the noise will certainly have some correlation over time.

In spite of the "simple reasoning," the rumors are true under most circumstances, or at least true to a very good approximation. When the rumors are true, wonderful things happen:

(a) digital signal processing systems are easy to design, and

(b) systems with quantization that are truly nonlinear behave like linear systems.

In order for the rumors to be true, it is necessary that the signal being quantized obeys a quantizing condition. There actually are several quantizing conditions, all pertaining to the probability density function (PDF) and the characteristic function (CF) of the signal being quantized. These conditions come from a "quantizing theorem" developed by B. Widrow in his MIT doctoral thesis (1956) and in subsequent work done in 1960.

Quantization works something like sampling, only the sampling applies in this case to probability densities rather than to signals. The quantizing theorem is related to the "sampling theorem," which states that if one samples a signal at a rate at least twice as high as the highest frequency component of the signal, then the signal is recoverable from its samples. The sampling theorem in its various forms traces back to Cauchy, Lagrange, and Borel, with significant contributions over the years coming from E. T. Whittaker, J. M. Whittaker, Nyquist, Shannon, and Linvill.

Although uniform quantization is a nonlinear process, the "flow of probability" through the quantizer is linear. By working with the probability densities of the signals rather than with the signals themselves, one is able to use linear sampling theory to analyze quantization, a highly nonlinear process.

This book focuses on uniform quantization. Treatment of quantization noise, recovery of statistics from quantized data, analysis of quantization embedded in feedback systems, the use of "dither" signals and analysis of dither as "anti-alias filtering" for probability densities are some of the subjects discussed herein. This book also focuses on floating-point quantization which is described and analyzed in detail.

As a textbook, this book could be used as part of a mid-level course in digital signal processing, digital control, and numerical analysis. The mathematics involved is the same as that used in digital signal processing and control. Knowledge of sampling theory and Fourier transforms as well as elementary knowledge of statistics and random signals would be very helpful. Homework problems help instructors and students to use the book as a textbook.

Additional information is available from the following website:

`http://www.mit.bme.hu/books/quantization/`

where one can find data sets, some simulation software, generator programs for selected figures, etc. For instructors, the solutions of selected problems are also available for download in the form of a solutions manual, through the web pages above. It is desirable, however, that instructors also formulate specific problems based on their own experiences.

We hope that this book will be useful to statisticians, physicists, and engineers working in digital signal processing and control. We also hope that we have rescued from near oblivion some ideas about quantization that are far more useful in today's digital world than they were when developed between 1955–60, when the number of computers that existed was very small. May the rumors circulate, with proper caution.