Ken Steiglitz

Isomorphism as Technology Transfer

EDITORS' INTRODUCTION

Ken Steiglitz was born on 30 January 1939 in Weehawken, New Jersey. He obtained the B.E.E. magna cum laude (1959), M.E.E. (1960), and Eng. Sc.D. (1963) degrees, all from New York University. Since 1963, he has been with the Department of Computer Science, Princeton University, where he is currently a professor of computer science. Dr. Steiglitz's research work has focused on digital signal processing (DSP) algorithms, digital filter design, and, more recently, on alternative models of computation, computing with solitons, and auction theory and applications. He authored the books *Introduction to Discrete Systems* (1974) and *A DSP Primer, with Applications to Digital Audio and Computer Music* (1996) and coauthored *Combinatorial Optimization Algorithms* (1982, 1998). He has received prestigious awards including the IEEE Signal Processing Society Award (1986), the IEEE Signal Processing Society's Technical Achievement Award (1981), the IEEE Centennial Medal (1984), and the IEEE Third Millennium Medal (2000). His hobbies include collecting ancient coins (especially from Roman Alexandria); researching numismatics; reading works by V. Nabokov, A.J. Liebling, and N. Gogol; and snatching a sleeper on eBay (if he isn't outsniped).

In this column, Dr. Steiglitz discusses the isomorphism (from the Greek words "iso" meaning *equal* and "morphosis" meaning *to shape*) between the analog and digital signal domains and its implications in terms of technology transfer. The discovery of the correspondence between these domains reshaped the way we think about the analog and digital worlds and was the basis of some of his later work in DSP audio and music synthesis. Trying to find more about what he likes in light of his life-long interest in this domain, we decided to ask him one simple question. The question was selected from the questionaire originally proposed in the French series "Bouillon de Culture" hosted by Bernard Pivot, which became a well-known coda used by James Lipton in the end of his interview show "Inside the Actor's Studio" on the American cable television channel "Bravo." So, "what sound or noise do you love?" we asked. "My wife's voice on the radio (she is a classical music announcer), Nikhil Banerjee's sitar, Django Reinhardt's guitar … " he answered.

—Adriana Dumitras and George Moschytz
"DSP History" column editors
adrianad@ieee.org, moschytz@isi.ee.ethz.ch

When I arrived at Princeton University in the fall of 1963, after having just completed my dissertation up the road at New York University under Prof. Sheldon S.L. Chang, I had no notion that I was carrying with me a seed of post-World-War II technology to the world of music, an unwitting agent of a kind of technical transpermia.

And it happened quite accidentally: I was wandering down the hall in an unfamiliar corridor of the new Engineering Quadrangle building and heard jarring noises emanating from an unmarked room. I stuck my nose in and saw that the cacophony was coming from loud speakers driven by a rack of electronics and a digital tape drive, which in those days was a refrigerator-sized device with vacuum columns to buffer the physical tape. The machine was manned by the composers Godfrey Winham and Jim Randall, and I learned that the roomful of equipment was a digital-to-analog (D/A) converter that had recently been donated to Princeton by Bell Laboratories, through Max Mathews, a pioneer in computer music. Having the local converter saved an 80-mile roundtrip to the Labs every time a composer wanted to hear any sound at all. Jim and Godfrey told me that the racket was caused by samples out of range at the output of a digital filter (a two-pole digital resonator), and they were struggling with the problem of controlling the level of the output signal.

I then told them that I had recently finished writing a dissertation on digital filters and that I might be able to help them. Today, this may seem commonplace, but in the early 1960s there was little (if any) published material on digital filters, and, in fact, this digital-to-analog converter was one of only a handful in the world that could operate at high audio frequencies. It is also true, by a turn of fate, that out of the millions of people who could have appeared at their door, I was one of only a handful who had ever even thought about filtering sound with a digital computer.

Their problem was simple: they weren't scaling the filter. It's a homework problem now in any introductory DSP course. I gave them a scaling constant, they coded it up, and their outputs then stayed in range. Although this was a very straightforward fix, it established me, quite unreasonably, as having godlike powers, as if I came bearing technology from another planet.

MUSICAL COLLABORATIONS

There followed years of enjoyable and fruitful collaboration with composers at Princeton, including especially Godfrey Winham, and later, graduate student Paul Lansky, now professor of music at the same institution. The converter chugged away for years, converting big reels of tape from university mainframes, like the IBM 7094 and 360/91, for composers at Princeton, as well as others throughout the country who, not having their own converters, mailed in tapes. My work with Winham included an algorithm for a digital pulse generator with arbitrary pitch, one that generates pitches "in the cracks." To explain, a pulse generator that interlaces ones with strings of zeros can generate only pitches with frequencies that are integral fractions of the sampling rate. These are not distributed densely enough in the musical scale to do music, and for a while it was a puzzle as to how to generate arbitrary pitches. The solution lies in exploiting an identity for the sum of sinusoids (for details see [1]).

My collaboration with composers continued with other projects, including adapting linear predictive coding (LPC) for the synthesis of singing. The idea is natural: LPC, a new tool in the 1970s, builds a speech model that separates the pitch and the formant structure. If speech is analyzed using LPC and then resynthesized using the right pitches, singing results. Godfrey Winham did some early experiments with the idea, and Paul Lansky went on to master the technique to make exciting music. He even played with transformations of the formant structure to change the perceptual size of instruments, making a virtual string ensemble from a scale played by a single violin [2].

This and my subsequent work in DSP was a natural continuation of my graduate work, which was so strongly influenced by Sheldon Chang, John Ragazzini, and their generation of researchers. By the 1950s, the early development of digital computers had engendered a great deal of interest in what was then called "sampled-data" control—that is, using digital computers in a control loop, at the time a new and radical idea. A transform theory for sampled signals had been developed, which became the z-transform. The transition from analog to digital ways of thinking was in the air.

ANALOG-DIGITAL CONNECTIONS

It was quite natural, therefore, that my research interests were drawn to this analog-digital connection. The theme of my dissertation was developing the view that digital signals can be thought of apart from the idea that they are "samples" of anything and that there is a correspondence between the two idealized (linear, shift-invariant) worlds of analog signal processing and DSP. This made possible a kind of technology transfer, especially in filter design. All the techniques for designing analog filters—for the famous Butterworth, Chebychev, and Cauer filters, for example—became immediately applicable to digital filter design. A half-century of brilliant work, including many elegant closed-form solutions, was retrieved for immediate use in the new context. The key idea was to map the frequency axis in the digital domain to the frequency axis in the analog domain in a one-to-one fashion, using the familiar bilinear transformation. Any book on DSP will cover this technique for designing closed-form infinite-impulse response digital filters, and my *DSP Primer* [3] is no exception. Soon after finishing my dissertation I became aware that Jim Kaiser, at nearby Bell Labs, was, independently and concurrently, using the same transformation for the same purpose. With his sharp insight and scientific ability he became a kind of mentor to me and to many others in the growing digital signal processing community.

The bilinear transformation has been around a long time and pops up throughout the sciences, usually in connection with mapping circles to circles or lines. Transmission line people use it in graphical form as a "Smith Chart" for impedance matching. I had come upon the bilinear transformation when I started working on designing digital filters in 1962, and it was my advisor, Sheldon Chang, who had scribbled it down for me one afternoon, with the offhand suggestion that maybe it might be of some use in filter design. Sheldon was, I believe, trained as a physicist, and has a corresponding deep and powerful intuition about how things work. In working under him I learned to pay attention to every word he said, a hard-learned lesson for many graduate students. As I followed this lead, it became clear to me for the first time that the transformation allowed an inter-conversion between analog filter and digital filter design.

We should take a moment to consider why the bilinear transformation plays such a central role in the connection between the analog and digital domains. To begin with, it is important to realize that if we want to establish a one-to-one mapping between analog and digital signals, sampling does not do the job. This is easy enough to see: many different analog signals have the same periodic samples. For example, we can interpolate sine waves of many different frequencies through the samples of one particular sine wave at one particular frequency. Thus, sampling can map more than one analog frequency to one digital frequency. (This fact lies at the root of aliasing problems.) It turns out that the simplest and most natural way to establish a *one-to-one* correspondence between the frequency axes in the analog and digital domains, the imaginary axis in the analog case, and the unit circle in the digital, is to use just the simple relationship $s = (z-1)/(z+1)$. It maps zero frequency in the digital domain ($z = 1$) to zero frequency in the analog domain ($s = 0$); it maps the highest digital frequency ($z = -1$) to the highest analog frequency ($s = \infty$); and it maps all the other frequencies to each other in a one-to-one fashion. The transformation between the frequency axes is usually encountered in the following form. The analog frequency variable is written as Ω, so $s = j\Omega$; the digital frequency variable is written as ω, so $z = \exp(j\omega)$; the sampling interval is represented by T; and the bilinear transformation is de-normalized by using $s = (2/T)(z-1)/(z+1)$. The analog and digital frequency variables are then related to each other explicitly as $\Omega = (2/T)\tan(\omega/2)$, which can be thought of as a nonlinear stretching of the ω-axis so it fits onto the Ω-axis.

technology, not only from analog to digital, but from control theory, fueled to a great extent by the postwar development of rocketry, through sampled-data control, to the exciting development of computer music in the second half of the 20th century.

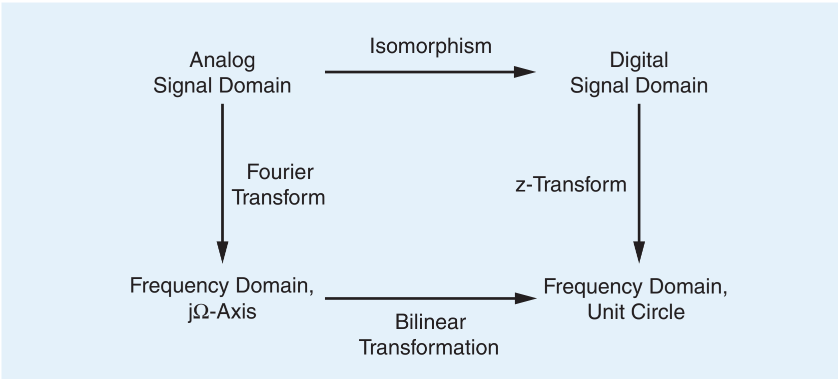

[FIG1] The diagram illustrates how digital and analog signal processing are connected by an isomorphism, induced by mapping the respective frequency domains using the bilinear transformation. The frequency variable in the analog case is usually denoted by Ω, and in the digital case by ω. The bilinear transformation connects these by $\Omega = (2/T)\tan(\omega/2)$, where T is the sampling interval. I make a distinction between a *transformation*, which maps variables, and a *transform*, which maps functions.

After seeing the close connection between the digital and analog filter design problems, I became absorbed by the larger idea of a formal correspondence between digital and analog signal processing, and went on to show that the bilinear transformation can be used to induce an *isomorphism* between the two signal domains: digital and analog signals are two reflections of the same thing, we are merely looking at the signals in two different coordinate systems (see Figure 1). The term *isomorphism* means, informally, that there is a correspondence between analog and digital signals that satisfies certain nice mathematical properties (specifically, it must be linear, reversible, and preserve the inner-product operation; see [4] for details). The isomorphism does work exactly as a change of coordinate system. We expand an analog signal in a basis that consists of a sequence of functions and use the coefficients of the expansion as the corresponding digital signal. The isomorphism has implications beyond filter design, including the equivalence of least-squares design problems in the two domains [4].

But the bilinear transform does much more than connect the analog and digital signal spaces: it also induces a mapping between linear, shift-invariant *filtering* operations. That is, if an analog signal x gets filtered by an analog filter to yield another analog signal y, the digital signal corresponding under the isomorphism to x gets filtered by a digital filter to yield the digital signal corresponding to y. And, furthermore, the digital filter that connects the signals in the digital domain is just the original analog filter with its frequency variable transformed by the bilinear transformation. This is how analog and digital signal processing can be connected perfectly: every filtering operation in one domain is mirrored by a corresponding operation in the other. The mathematical foundation provided by the isomorphism shows us that the rich intuition we bring with us from the analog realm of linear signal processing, with all the power of frequency domain interpretation that we take for granted, applies equally well to the digital realm.

In historical context, I like to think of this connection as enabling a transfer of

REFERENCES

[1] G.C. Winham and K. Steiglitz, "Letter-to-the-Editor: Input generators for digital sound synthesis," *J. Acoust. Soc. Amer.*, vol. 47, no. 2, (Part 2), pp. 665–666, Feb. 1970.

[2] P. Lansky and K. Steiglitz, "Synthesis of timbral families by warped linear prediction," *Comput. Music J.*, vol. 5, no. 3, pp. 45–49, fall 1981. Reprinted in *The Music Machine*, C. Roads,Ed. Cambridge, MA: MIT Press, 1989.

[3] K. Steiglitz, *A DSP Primer, with Applications to Digital Audio and Computer Music*. Menlo Park, CA: Addison-Wesley, 1996.

[4] K. Steiglitz, "The equivalence of digital and analog signal processing," *Inform. Contr.*, vol. 8, no. 5, pp. 455–467, Oct. 1965. Reprinted in *Digital Signal Processing*, L.R. Rabiner and C.M. Rader, Eds. New York: IEEE Press, 1972.

SP