Lawrence Rabiner

The Chirp z-Transform Algorithm—A Lesson in Serendipity

It was the fall of 1968. I had finished my Ph.D. degree on speech synthesis at MIT in May of 1967 and started working as a member of the technical staff in Jim Flanagan's group at AT&T Bell Laboratories in May of 1967. My initial assignment was working on problems of speech analysis and synthesis using a serial formant speech synthesizer. Ron Schafer joined Jim Flanagan's group at Bell Laboratories in the spring of 1968, having completed his Ph.D. at MIT in the area of homomorphic signal processing and its applications to speech processing. Ron and I started working together almost immediately on the problem of formant estimation using homomorphic methods. The goal was to compute a smooth speech spectrum using homomorphic filtering methods and then to estimate the relevant formant frequencies by peak picking the homomorphically smoothed spectrum. Most of the time this algorithm worked well; however once in awhile the formants were either too close to each other to resolve reliably or had bandwidths that were just too high to find from simple peak picking methods. Ron and I investigated a number of speech enhancement methods with the goal of making formant estimation be more reliable and robust. None of the proposed methods worked as well as we would have liked.

Next came the role of serendipity. I was attending an IEEE conference in New York City in the fall of 1968 when I ran into Charlie Rader from MIT Lincoln Laboratory. Charlie had become a good friend through our joint association with Ben Gold of Lincoln Laboratory and through our joint membership in the IEEE Signal Processing Group, a subgroup of IEEE's group on audio and electroacoustics. Charlie asked me what problems I was working on, and I brought up the issue of formant estimation from homomorphically smoothed spectra and the problems that Ron and I faced as we tried to enhance the speech signal by appropriate signal weighting to reduce the formant bandwidth. Charlie asked me to define what I would consider to be the ideal solution to this problem, and after a bit of thought I told him that what we needed was a signal enhancement method that changed the unit circle in the z-plane (the usual curve for evaluating the spectrum of a real signal) into a spiral curve that increased gradually from the origin of the unit circle to a circle that curved gradually into the inside of the unit circle. Charlie told me that he had just heard a lecture at Sylvania Labs where Leo Bluestein had been working with chirp signals to convert a discrete Fourier transform (DFT) for nonpower of two signals into an efficient convolution that could be computed using power of two DFTs. Charlie realized imme-

The author of this column is Dr. Lawrence (Larry) Rabiner. He was born in 1943 in Brooklyn, New York, and completed his B.S., M.S. (1964), and Ph.D. (1967) degrees at the Massachusetts Institute of Technology. His career at AT&T Laboratories—Research, New Jersey, spanning over 40 years, focused on DSP research with application to speech processing and speech recognition. He coauthored *Theory and Application of Digital Signal Processing* (1975), *Digital Processing of Speech Signals* (1978), *Multirate Digital Signal Processing* (1983), and *Fundamentals of Speech Recognition* (1993). Dr. Rabiner received numerous prestigious awards, most recently the IEEE Kilby Medal (1999) and the IEEE Millennium Medal (1999). Called affectionately "Mr. Energy" by his collaborators for his insatiable love of technology and his desire to learn about new areas, he appreciates (in his own words), "attention to detail, clarity of thought, clarity of vision, willingness to laugh at yourself when things go wrong." He is straightforward, enthusiastic, and passionate, a complex personality so difficult to define, yet so easy to like and admire right away. His hobbies include racquetball, bicycle riding, digital photography, and stamp collecting. Dr. Rabiner shared with us some of his thoughts—one is summarized by his favorite quotation ("When you come to a fork in the road—take it" —Yogi Berra), another one by a lesson in serendipity.

1053-5888/04/$20.00©2004IEEE

diately that a similar approach could be used to compute the spectrum of a signal on a curve that spiraled in gradually from the unit circle. Furthermore, with a bit of work, it became clear that the resulting spectral analysis could be realized on an arbitrary spiral in the z-plane, i.e., one that began at any arbitrary point (magnitude and phase) and ended at another arbitrary point, with arbitrary resolution between the beginning and ending points of the curve. Furthermore it also became clear that such an off-circle spectrum could be computed efficiently using DFT convolution techniques. Charlie and I spent the better part of the next couple of hours outlining the main features of what ultimately became known as the chirp z-transform (CZT) algorithm.

I returned to Bell Laboratories the next day and spent the better part of the next several weeks working with Ron Schafer and developing a general structure for efficiently implementing the CZT algorithm and using it as part of the speech spectral analysis system that was developed to estimate formants of speech. The algorithm turned out to work flawlessly for the intended application and ultimately was extended to a range of applications in the area of spectrum analysis. Furthermore, a highly efficient method for implementing a DFT for nonpowers of two (the original Bluestein motivation for using chirp signals) was created and used widely in several "pitch-synchronous analysis" methods developed at Bell Labs and at other research labs.

Being in the right place, at the right time, with the right person is a great formula for success in almost any area, but it proved essential for the creation of the CZT Algorithm on a quiet fall day in New York City. Having serendipity work for you never hurts.

The CZT

Unlike the discrete Fourier transform (DFT), which computes the z-transform of a finite duration signal around the unit circle, the chirp z-transform (CZT) computes the z-transform of the finite duration signal along a general spiral contour in the z-plane. Such a contour can be described by the spiral curve:

$$Z_m = AW^m, \quad \text{with} \quad m = 0, 1, \ldots, M-1$$

where A is the contour starting point, W determines the rate at which the contour spirals in (for $|W| > 1.0$) or out (for $|W| < 1.0$) from a circle of radius A_0 (where A_0 is the magnitude of A), and M is the length (number of points) of the transform.

Example: For $M = 101$, $A = 0.98$, and $W = 1.00625 \exp(-0.02\, j\pi)$ the spiral contour shown below is obtained:

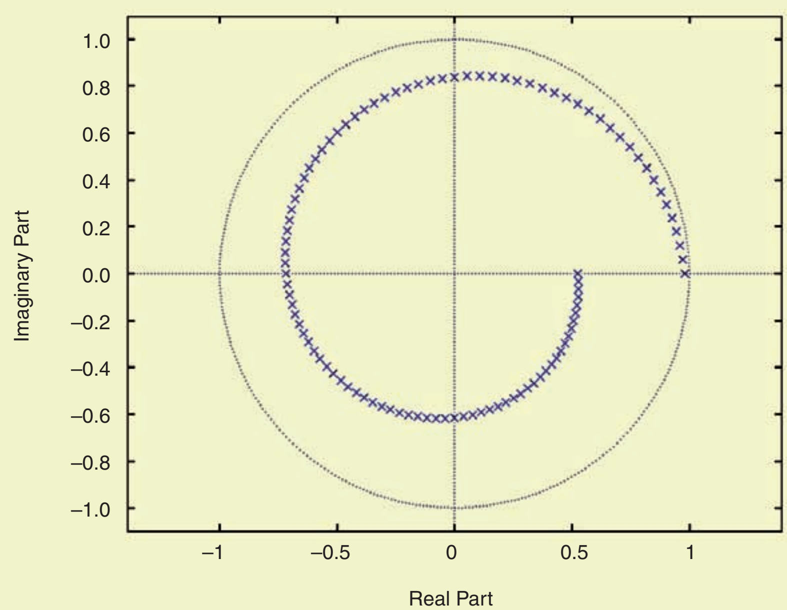

More details on the CZT can be found in *Theory and Applications of Digital Signal Processing* by L.R. Rabiner and B. Gold (Englewood Cliffs, NJ: Prentice-Hall, 1975), pp. 393–399.