Session 2532

A Unified Vector Space Approach to Teaching the Fourier Transform

Andrew Sterian
Grand Valley State University

Abstract

This paper is concerned with the approach to teaching an introductory course on Fourier theory for engineers. Commonly called *Signals and Systems* (or some variation), this course generally introduces four transforms: the Fourier Transform, the Fourier Series, the Discrete-Time Fourier Transform (DTFT), and the Discrete Fourier Transform (DFT). Our concern is that the method and order of topic presentation (as reflected in popular textbooks) creates unnecessary difficulties for students. We propose spending less time on the transforms themselves and more time at the beginning of the course in presenting a finite-dimensional vector space framework. The DFT then becomes a natural application of this framework: the projection of a signal onto a complex exponential basis. The remaining three transforms follow with the same interpretation, differing only in the domain of application. Thus, students are presented with a rigorous but tractable development (the DFT) that supports all four transforms with a common foundation.

1. Introduction

Electrical engineering curricula traditionally include an introductory course on signals and systems as a foundation for subsequent courses in communications, control, and signal processing. A large component of this course (in addition to LTI systems, convolution, the impulse response, etc.) is the presentation of the Fourier Transform in its four different forms (Fourier Transform, DFT, DTFT, and Fourier Series). Students often have difficulty with this material because it is highly abstract and because the four transforms are tantalizingly similar yet have subtly different properties and domains of application. These differences and similarities are sources of confusion because students are usually not given an underlying foundation to bind the transforms together. In addition, students are expected to accept on faith that the various transform definitions are correct. Computer demonstrations of harmonic sinusoids summing together to eventually describe a square wave are reassuring and compelling, but there is still a residual feeling that there is magic involved.

Our approach to teaching the Fourier transforms is founded upon basic concepts of abstract linear vector spaces. After an introduction to these basics, we show that the complex exponentials and shifted unit sample functions are equally valid bases for finite-length sequences. The DFT is then derived as a simple change-of-basis transformation from the unit sample basis onto the basis of complex exponentials. The finite-dimensional nature of this development is mathematically tractable for the students yet rigorous enough that they are convinced of its validity. Thus, the DFT is not just a formula that must be accepted as true.

The change-of-basis interpretation of the DFT easily extends to the remaining three Fourier transforms. With this unified interpretation, the students accept the presentation of these three transforms without proof (hence without entering the realm of infinite dimensional vector spaces). Thus, we propose that “a Signals and Systems course” begin with the fundamentals of linear vector spaces (possibly as a review), followed by the DFT, and only then followed by the remaining three Fourier transforms. Based upon student feedback from a Signals and Systems course taught in the summer term of 1999 at GVSU, we believe that this method of presentation is an improvement over the prevailing approach.

Our method of presentation is markedly different from current approaches. For example Oppenheim and Willsky[5] begin with Fourier series and progress to the Fourier Transform followed by the Discrete-Time Fourier Transform. The DFT appears as the discrete-time version of the Fourier series but without a vector space interpretation. McGillem and Cooper[4] also begin with the Fourier series and then progress to the Fourier transform, followed by a few pages on the DFT. Lathi[2] begins with the Laplace transform, followed by the Z transform, Fourier series, the Fourier Transform, and then the DFT and DTFT.

As mentioned above, we spend a significant amount of class time on linear vector spaces, which leads directly to the DFT. The DFT, then, is our entry point to the instruction of Fourier transforms. In addition to the mathematical tractability of proving the DFT’s validity, as mentioned above, this transform lends itself well to computational exercises (on a computer) with exact answers. Attempting to demonstrate the Fourier Transform of a sinusoid on a computer, in comparison, leads to the problem of explaining why the answer doesn’t look like an impulse. This is a useful question to address, but not at the beginning of a course when the students have not yet been oriented to the material.

In the subsequent sections we describe the order of topics that we follow to lead us to the DFT and then beyond to the remaining three transforms. The field of linear vector spaces can occupy several semesters of instruction, so we must choose carefully the topics that will support our eventual goal.

2. Linear Vector Spaces

While students may have considerable experience with the notion of “vectors” as directed line segments representing forces or velocities, the concept of abstract linear vector spaces (LVS) is not taught uniformly at the undergraduate level. Our first step, then, is to expand the range of what the students accept as a “vector” to include the abstract. A particular lecture technique has proven effective: the students are presented with a list of items and are asked to vote upon which items are vectors. The list of items includes the usual directed line segments and a one-column matrix, but also includes a 2x2 matrix, a real number, a sequence of real numbers, a continuous function such as $\cos(x)$, and a nonsense item such as a happy face. After voting for the first two items and rejecting the rest, the students are quite puzzled when they are told that all of these items are vectors...given the appropriate definition of a vector *space*. This exercise serves to challenge the students' preconceptions and motivates the definitions and axioms that define an abstract linear vector space[3] (e.g., closure, associativity, etc.)

2.1 Basis, Inner Product, and Norm

The concepts of span, linear independence, basis, dimension, norm, and inner product are presented next. To illustrate these concepts, we frequently make use of both the vector space of continuous polynomials over $[0, 1]$ with the inner product:

$$\langle p(x), q(x)\rangle = \int_0^1 p(x)q^*(x)dx \tag{2.1}$$

and the space of finite-length discrete-time sequences of complex numbers with the inner product:

$$\langle p[n], q[n]\rangle = \sum_{n=0}^{N-1} p[n]q^*[n] \tag{2.2}$$

These are both familiar domains for students and are easily shown to be vector spaces, even as assigned homework problems. Using two different vector spaces to illustrate the concepts of basis, etc. keeps the students aware that we are discussing abstract concepts applicable to any vector space.

These concepts require several lectures to present (approximately a quarter of a semester, including the projection theorem as described below) but they form the essential foundation for the vector space interpretation of the Fourier transforms. Thus, we invest time up front but reclaim this time later in the course when presenting the four Fourier transforms.

2.2 The Projection Theorem

After presenting the above concepts, we introduce the projection theorem. Again, an effective exercise is as follows: given a line in two dimensions and a reference point not on that line, students are asked to identify the closest point on the line to the reference point. Students have no trouble with the answer: the closest point on the line is given by the shortest line segment that joins the line with the reference point. This line segment is easily accepted as being perpendicular to the line. We then state another problem: find the third-order polynomial that most closely approximates e^x over $[0, 1]$. The fact that this is the exact same problem, but without the convenient Cartesian interpretation, prepares the students for the projection theorem.

The projection theorem (which states that the closest point in a subspace to a point outside the subspace is given by finding an error vector that is orthogonal to the subspace) is an abstract version of the intuition that guided students to solve the Cartesian closest-point-to-a-line problem. Here, we tie together all the concepts of basis, inner product and norm to reinforce and demonstrate the usefulness of the LVS foundation.

2.3 Orthonormality

The final piece of the LVS foundation uses the concept of inner product as an abstraction of angle (as demonstrated by the projection theorem) to motivate the benefit of orthonormal bases. On this type of basis, the projection of each vector onto the subspace is given by the inner product between the vector and each basis element. This naturally leads to the concept of a change of basis from one orthonormal basis to another. At this point, the students are thoroughly convinced that they are in the wrong class.

3. The DFT

Only a few small steps remain before the students are appeased. We trivially demonstrate that the N shifted-unit-sample functions $\delta[n-k]$ for $k = 0 \ldots N-1$ are an orthonormal basis for the vector space of N-point discrete-time complex-valued sequences. Any N-point sequence x can then be represented by a linear combination of these vectors using a vector of N coefficients X so that $x = IX$ where I is the identity matrix and is formed from the column-wise arrangement of the shifted-unit-sample functions. In this case, $x = X$. In this finite-dimensional case, the question of the inverse problem is easy to answer. That is, given a vector x, what should the coefficients be so that the linear combination of vectors IX gives the vector x? Inverting this matrix equation gives $X = I^{-1}x = x$.

The students may well be amused at the trivial nature of this development, however we have presented an important concept relating to representation: given a set of vectors arranged column-wise in a matrix V, it is true that V is a basis (by definition) if and only if we can solve the equation $X = V^{-1}x$ for any vector x. This equation yields the unique coefficients of the linear combination necessary to represent x. Clearly, we require that V be non-singular. In this development, we have unified LVS concepts with familiar matrix manipulations.

Finally, we consider the matrix $V = [\boldsymbol{e_0}|\boldsymbol{e_1}|\ldots|\boldsymbol{e_{N-1}}]_{N \times N}$ where each $\boldsymbol{e_k} = e^{j\frac{2\pi kn}{N}}$ for $n = 0 \ldots N-1$. If we can show that V is non-singular and the columns of V form an orthonormal basis (both are easily shown as homework problems) then we have shown that the vectors $\boldsymbol{e_k}$ are a valid orthonormal basis for N-point complex-valued sequences.

And at last, we return to the point of the course: with the definition of V as above, the equation $X = V^{-1}x$, which computes the coefficients of the expansion of any vector x onto the basis vectors $\boldsymbol{e_k}$, is simply the Discrete Fourier Transform (DFT) of the discrete-time sequence x.

Note that we have achieved the following:

- The DFT is presented in an LVS framework as a change-of-basis transformation onto the basis of complex exponentials. It is not simply presented as a formula that must be memorized.
- The development of the DFT as presented above rests only upon the fundamentals of linear vector spaces and matrix algebra. Both are topics that are tractable for undergraduate engineering students.
- By working backwards, the students can rigorously show that the DFT is an invertible transformation that represents a discrete-time sequence x as a series of coefficients X where each coefficient weights a complex exponential basis vector. Thus, the concept of representing any signal as a sum of sinusoids is given concrete form. By convincing the students of this concept's validity in the relatively simple world of finite-dimensional vector spaces, we can appeal to this result in the remainder of the course without unnecessary complexity.

The above interpretation of the DFT is not new; standard graduate-level courses on signal processing[1] present the DFT as a matrix transformation. By carefully introducing this interpretation at the undergraduate level, however, we believe that the presentation of a standard Signals and Systems course can be improved in the manner detailed above.

4. The DTFT, FT, and FS

At this point, the majority of the preparatory work in the course is done. The DTFT, for example, follows naturally by letting N, the length of the discrete-time sequence, approach infinity. We can now leave the LVS formalisms and present the four Fourier transform definitions in their conventional forms as summations or integrals. All of these forms can be readily compared to previously-presented inner products, thus reassuring the students that we are still simply projecting signals onto complex exponentials. Formal arguments and theories of infinite vector spaces are not required.

Figure 4.1 below additionally serves to highlight the similarities and differences between the four Fourier transforms. After convincing the students that all four transforms project a signal onto a basis of complex exponentials, the diagram is a reminder that the four transforms exist in different domains. Hopefully, this will serve as a good starting point for answering the "When do I use which transform?" question that invariably arises near the end of the semester.

5. Evaluation

A Signals and Systems course as presented in this paper was taught at GVSU in the summer of

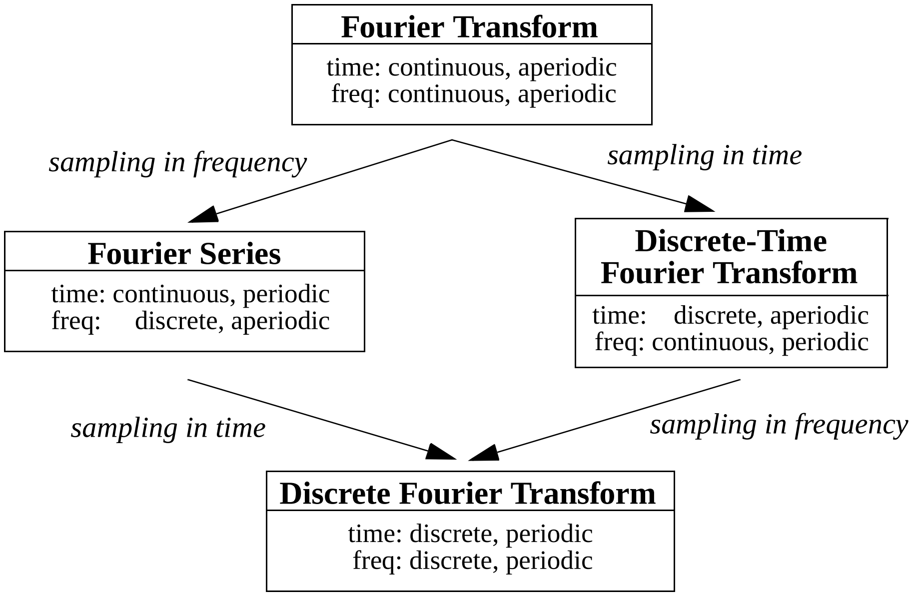

Figure 4.1: This figure is used to present a high-level view of the relationships between the four Fourier transforms and their domain of application. The DTFT, for example, can be interpreted as the time-sampled Fourier Transform as shown in several texts. This sampling converts one domain from continuous to discrete and induces periodicity in the dual domain. In addition to the vector space interpretation as discussed in this paper, the above diagram helps students see both the similarities and the differences between the four transforms.

1999 to 10 electrical engineering students in a 13-week course. Initial feedback indicated positive results: the course material was seen as a cohesive whole rather than a sequence of several, vaguely similar transforms. The presentation of linear vector space theory took longer than expected, due to lack of exposure in prior courses. Because of this delay, the presentation of the last topic in the course, the Z transform, was hastened somewhat. It should also be noted that the Laplace transform, traditionally taught in a Signals and Systems course, is thoroughly covered in a different course at GVSU (Electronic Circuits) hence was not taught in our course, freeing additional time for the vector space topics.

Bibliography

1. Kay, S.M. Modern Spectral Estimation. Prentice-Hall, 1988.
2. Lathi, B.P. Linear Systems and Signals. Berkeley-Cambridge Press, 1992.
3. Luenberger, D.G. Optimization by Vector Space Methods. John Wiley & Sons, 1997.
4. McGillem, C.D., Cooper, G.R. Continuous and Discrete Signal and System Analysis. Oxford University Press, 1991.
5. Oppenheim, A.V., Willsky, A.S., Young, I.T. Signals and Systems. Prentice-Hall, 1997.

ANDREW STERIAN
Andrew Sterian is currently an Assistant Professor in the Padnos School of Engineering at Grand Valley State University. He received his B.A.Sc. in Electrical Engineering from the University of Waterloo, Canada and the M.S.E. and Ph.D. in Electrical Engineering from the University of Michigan, Ann Arbor. He has taught courses in signal processing, digital systems, and microcontrollers.