



ELECTRONOTES 227

Newsletter of the Musical Engineering Group

1016 Hanshaw Road, Ithaca, New York 14850

Volume 23, Number 227

February 2016

FFT OF CHIRPS – LINEAR FM

-by Bernie Hutchins

INTRODUCTION

In the previous issue (EN#226 [1]) we looked at the FFTs of linear and exponential ramps. These proved interesting in part because they were NOT objects for which we immediately think of the FFT (the DFT) as providing an obvious analysis. We generally think of using an FFT for something which looks as though it is composed of sinusoidal components. In fact, we know that the FFT provides an analysis of a length N time sequence exactly as though a signal were composed (generally) exclusively of DFT harmonics ($2\pi k/N$), where k is an integer from 0 to $N/2$ (N even) or $(N-1)/2$ (N odd). For the most part, we have no difficulty of considering components for which k is not actually an integer, with the result understood to “leak” into nearby values of k [2]. But, such frequencies were still constant in time. Not that we can’t take the FFT of any N samples!

One of the joys of digital signal processing is that of relating a new situation to well-known results. Indeed, we often look at the analysis of a previously untried signal by stepping carefully through simpler familiar examples as a means of calibrating both our software and our minds. The satisfaction is when these familiar steps still work (!) and the delights are when we think ahead successfully to what should happen in the new case. The absolute joy however is when something totally unanticipated happens. That is, we have removed the “bonehead” errors and bugs and something strange is still going on. (This “joy” assumes we are under no immediate pressure to finish a project!).

“How wonderful that we have met with a paradox. Now we have some hope of making progress.”

Niels Bohr

So we most often think of taking an FFT of a signal composed of certain frequency components where the component frequencies are fixed throughout all N samples. This is not to say that we don't think of a "spectrum" or signal where frequencies are changing. Indeed we often specifically write down analytic formulas for spectra of "frequency modulated" signals in terms of sidebands, and the corresponding results can be demonstrated by using the FFT as a calculation device. We also often display what is apparently a spectrum of a speech or music signal where component frequencies change. These notions arise mainly from what we call "spectrograms" which display time varying spectra. Typically we do this by calculating a FFT as a running property (of time arrayed "frames") and these are then (effectively) tipped up and displayed side-by side so that the horizontal axis is time, the vertical axis is frequency, and the magnitude of a particular component is related to darkness (out of the page). These are usually useful in cases where the frequencies vary slowly from frame to frame (like is very often the case with speech and music).

FINDING A PUZZLE

Let's suppose we are asked to guess what an FFT of a linear FM signal would look like. Such a signal is often called a "chirp" almost certainly because it resembles a call of a songbird. It would have (for example) $N=2000$ time samples beginning at a frequency $k=10$ and running to $k=14$. That is, $k=10$ would mean that the sequence would have 10 full cycles in the full length of 2000, and $k=14$ would mean that the sequence would have 14 full cycles. For example, in Matlab code the signal sequence for $k=10$ would be:

```
n=0:1999;  
x10=cos(2*pi*10*n/2000);
```

The resulting length $N=2000$ signal would have a length $N=2000$ FFT (DFT) which happens to be real (because of the cosine used) and all the signal energy would be at $k=10$ and $k=2000-10 = 1990$. Usually we would only plot the lower half of the FFT since the magnitude of the upper half is symmetric. Further, in this example, even for the lower half we have all zeros except at $k=10$, so only need plot a restricted range (Fig. 1a).

The case where we have exactly an integer number (k) of cycles in length N is classic in that it has non-zero energy at only k (and at $N-k$). This is often a "calibration" plot. Just to make sure we are thinking correctly. In as much as we are supposing a time-varying frequency from $k=10$ to $k=14$, we would envision only isolated spikes at $k=10, 11, 12, 13,$ and 14 in the cases where we have fixed frequencies such that we have 10, 11, 12, 13, and 14 full cycles of a single FIXED frequency. As a second example, the case of $k=11$ is shown in Fig. 1b. The result is similar for $k=12, 13,$ and 14 but we won't show these. These exact results are lovely and often essential stepping-stones to a less familiar investigation.

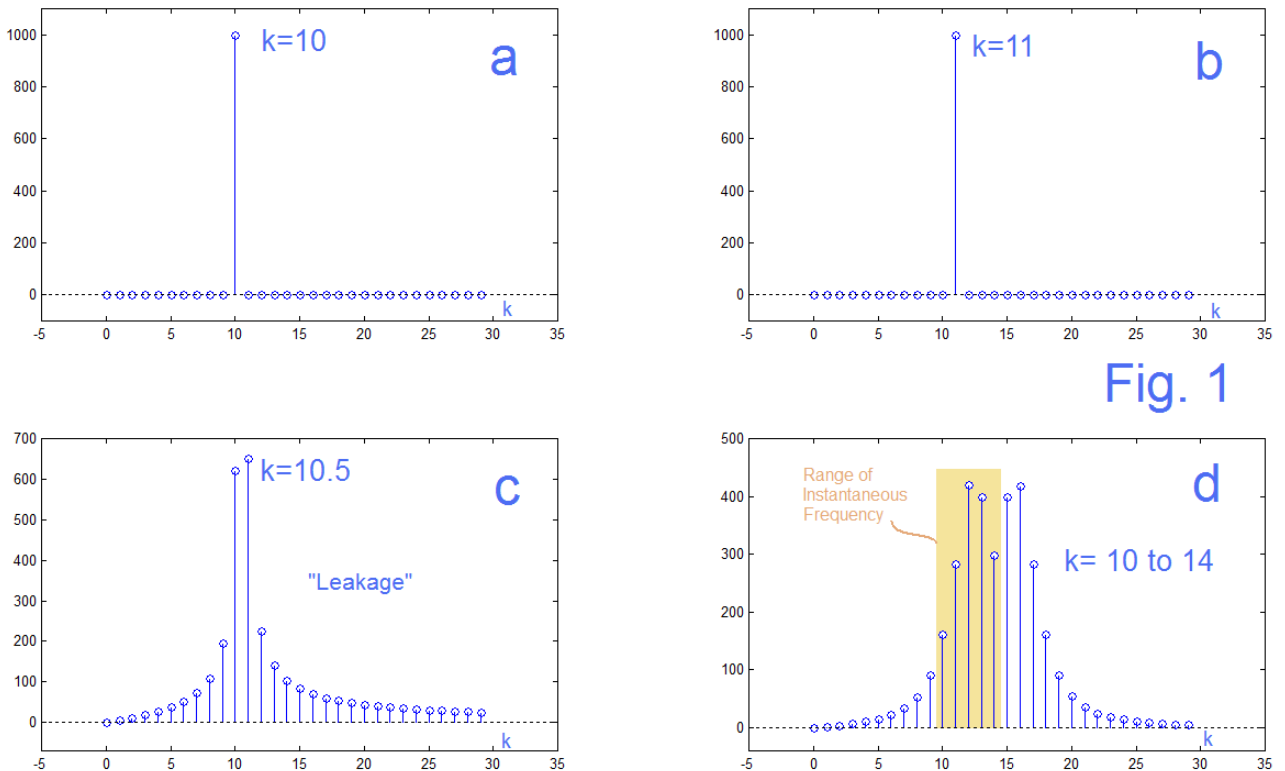


Fig. 1

The next stepping-stone or calibration step is to use a frequency that is a non-integer. That is, the frequency is between fully-resolved FFT harmonics. This would be the general case, the exact integers being contrived (but essential) examples. For example, suppose we choose a frequency corresponding to 10.5? Of course we could resolve this with $2N$ samples rather than just N , but eventually, we can't accommodate in this way. What happens with $k=10.5$ is shown in Fig. 1c. The FFT reports out energy that "leaks" into adjacent available values of k (10 and 11 in particular) [2]. This is at once logical, well-understood, and useful. This is old stuff. NOTE that the frequency $k=10.5$ is fixed across all $N=2000$ samples here. It does not move from $k=10$ to $k=11$.


So what happens when the frequency moves during the $N=2000$ samples? We expect something different, but perhaps, not too different. What a surprise (Fig. 1d). Perhaps we supposed, based on Fig. 1c, that the spectrum would sort of fill in between $k=10$ to $k=14$, perhaps mainly from $k=9.5$ to $k=14.5$, for example (the light tan of Fig. 1d). If you get a stranger result, you usually start with the assumption that you have likely done something wrong! A disconcerting plot such as Fig. 1d typically appears when you program does eventually run to completion (no beeps for syntax errors and typos – usually a half dozen such!) but the plot seems off. Often we get to slap ourselves one additional time and put in (or remove) a factor of 2, of 2π , something like that. Then we can get to the business of making the plot "pretty" and interpreting it. Here however, there seems to be nothing wrong – except an error as to what we were SUPPOSED to get. This is at once a bewildering situation and an exciting one. Somewhere a bell is going off saying "This is a modulation problem, and modulation problems always fool us at first." First, let's see what we checked, and the hints these offered us.

The first thing I tried in order to see what was going on was to look at the signal in the time domain. What signal? Well the one I wrote as:

```
n=0:1999;  
f10to14=10+[0:.002:3.998];  
x10to14=cos(2*pi*f10to14.*n/2000);
```

this seemingly being the straightforward extension of the fixed frequencies to the case where the frequency (f10to14) shifts linearly with time from 10 to 14. We note here that this was the usual dodge of trying to sneak in an “instantaneous frequency” (the time derivative for phase) as an ordinary frequency [3, 4]. If we do not do this right, it is the case that we don’t get the right answer, although the general features may well appear. It turns out that this “linear FM” or “chirp” is certainly non-trivial.

If we search the Internet we find a fair number of folks computing the FFT of a chirp – but this claim of computing has two differing notions, so it seems. One is to just set up an example time sequence (as I have done) and then compute the FFT as an example. Plug in and grind! Pretty easy. A second notion of computing is to work analytically to obtain a general (closed form) formula. Hard in this case. It seems that the math was worked out by about 1960, and in various journals (such as B.S.T.J.) but hard to find online. One recent result is as an appendix to a thesis [5].



I do not propose to repeat this derivation here. For one thing, it is difficult and I have not even tried to study it. Seeking an excuse, it seems possible to resort to quoting Richard Hamming

"The purpose of computing is insight, not numbers."

Richard Hamming, Numerical Methods for Scientists and Engineers (1962) Preface

This does in fact represent the theme here. In fact, a rigorous mathematical development too often does not yield any physical insight. Again, we cite Hamming [6]. The goal is to see what is happening.

The first "ouch" in our first look here is Fig. 1d where the spectrum seems to reflect about (be centered at) a frequency $k=14$ while we were thinking of $k=14$ as the highest frequency we input. Fig. 2 is our time domain plot and it too shows strangeness. Here we plot for the $n=2000$ samples the chirp (blue) along with the green ($k=10$) and red ($k=14$) end frequencies. For convenience, the peaks of all three waveforms (same color code) at the top as stars.

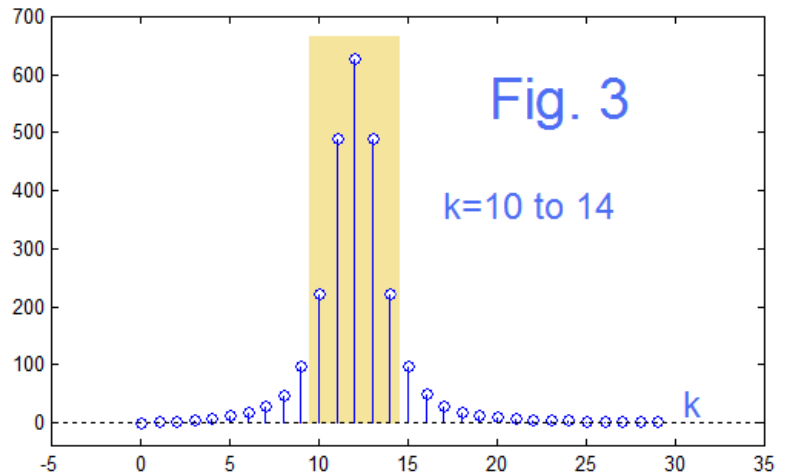
At first blush, things look proper. Indeed the blue waveform is speeding up left to right. Comparing, in the upper left corner of Fig. 2, the leftmost two blue stars to the leftmost two green stars ($k=10$), we see the blue getting a noticeable head start. On the far end, however, the rightmost two blue stars are closer together than the rightmost two red stars ($k=14$). Were you guessing it would be behind – I was? So, this is consistent with the spectral content exceeding the $k=14$ which we were supposing was the high limit. That's the second "ouch". Further, there are of course 10 cycles of the green and 14 cycles of the red. Even though it is not an actual cosine, if we count the "cycles" of the chirp (blue) we see 14! Ouch number 3. Weren't we kind of expecting 12?

A FACTOR OF 2, IF YOU PLEASE


It looks here like we messed up by a factor of 2 somewhere. A further hint is that this factor of 2 does not seem to be any overall multiplier (for example) but rather has to do with the width of the modulated region – the amount it "wiggles". This is a familiar error with modulation work associated with (incorrectly!) assigning instantaneous frequency where phase (integration) is what is required. This correction we will approach by first adjusting our Matlab code, and then doing the correct math. Here is the revised code corresponding to the lines at the top of page (4). I still call the result `x10to14`, because that's what it is. If some change was needed, why not rename the first `x10to18`? Here is what I used:

```
n=0:1999;  
f10to14=10+[0:.001:1.999];  
x10to14=cos(2*pi*f10to14.*n/2000);
```

Making this change, we can find a replacement for Fig. 1d as Fig. 3, which is immediately more agreeable to our expectations. We note that it is centered about $k=12$ (we are sweeping from 10 to 14), it has single peak, and is fairly well contained outside of $k=10$ to 14. It looks like what we might expect for varying a frequency about a $k=12$ center.




Next we look at the time domain, replacing Fig. 2 with Fig. 4, and again we see something a lot more likable. Immediately instead of 14 full cycles for the blue modulation waveform we see just the 12 we were kind of looking for. Looking closely at Fig. 4, we see that in the upper left corner, the second blue peak (blue star) is just slightly ahead of the green $k=10$ star just above it, so it is slightly faster. Its instantaneous frequency is above $k=10$. On the upper right of Fig. 4, the second from the last blue star is slightly behind the $k=14$ red star almost directly above it, the blue is lower in frequency than $k=14$.



So – things seem to have been fixed. We see that a factor of 2 on the spread does the job. What remains to be explained is why this happened. Here we make the distinction between accumulated phase and instantaneous frequency, a confusing element not

uncommon in modulation problems [3, 4]. The essential picture is the simple rotating vector of Fig. 5. The modulated signal is taken to be the cosine of the time-varying phase $\phi(t)$ – the x axis in the diagram. Note that the phase angle $\phi(t)$ is NOT in general a linear function of frequency. In the case of a constant frequency (no modulation), we would have a familiar relationship $\phi(t) = \omega t$, where ω is constant. In the modulation case, ω is not a constant, and is usually called an instantaneous frequency $\omega_i(t)$.



Thus $\omega_i(t) = d\phi(t)/dt$, and conversely:

$$\begin{aligned} \phi(t) &= \int_0^t \omega_i(\tau) d\tau \\ &= \int_0^t \left(\omega_0 + \frac{\omega_1 - \omega_0}{T} \tau \right) d\tau \\ &= \omega_0 t + \frac{\omega_1 - \omega_0}{2T} t^2 \end{aligned}$$


And thusly the 2 (or 1/2 if you prefer) appears as a consequence of the integration. Here the instantaneous frequency starts at the lower frequency of the sweep (ω_0) and sweeps up to ω_1 after a sweep time T (inside integral). The modulated signal itself is then just:

$$x(t) = \cos \phi(t)$$

A quick additional test of the idea that we now have it right would be to sweep the modulation in the reverse direction as:

```
n=0:1999;
f14to10=14-[0:.001:1.999];
x14to10=cos(2*pi*f14to10.*n/2000);
```

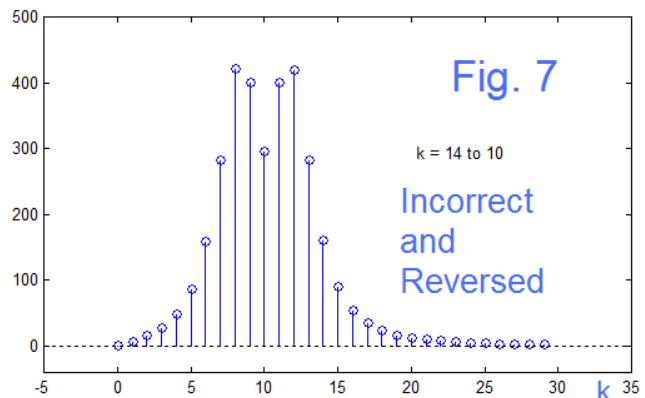
Running the FFT and the time domain plots for this change give us exactly what we would expect (but worth verifying). The FFT is absolutely identical to Fig. 3, the forward sweep. The time domain (plotted in Fig. 6 here) is just the time reversal of Fig. 4.

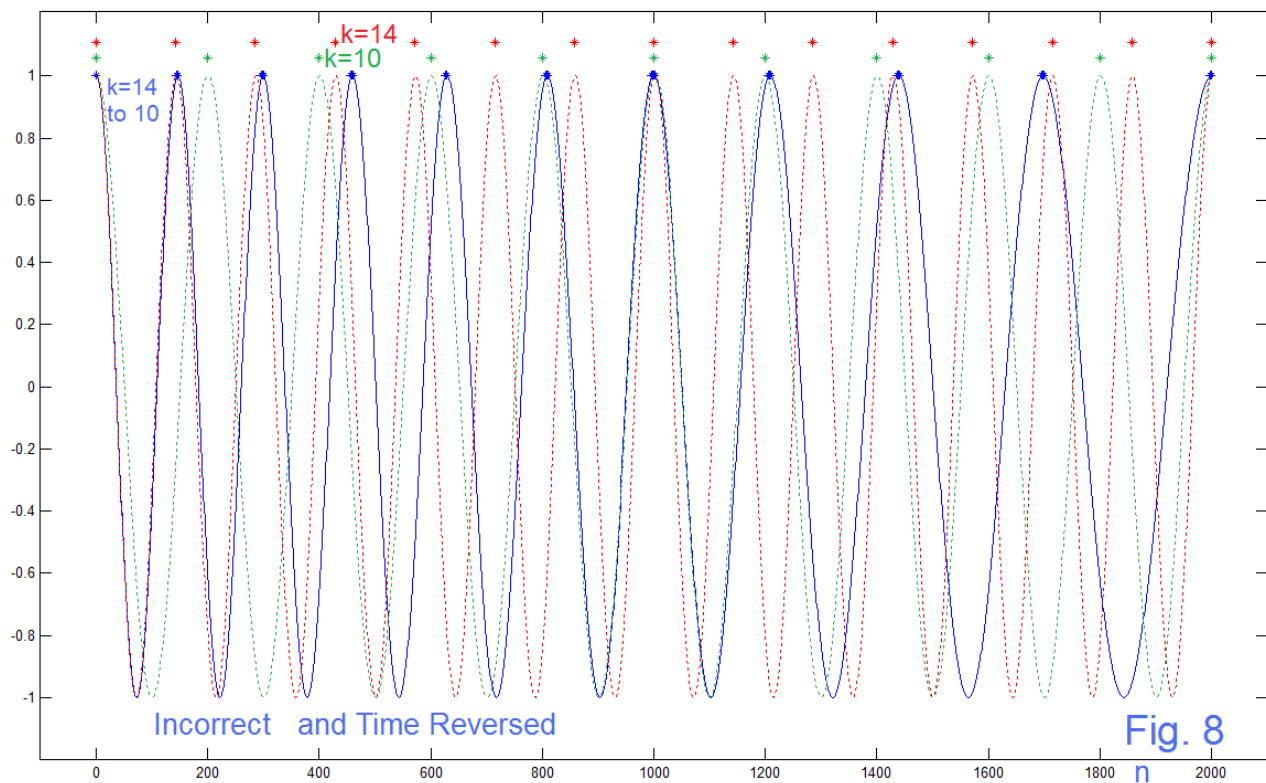


This result for Fig. 6 is a comparison of the correct setup of the phase, and we find that the time reversal did not matter. What if we did the time reversal on the incorrect phase of Fig. 1 and Fig. 2.

```
n=0:1999;
f14to10=14-[0:.002:3.998];
x14to10=cos(2*pi*f14to10.*n/2000);
```

In this case, the time reversal gives a shifted FFT (Fig. 7) and now one that is centered about $k=10$ instead of $k=14$. In fact, the time domain view (Fig. 8) shows





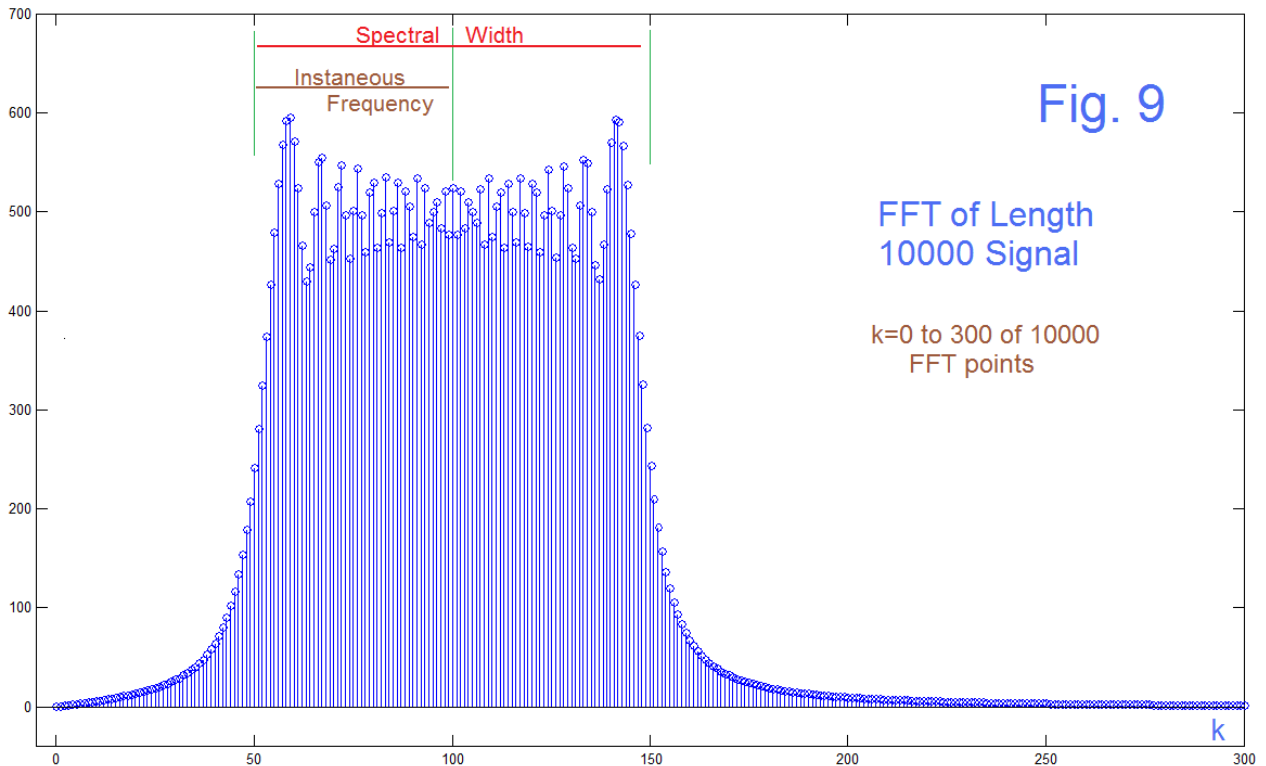
that the modulated signal has 10 actual “cycles” (just as Fig. 2 had 14 cycles). With the correct phase setup (Fig. 4) we had 12 cycles. Thus Fig. 7 and Fig. 8 are provided as evidence (at least as a curiosity) in support of the math of equations on page (7).

AN ADDITIONAL EXAMPLE

Above we have used mainly frequencies of $k=10$ to $k=14$. These are DFT normalized frequencies. For our next example we want to use a slightly larger sweep range and will use (equivalently) $k=10$ to $k=20$ for the instantaneous frequency. This will mean that the spectrum will sweep from $k=10$ to $k=30$. At the same time, we are increasing the length of the signal from $N=2000$ to $N=10000$. So values of k in the length 10000 FFT will be 5 times larger. That is the code lines becomes:

```
n=0:9999;
f10to20=10+[0:.001:9.999];           % 10 to 20
x10to20=cos(2*pi*f10to20.*n/2000);   % frequencies scaled by 5 (10000/2000)
```

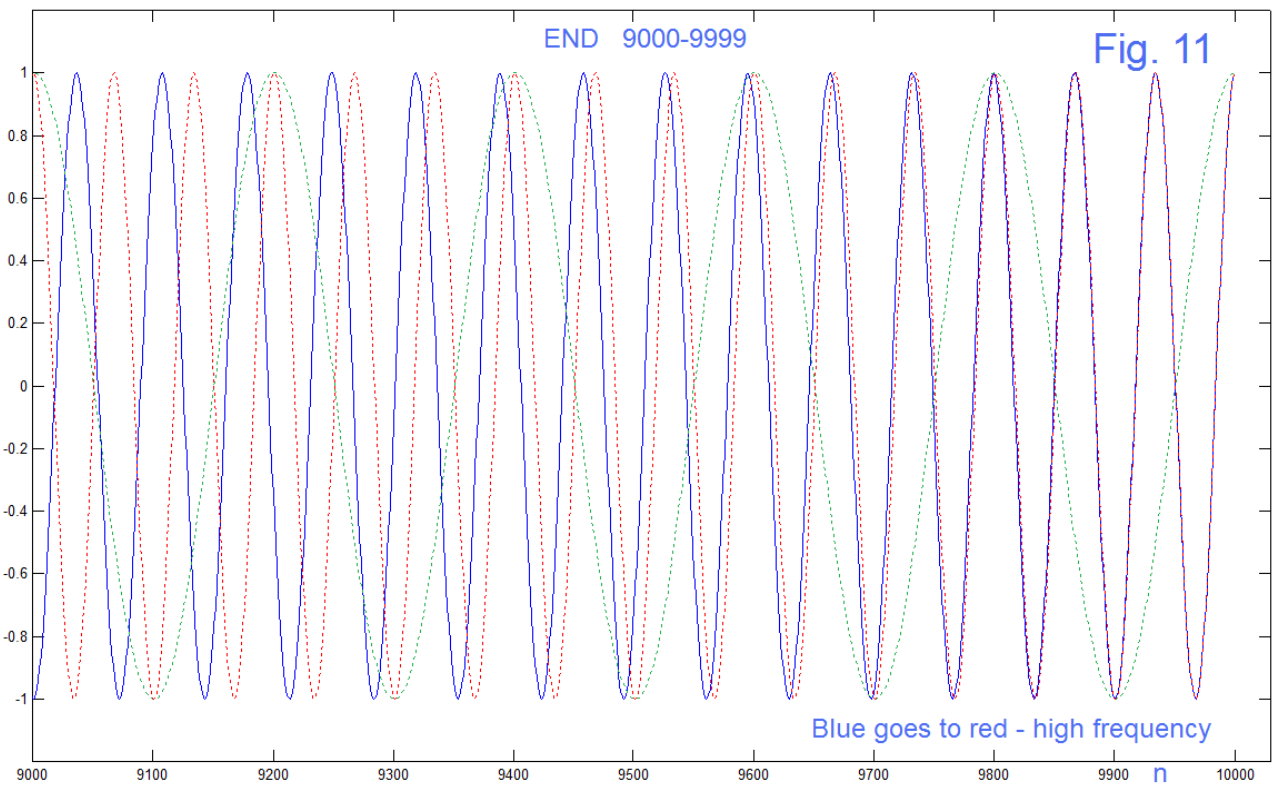
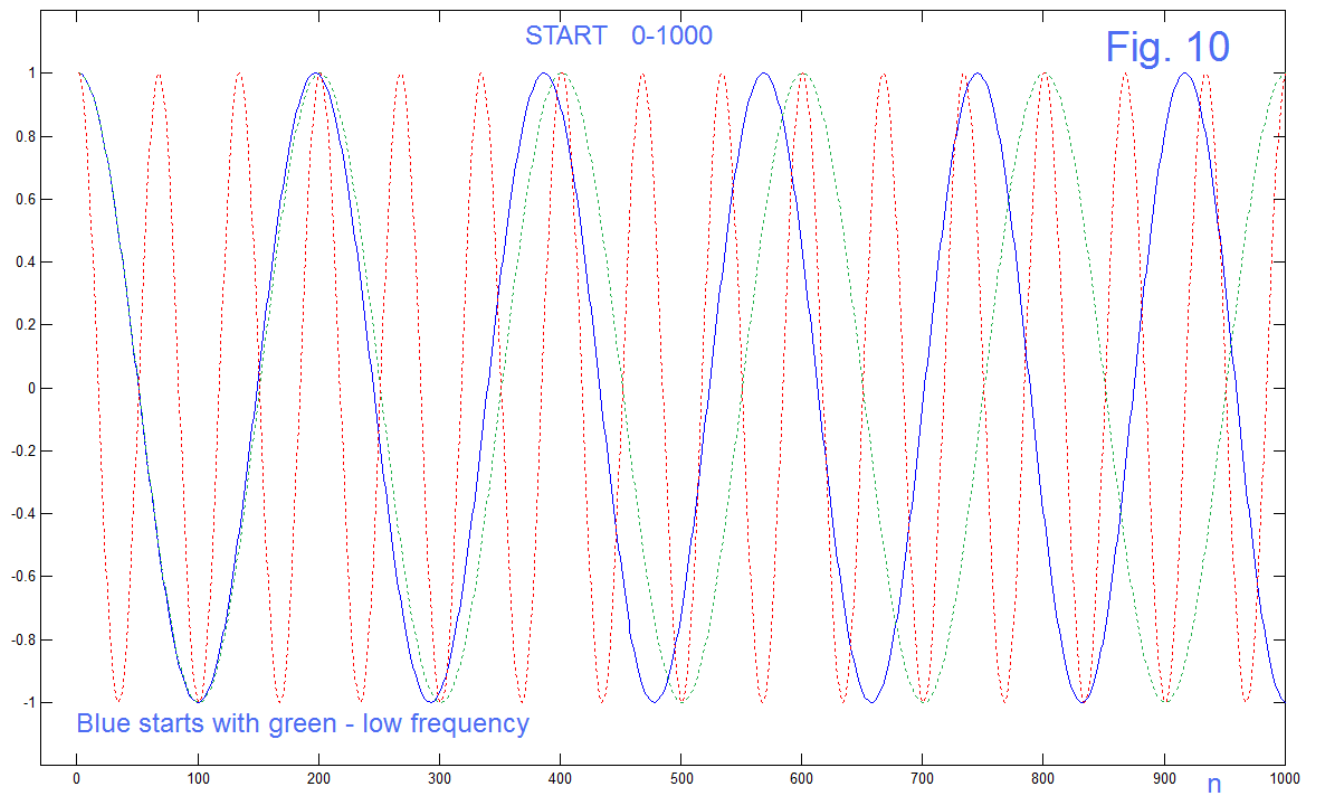
The goal here is to see if the spectrum has multiple “bumps” instead of just the central bump as in Fig. 3, or multiple bumps as in Fig. 1d (and in various Internet examples). While we were pleased to have the correct center as in Fig. 3, it was not clear if any multiple bumps were the result of an error. In actuality, the multiple bumps could not have been due to an error. Nothing was wrong – just the scaling was off.



The time-domain signal is length 10000 and consists of a full 100 “cycles” or 100 full rotations in Fig. 5. It is thus quite difficult to plot on a single figure, and accordingly we will settle for the first 1000 samples (Fig. 10) and the last 1000 samples (Fig. 11) plotted as a continuous curve. This shows us what we need to know. Similar to the previous time-domain plots, we have fixed-frequency plots corresponding to the beginning and ending frequencies. The green curve is the low end (50 full cycles) while the red curve is the high end (150 cycles). These are target values for the ends of the sweep.

Accordingly we see in Fig. 10, the left side, that the blue curve (the modulation or chirp) pretty much matches the green curve ($k=50$). In Fig. 11, on the right side, the blue curve has speeded up to three times as fast and pretty much matches the red curve ($k=150$). This (upper right corner of Fig. 11) is perhaps also an excellent point to notice that the red curve is indeed three times faster than the green one. Do not be misled by the factor of 5 that multiplies k as we increase the number of samples (without changing actual frequencies) as the signal length increases from 2000 to 10000. It’s like changing the sampling frequency where we had plenty of excess samples already.

Fig. 9 reminds us that the instantaneous frequency only varies from $k=50$ to $k=100$. We also see that there is structure (bumps) in the FFT magnitude plot of Fig. 9. So the general idea of a rough rectangular “painting” out by the sweeping frequency is found. If the bumps still bother you (as they do me) keep in mind a comparison with FM results that correspond to structured “sidebands”.



ADDENDUM ON FACTOR OF 2

Above we chose to dismiss a factor of two by pointing out (correctly) that it was a familiar factor that comes up in the integration of a linear function. In a sense, this helps a great deal because it “explains” that we have not made an obvious mistake. On the other hand, things like integral calculus are so familiar we accept a result with no further question as to why the explanation, as now displaced, is true in turn. This we will get to monetarily – it is essentially the $(1/2)$ in the area of a triangle. But first, a related amusing story.

Back in the Second War, there was a need for early detection and targeting – i.e., for what came to be known as radar. It was clear that if one hoped to electronically “paint” and resolve an object perhaps the size of an aircraft, one needed high frequencies (microwaves) and power, and if possible, a physically smallish piece of equipment. Short wavelengths to resolve, high power to assure that a tiny fraction of a beam as reflected was still enough to be meaningfully received. Microwaves (centimeter wavelengths) were then available from “klystrons” which only produced tiny amounts of power. Useful on a physics bench. The British (and others) were responsible for discovering a high power microwave device called a “Magnetron” (you have one in your home today – it is the heart of your microwave oven). These had to be developed and manufactured, and the British turned to the Americans as having all the necessary resources overall. Hence, the technology and samples were transferred, eventually to the newly formed “Rad Lab” at MIT. No one really understood how the magnetron worked. The great physicist I. I. Rabi took charge [7]:

"It's simple," Rabi told the group. "It's just a kind of whistle."

"Okay, Rabi," said Edward U. Condon, another Berkeley-trained physicist brought in from Westinghouse Research Laboratory, "how does a whistle work?"

How marvelous that physics and engineering are such that intuition is often so valid and can give a useful boost to understanding a mystery. And what a charming story.

Moving back now to calculus, likely we all remember the presentation of integrals as involved formulas and methods, with the underlying notion that it was a way to calculate an “area under a curve”. I don’t remember much about my impression when this was first presented, but clearly it was a relief to have a physical interpretation (and a useful application) of mandated math instruction. This obviously permitted us calculate areas of math objects (like $y = x^2$) which exceeded that methods of geometry. I don’t remember if we applied this to familiar geometric objects (triangles, rectangles, trapezoids, semi-circles, etc.). I presume we did – perhaps as part of derivations and proofs. Certainly interpreting the integral as an area that could be numerically approximated by the sums of very small areas of rectangles and/or trapezoids was central to numerical integration, which became most of the story in many areas of engineering. So, how often do we go back and relate simple integrals to geometric shapes? Possibly we long-abandoned any such concern. The $(1/2)$ in integrating a linear function is just THERE – has been forever.




Fig. 12

$$\begin{aligned}
 x(t) &= mt + b \\
 &= \frac{x_1 - x_0}{t_1 - t_0} t + b \\
 b &= \frac{x_1 t_0 - x_0 t_1}{t_0 - t_1} \\
 a(t) &= \int_{t_0}^{t_1} x(t) dt \\
 &= \frac{x_1 - x_0}{t_1 - t_0} \frac{t^2}{2} \Big|_{t_0}^{t_1} + b \Big|_{t_0}^{t_1} \\
 &= \left(\frac{(x_1 - x_0)(t_1 + t_0)}{2} \right) + b(t_1 - t_0) \\
 &= x_0(t_1 - t_0) + \left(\frac{1}{2} \right) [(t_1 - t_0) m (t_1 - t_0)]
 \end{aligned}$$

Fig. 12 is a simple figure showing the area under a finite portion of a linear function. We want here to relate this to the accumulated phase $\varphi(t)$ as in Fig. 5, where $d\varphi(t)/dt$ is the instantaneous frequency, a linear function in our chirp examples. This is obviously a “trivial” problem, and is so easy it is easily messed up! The equations to the right of the figure are standard. To define a linear curve (a straight line) we require only two points. Or alternatively a single point and the slope, of which the slope (m) and vertical intercept (b) are universally remembered. The equation of the curve $x(t)$ and the calculation of the intercept b are the first two equations. Below those two, we show the standard (definite) integration to the area, bringing in the factor of $(1/2)$ in this manner. If we substitute the intercept b into the result of the integration, we get the area in terms of just the time points themselves. The bottom line in Fig. 12 bypasses the integration and just writes down the area (yellow trapezoid) in terms of the rectangular area with the triangle added on [hence the $(1/2)$ comes in as half the product of the base times the height.] This form should be and is the same as the integration result. To see this note that when we substitute the slope m into the bottom equation a factor $(t_1 - t_0)$ cancels. The result that is common to the integral approach and to the geometric approach is:

$$a(t) = [x_0 t_1 - x_0 t_0 - x_1 t_0 + x_1 t_1] / 2$$

which happily involves only the two endpoints and clearly, the factor of $1/2$.

STILL WORK TO DO?

The only (!) thing that is lacking (after looking at this for a month) is a clear intuitive understanding of why the spectrum is so wide as compared to the instantaneous frequency range. Likely there will be more investigation coming. So it is tempting to just add “Part 1” to the title. [But the last time I did this – the Finite Word-Length material – there was a delay of eight years. That Part 2 consisted of mostly hand-written notes which I did use to teach. I had to retire before I got time to assemble the posting. At the very least, you need to have some notion of the direction you still want to go in order to call something Part 1.]

So the title above is “FFT OF CHIRPS – LINEAR FM” and that’s it. Here we have not said we have looked at the “Fourier Transform” (FT), or found the “spectrum” of a chirp. The FFT is a Fast Fourier Transform but of course, it is a fast “Discrete Fourier Transform.” We specifically did not calculate the FT (in the sense of a single continuous-time chirp). It is possible to estimate a FT using the FFT [2]. Largely this is a matter of isolating the non-zero part of a signal (with plenty of extra samples) to be analyzed by embedding it in a much longer sequence of zeros. This is worth looking at.

A second type of follow-up might well be to consider just what the FFT is. It represents a sequence $x(n)$ where $x(n)$ is PERIODIC in length N . So when we look at the FFTs as above, these correspond not just to the chirps (plotted as blue) but to the periodic continuation (repeat forever) of $x(n)$. That’s an interesting notion, but its larger impact is that this fact in turn means that the linear modulating signal is periodic. This modulating signal is thus a SAWTOOTH, and we in turn regard the sawtooth as being composed of sinewaves (Fourier Series). Well, not exactly, but the FFT fits (forces) the sampled version of the signal into a finite set of “DFT harmonics”. This is a sawtooth-like waveshape with some smoothed off edges as well as a wavering ramp. It can be made an excellent approximation.

The final key to the puzzle is that we know [3] well how to analyze the case of a sinusoidal modulating signal and (I think) therefore of a sum of sinusoidal components. [sidebands determined by Bessel functions, etc.] If we can obtain the FFT result via this alternative route, that would be encouraging. That is we consider the question that the FFT spikes are just a superposition of sidebands. Hopefully this would also yield more insight.

Thus having set at least three tasks (a zero-padded FFT, comparing a Fourier Series to the FFT of a sawtooth, and the review and application of ordinary FM) we note that the work here is self contained, even as it suggests additional work.

REFERENCES:

- [1] <http://electronotes.netfirms.com/EN226.pdf>
- [2] "Practical Aspects of Obtaining Spectra from FFT's", *Electronotes*, Vol. 18, No. 181, April 1993
- [3] (a) See for example "Frequency Modulation" in a very early issue of *Electronotes*, Vol. 2, No. 13, Oct. 10, 1972; (b) Numerous books contain excellent derivations such as S. Goldman's *Frequency Analysis, Modulation and Noise*, Dover (1948); (c) Also be sure to see John Chowning, "The Synthesis of Complex Audio Spectra by Means of Frequency Modulation," *J. Audio Eng. Soc.*, Vol. 21, No. 7, Sept 1973, pp 526-534
https://ccrma.stanford.edu/sites/default/files/user/jc/fmsynthesispaperfinal_1.pdf
http://people.ece.cornell.edu/land/courses/ece4760/Math/GCC644/FM_synth/Chowning.pdf
Also see the interesting YouTube: <http://www.youtube.com/watch?v=w4g92vX1YF4>
- [4] B. Hutchins, "The Frequency Modulation Spectrum of an Exponential Voltage-Controlled Oscillator", *J. Audio Eng. Soc.*, Vol. 23, No. 3, April 1975, pp 200-206
<http://electronotes.netfirms.com/AES3.PDF>
- [5] R.J. Bullock, "Motion Compensation for Aircraft-Borne Interferometric SAR", (Thesis Univ. College London) Appendix D, "Derivation of Linear FM Spectrum Using the Principle of Stationary Phase" Pg. 287 (2003) [Or search for "spectrum of chirp"]
http://discovery.ucl.ac.uk/1470281/1/RJ_Bullock_Thesis.pdf
- [6] Richard Hamming quotes. https://en.wikiquote.org/wiki/Richard_Hamming
- [7] R. Buderer, *The Invention that Changed the World*, Simon & Schuster (1996), pg 49

PROGRAM

For a long time in our publishing efforts, much of what appears in the text, and particularly as figures that are essential to the text, is computer generated. We have emphasized that if we add the computer code to the presentation, it serves as a backup and generally, a completely unambiguous version of what was actually done. That is, the interpretation of the code may require some additional study (although often it is easily understood), but it is never ambiguous. To the extent that we can publish runnable code as actually used, we have near perfect documentation. To the extent that the code sections are organized as the text is, things are particularly agreeable. To the extent that the code is simple minded, the lack of sophistication may be deemed a virtue. In cases where we feel we have succeeded (as here) in mirroring the actual presentation in the code, we justify printing it. At the least, it should (also) be available online as a last resort, which is the case here <http://electronotes.netfirms.com/EN227MFile.txt>

```

% en227.m    (Matlab)
% All figure numbers in program match those of EN227 text
%
% Fig 1 and 2 of EN227
% frequency 10 to 14 as direct argument
n=0:1999;
x10=cos(2*pi*10*n/2000);
x11=cos(2*pi*11*n/2000);
x12=cos(2*pi*12*n/2000);
x13=cos(2*pi*13*n/2000);
x14=cos(2*pi*14*n/2000);
x10p5=cos(2*pi*10.5*n/2000);
f10to14=10+[0:.002:3.998];
%
% Alternative Overwrites
%f10to14=10+[0:.002:3.998];
%f10to14=10+[0:.001:1.999];
%f10to14=9.5+(5/4)*[0:.002:3.998];
x10to14=cos(2*pi*f10to14.*n/2000);
%
X10=fft(x10);
X11=fft(x11);
X10p5=fft(x10p5);
X10to14=fft(x10to14);
X10(1:20);
X11(1:20);
X10p5(1:20);
X10to14(1:20);
%
figure(1)          % FFTs
MX10=abs(X10);
MX11=abs(X11);
MX10p5=abs(X10p5);
MX10to14=abs(X10to14);
%
subplot(221)
stem([0:29],MX10(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 10')
%
subplot(222)
stem([0:29],MX11(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 11')
%

```



```

subplot(223)
stem([0:29],MX10p5(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -70 700])
title('k = 10.5')
%
subplot(224)
stem([0:29],MX10to14(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -40 500])
title('k = 10 to 14')
%
%
figure(2)          % Time Domain
plot(x10to14)
hold on
plot(x10,'g:')
plot(x14,'r:')
plot([0:2000/10:2000],1.05*ones(1,11),'g*')
plot([0:2000/14:2000],1.1*ones(1,15),'r*')
%
mm=0;
mmm=[];
for nn=1:2000
    if x10to14(nn)>.998
        nn-1;
        mm=mm+1;
        mmm(mmm)=nn-1;
    end
end
mm;
plot(mmm,ones(1,mm),'b*')
hold off
axis([-100 2100 -1.2 1.2])
title('Sweep k=10 to 14, 10=green 14=red')
%
%
%
% Fig 3 and 4 of EN227
% integrated instantaneous frequency
%

```

```

n=0:1999;
x10=cos(2*pi*10*n/2000);
x11=cos(2*pi*11*n/2000);
x12=cos(2*pi*12*n/2000);
x13=cos(2*pi*13*n/2000);
x14=cos(2*pi*14*n/2000);
x10p5=cos(2*pi*10.5*n/2000);
f10to14=10+[0:.001:1.999];
%
x10to14=cos(2*pi*f10to14.*n/2000);
%
X10=fft(x10);
X11=fft(x11);
X10p5=fft(x10p5);
X10to14=fft(x10to14);
X10(1:20);
X11(1:20);
X10p5(1:20);
X10to14(1:20);
figure(3) % FFTs
MX10=abs(X10);
MX11=abs(X11);
MX10p5=abs(X10p5);
MX10to14=abs(X10to14);
%
subplot(221)
stem([0:29],MX10(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 10')
%
subplot(222)
stem([0:29],MX11(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 11')
%
subplot(223)
stem([0:29],MX10p5(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -70 700])
title('k = 10.5')
%

```

```

subplot(224)
stem([0:29],MX10to14(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -40 700])
title('k = 10 to 14')
%
%
figure(4)      % Time Domain
plot(x10to14)
hold on
plot(x10,'g:')
plot(x14,'r:')
plot([0:2000/10:2000],1.05*ones(1,11),'g*')
plot([0:2000/14:2000],1.1*ones(1,15),'r*')
%
mm=0;
mmm=[];
for nn=1:2000
    if x10to14(nn)>.998
        nn-1;
        mm=mm+1;
        mmm(mm)=nn-1;
    end
end
mm;
plot(mmm,ones(1,mm),'b*')
hold off
%
axis([-100 2100 -1.2 1.2])
title('Sweep k=10 to 14, 10=green 14=red')
%
%
%
%
% Fig. 5 Rotating Vector Picture
theta=0:0.1:360;
cx=cos(2*pi*theta/360);
cy=sin(2*pi*theta/360);
figure(5)
plot(cx,cy,'b')
hold on
plot([-2 2],[0 0],'k:')
plot([0 0],[-2 2],'k:')
plot([0 cos(pi/3)],[0 sin(pi/3)],'r')
plot(0.3*cx(1:600),0.3*cy(1:600),'g')
axis('equal')
axis([-1.3 1.3 -1.3 1.3])
hold off
%
%
```

```

% Fig 6 and 7 of EN227
% Reverse Direction Integrated Instantaneous Frequency - As Check
%
n=0:1999;
x10=cos(2*pi*10*n/2000);
x11=cos(2*pi*11*n/2000);
x12=cos(2*pi*12*n/2000);
x13=cos(2*pi*13*n/2000);
x14=cos(2*pi*14*n/2000);
x10p5=cos(2*pi*10.5*n/2000);
f14to19=14-[0:.001:1.999];
%
x14to10=cos(2*pi*f14to19.*n/2000);
%
X10=fft(x10);
X11=fft(x11);
X10p5=fft(x10p5);
X14to10=fft(x14to10);
X10(1:20);
X11(1:20);
X10p5(1:20);
X14to10(1:20);
%
figure(6) % FFTs
MX10=abs(X10);
MX11=abs(X11);
MX10p5=abs(X10p5);
MX10to14=abs(X14to10);
%
subplot(221)
stem([0:29],MX10(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 10')
%
subplot(222)
stem([0:29],MX11(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 11')
%
subplot(223)
stem([0:29],MX10p5(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -70 700])
title('k = 10.5')
%

```

```

subplot(224)
stem([0:29],MX10to14(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -40 700])
title('k = 14 to 10')
% Overwrite (Discard) Above Fig. 6
%
figure(6)
clf
plot(x14to10)
hold on
plot(x10,'g:')
plot(x14,'r:')
plot([0:2000/10:2000],1.05*ones(1,11),'g*')
plot([0:2000/14:2000],1.1*ones(1,15),'r*')
%
mm=0;
mmm=[];
for nn=1:2000
    if x14to10(nn)>.998
        nn-1;
        mm=mm+1;
        mmm(mm)=nn-1;
    end
end
mm;
plot(mmm,ones(1,mm),'b*')
hold off
%
axis([-100 2100 -1.2 1.2])
title('Sweep k=14 to 10, 10=green 14=red')
%
% WRONG WAY TIME REVERSED
%
%
% SECOND RUN - REVERSE
%
%
n=0:1999;
x10=cos(2*pi*10*n/2000);
x11=cos(2*pi*11*n/2000);
x12=cos(2*pi*12*n/2000);
x13=cos(2*pi*13*n/2000);
x14=cos(2*pi*14*n/2000);
x10p5=cos(2*pi*10.5*n/2000);
f14to10=14-[0:.002:3.998];
%
x14to10=cos(2*pi*f14to10.*n/2000);

```

```

X10=fft(x10);
X11=fft(x11);
X10p5=fft(x10p5);
X14to10=fft(x14to10);
X10(1:20);
X11(1:20);
X10p5(1:20);
X14to10(1:20);
%
figure(7)
MX10=abs(X10);
MX11=abs(X11);
MX10p5=abs(X10p5);
MX14to10=abs(X14to10);
%
subplot(221)
stem([0:29],MX10(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 10')
%
subplot(222)
stem([0:29],MX11(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -100 1100])
title('k = 11')
%
subplot(223)
stem([0:29],MX10p5(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -70 700])
title('k = 10.5')
%
subplot(224)
stem([0:29],MX14to10(1:30))
hold on
plot([-10 50],[0 0],'k:')
hold off
axis([-5 35 -40 500])
title('k = 14 to 10')
%

```

```

figure(8)
plot(x14to10)
hold on
plot(x10, 'g: ')
plot(x14, 'r: ')
plot([0:2000/10:2000], 1.05*ones(1,11), 'g*')
plot([0:2000/14:2000], 1.1*ones(1,15), 'r*')
%
mm=0;
mmm=[];
for nn=1:2000
    if x14to10(nn)>.998
        nn-1;
        mm=mm+1;
        mmm(mm)=nn-1;
    end
end
mm;
plot(mmm, ones(1,mm), 'b*')
hold off
%
axis([-100 2100 -1.2 1.2])
title('Sweep k=14 to 10, 10=green 14=red')
%
%
%   EXPERIMENTATION SECTION
%
%
% Fig 9 and 11
% integrated instantaneous frequency (longer)
%
n=0:9999;
x10=cos(2*pi*10*n/2000);
x11=cos(2*pi*11*n/2000);
x12=cos(2*pi*12*n/2000);
x13=cos(2*pi*13*n/2000);
x30=cos(2*pi*30*n/2000);
x10p5=cos(2*pi*10.5*n/2000);
f10to20=10+[0:.001:9.999];
%
x10to20=cos(2*pi*f10to20.*n/2000);
%
X10=fft(x10);
X11=fft(x11);
X10p5=fft(x10p5);
X10to20=fft(x10to20);
X10(1:20);
X11(1:20);
X10p5(1:20);
X10to14(1:20);

```

```

figure(9)      % FFT of longer sequence
%
MX10to20=abs(X10to20);
stem([0:399],MX10to20(1:400))
hold on
plot([-10 1000],[0 0],'k')
hold off
axis([-5 300 -40 700])
title('Experimental FFT - Longer')
%
figure(10)     % left side
plot(x10to20)
hold on
plot(x10,'g:')
plot(x30,'r:')
hold off
%
axis([-30 1000 -1.2 1.2])
title('Experimental - Longer (Left Side)')
%
figure(11)     % Right side
plot(x10to20)
hold on
plot(x10,'g:')
plot(x30,'r:')
hold off
%
axis([ 9000 10030 -1.2 1.2])
title('Experimental - Longer (Right Side)')

%
% Formula test
t0=2.3
t1=5.7
x0=4.1
x1=7.3
m=(x1-x0)/(t1-t0)           % Slope
b=(x1*t0-x0*t1)/(t0-t1)    % Intercept
a1=(x1-x0)*(t1+t0)/2 + b*(t1-t0) % Area by Integration
a2=x0*(t1-t0)+(t1-t0)*m*(t1-t0)/2 % Area by Rect + Tri
a3=(x0*t1-x0*t0-t0*x1+t1*x1)/2 % Area by Common Formula

% Check 100 full cycles - Figure Not Used in Text
figure(13)
plot(x10to20)
axis([-100 10100 -1.1 1.1])

```

Matlab file online at Electronotes site as <http://electronotes.netfirms.com/EN227MFile.txt>