8 KLSPEC: A SPECTRAL ANALYSIS PROGRAM

The KLSPEC speech analysis program accepts user commands to read in waveform files and generate spectral displays of various types. More than one waveform file can be analyzed at one time; for example, natural and synthetic utterances can be compared. It is useful to have printouts made by the program SPECTO available before you while using KLSPEC in order to more easily find your way through the utterances.

Spectral representations that can be obtained include (1) the DFT magnitude spectrum, (2) a spectrogram-like spectrum, (3) a critical-band spectrum that employs a Mel frequency scale and filter bandwidths that increase with increasing frequency, and (4) a linear prediction spectrum. Voice fundamental frequency is displayed with the DFT spectral representation, while estimates of formant frequency positions are displayed with the spectrogram-like spectral representation and the lpc representation, and spectral peak locations/amplitudes are displayed with the critical-band representation. The dft spectrum can be displayed below the lpc spectrum or below the spectrogram-like spectrum by special commands if one wishes to know how good a fit there is between the smoothed representation and the actual spectrum.

Permanent hard copy of any display can be obtained by executing the appropriate user command.

8.1 STARTING KLSPEC

The program is evoked simply by typing:

$ *klspec<CR>*

The program begins by asking for the first name of a waveform file to be analyzed. A sample of a complete dialog for analyzing the waveform file 'foo' is:

$ *klspec<CR>*

```
KLSPEC   Version _   D.H. Klatt   Date-created

First name of input waveform file 1: foo<CR>
```

```
Begin reading waveform file  foo  into core buffer 1 ...
Waveform contains xxxxx samples (at xxxxx samples/sec)

Assume VT1xx terminal at Vax-x workstation

Ready for first command
```

After this waveform is read into core, the program types a few lines of information and then waits for the user to issue an analysis command. Legal analysis commands are listed in Table 15, and are described in Section 8.2.

8.1.1 Optional Arguments

The program KLSPEC can take arguments of two types. The first type is an argument preceded by a minus sign, which identifies the type of terminal on which the spectral plots are to be displayed. The computer knows what type of terminal is usually connected to each computer line, so terminal specification is not necessary unless a terminal has been moved from its normal location.

$ *klspec<CR>*	(plot parameter data on default terminal)
$ *klspec -vt100<CR>*	(request VT-100 terminal, no plots)
$ *klspec -vt125<CR>*	(request VT-125 terminal)
$ *klspec -vs100<CR>*	(plot parameter data on a VS-100 workstation)

The second type of argument, identified by the absence of a minus sign, is the name of a waveform file. Up to ten waveform names can be strung together, separated by spaces. However, normally, one or two names are given, as analysis of one, and comparison between two waveforms, are the most common activities:

$ *klspec foo<CR>*	(analyze waveform foo)
$ *klspec foo bar<CR>*	(analyze waveforms foo and bar)

Arguments are separated by spaces. If information such as the format of the arguments to KLSPEC is needed, but this manual is *not* at hand, type '?' in place of an appropriate command or argument to get a list of legal options at that point:

$ *klspec ?<CR>*

8.2 ANALYSIS COMMANDS

When the flashing square appears at the left of the screen, the user may request any of the actions listed in Table 15 by typing a single character, not to be followed by a carriage return.

8.2.2 'q', Quit

The only correct way to quit the KLSPEC program is by typing 'q'. Otherwise, the terminal may be left in a strange state.

If an emergency requires use of '↑C' to get out of a program, and the terminal is not using the full screen, hit the 'reset' key followed by the '0' key. Wait until the self-test is over, and then type '<CR>'. The usual '$' system prompt should appear. If not, it is possible that the terminal baud rate is not correct and should be changed to 4800 or 9600, whichever it isn't.

8.2.3 'd', DFT Spectrum Requested

A 512-point discrete Fourier transform is computed of a selected chunk of waveform that is first differenced, and multiplied by a Hamming window. The magnitude, in dB is displayed of the resulting 256 dft samples. An example of a dft spectrum is shown in Figure 17. This spectral representation is useful for analyzing periodic speech sounds in which there is an interest in the amplitudes of individual harmonic components.[26]

The waveform is displayed below the spectrum. This takes some time on a VT-125 terminal. To speed up the process, fewer waveform points can be displayed by increasing 'wi', waveform increment' from its default value of 2, using the 'C' change-parameter-value command. If 'wi' is set to zero, the waveform is not displayed at all, and the display is maximally fast.

The program computes and displays fundamental frequency if it is determined that local spectral maxima occur at fairly regular intervals. The algorithm used to estimate fundamental frequency is a preliminary simple-minded implementation of the Goldstein algorithm.[27] The fundamental is computed by collecting frequencies of local maxima in the dft spectrum (presumably each local maximum is a harmonic of f0). Only peaks below 3000 Hz contribute to this pool, and the fundamental is judged to be the f0 that accounts for the most peaks as harmonics. If there is little low-frequency energy present in the spectrum, or if the distribution of differences is too dispersed in frequency, the spectrum is termed unvoiced and no fundamental is displayed. The algorithm is susceptible to errors, particularly when f0 is changing within the analysis window.

---

[26] The true harmonic amplitudes would be obtained only if the first difference preemphasis were deactivated, using the 'C' command to set 'fd' to zero.

[27] Duifhuis, H., Willems, L.F., and Sluyter, R.J. (1982). "Measurement of Pitch in Speech: An implementation of Goldstein's Theory of Pitch Perception", J. Acoust. Soc. Am. 71, 1568-1580.

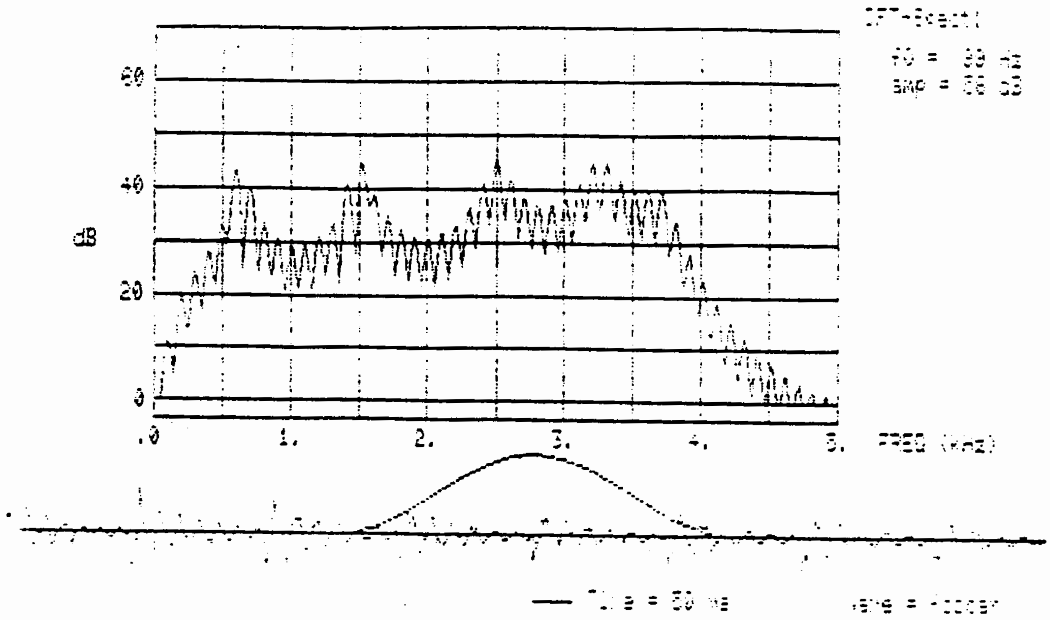

Figure 17: KLSPEC DFT magnitude spectrum of a synthetic vowel, produced with the 'd' command.

The duration of the Hamming window is set to a default of 299 samples,[28] which is 30 ms at a sampling rate of 10,000 samples/sec. The effective (3 dB down) duration of the window is about half this value, but it is still relatively long compared with default windows used for spectrographic and critical-band spectral representations, so as to be able to track low fundamental frequencies. The window duration used in computing the dft is called 'wd' (for 'window-dft') and is user adjustable by typing the 'P' and 'C' commands described below. A shorter window, e.g. 'wd'=200, might profitably be used for analysis of female and children's voices to minimize the possibility that f0 is non-constant during the analysis window.

The dft magnitude remains as the default spectral display type until either 's', 'c', or 'l' commands are used to change the default display type.

---

[28] The remainder of the 512 samples sent to the dft routine are, of course, set to zero.

8.2.4 's', SPECTROGRAM-like Spectrum Requested

An approximation to the filter set used in a broadband spectrogram display has been created by windowing a selected portion of waveform (window duration for spectrogram, 'ws', has a default value of 256 samples), computing a 256-point dft, and forming a weighted sum of adjacent dft sample energies for each of 128 spectrogram-like 'filters'. An example of this display is shown in Figure 18.

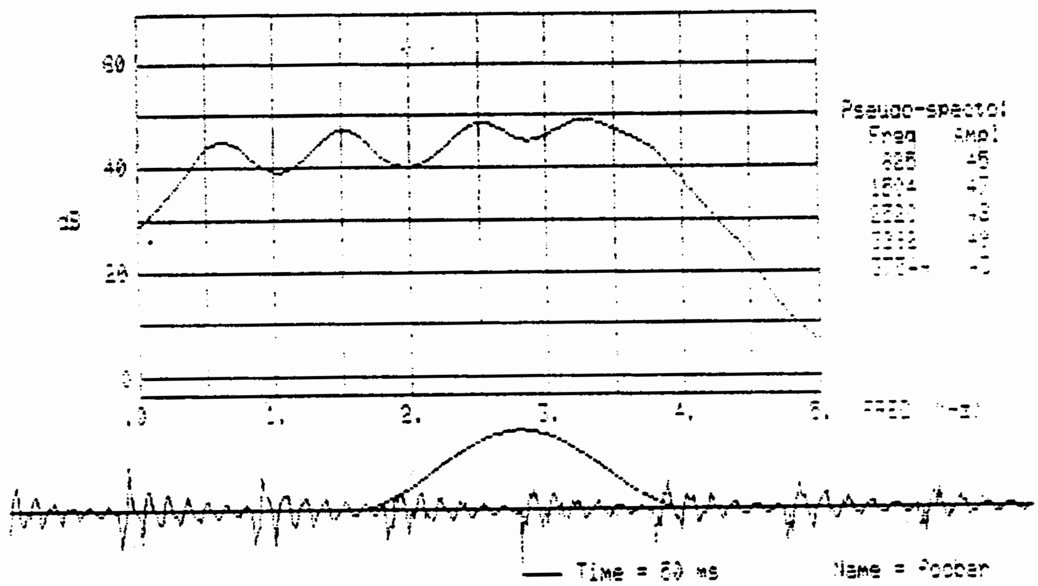

Figure 18: KLSPEC Spectrogram-like spectrum of a synthetic vowel, produced with the 's' command.

Each filter has a frequency-domain shape that is approximately Gaussian, and has a 3-dB bandwidth, 'bs' -- bandwidth specto display, that has a default value of 300 Hz. When analyzing spectra obtained from speakers with a high fundamental frequency, it is a good idea to increase this analyzing bandwidth to 400 Hz or more, using the 'P' and 'C' commands described below.

The duration of the Hamming window, 'ws' -- window specto display, can also be adjusted using the 'C' command. A shorter window, which is perhaps desirable for analysis of female voices, requires padding the windowed waveform with zeros so as to provide the dft subroutine with 256 points (so that 128 spectrogram-like filter outputs are still computed). On the other hand, if the requested window duration is increased above 256 samples, as would of course be desirable if the sampling rate of the waveform being analyzed were e.g. 16.000 samples/sec. then the present algorithm uses a folding procedure to fit the long waveform chunk into a form suitable for a 256-point dft calculation. (FIX) Users should consult a text on digital signal processing to get a feeling for the unfortunate consequences of this.

Local maxima in the display are often indicative of the frequency positions of formants. The program estimates and displays the locations of prominent spectral peaks. The interpolation algorithm used to estimate formant frequency positions provides an improvement in accuracy over the 40-Hz resolution implied by a 128-sample spectrum covering 5 kHz, and is quite accurate when formant peaks are well-separated. When two formants are hiding under a single broad maximum, the program tries to find the hidden formant by looking for a local minimum in the derivative of the spectrum with respect to frequency. This is how the fifth formant frequency was estimated in the figure. However, this simple-minded initial version cannot always resolve two formants hiding under a single peak. The algorithm also attempts to ignore small local maxima in the valleys between formant peaks, but again, it is not always successful.

When in pseudo-spectrogram mode, the spectrum computed previously is displayed faintly. This facilitates detection of spectral changes, and also permits direct comparison of spectra from two waveform files, using the 'x' command described below.

8.2.5 'S' ~~[illegible]~~, SPECTROGRAM-like Spectrum and LPC Spectrum Requested

The ~~four-keystroke~~ command 's' ~~'(esc)' followed by 'd', 's', and another '(esc)'~~ requests that both the SPECTO and DFT spectra be computed and superimposed, and that both formant frequencies and fundamental frequency be estimated. ~~(The '(esc)' is a meta-command key that permits stacking up several commands before the next display is created, see Section 8.3).~~ The cost, of course is that this amount of computation takes a little longer. One may wish to locate a position in the waveform more quickly using the 's' command, and then switch to the ~~[illegible]~~ command for plotting purposes. This command is useful for determining the relation between a smoothed peak location and the individual harmonics of a voiced speech sound in order to verify the reasonableness of formant estimation. An example is shown in Figure 19.

8.2.6 'c', CB (Critical-Band) Spectrum Requested

The critical band is a concept originating in auditory psychophysics, and it refers to the inability of the ear to resolve the frequency locations of individual sinusoidal frequency components if they are within a critical frequency interval. There are many different estimates of the critical bandwidth as a function of center frequency of an auditory filter, but a typical quoted bandwidth relationship as a function of frequency is 80 Hz at low frequencies, 160 Hz at 1000 Hz, and 0.16 of the filter center frequency above 1000 Hz.

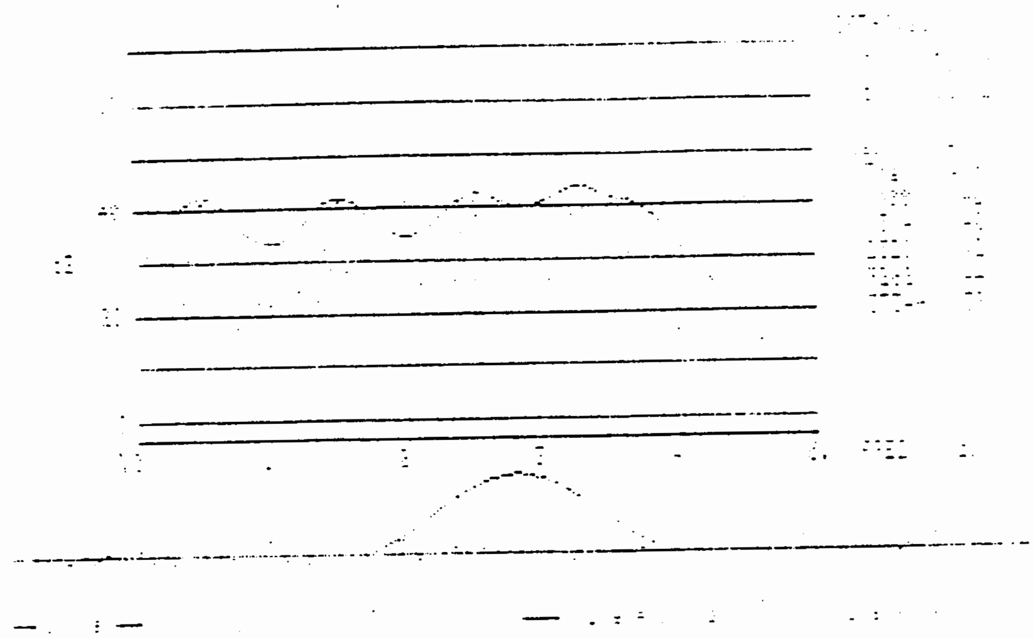

Figure 19: KLSPEC Spectrogram-like spectrum and superimposed DFT spectrum of a synthetic vowel, produced with the '(esc)ds(esc)' command.

The critical-band filter set employed in KLSPEC resembles these characteristics, but has been redesigned to have rather wider filters, particularly at low frequencies, so as to not resolve individual harmonic components (see Table 16). In this way, local maxima usually reflect formant locations rather than harmonic locations. The lowest frequency filter is centered at 200 Hz -- energy below about 80 Hz is thus excluded from the critical-band representation -- the reason being that these very low frequencies typically contain mostly background noise.

Each critical-band filter has been approximated by windowing a selected portion of waveform (window duration for critical-band spectrum, 'wc', has a default value of 256 samples), computing a first difference preemphasis, computing a 256-point dft, and forming a weighted sum of adjacent dft sample energies for each of about 36 critical-band 'filters'. The weighted sum is computed only for dft samples within 40 dB of the peak in order to speed up calculations. An example of this display is shown in Figure 20.

A default window duration of 25.6 ms and a Hamming shape were chosen as a compromise between the desirability of long windows to be insensitive to locations of pitch periods within the window and the desirability of short windows to characterize rapidly

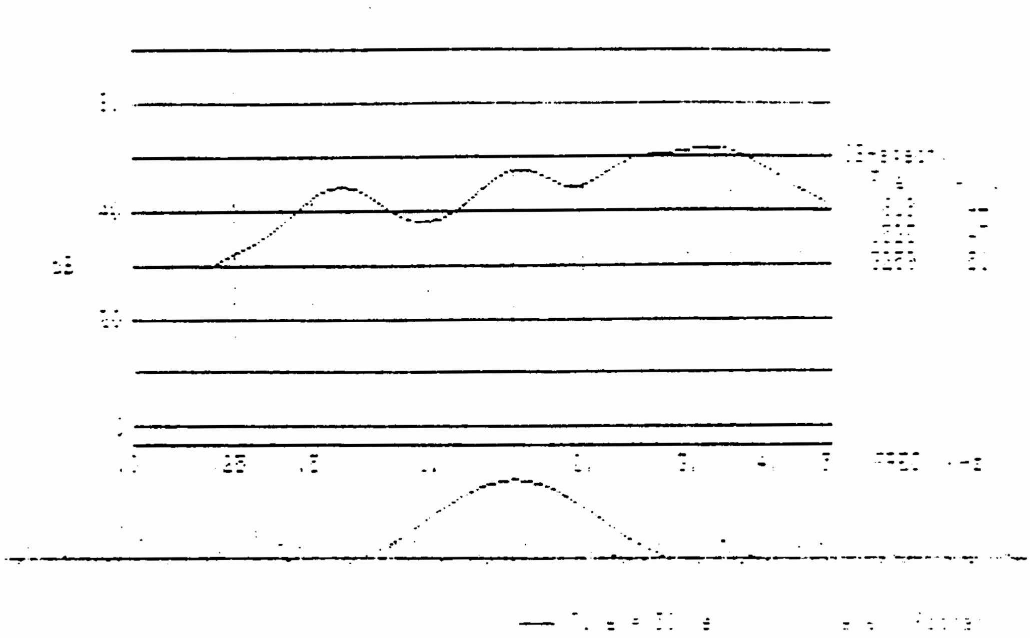

Figure 20: KLSPEC critical-band spectrum of a synthetic vowel, produced with the 'c' command.

changing spectral events. The window has an effective duration (time between 3 dB-down points) of about 10 to 12 ms, which is not unlike the duration of perceptual analysis that is implied by several psychophysical experiments (such as gap detection). However, the auditory system appears to have better temporal resolution at higher frequencies, a characteristic that cannot be duplicated when using a windowed dft to create a set of pseudo-filters, as we have done. The 'C' change-parameters command can be used to change the critical-band analysis window duration parameter 'wc'. If the duration is increased above 256 samples, as would of course be desirable if the sampling rate of the waveform being analyzed were e.g. 16,000 samples/sec. then the present algorithm uses a folding procedure to fit the long waveform chunk into a form suitable for a 256-point dft calculation. (FIX) Users should consult a text on digital signal processing to get a feeling for the undesirable negative consequences of this.

The first difference boosts the relative energy in the higher frequencies and attenuates very low frequency components in a manner not unlike the equal loudness contour at 40 dB SPL. Speech spectra look good with this preemphasis because it is easier to see higher-frequency peaks that have perceptual importance, and the first harmonic does not look like a spurious formant,[29] but one must always remember that the true spectrum is tilted down with respect to what is displayed. To see the true spectrum, use the 'C' change-parameter command to set 'fd' to zero.

Filters are spaced at equal intervals on the technical mel frequency scale, which relates frequency *f* to pitch *m* in mels by the equation:

$$m = k * \log[\, 1. + (\, f / 700.\,)\,]$$

where k is set so that f = 1000 results in m = 1000. The effect of this transformation can be seen in the horizontal axis of the critical-band spectrum displayed in Figure 20 -- low frequencies take on greater visual important than higher frequency components in a manner analogous to their perceptual importance, as measured e.g. by intelligibility of bandpass filtered speech. Filter bandwidths change by a similar function such that adjacent filters overlap at about their 1 dB down points. The frequency at which the critical-band spectrum frequency axis goes from being linear to logarithmic, 700 Hz, is actually user adjustable, being the changeable parameter 'fl' in Table 17.

To increase or decrease the total number of filters used, simply change the parameters 'f1' and 'f2'[30] which are the center frequencies of the first and second filters. Since filter center frequencies follow the technical mel scale, increasing 'f2' will reduce the total number of filters and space them out more widely (but will not change their bandwidths). Due to program space limitations, it is not possible to increase the number of filters (by decreasing 'f2') beyond a certain limit -- the limit can be found by trial and error.

The default filter design values result in a set of 36 filters spanning the frequency range from 200 to 4850 Hz. If fewer than 36 filters were used, as is often the case in speech analysis/recognition systems, formant frequency estimation would not have the resolution that is attainable with the current interpolation scheme. Further evidence for the desirability of many overlapping critical-band filters comes from experiments that have shown that a critical-band vocoder degrades in intelligibility if fewer than 30 or so filters are employed.[31]

---

[29] This representation might be closer to the 'Central Spectrum', i.e. the spectrum representation presented to the central nervous system by the peripheral auditory system, but not enough is presently known to be sure.

[30] See the 'P' and 'C' commands described below.

[31] Klatt, D.H., Seneff, S. and Zue, V. (1982). "Design Considerations for Optimizing the Intelligibility of a DFT-Based, Pitch-Excited, Critical-Band Spectrum Speech Analysis/Resynthesis System", Speech Communication Group Working Papers, Vol. 1, 31-46.

Each filter has a frequency-domain shape that is approximately Gaussian, and has a 3-dB bandwidth that has a default value as given in Table 16. The bandwidth of each filter is determined by the parameters 'b0' and 'b1', where 'b1' is the nominal bandwidth of the filter centered at 1000 Hz. All other filter bandwidths are given values that are either proportional to filter center frequency (with the same proportionality as 'b1'/1000) or fixed at 'b0' Hz, whichever is greater. For the default values of 'b0' and 'b1', filter bandwidths are never smaller than 210 Hz, are 250 Hz at 1000 Hz, and increase thereafter proportional to filter center frequency.

When analyzing spectra obtained from speakers with a high fundamental frequency, it is a good idea to increase these analyzing bandwidths at low frequencies to 350 Hz or more, using the 'C' command described below to set 'b0'. To increase the bandwidths at high frequencies, as for example might be desirable for more smoothing of noise spectra, the parameter 'b1' would be increased.

Local maxima in this display are sometimes indicative of the frequency positions of formants. The program estimates and displays the locations of prominent spectral peaks. However, at higher frequencies, the analysis bandwidth is too great to resolve individual formants, and peak frequencies displayed are often the average of several formants contributing to a broad energy concentration. If formant data is desired, use the 's' command.

When in critical-band spectrum mode, the spectrum computed previously is displayed faintly. This facilitates detection of spectral changes, and also permits direct comparison of spectra from two waveform files, using the 'x' command described below.

8.2.7 'l', LPC, Linear-Prediction Spectrum

The linear prediction spectrum has been used a great deal recently in speech research. An LPC spectral representation is provided here for completeness, but LPC has several limitations so that its use is not advocated. In particular, bandwidth estimates are very sensitive to the location of harmonics relative to formant frequency, spectral zeros are poorly represented, and fricative spectra are represented as having multiple peaks even when the dft indicates that the true spectrum is relatively smooth. An example of the LPC spectral display for a very regular synthetic waveform (for which LPC works very well) is shown in Figure 21.

The number of points in the Kaiser time weighting window, 'wl' -- window lpc display, has a default value of 256, but can be changed with the 'P' and 'C' commands.

The number of LPC coefficients, 'nc', is assigned a default value of 14, which is a good choice for analyzing the speech of an adult male at 5 kHz. To reduce the value of 'nc'

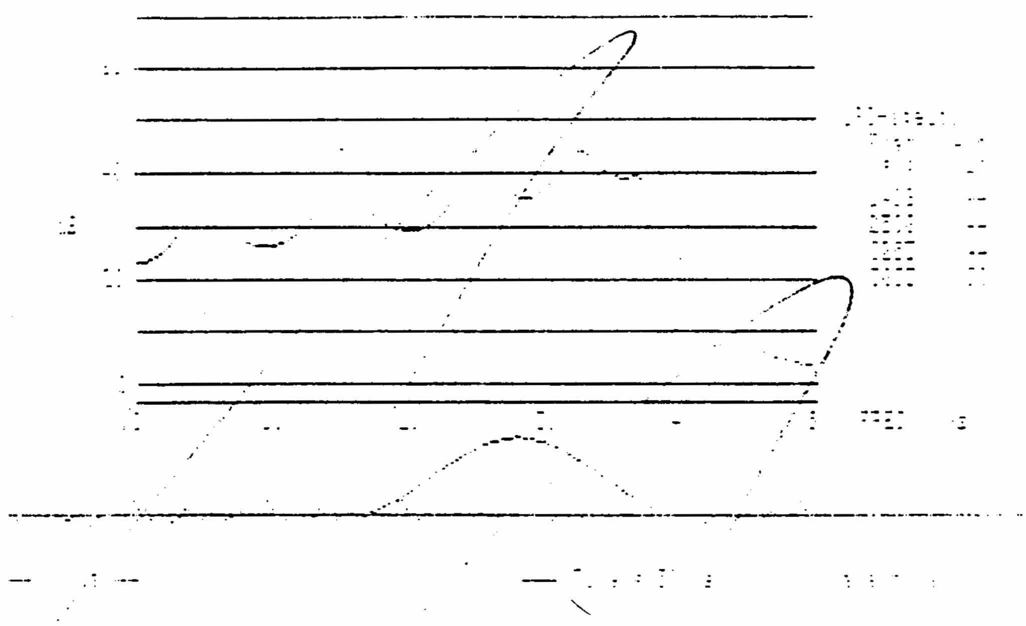

Figure 21: KLSPEC LPC linear-prediction spectrum of a synthetic vowel, produced with the 'l' command.

to e.g. 12 in order to study female speech, the 'P' and 'C' commands may be used.

8.2.8 '(esc)dl(esc)', LPC and DFT Spectrum Requested

The four-keystroke command '(esc)' followed by 'd', 'l', and another '(esc)' requests that both the LPC and DFT spectra be computed and superimposed, and that both formant frequencies and fundamental frequency be estimated. (The '(esc)' is a meta-command key that permits stacking up several commands before the next display is created, see Section 8.3). The cost, of course is that this amount of computation takes a little longer. One may wish to locate a position in the waveform more quickly using the 'l' command, and then switch to the '(esc)s' command for plotting purposes.

For regular synthetic vowel-like waveforms, such as in Figure 22, the LPC fit to the DFT is quite good. If LPC analysis is to be used, it is strongly suggested that the DFT spectrum be superimposed, so that the deficiencies of the LPC match to fricatives and spectral zeros is appearant.

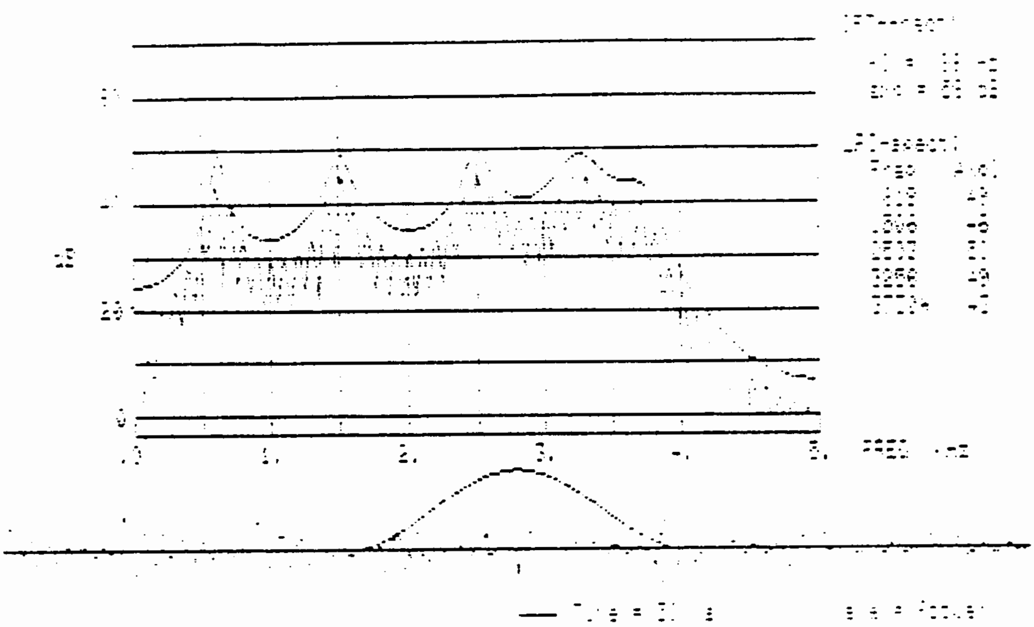

Figure 21: KLSPEC LPC spectrum and superimposed DFT spectrum of a synthetic vowel, produced with the '/' command.

8.2.9 'g', GRAPH on LA-50 Hard Copy Unit

Any of the above four spectral displays may be produced by a command listed above, and then copied to the LA-50 hard-copy unit using the 'g' command.

If the LA-50 should run out of paper in the middle of a plot, simply insert more paper, as you would for an ordinary typewriter, and then hit the *ready* button on the front panel of the LA-50.

8.2.10 '(esc)', Stack up Commands

Since the VT-125 is rather slow, it can be frustrating to go through a sequence of commands in order to achieve some spectral representation because a full plot is generated after each command. To inhibit plotting temporarily while a set of commands is stacked up and executed without plotting, one simply uses the '(esc)' meta-command. For example, to (1) move to the other waveform buffer, (2) select a different spectral representation, and (3) select a time at which to do the analysis, the following typed string will accomplish all three with a minimum delay:

```
(esc)xdt250<CR>(esc)
```

The first '(esc)' sets the program into stacking mode, and the second '(esc)' exits from this mode, simultaneously requesting a plot. Note that a waveform selecting command must always be given before a time setting command. Otherwise, order of commands within a stack is free.

8.2.11 'p', PLAY Current Waveform

The entire contents of the waveform buffer currently being displayed is played out through the d/a converter whenever the 'p' command is issued. The requested material is played once, and control is returned to the command processor.

8.2.12 'O', OUTPUT Filter Characteristics

A printout of the center frequencies and bandwidths of either the spectrogram-like filter set or the critical-band filter set in use at the time can be obtained with the 'O' command (FIX). An example, which describes the default characteristics of the critical-band filter set, is shown in Table 16.

FILTER CHACTERISTICS:

n	FREQ	BW	n	FREQ	BW	n	FREQ	BW
1	200	210	13	979	250	25	2432	622
2	248	210	14	1068	273	26	2599	665
3	299	210	15	1163	297	27	2775	710
4	352	210	16	1262	323	28	2960	757
5	408	210	17	1367	349	29	3156	807
6	467	210	18	1477	378	30	3361	860
7	529	210	19	1593	407	31	3578	915
8	595	210	20	1715	439	32	3806	974
9	664	210	21	1844	472	33	4046	1035
10	737	210	22	1980	506	34	4299	1100
11	813	210	23	2123	543	35	4566	1168
12	894	228	24	2273	581	36	4847	1240

Table 16: Center frequencies and bandwidths of the critical-band filter set, as obtained using the 'O' command.

8.2.13 'o', OUTPUT Printout of Waveform Values in Window

If numerical values of waveform samples are desired (such information is rarely needed), the samples within the analysis window are printed out in decimal when the 'o' command is given.

8.2.14 'P', Print Changeable Parameters

A complete printout of analysis parameters that can be changed during an analysis session is obtained by typing 'P'. These constants are listed in Table 17. The command used to change the value of any one of them, 'C', is described next.

2-CHAR SYMBOL	CURRENT VALUE	DESCRIPTION
sr	10000	Sampling rate assigned to current waveform
wc	256	Number of samples in Hamming time window for CB
wd	299	Number of samples in Hamming time window for DFT
ws	256	Number of samples in Hamming time window for SPECTO
wl	256	Number of samples in Kaiser time window for LPC
nc	14	Number of coefficients in LPC analysis
fd	1	First difference preemphasis if =1, not if =0
b0	210	Bandwidth of CB filter at lowest center frequency
b1	250	Bandwidth of CB filter at 1000 Hz
bs	300	Bandwidth of Spectrogram-like filter set
f1	200	Center frequency of first CB filter
f2	248	Center frequency of second CB filter
fl	700	Freq at which CB spect freq axis goes from linear to log
wi	2	Wave incr: display with wave point for speed, none if=0

Table 17: KLSPEC changeable analysis parameters.

Each of these parameters has built-in limits to prevent the user from doing things that are not likely to work.

The parameter 'wi' controls how (or whether) the waveform is plotted at the bottom of the screen. Because the VT125 is a relatively slow terminal, there may be situations where you wish to speed it up, either by deleting the waveform portion of the plot ('wi' = 0), or by degrading the waveform plot by skipping over points. For example, if 'wi' is set to 4, only every 4th point is displayed, and it takes only one quarter of the time to display the waveform.

8.2.15 'C', CHANGE Value of an Analysis Parameter

To change the default value of an analysis parameter, the following dialog is entered:

```
> C
    Par: sr<CR>
        Change value from 10000 to: 16000<CR>
>
```

The 'C' request causes the program to prompt for a parameter, which is identified by typing the appropriate 2-character symbol from Table 17. If a legal parameter is requested, the program then prompts for a value. If the value typed is within expected limits, it is silently entered into the parameter table, and the program is ready for a new user command.

However, if the value typed by the user is outside these limits, the program responds as in the example below:

```
> C
    Par: sr<CR>
        Change value from 10000 to: 0<CR>
    Requested value for sr less than min of 5000, try again
        Change value from 10000 to: 6000<CR>
>
```

8.2.16 'r', READ Another Waveform File

Up to ten waveform files can be read into core at once. However, for best performance of the program (speed of response to user commands) it is best to limit analysis to no more than about a half-minute (30 seconds) of waveform files.

To read a new file into core, type:

```
> r
    First name of input waveform file 2: bar<CR>
>
```

The program indicates the number of the file read, the sampling rate stored in the file header, and the duration of the waveform file, in samples. If the system cannot find a waveform file with the specified name, the program aborts and you have to start over (FIX).

Reading a waveform file automatically makes it current (the waveform being analyzed and displayed). To return to view/analyze one of the previous waveform files, use the '-', 'n' or 'x' commands described below.

8.2.17 '+', INCREMENT Waveform Counter, View Next Waveform

To change from viewing one waveform file to another, the '+', '-', 'n' and 'x' commands are used. One can only move the waveform counter up or down by one position at a time with the '+' and '-' commands, but 'n' allows skipping to any nonadjacent waveform buffer. The '+' increment pointer command causes the next waveform to become current. If the pointer is already set to the maximum (e.g. there are 5 files and the pointer is sitting at 5 already), then the '+' command has the effect of moving the counter to 1, i.e. wrap-around arithmetic is used.

If the waveform time cursors are locked together (see 'L' command described below) the program tries to find a cursor time in the new waveform that corresponds to the last viewed time in the previously displayed waveform. However, if that would put the time cursor off the end of the file, a maximum legal time is selected instead, and this new selected cursor position is made the default for all other waveform files.

8.2.18 '-', DECREMENT Waveform Counter, View Last Waveform

The '-' decrement pointer command causes the last waveform to become current. Wrap-around arithmetic is used, so if the counter is at 1 when a '-' command is given, the counter is set to its maximum value.

8.2.19 'n', NUMBER of Waveform to be Displayed

The program will ask for the number of the waveform buffer to be displayed. In this way, one can move between any pair of waveform buffers and compare spectra.

8.2.20 'x', EXCHANGE Waveforms

If there are two waveforms in core, and only if there are two waveforms in core, the 'x' command causes the other waveform to become current. Otherwise, 'x' has no effect.

If the two waveform time cursors are locked together (see 'L' command described below) the program tries to find a cursor time in the new waveform that corresponds to the last viewed time in the previously displayed waveform. However, if that would put the time cursor off the end of the file, a maximum legal time is selected instead (and this new selected cursor position is made the default when returning to display the other waveform file).

8.2.21 'f', FORWARD Movement of the Waveform Cursor 10 ms

The waveform displayed at the bottom of the screen can be moved forward or backward with respect to the fixed Hamming window that is displayed faintly. On a VT125 terminal, this movement is not continuous, but is in steps of +10 ms, -10 ms, +50 ms, or -50 ms. In addition, the 't' command can be used to position the waveform within 1 ms of any desired sample.

The time displayed corresponds to the waveform sample point at the center of the analysis window. To move forward 10 ms, type 'f', and wait for the picture to be repainted.

8.2.22 'F', FORWARD Movement of the Waveform Cursor 50 ms

To move forward 50 ms, type 'F', and wait for the picture to be repainted.

8.2.23 'b', BACKWARD Movement of the Waveform Cursor 10 ms

To move backward 10 ms, type 'b', and wait for the picture to be repainted.

8.2.24 'B', BACKWARD Movement of the Waveform Cursor 50 ms

To move backward 50 ms, type 'B', and wait for the picture to be repainted.

8.2.25 't' TIME of Desired Waveform Cursor Position

Upon receiving a 't' command, the program asks for a time in msec at which to position the waveform with respect to the center of the Hamming window. Times outside the range of possibility are truncated silently to the nearest legal time.

8.2.26 'L' LOCK Waveform Cursors Together

Normally, when a waveform cursor moves, by execution of a 'f', 'b', 'F', 'B', or 't' command, only the currently displayed waveform cursor changes. However, sometimes, it is desirable to e.g. align the beginning times of two waveforms to burst onset, lock the cursors together with the 'L' command, and then step forward in time, observing and comparing spectra with the 'x' command.

The program is moderately clever about handling situations where the cursors are locked, and one file is longer than the other.

8.2.27 'U', UNLOCK Waveform Cursors

If the cursors are locked together with an 'L' command, but get out of synch because one waveform has a longer phonetic segment than the other, the 'U' command can be used to detach the cursors so that they can be realigned, or so that some other activity can be pursued.

8.3 MULTIPLE COMMANDS USING THE (esc) KEY

Since the VT-125 is not very fast in plotting spectra and waveforms, it can be rather painful to make multiple changes such as selecting another waveform, a different time, and a different spectral representation, all at once. The (esc) key has been provided to solve this problem. All commands given after an (esc) keystroke are remembered but not performed until another (esc) keystroke.

For example, the typed sequence '(esc)dsxt450(esc)' selects both dft and specto plotting, moves to the other of two waveform files, and selects the time 450 ms for plotting.

NEWKLSPEC: Some Minor Modifications in KLSPEC

T. V. Anantha
3rd Feb 1992

Incorrect File Name: By mistake if you type in a wrong file name, the program quits not giving you a chance to correct the mistake. Now this doesn't happen. The program gives you a chance to reenter the file name.

Command 'r' : This command lets you load several signal files and compare them etc. However, the way it is being used is to effectively run KLSPEC in a loop instead of quitting and reentering with the new file name. After a certain number of loops you are thrown out of the program and again you enter with a new name. The command is not meant to be used that way.

New command 'v':
Case.1.: When you are working with **one file** at a time and want to load a new file: Type in 'v' on prompt '>' like,
>v
This will replace the existing file with the new file. In this way there is no limit to the number of files you can go on reading.
Case.2: When you are working with **more than one file**:
To replace the currently active file with a new file, use command 'v' . By typing 'x' you can set the active window. By setting the active window repeatedly and typing in 'v' commands you can replace old set by a new set of files.

New Command 'k':
This command is an alternative to 'a' command. This command implements spectral averaging method suggested by Prof. Ken Stevens. Begin and end time locations for spectral average are:

Begin at cursor location - kn/2
End at cursor location + kn/2

The window size is still determined by WD parameter.
What is kn ?
Initially the default value for kn is 200 samples. Use 'C' command to change kn. Note that when kn = WD, there is a single DFT spectrum, hence the lower limit for kn is WD. Maximum limit for kn is 512, the maximum DFT size used in KLSPEC.

KLSPEC92: Some Minor Modifications in KLSPEC

T. V. Anantha
3rd Feb 1992

__Incorrect File Name__: By mistake if you type in a wrong file name, the program quits not giving you a chance to correct the mistake. Now this doesn't happen. The program gives you a chance to reenter the file name.

__Command 'r'__ : This command lets you load several signal files and compare them etc. However, the way it is being used is to effectively run KLSPEC in a loop instead of quitting and reentering with the new file name. After a certain number of loops you are thrown out of the program and again you enter with a new name. The command is not meant to be used that way.

__New command 'v'__:
Case.1.: When you are working with **one file** at a time and want to load a new file: Type in 'v' on prompt '>' like,
>v
This will replace the existing file with the new file. In this way there is no limit to the number of files you can go on reading.
Case.2: When you are working with **more than one file**:
To replace the currently active file with a new file, use command 'v' . By typing 'x' (for two files) or 'n' (for more than two files) you can set the active window. By setting the active window repeatedly and typing in 'v' commands you can replace existing set by a new set of files.

__New Command 'k':__
This command is an alternative to 'a' command. This command implements spectral averaging method suggested by Prof. Ken Stevens. The consequitive blocks are imsec apart for both 'a' and 'k' commands.
Begin and end time locations for spectral average are:

Begin at cursor location - kn/2; End at cursor location + kn/2
The window size is still determined by WD parameter.

__What is kn ?__
Initially the default value for kn is 200 samples. Use 'C' command to change kn. Note that when kn = WD, there is a single DFT spectrum, hence the lower limit for kn is WD. Maximum limit for kn is 512, the maximum DFT size used in KLSPEC.

Notes on Using KLSPEC93
10-March-93

Including a spectrogram in the klspec93 display:

KLSPEC93 uses an initialization file (klspec93.ini) to locate the klspec93 windows to accommodate a spectrogram. If you include a line that begins with a 'S' or 's' in this file the program will display a spectrogram. There is no way to request a spectrogram once the program has started without one, or to delete the spectrogram if you started KLSPEC93 with one. The default klspec93.ini causes a spectrogram to be drawn.

Specification of window locations in klspec93.ini:

The location of the spectrogram on the screen is not included in this file. Changing the default location of the spectrogram window is not trivial and is not recommended. If you're dying to change the location of the spectrogram, refer to the spec_kl_defs.readme and the notes at the bottom of the spec_kl_defs.dat. The latter is the file of default settings that the spectrogram utility reads on startup. The waveform, spectrum windows, and even the text window, are located at startup based on the information in klspec93.ini. Although the locations of the graphics windows may be changed and saved when exiting KLSPEC93, the text window position cannot be saved (due to preexisting program structural considerations.)

Changing the size of the spectrogram window:

Generally it's not expected that the user will resize the spectrogram window. The window size is controlled to a large extent by the algorithm responsible for the calculation of the display and for efficiency reasons is not infinitely variable.

It is possible do create a VERY large spectrogram window suitable for group viewing by resizing the window to its maximum vertical size on the screen.

Two new spectrogram commands 'j' and 'u': (Nothing mnemonic here)

The 'u' command will generate a new spectrogram display based on the current location of the marker in the active waveform window. The algorithm used is clever enough to detect if the marker is still located in the region of the waveform used for the current spectrogram display, and will NOT recalculate the spectrogram in this case. This is done to avoid the delay in recalculating and displaying the spectrogram.

If multiple waveforms are being analyzed (using the 'r' command,) the 'j' command will generate a new spectrogram for the 'active' waveform window, centered around the current marker location in the active waveform window.

A message is displayed in the command window if the marker is placed outside of the region last used as the basis for the spectrogram display and the user is reminded to use the 'u' command to recalculate the spectrogram display.

Note: None of the parameters used in the klspec93 'spectrum' window affect the spectra in the spectrogram. Please see the spec_kl_defs.dat and spec_kl_defs.readme files for more information.