

## Chapter 18

# HADES: A Case Study of the Development of a Signal Analysis System<sup>1</sup>

Philip E. Rubin  
Haskins Laboratory

### 1. Introduction

Research on speech and language, and the development of useful technology based on this research, depends on increasingly sophisticated tools, and most particularly on signal analysis systems. We are using ever more powerful tools to display, measure, edit and analyze time domain signals including both acoustic productions and physiological activity. The evolution of speech and signal analysis systems at Haskins Laboratories over the past 15 years will be described. The details of our systems are dependant on our particular computing environment, but many of the issues that are discussed are quite general and likely arise in some form whenever an attempt is made to create research tools for general use.

### 2. Background

The development of speech analysis systems at Haskins Laboratories has been influenced by the nature of the institution, the interests and needs of the researchers associated with it, and the available hardware. Haskins is a non-profit research institution affiliated with academic institutions in the Northeast and internationally, including Yale University, the University of Connecticut, City College of New York, Wesleyan University, Wellesley College, and the Hebrew University. The Laboratories provide a central facilities resource for research in areas related to speech and language,

---

<sup>1</sup>The development of many of the systems mentioned in this paper and the preparation of this document were supported, in part, by a variety of grants and contracts to Haskins Laboratories including NIH Grants HD-01994, MS-13870, NS-13617, and NIH Contract N01-HD-5-2910.

including speech perception, speech production, phonetics, phonology, reading, motor behavior, and cognitive and ecological psychology. The Laboratories serve a small group of full-time research and technical support staff and a larger group of part-time researchers, graduate students, undergraduates, and postdoctoral students from many institutions.

Since the mid-1970s software efforts have concentrated on the development of special-purpose tools for a VAX -11/780, and more recently the associated VAX minicomputers that are part of our local area VAX cluster. A major goal of new software development has been to exploit the power of the VAXstation workstations while ensuring that new tools are integrated with our day-to-day computing environment. In addition, our systems must provide the power, consistency, features, and speed demanded by long-term users, while still being accessible to visitors who need to get their work done quickly and with minimal support from the relatively small technical staff.

### 3. Early Developments—WENDY and SPA

The design of new display and analysis systems has been influenced by systems that were developed in the early 1970s, many of which are still in use at Haskins and other institutions. One of the most influential programs, still in daily use, is the Waveform Editor aNd DisplaY program (WENDY). The program is used to edit 12-bit sampled data (PCM)<sup>2</sup> files (see O'Shaughnessy, Chapter 2, and Gopal, Chapter 8) that conform to a local standard. While the system is most commonly used for speech signals, it is readily adapted to physiological records, model-generated data, or the

<sup>2</sup> Haskins Laboratories uses the Pulse Code Modulation (PCM) method of digitizing analog signals, which consists of taking amplitude samples at frequent, regular intervals, and representing the continuously varying signal as binary digital numbers (Whalen, *et al.*, 1990). The number of intervals per unit of time is known as the sampling rate. At Haskins, sampling rates of 10,000 and 20,000 Hz (samples per second) are used for speech data, while physiological signals are samples at a wide variety of rates, including non-integer sampling rates. The resolution within the voltages that the system can encompass (known as the dynamic range) is specified by the number of bits used for the encoding. In the case of the Haskins system, the approximate range of -10V to +10V is coded into 12 bits. The Haskins systems, like many others, avoids having a sign bit by adding a dc offset half as large at the dynamic range. For a 12 bit system, the original values (-02048 to +2047) are internally stored as values in the range of 0-4095; the Haskins systems represents these values as 16-bit two's complement numbers with the highest four bits reserved as a control field. The Haskins PCM file format saves sampled data in a binary file (64 words per record), with a 4-record header block that supplies additional information about the file (Whalen, *et al.*, 1990). Segment labels can either be stored in trailer blocks (old label format) or as separate ASCII files (new label format).

output of synthesizers. Signals can be displayed as waveforms, amplitude and duration measurements can be made, and signals can be combined or edited as needed. The program also includes a procedural language for performing iterative operations and defining macros to repeat complex or frequently used operations. WENDY was designed to run on a single central CPU (originally a PDP 11/45 and later a VAX 11/780) that supported a number of Tektronix 4010/4014 compatible display terminals.

Figure 1 shows a typical display from the program. In the top panel is the amplitude waveform of an adult male saying "The cow chewed its cud." The bottom panel is an expanded section of a portion of this signal. Amplitude waveforms are shown in "display ports"; WENDY can accommodate up to eight simultaneous display ports.

WENDY is well adapted to performing its main function: displaying and editing signal waveforms. It can handle multiple signals, and the internal procedural language provides considerable flexibility for the manipulation of these signals. The program can be run from a variety of Tektronix-compatible terminals, including Digital VT240, GraphOn, and microcomputer-based Tektronix emulators. However, the program is not easy to learn. The command set is unique to this program, and is not self-documenting, and the functions of many commands are not intuitively obvious. The LISP-like procedural language is powerful and flexible once mastered, but is sufficiently difficult to learn that many users never take full advantage of it. The storage-mode display technology results in long display times, with many total screen redispays. The program is also not well integrated with other forms of signal analysis used at the Laboratories.<sup>3</sup>

Shortly after the development of WENDY, a program called SPA (SPectral Analysis) was designed to provide FFT-based spectral analysis (see O'Shaughnessy, Chapter 2). SPA was intended to supplement ILS (Interactive Laboratory System), a commercial package from Signal Technology Incorporated of LPC-based analysis and filtering routines. SPA is a menu-driven program using the standard IEEE routines (IEEE, 1979) that is powerful, and much more accessible to the occasional user than ILS. Table 1 illustrates the SPA command menu (which is based on a Laboratory-wide applications interface shell called HUI—Haskins User Interface). Sampled

<sup>3</sup> An exception to this is WES, which is a modified version of WENDY that runs as an integrated component of our PSP Physiological Signal Processing system, allowing display, edition and measurement of waveforms from within PSP.

data files can be analyzed, displayed, and saved in a Discrete Fourier Transform (DFT) format. Graphical interactions with the spectral data are, once again, limited by the inadequacies of the storage mode displays.




Figure 1: WENDY waveform display. Top panel—the utterance “The cow chewed its cud.” Bottom panel—expanded view of the portion of the waveform that contains the word “cow.”

Table 1: Menu of SPA commands.

MENU OF SPA COMMANDS:

SPE	SPECTral analysis
DIS	DISplay spectrogram
DSP	Display SPECTral Cross-section
PRI	PRInt spectral values
AVE	AVERage magnitude spectra
RATE	decimation and interpolation (rate)
RMS	RMS signal energy
S3D	diSplay 3-Dim. spectrum
WAVE	WAVEform display
SET	SET variables, flags
SHOW	SHOW variables, flags
VAX	VAX DCL command
MACRO	MACRO definition
SUM	SUMmary of SPA
HELP	HELP
MENU	MENU of SPA commands
EXIT	EXIT SPA

Figure 2 shows a pseudo-spectrogram of the first 275 msec. of the utterance, “The cow chewed its cud.” The figure is of low quality because spectrogram displays require gray-level representation: time is on the horizontal axis, frequency on the vertical, and intensity is represented by darkness. Figure 2 represents the best that can be done on a monochrome display without resorting to simulating a grey-scale using custom dot patterns. Even such a crude display is often useful (e.g., choosing points for drawing spectral cross-section information), and when additional resolution is needed the user can always use an analog spectrograph.

#### 4. HUI—The Haskins User Interface Shell

During the mid-1970s, a wide variety of application programs were developed at the Laboratories. There were PCM file format support programs such as AFM (Arithmetic File Manipulation), CPC (viewing PCM data), INPUT (converting analog signals to PCM), and OUTPUT (PCM to analog). There were also programs to synthesize speech and other sounds, including SYN (acoustic synthesis and synthesis-by-rule), ASY (articulatory synthesis), and SWS (sinewave synthesis). PSP (Physiological Signal Processing), ACT and ACE were developed to display and analyze physiologi-




Figure 2: SPA pseudo-spectrogram.

cal time domain data. Because the user population was large and varied, and the number of application programs was growing quickly, we decided to create a simple, generic application shell that we called the Haskins User Interface, or HUI. HUI provides a flexible and consistent user interface, and gives programmers a framework and collection of tools for developing new applications. The basic HUI commands, those common to all applications, are shown in Table 2. HUI also supports a standard command-line parser to ensure commonality of command format. Except for enhancements to the MACRO function, HUI has been stable for many years.

Table 2: Menu of basic HUI commands.

#### MENU OF HUI COMMANDS:

EXIT	Exit from program
HELP	Display help for commands
MACRO	Define a macro
MENU	Display a list of commands
NEWS	Display the latest program news
SET	Set program flags
SHOW	Show program flags
VAX	Execute a DCL command

## 5. A New Approach to System Development

The design of speech and signal analysis systems at Haskins, and at many other institutions, has changed radically in recent years because of the ready availability of powerful engineering workstations and raster graphics displays. Our earlier systems were severely restricted by relatively scarce and expensive processing capacity and by the limitations of storage mode graphics. During the latter half of the 1970s we expended considerable effort in creating our own workstation, known as the Digital Pattern Playback (DPP) (Nye, *et al.*, 1975, Nye, *et al.*, 1977). This system combined custom-designed software with special-purpose hardware including a DEC GT40 display processor connected to a PDP 11/45 minicomputer, a Hewlett Packard gray-scale raster display screen, RAM accessible to both processors via a bus window, a hardware spectrum analyzer, a 40-channel speech vocoder, potentiometer input control devices, a GT40 lightpen, and a Summagraphics writing tablet.

The user interface was based on a primitive windowing system with lightpen-controlled menus. Speech could be input directly or from disk files, a scrolling waveform could be displayed on the GT40, frequency-band analysis was done by hardware, and one or two gray-scale spectrograms could be displayed on the HP screen. The user could make a variety of waveform and spectral measurements and do spectral editing with subsequent vocoder resynthesis. This system had many of the features that are common today, including using special-purpose hardware to relieve the CPU of processing-intensive tasks, and a windowing system that supported both gray-scale and vector graphics. However, the system was complex, very specialized, and difficult to maintain. The Digital Pattern Playback foreshadowed many features that have now become commonplace, and was an important step from time-sharing systems to modern distributed workstations.

The desire to modernize our approach to signal display and analysis has been spurred by a variety of factors. These include a change in the mid-80s to a distributed processing workstation computing environment (Talkin, 1989), the availability of affordable new display technologies, and the rapid development and acceptance of graphical user interfaces (GUIs such as the operating environments on the Xerox Star, the Apple Lisa and Macintosh, Suns, etc., and more recently, *Windows*, *Open Look* and *Motif*). While our newest systems are a logical continuation of what we tried to do with DPP, we have profited from such pioneering systems as enthusiasm and excitement generated by low-end solutions such as the *MacRecorder/Sound Edit* system on the Macintosh (Farallon, 1989).

A number of considerations influenced the direction of our newest systems. Of most importance, we needed the power of a VAX or a Sun workstation combined with a graphical user interface. The complexity and sheer number of applications made the consistency and ease of use of GUIs imperative, but at the time micro-computer based applications lacked the support, hardware integration, and raw processing power necessary for modern signal-processing tasks.

The need for VAX compatibility precluded the use of products from other workstation vendors such as Sun, Apollo, MASSCOMP, or Lisp Machine. We were also faced with a proliferation of new types of data (e.g., magnetic resonance images, microbeam and X-ray data, model-generated values) that made it necessary to display significantly more data channels than we previously had, and to maintain internal consistency between applications and Haskins-specific data types. Finally, since we have formal and informal ties to many institutions, it was essential to both generalize and standardize our data types to facilitate file interchanges with individuals and laboratories world-wide.

As is true for any institution, we had to work within our own limitations. Haskins is not a software company—our primary mission is research—and we generally cannot support speculative projects; nor can we afford to do system development just because “someone should do this.” However, the arrival of our first VAXstation II/GPX system in 1986 forced us to embark on some pioneering efforts, largely because the graphics windowing software system (VWS) was new to us and to DEC. While we realized that our initial efforts might not result in a completely successful system, we believed that we would gain essential knowledge and hoped to develop a set of software tools that future applications could use.

## 6. Initial Steps to a Modern Design: *SPEED*

*SPEED* was conceived as a prototype signal processing, display, and editing program that, while limited in functionality, would be easy to learn and to use. Our initial requirements for the user interface included:

- Simplified command set
- Integration of waveform and spectral functions
- Graphical User Interface (GUI)
  - Menu Bar with command-key equivalents
  - Windows for displays
  - Mouse control

Widgets: close and resize boxes, options; elevators and scroll bars, on-screen tools  
 Control Panel (toolbox with icons)  
 Dialog boxes for options  
 File finder

- Advanced Graphics features
  - Scrolling waveform displays
  - On-screen readouts
  - On-screen option selection
  - On-screen cursors for measurement
  - Graphical text editor
- Cut-and-paste editing of signal sections
- Beginnings of limited analysis and modification primitives

*SPEED* is currently in use at a number of sites around the world. The most recent version provides simple waveform editing, display, and measurement, and limited spectral analysis, in a modern GUI environment of menu bars, pop-up menus, windows, icons, and dialog boxes. The program displays data in the Haskins PCM format, and performs Discrete Fourier Transform analyses of the data. Up to four PCM files can be displayed at once, each in a “main” and an optional “scrolling” window. Basic waveform editing operations are available, including cut, copy, paste, save section, invert, rectify, set to silence, reverse, etc. Only one DFT file can be displayed, but several views are possible: the main spectrogram; up to four spectral cross-sections, providing a frequency-magnitude plot at a given time frame; and a “waterfall” spectrogram with overlapping cross-sections for a series of frames.

Figure 3 shows an example of the basic *SPEED* displays. The menu bar is at the top of the screen, with individual commands accessed using pull down menus. The menu structure is simple:

<b>SPEED</b>	Help, Quit
<b>Display</b>	Wave, Display DFT
<b>Analysis</b>	Spectral
<b>Windows</b>	Free, Aligned, Hard Copy, TIFF (Tag Image Format)

The “Display Wave” command invokes a “file finder” dialog box that the user can use to locate a file by entering a name or selecting a name from a list. Once selected, the file is opened and displayed in a waveform window that has a graphical display of the waveform and an array of numerical

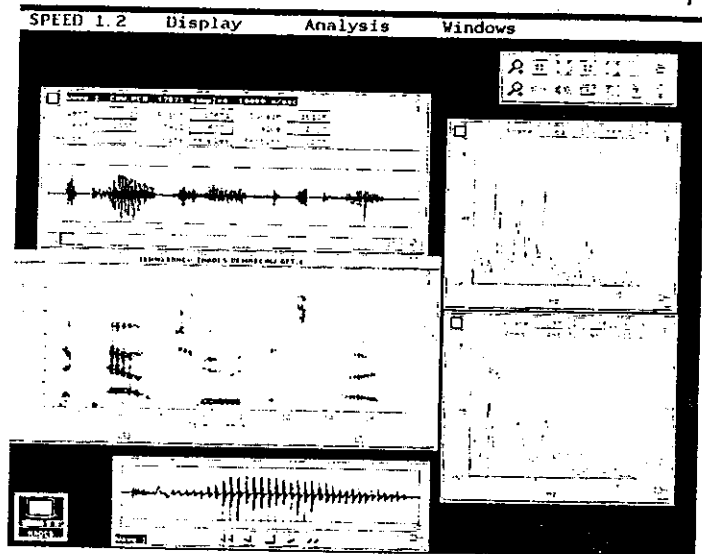


Figure 3: *SPEED* displays

values for the various display and editing parameters. *Left* and *Right* show the start and end of the waveform displayed, and *Head* and *Tail* point to a marked *Section*, which is the part of the waveform to which all editing commands are applied. *Cursor* shows the position of the graphics cursor, and *Value* shows the magnitude of the wave at the cursor position. Values can be changed by typing into fields or by graphical operations: placing the cursor by clicking on the waveform rectangle, moving the head by dragging a handle (a small triangle at the bottom of the display), and so on.

At the top right in Figure 3 is the control panel, a small window with two rows of icons, which is similar in function to the tool palette in programs like MacPaint. The controls affect the active waveform window. From left to right the icons are:

Top row: magnification increase; zoom-in display; zoom-out display; zoom-in section display; zoom-out section display; section control; open/close scrolling section window.

Bottom row: magnification decrease; play entire file; play section; cut-and-paste editing functions; special effects such as reverse, set to silence, rectification; labeling options dialog; general options dialog.

The waveform at the bottom of Figure 3 is a scrolling window that displays the section marked in the waveform window. The scrolling window has tape-recorder-like controls that provide a convenient way of tracking an enlarged portion of the master file or for making fine adjustments in section size and location.

Spectrograms (the middle window in Figure 3) are generated by Discrete Fourier Transform analysis of the sampled data files using the IEEE package of signal processing routines (IEEE, 1979). A dialog box guides the user through the selection of such analysis options as source data file, analysis type, window type and size, starting and ending portions of the waveform, and so on. Up to four spectral cross-section windows are visible at a time, two are shown on the right-hand side of Figure 3. Spectral cross-sections show frequency-magnitude information for a single frame that can be moved by dragging markers (the "houses" labelled "1" and "2" at the bottom of the Spectrogram window) or by entering a time value in the cross-section window. Although not illustrated, part or all of the spectrogram can be displayed in "Waterfall" form, a projection of a three-dimensional plot of successive cross sections.

Dialog boxes proved to be an essential tool for the *SPEED* interface, and contribute greatly both to the power and ease-of-use of the system. A judicious use of dialog boxes makes the system nearly self-documenting, provides useful default values for novice users, and imposes negligible time penalties on experienced users. We have also found it useful to include buttons that set values for a group of choices. For example, the Wide Band and Narrow Band buttons will set the Window Width, Skip Between Frames, and # Spectral Values variables to values appropriate for wide and narrow band spectral analyses.

## 7. The HADES Prototype: An Advanced Design

*SPEED* was intended to provide a low-end VAXstation speech analysis and display system that used the full power of VAX workstations. It was, however, conceived as a prototype and never intended to be a large-scale system. Our next step was to be a system that had an extensible programming language with flexible signal processing primitives, and that was fully integrated with other Haskins software tools. The Haskins Analysis Display and Experiment System (HADES) would provide a vehicle for standardizing and consolidating many of the signal processing tools (e.g., PSP, ACT, ACE) that had been developed at the Laboratory over the years. Our initial design requirements included:

- The development of a single system to handle data acquisition, in an experimental context, with subsequent display and analysis.
- Multiple waveforms in multi-panel windows, for manipulation of multi-channel data (e.g. physiological measurements, dichotic speech).
- A display primitive for multiple waveforms within a single window, using different colors or line types to differentiate the individual signals.
- The implementation of a procedural programming language within the program.
- Consolidation of a variety of signal-processing tools from other programs.
- A command window for entering text commands and/or command strings.
- Handling both existing Haskins label structures and the new Haskins label structure.
- Improvements in the quality of the spectrographic display.
- Redesign of the spectral analysis section so that:
  - a. Analysis is performed on signals or signal sections.
  - b. Analysis results can be stored in a buffer instead of in a file.
- A journaling feature for recording data readings, such as spectral peaks, duration measurements, etc.
- Spectral averaging: both non-destructive display (like cross-section) and actual data averaging.
- Specialized display primitives for frequently occurring activities and for new data types (such as images).

At the heart of HADES are three functions: acquisition, display, and analysis. The proposed final system will include acquisition of multiple streams of data in real-time, experimental set-up and interactive control of experimental conditions and peripheral devices, standardized data-conditioning and signal-processing routines, multiple display tools for qualitative data exploration, and general-purpose analysis procedures. The internal procedural language (SPIEL) supports the creation of specialized analysis macros that can be stored and manipulated as text files or executed from a

command-line entry. All data displayed by or stored in HADES are accessible to SPIEL as variables, and all operations are available as commands.

A typical HADES display is shown in Table 3. The menu bar commands are substantially more elaborate than in SPEED and in particular the Display options are greatly augmented:

WAVE	sampled data format multi-panel waveform display
OVERLAY	multiple overlaid waveform display
LISSAJOUS	plot X (horizontal axis) versus Y (vertical axis)
PHASE PLOT	phase plane: value on X versus delta value on Y
TIFF	tag image file format for raster images
SPECTROGRAM	gray-scale spectrogram
CROSS SECTION	spectral cross-section
3D	3D waterfall spectral display

Table 3: HADES 0.1 menu structure.

Hades 0.1		Data	Display	Analysis	Misc.		
Flags	•F	Acquire	•F	Wave	•W	Spectral	Show Journal
Edit	•E	Save	•E	Overlay			Save Journal
Show control panel	•I	List	•I	Lissajous			Experiment
Hardcopy		Open	•O	Phase plot			Layout
News		Play	•P	TIFF			
Help	•?	Spectrogram					
Quit	•Q	Cross section					
		3-D spectrogram					
		FBA					

The HADES control panel (bottom of Figure 4) has 21 icons and permits much greater control over multiple window displays and reduces the number of mouse movements needed to set values. The spectrographic display is much better than that provided by SPEED (compare Figure 3). Examples of some of the additional display tools can be seen in Figures 5 and 6. A multiple waveform display window; with two panels is shown at the top of Figure 5. The top panel (PXL A) is a model-generated signal showing values of lip aperture; the bottom panel (PXL H) shows lip height. The small window at the bottom left is a phase display for the PXL A data. The small window at the bottom right plots PXL A data on the horizontal axis and PXL H data on the vertical axis (a LISSAJOUS display). An

additional display type can be seen in Figure 6. At the top are waveforms representing speech and physiological records of the positions of speech articulators derived from microbeam analysis. Below the waveforms are two windows containing gray-scale images. At the right is a midsagittal MRI view of a head, showing a portion of the vocal tract. The program can display and save files in TIFF format (Davenport and Vellon, 1987). Using a public-domain format for digital images allows free interchange between Macintosh, AT-compatible, and VAXstation-based programs.

The text window at the bottom of Figure 5 is the HADES command window. This feature is a significant departure from SPEED, and was included to allow for preferred user workstyles. Different interface designs are needed for users who need easy access to program tools with minimal training and users who wish to bypass the graphical interface to get at the raw processing power of the system. The command window also lets users run non-graphics HADES commands from text only terminals.




Figure 4: HADES 0.1 displays.

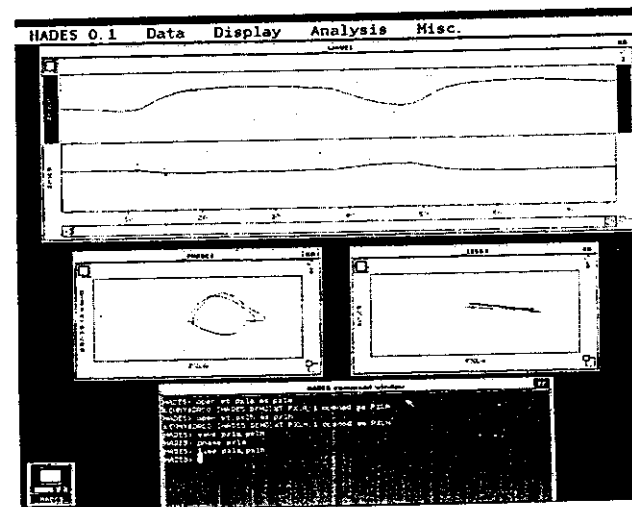


Figure 5: Additional HADES 0.1 displays.

## 8. HADES 0.8: The Beta Version

As is typically the case when moving from a conceptual prototype to a release version, many of the low level routines were redesigned to increase robustness, system performance, ease of maintenance, and extensibility. Improvements in system functionality are readily seen by looking at the menu structure summarized in Figure 7. The new structure includes hierarchical menus, EDIT and PLAY functions as menu selections with command key equivalents, and additional FLAG, DATA, and LABEL options. We also eliminated the control panel, some of its functionality has been incorporated in the main menu while other tools appear at the bottom of the temporal displays (see Figure 8). Breaking with the Macintosh tradition, all three buttons of the VAXstation mouse are used. For example, clicking on the Selection Zoom icon with the right button increases the selection size, pressing the left button decreases it, and clicking with the center button selects the entire display. The readout portion of what was the control panel has been enhanced and is shown in its new form at the bottom of Figure 8. The readout panel now contains information about signal values (minimum and maximum values) and, on occasion, analysis results.






Figure 6: HADES 0.1 multiple waveform and TIFF displays.

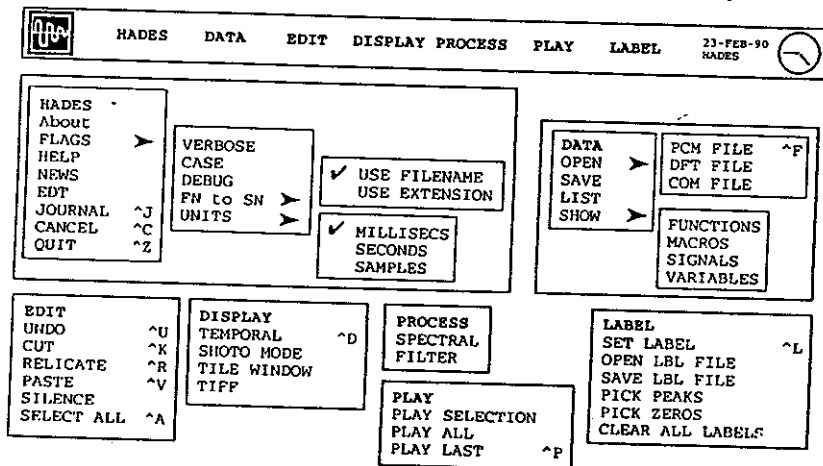


Figure 7: HADES 0.8 menu

What was formerly considered to be a waveform display window has been reconceptualized as a temporal display window. This reformulation provides for easier synchronization of data, as shown by the waveform and related spectrogram in Figure 8. Time domain waveforms have also been integrated and simplified as is shown in Figure 9, in which the two panels of WINDOW1 have been combined into a single panel in WINDOW2 by using the OVERLAY option. The VERTICAL DISPLACEMENT icon could be used to reposition the panels for comparison purposes. Many other small changes in the interface have been made in response to user testing or user requests.

## 9. Moving to Release

Few interface changes are planned as we move from Beta test to the release version of HADES. It is, however, anticipated that the signal processing portion of SPIEL will be enhanced and refinements will be made in the facilities it provides for designing and conducting experiments. Considerable care was taken to maintain a high level of modularity in the design, and many new primitives will be imported from existing programs. In future releases of HADES we expect to pay particular attention to the issues associated with file interchanges both between other Haskins applications and systems at other institutions (see Mertus, 1989, for a discussion of the issues associated with standardization of file formats). We are also evaluating the possibility of migrating our system to the DECwindows (or MOTIF) implementation of the X windows systems (Nye, 1988).

## 10. Conclusion

The evolution of the signal processing environment at Haskins is a good example of the many influences including political, scientific, academic, human, and technological that contribute to final system design. Developing a powerful but accessible signal processing system has forced us to balance what often appear to be conflicting design requirements. Our systems must always meet the needs of a diverse user population so we must constantly be aware of the subtle interplay between power and ease of use. While we try to build on the work of others, Haskins is unique and new tools must be integrated into our environment. User interfaces should simplify the user's task, but can also enrich their understanding of the problem domain and provide ever more sophisticated and powerful tools. While we need to provide continuity, speech research is a dynamic, rapidly changing field and we must provide flexible open-ended systems that gracefully respond to changing user needs. We need to be constantly aware of "real-world" constraints on our development projects




Figure 8: HADES 0.8 waveform and spectrogram

including not only resource limitations, but also changes in user workstyles, the need to maintain continuity despite changes in programming staff, and institutional decisions such as funding, changes in priorities, and delays due to competing obligations. Finally, we must contend with a need to support projects that are likely to have immediate benefits, while still reserving resources for exploratory work so we will be prepared to take advantage of technological and scientific innovations.




Figure 9: HADES 0.8 multiple waveform and display comparison.

## 11. References

- Cyphers, D. S. (1985). *Spire: A Speech Research Tool*. S. M. Thesis, Massachusetts Institute of Technology.
- Davenport, T., & Vellon, M. (1987). Tag Image File Format—Rev. 4.0. *Aldus/Microsoft Technical Memorandum*, 4-31.
- Farallon Computing, Inc. (1987). *MacRecorder User's Guide*. Berkeley, CA: Farallon Computing, Inc.
- Gayvert, R. T., Biles, J. A., Rhody, H., & Hillenbrand, J. (1989). ESPRIT: A signal processing environment with a visual programming interface. *Journal of the Acoustical Society of America*, 85(Sup1), S57.
- Hancock, B. (1989). DECwindows: X'ing with a vengeance. *DEC Professional*, 8(2), 70-84.
- Henke, W. L. (1987). MITSYN Languages: A Coherent Family of High-Level Languages for Time Signal Processing. *Language Reference Manual, V6.0*. Belmont, MA: WLH.

IEEE Acoustics, Speech, and Signal Processing Society (1979). *Programs for Digital Signal Processing*. New York: IEEE Press.

Jamieson, D. G., Nearey, T. M., & Ramji, K. (1989). CSRE: A Speech Research Environment. *Canadian Acoustics / Acoustique Canadienne*, 17(4), 23-35.

Maverick, V., & Rubin, P. (1988). *SPEED—Signal processing, editing and display*. Internal Haskins Laboratories documentation note.

Mertus, J. (1989). Standards of PCM Files. *Behavior Research Methods, Instruments, & Computers*, V21(2), 126-129.

Nye, A. (1988). *Xlib Programming Manual Volume One*. MA: O'Rielly & Associates, Inc.

Nye, P. W., Cooper, F. S., & Mermelstein, P. (1977). Interactive Experiments with a Digital Pattern Playback. *Haskins Laboratories Status Report on Speech Research, SR-49*, 86-96.

Nye, P. W., Reiss, L. J., Cooper, F. S., McGuire, R. M., Mermelstein, P., & Montlick, T. (1975). A Digital Pattern Playback for the Analysis and Manipulation of Speech Signals. *Haskins Laboratories Status Report on Speech Research, SR-44*, 95-107.

Scully, W. (1988). *User's Guide to the HUI system (Haskins User Interface, V2.0)*. Haskins Laboratories internal documentation note.

Szubowicz, L. S. (1982). *WENDY Version 1.6 Reference Manual*. Haskins Laboratories internal documentation note.

Talkin, D. (1989). Looking at Speech. *Speech Technology*, V4(4), 74-77.

Whalen, D. H., Wiley, E. R., Rubin, P., & Cooper, F. S. (1990). The Haskins Laboratories' Pulse Code Modulation (PCM) System. Haskins Laboratories internal documentation note.

Zue, V. W., Cyphers, D. S., Kassel, R. H., Kaufman, D. H., Leung, H. C., Randolph, M., Seneff, S., Unverferth, J. E. III, & Wilson, T. (1986). The Development of the MIT LISP-Machine Based Speech Research Workstation. Proceedings of the ICASSP 86,

International Conference on Acoustics, Speech, and Signal Processing, held in Tokyo, Japan, April 8-11, 329-332.

## 12. Acknowledgments

The systems described in this paper have been worked on by a large number of individuals at Haskins laboratories over a period of 15 years. There is not room to give proper credit to everyone involved with these projects, however, I would like to give special thanks to the following individuals: Leonard Szubowicz, Patrick Nye, Marion MacEachron, Mark Tiede, Vance Maverick, Michael D'Angelo, Vincent Gulisano, William Scully and Ed Wiley.

## Library of Congress Cataloging-in-Publication Data

Syrdal, A. (Ann K.)

Applied speech technology / A. Syrdal, R. Bennett, S. Greenspan.

p. cm.

Includes bibliographical references and index.

ISBN 0-8493-9456-2

I. Speech processing systems. I. Bennett, R. (Raymond W.)

II. Greenspan, S. (Steven L.) III. Title.

TK7882.S65S97 1994

006.4'54—dc20

94-26882

CIP

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

Neither this book nor any part may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage or retrieval system, without prior permission in writing from the publisher.

All rights reserved. Authorization to photocopy items for internal or personal use, or the personal or internal use of specific clients, may be granted by CRC Press, Inc., provided that \$.50 per page photocopy is paid directly to Copyright Clearance Center, 27 Congress Street, Salem, MA 01970 USA. The fee code for users of the Transactional Reporting Service is ISBN 0-8493-9456-2/95/\$0.00+\$.50. The fee is subject to change without notice. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

CRC Press, Inc.'s consent does not extend to copying for general distribution, for promotion, for creating new works, or for resale. Specific permission must be obtained in writing from CRC Press for such copying.

Direct all inquiries to CRC Press, Inc., 2000 Corporate Blvd., N.W., Boca Raton, Florida 33431.

© 1995 by CRC Press, Inc.

No claim to original U.S. Government works

International Standard Book Number 0-8493-9456-2

Library of Congress Card Number 94-26882

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

Printed on acid-free paper

Edited by  
**A. Syrdal**  
**R. Bennett**  
**S. Greenspan**

# **Applied Speech Technology**



CRC Press

Boca Raton Ann Arbor London Tokyo