
10; PAR$("CONGRATULATIONS",A$,P$)&PAR$
I need a set of LPC parameters of binary numbers fed to a group of reflection
Hi
Hallo
You talk to much, you
I can Speak English. I learned it from a Book.
to much, you worry me to death.
I generate an analog signal, then I attenuate it, filter it through a high-pass and low-pass filters at 100 Hz to 3000 Hz and amplify it.

The new HP speech synthesis module for Series 80 was designed for both the serious software writer and the hobbyist.

SERIES 80 TALKS BACK

BY LOREN KOEHLER

Why would one want a talking computer? We have cars that remind you when you've left the key in the ignition and phones that alert you when your receiver is off the hook. So a computer that talks is logical in view of the growth of personal computer use. And talking is probably the most natural means of communication.

It's been said that "user friendly" really means idiot proof. But computer users are ordinary people—people who make mistakes, need instructions, want to be warned when the machine's about to break down or when the job's done, and who need praise. What better way to do these things and improve the quality of interaction between people and their computers than with a human voice. That's what speech synthesis can do.

We'll talk about a software/hardware approach to speech synthesis using the new HP 82967A Speech Synthesis Module for HP Series 80 personal computers. This module was designed to meet the needs of both the serious software writer and the hobbyist. Enough software and documentation is available for a wide variety of speech synthesis tasks.

Two popular approaches to speech synthesis are 1) phonetic synthesis, and 2) analysis/synthesis. The HP 82967A makes use of an analysis/synthesis technique called linear predictive coding (LPC). We'll discuss LPC analysis/synthesis here.

The phonetic approach allows your computer to verbalize the full set of words in a given language. LPC-type speech can produce very high quality English and foreign-language speech. But what's really needed is a speech synthesis approach that has the word generation flexibility of phonetic text-to-speech systems without sacrificing the high quality and multiple-language capability of LPC. And at a reasonable price.

ILLUSTRATION BY ALAN COBER

Generating a word takes just a few minutes thanks to a powerful and friendly user interface.

The quality of voice output has a direct bearing on how it can be used. Quality, in this discussion, will be divided into two categories: intelligibility and how closely the synthetic speech sounds like the human voice.

Intelligibility means how easy it is to distinguish what is being said. Some speech systems are very intelligible, but require great motivation from the user to listen to the speech for any length of time. This isn't surprising because it's sometimes tiring to listen to other people, especially if they talk in a monotone. The second quality, human-likeness, is very subjective—what sounds pleasing to you might sound pallid to your neighbor. The important point is that as synthetic speech sounds more like the real thing, the motivation required to listen to it becomes lower.

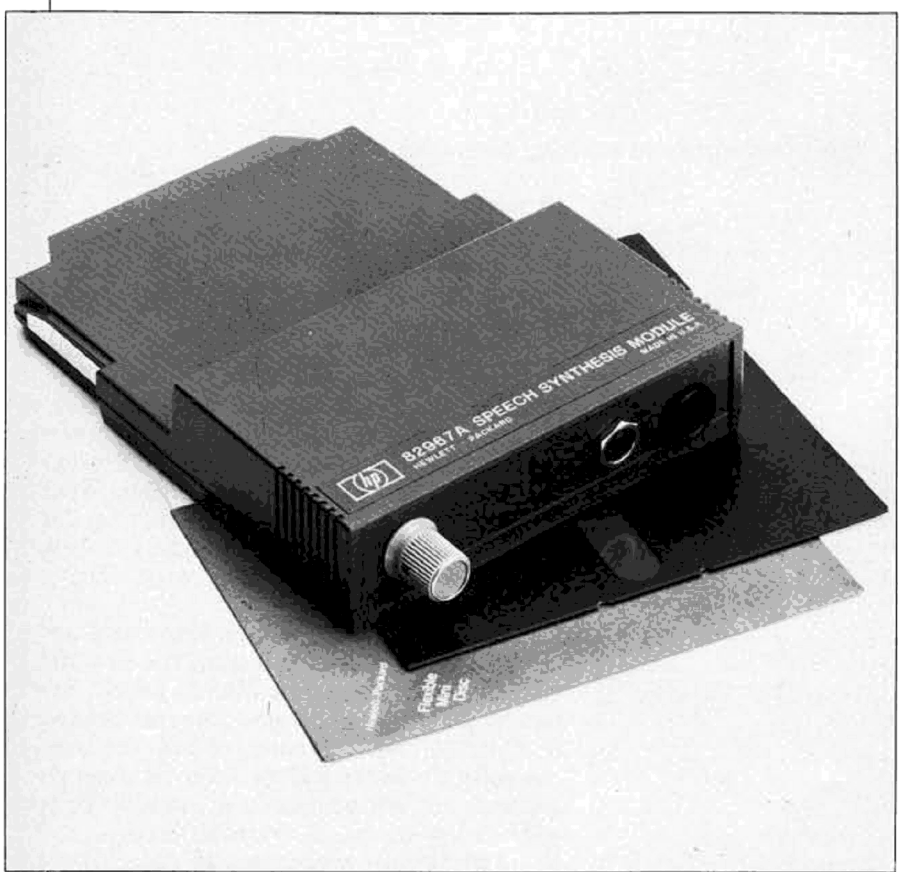

A key aspect of the use of speech synthesis is how well the speech output hardware and software are integrated with the system. Is the vocabulary sufficient for your needs? How easy is it to incorporate voice output into your programming?

A talking computer can be a novelty for a while, but using speech as a computer tool requires that the speech execution be incorporated into the high-level language of the computer in such a way that it will work well with the other I/O devices in real time and be easy to use.

Speech synthesis can be based on either a digitized waveform of the original speech or a parametric model representing the human vocal tract. Waveform representation may sound very good but it has the inherent disadvantage of operation at a high data rate—requiring more memory or producing few words. Parametric modeling requires less memory but involves more analysis to derive words.

PHONETIC SYNTHESIS

The speech hardware used with phonetic software ranges from a formant synthesizer implemented on a digital signal processing chip to a library of digitized phonetic sounds and a digital-to-analog converter. A phoneme synthesizer can be implemented entirely in the hardware—a good example is the Votrax speech chip. Phonetic synthesis can also be produced by an LPC synthesizer and a set of phoneme-length segments derived from actual speech. Both formant and LPC-type synthesizers are based on a parametric model of the vocal tract.

Phonetic synthesis uses a set of pre-defined phonetic sounds. Words are formed by concatenating the phonetic sounds according to a set of rules. The phonetic synthesis approach has two main features: 1) a low data rate—approximately 3 to 5 bytes per phoneme and up to 100 bytes per second to construct a word (most people verbalize one word in two-thirds of a second) and 2) verbalization of the words in a given language either through the use of phonetic spelling or a text-to-speech program using the given set of sounds and rules. The disadvantage of phonetic synthesis is that many of the qualities of the human voice we take for granted—emphasis, intonation, pitch—may be missing.

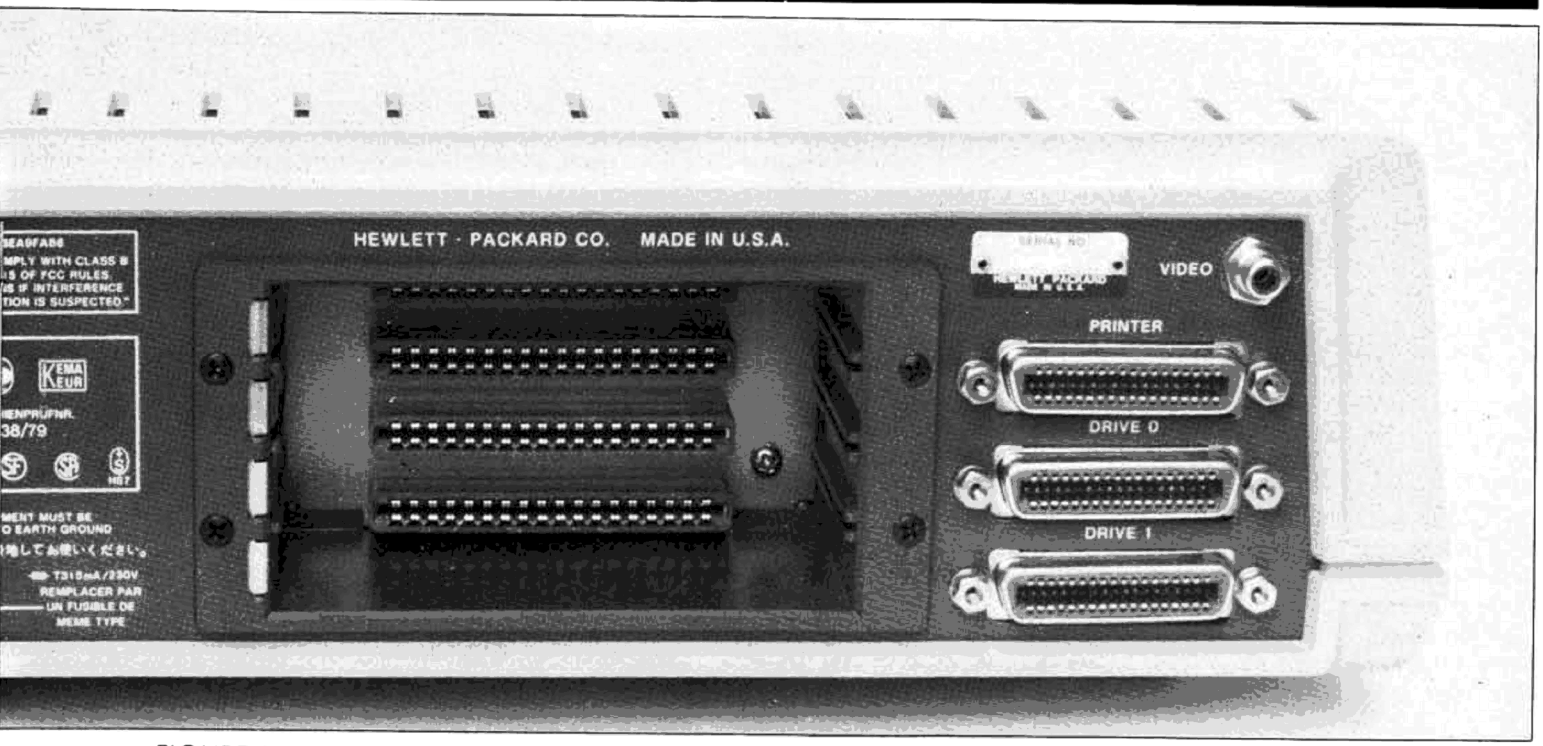

FIGURE 1. **Back view of HP Series 80 personal computer showing I/O ports where speech synthesis module is inserted.**

ANALYSIS/SYNTHESIS

Using this technique, a recording is made of a person speaking an utterance. An utterance is a human sound, a word, or multiple words spoken together, each one of which is analyzed. The recorded utterance is mathematically analyzed using the LPC technique. This yields information about the word in a condensed form, reducing one second of speech from 8,000 bytes to approximately 150 to 250 bytes. To synthesize the analyzed word, an appropriate speech chip must be used—for example, the Texas Instruments TMS 5220. The analysis/synthesis approach is important because it can produce speech of high quality and will work with multiple languages without altering the speech chip or analysis software.

Text-to-speech is a piece of software that works in conjunction with a speech synthesizer to convert a stream of ASCII words into synthesized words. Usually, text-to-speech is only associated with a phonetic synthesizer because of the large number of words required to make text-to-speech viable. Text-to-speech programs range in size from a few thousand bytes to more than 100,000 bytes, depending on the sophistication of the program and the phonetic rules set for generation and/or concatenation of the phonetic sounds.

HP 82967A SPEECH SYNTHESIS MODULE

The speech synthesis software/hardware approach we'll talk about is based on a new speech card for HP Series 80 personal computers.

The HP 82967A has four key elements:

1. Hardware
2. Speech Output
3. Vocabulary
4. Software
   a. Editors and binary programs
   b. Utilities

The new module is an LPC-type of speech synthesis system, based on the Texas Instruments TMS 5220 speech chip. Software plays a major role in making the module useful—even with the most elegant hardware in the world, a computer without the necessary software is useless.

HARDWARE

The HP 82967A is an I/O module that plugs directly into the back of any HP Series 80 computer (see Figure 1). The module is fully enclosed in a plastic case for easy attachment to the computer.

When the computer wants a word to be verbalized, it outputs the LPC parameters associated with the word to a 90-byte FIFO contained in the RAM of an 8049 microprocessor on the speech card. The parameters reside in the FIFO until required by the TMS 5220 speech chip. The chip takes the speech parameters and generates an analog signal representing the desired word. The analog signal is in succession attenuated, filtered and amplified. (Attenuation is the process of reducing the analog signal level, allowing programmable software volume control of the speech output. Filtering to remove unwanted noise is done with high-pass and low-pass filters rolling off at 100 Hz and 3000 Hz, respectively.)

The speech card, with its manual volume control knob at the back, was designed to take advantage of the key attributes of the computer—memory, programmability, mass storage —without compromising quality, in order to keep the cost down.

There are two 8-ohm speaker jacks on the back of the card—an RCA-type jack and a ¼-inch phono jack. The output jacks may be used with a speaker or headphones but not both at the same time. Using the RCA jack you can easily connect the speech card to the internal speaker of the HP 86 monitor (see Figure 1). With the volume control turned down all the way, the volume level is set

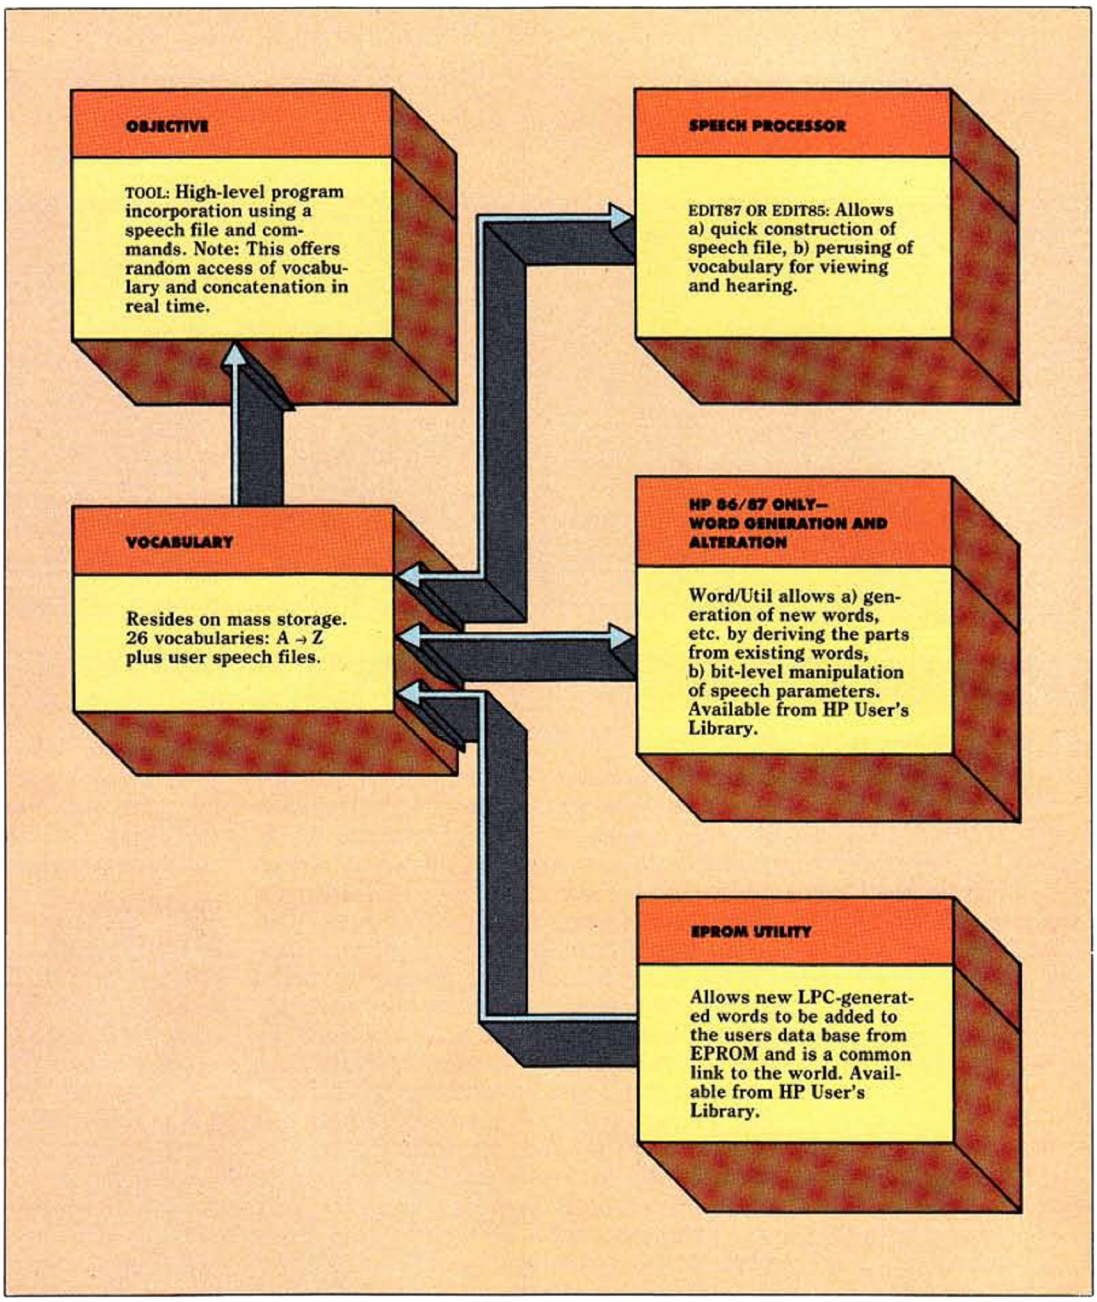

FIGURE 2. **Block diagram of separate parts of software process.**

correctly for use with headphones.

As we've said, the speech module uses an LPC-type speech chip. When speaking a word, the speech chip requires a set of LPC parameters (binary numbers) to be fed to it every 25 milliseconds until the word is completed. (The LPC parameters are digital numbers that represent the characteristics of a word.) The parameters represent 1) energy, 2) pitch, 3) a repeat bit that tells the speech chip whether to reuse parameters it already has, and 4) a group of parameters called reflection coefficients that change the characteristics of the speech filter in the synthesizer over time.

It's worth noting that these parameters are, in most cases, mathematically derived from some original audio recording of the word or sound using an LPC-type analysis, the purpose of which is to reduce the amount of information required to represent the word or sound. This is the source of both the advantage and the disadvantage of LPC-type speech for end users: The advantage is that a good LPC analysis can produce splendid results. The disadvantage is the high cost of the hardware and software to do a good LPC analysis.

When whole words and/or phrases are produced with an LPC analysis, many of the important qualities of the original human utterance such as pitch, emphasis, and accent are preserved.

Next, we'll talk about the goal of the HP 82967A speech synthesizer software/hardware and work back to its separate processes (see Figure 2).

SPEECH OUTPUT AS A TOOL

Making speech on the personal computer a meaningful tool is accomplished by allowing easy incorporation of speech into the user's high-level language program while not sacrificing speech quality. Since speech synthesis on a computer is not an end in itself, but an enhancement, it requires very little additional effort to write the necessary program.

Three things are needed to incorporate speech into an HP BASIC language program:

Talking is probably the most natural means of communication, so a computer that talks is logical considering the growth of personal computer use.

1) a command to execute or speak the desired utterance or phrase; 2) a speech file that can be placed in the memory of the computer, containing the words, sentences or sounds you want to verbalize; and 3) a command that will select the parameters for the desired utterance or phrase from the speech file so that it in turn can be spoken.

In the programming example, a speech file was loaded from mass storage under the file name "WORDS." A speech file is a data file of vocabulary used in an application, whether the application is a warning system or a spelling-bee program. In essence it is a subset of your vocabulary data base. Using a speech file allows a minimum of computer memory to be used and reduces the transfer time required to move the data file from mass storage to RAM.

A speech file may consist of one phrase or a thousand phrases. A phrase can be whatever you want—a word, sentence or sound. It can even be the preamble to the Constitution.

The file may reside on a floppy disk, cassette tape (for the HP 85) or in EPROM, using an HP EPROM module.

VOCABULARY

The HP 82967A is shipped with a vocabulary on a disk of approximately 1,500 utterances composed of English words, variations and sounds. All the vocabulary, which has a common male voice, is broken into 26 vocabulary data files, A through Z, with all the A words residing in the A vocabulary data file, etc. The average utterance in the vocabulary requires 95 bytes of data. However, some utterances are much larger—a bugle call is the largest at 1,469 bytes.

Variations of some words are included to aid in sentence construction since people pronounce words differently and at different rates, depending on the preceding and following words in a sentence and what the speaker really wants to communicate. To denote variations of a word in the data base, each variation is followed by a number—for example, warning, warning1, warning2.

Prefix and suffix sounds are denoted by '—for example, the plural sound sss is denoted sss'.

Musical sounds are denoted by a right-hand parenthesis—for example, bugle call), whistle).

Speech data files have a unique structure to facilitate the handling of the words by a personal computer. Although the user doesn't need to understand the data file structure to utilize speech, the file structure is described for completeness.

When a speech data file is generated, the first ASCII string stored in the data file contains information regarding all the words in the data file; this string is called the name string. The information for each phrase includes the name length, the number of concatenations in the phrase, the relative location in the data file of the speech data with respect to the other

The following example demonstrates programming with speech in HP Basic:

```
10 Dim A$[1000],P$[2000]
20 LOADBIN "SPKB87" ! load a binary program that contains the
30                  ! speech commands
40 !
50 ! next we will load a speech file
60 !
70 ASSIGN#1 TO "WORDS" ! WORDS - name of speech file
80 DLOAD#1;A$,P$       ! ASCII goes into A$ & data into P$
90 ASSIGN#1 TO *       ! close the mass storage data file
100 !
110 ! Let's set the volume from the keyboard
120  SVOL 10;15 ! 10 is the speech card address & 15 is the
                  maximum volume setting
130 !
140 ! now we are ready to pick 2 phrases and speak them
150 ! sequentially
160 SPEAK 10;PAR$("CONGRATULATIONS",A$,P$)&PAR$("ACE",A$,P$)
170 !
180 ! or you could do it as follows
190 !
200 S$="CONGRATULATIONS"
210 Q$="ACE"
220 R$=PAR$(S$,A$,P$)&PAR$(Q$,A$,P$)  !(locate the required
221 !                                   speech data using an
222 !                                   ASCII name in a string
223 !                                   variable)
230 SPEAK 10;R$ ! speak the words - 10 is the address the speech
240 !                              card resides at
250 !
260 ! CONTINUE WITH YOUR PROGRAM
```

When whole words and phrases are produced with an LPC analysis, many of the important qualities of the original human utterance are preserved.

phrases, a search tag and the phrase name. The first two bytes of each set of speech data parameters give the length in bytes of that set of parameters (see Figure 3).

The vocabulary disk is initialized with 25-byte record boundaries to minimize wasted space on the disk. After the ASCII name string is placed on the disk, the speech data for each phrase in the data file is stored out as a string, always starting on the next available record boundary. When speech data for a phrase is wanted from the disk, the ASCII name string is searched to find the placement of the speech data for the phrase within the data file.

The routines necessary to store and read from the speech data files are contained in the EDIT87 and EDIT85 programs shipped with the speech card. The routines are written in BASIC. When the EDIT87 or EDIT85 programs are being run, the program may be paused and A$ (ASCII information string) or P$ (speech data) dumped on the screen for observation.

SOFTWARE

A speech editor called EDIT for both the HP 86/87 and the HP 85/83 is shipped on the vocabulary disk (see Figure 4). The HP 86/87 editor is called EDIT87; the HP 83/85 editor is EDIT85.

The first thing a person likes to do after buying a speech card is hear it talk. Do this by loading EDIT87 or EDIT85 and pushing the run key. You can then make the computer talk to your heart's content.

EDIT is a program that allows you to peruse the vocabulary data base, selectively viewing and listening. It permits you to compose a data file, referred to as a speech file, consisting of the sentences, words and sounds you want to use in your applications program. EDIT assists a user in composing a speech file in much the same way as a word processor assists in writing text. Therefore, it's useful to think of EDIT as a speech processor.

Here's the softkey menu for the major functions of EDIT.

VIEW—allows viewing of names in a data file or all data files.

LISTEN—allows the selection of a word or words concatenated into a sentence for listening. For example,

HELLO_ACE_THIS_IS_COMPUTER_
SPEECH SYNTHESIS.

The underline is typed in after each desired phrase. The name SPEECH SYNTHESIS is one phrase. To add pauses between the phrases, add more underlines. You can use LISTEN to hear vocabulary in your data base on mass storage and EDIT for vocabulary in RAM.

GET—allows you to get a whole vocabulary file and place it in the editor for alteration, addition, etc.

INCLUDE—allows you to pull selected phrases—whether they are words, sentences or sounds—into the editor from the mass storage device.

SAVE—allows you to store all words, sentences and sounds in the editor on a mass storage device. The saved file is your speech file, used in your applications program.

DELETE—allows deletion of phrases from the editor.

RENAME—allows changing of the name(s) associated with the words, sounds and phrases. For example, "HELLO HEWLETT PACKARD" could be renamed "GREETING."

There are various techniques for phrase construction.

Word Selection: Note that many words in the data base have existing variations in pronunciation. Using the correct variation of each word in phrase construction is the most important consideration in producing good quality speech. The additional time required to listen to each variation until the correct fit is selected is a good investment. Listening time is minimized by adding the variations of the words into memory from the vocabulary disk and listening to the concatenated words from RAM instead of the disk.

Renaming: Names associated with words can be numeric as well as alphabetic. This is sometimes useful for incorporation of speech into a basic program. A simple example is verbalizing numbers:

Name of word: 60
Synthetic word: "sixty"

File Order and Sort: Sometimes it is desirable to place the existing sequence of phrases in a speech data file in alphabetic order. The EDIT programs will do a sort for you as an option—when you do a GET of a whole data file from disk—by typing in the desired "dictionary name" followed by ",S".

Example: "WORD,S"
where WORD is the file name.

Vocabulary Listing: Nobody wants to waste time trying to figure out how much vocabulary is available in a data base. The following resources eliminate guesswork:

1. The manual has a listing of every word and its byte count.
2. The manual also contains a tearout

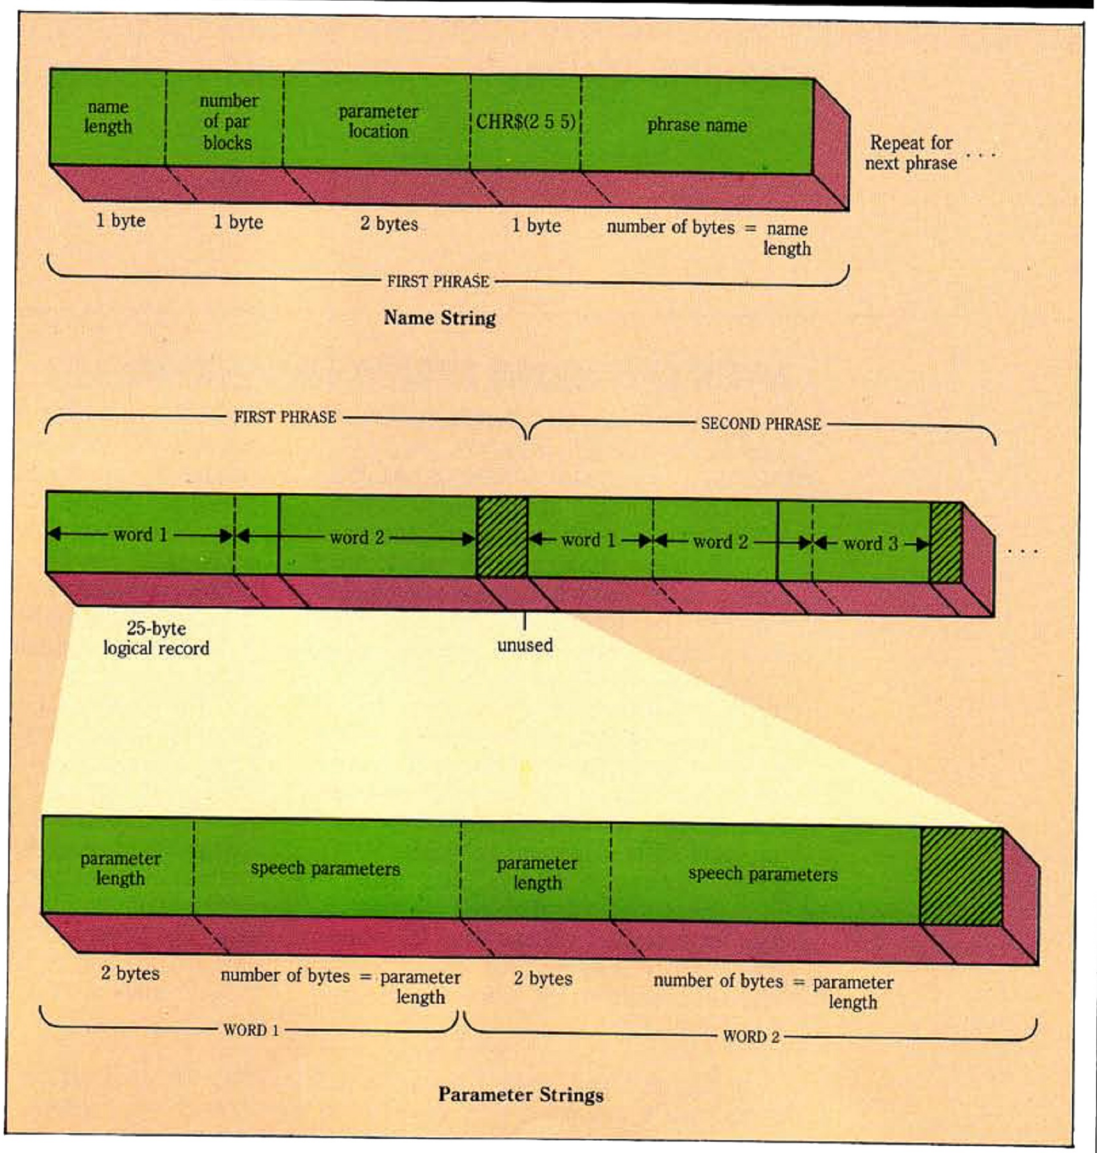

FIGURE 3. **Structure of a speech parameter string.**

card listing all 1,530 words in the data base provided (see speech vocabulary table).

3. Using the EDIT87 or EDIT85 programs you may view the existing names of any or all speech data files. You can also print out the listing that appears on the screen.

4. If all else fails, the EDIT program will tell you that the non-existent word you are trying to use is missing.

Speech synthesis is useful as long as the words you want are in the vocabulary data base. But what if you want to say "Thomas Edison" or "California" and they're not in the vocabulary?

Word/Util is a program that allows you to create words and sounds you need but don't have. It was written for use on the HP 86 and HP 87 only. The program is softkey-driven, with binary programs totaling over 28K bytes and several BASIC subprograms.

When you want to create a word, you extract the parts of the word from the existing vocabulary data base of words and patch them together. The resulting word will probably meet your expectations for quality. The process of generating a word takes just a few minutes in many cases, thanks to a friendly and very powerful user interface. To quickly locate the words you want to use, Word/Util will dump any existing vocabulary data file catalog on the screen in an orderly fashion.

It's worth noting about the process of word generation that the original vocabulary data base is large enough to yield the required sounds, so that almost any word you can think of can be generated. Since the entire data base has only one speaker, the quality of the newly generated word is much better than if the parts of the word had been derived from several different speakers.

Word generation by patching groups of sounds together is not unique. The phonetic synthesis approach of some popular speech systems on the market allows you to misspell a word or spell it phonetically to derive the correct pronunciation. What is really occurring is the concatenation of phonetic sounds to generate the desired word.

There are three distinctions between Word/Util and strict phonetic word generation systems:

1. Word/Util allows the user generating a word to select the parts of the new word in as complete a form as possible. Instead of thinking in terms of phonetic parts, much of the time the user can think in terms of syllables.

2. The ability to work at the syllable level when constructing a new word can be faster than at the phonetic level.

3. The results sound better.

BASIC LANGUAGE SPEECH COMMANDS

The following speech commands were used in the example program:

DLOAD example DLOAD#1;A$,P$

This command allows you to bring in a speech data file from mass storage. DLOAD takes all the speech data for each phrase in the speech data file and places it in P$. The information contained in the ASCII name string—word location, name length, etc.—is placed in A$.

SVOL example SVOL 10;15

This command allows the BASIC language program to adjust the speech volume to any one of sixteen possible levels ranging from 0 to 15, where 0 is no volume and 15 is maximum volume.

PAR$ example PAR$("ACE",A$,P$)

This command retrieves the speech data for a desired phrase from P$ when given an ASCII name (line 160) or a string variable containing the same ASCII name (line 220), as long as the phrase is in your speech file. Remember, P$ contains all the speech data for all the phrases in your speech file and A$ contains the ASCII name information. Par$ uses A$ as a lookup table into P$.

SPEAK example SPEAK 10;R$

This command feeds the speech data to the speech card and tells the card to invoke speech.

The number 10 is the I/O address given to the speech card as a factory setting. If this address needs to be changed, it can be done by flipping a switch on the card.

These functions could have been accomplished using BASIC language subroutines and existing commands, with the disadvantage of a less friendly user interface. The speech commands are part of the SPKB87 and SPKB85 binaries used on the HP 86/87 and HP 85/83, respectively.

If the resulting word doesn't have the characteristics you want, or if you want to do something really exotic, Word/Util will allow you to quickly manipulate, add or delete LPC parameters the utterance is composed of.

Word/Util has been called an "experimenter's dream" by its writers. It's available from the HP personal computer user's library at a nominal cost.

Speech parameters generated for the TMS 5220 speech chip should look identical for any unique utterance whether the speech chip is in a Hewlett Packard product or a Texas Instruments product. However, the way the speech parameters are handled and stored in their respective systems will most likely be entirely different.

Custom LPC-analysis word generation is available from different design support organizations such as Texas Instruments Regional Technology Centers around the world. Equipment that will allow users to do custom LPC-analysis word generation is slowly becoming available.

The most common medium to send and receive speech parameters on is an EPROM, unless you are dealing with hundreds of words. In the case of the HP 82967A speech synthesis module discussed here, receiving an EPROM containing the speech parameters is not enough. The speech parameters must be put into the correct data file structure.

To establish a friendly link between the data file structure and the EPROM containing speech parameters, an EPROM utility was written. The utility is easy to use. All you need is an HP personal computer system, the EPROM and an EPROM module made for the HP Series 80 personal computer line.

The user inserts the EPROM into the card and the card into the back of the computer, loads the EPROM utility software, then supplies a little information to the program (when prompted) about the names of the words in the EPROM and the starting addresses within the EPROM. This information should be supplied with the EPROM containing the words. The new words are placed in mass storage by the utility and are ready to be used.

The most common reasons for generating new LPC words are 1) for a foreign language; 2) for a different voice—for example, a female voice—or a special sound such as the chimes of a grandfather clock; and 3) for multiple words such as "the weather for today." (The speech resulting from multiple-word generation sounds like a recording.)

The HP 82967A was designed as a tool with both local and remote applications. Like any tool, getting the most from it requires time and attention; however, the results more than justify the user's efforts.

Loren Koehler is an R&D engineer at HP's Portable Computer Division, Corvallis, OR.

Word/Util is a program for the HP 86 and 87 that lets you create words and sounds you need but don't have in the vocabulary data base.

SPEECH VOCABULARY

The tear-out card in the manual contains a list of the vocabulary supplied with the HP 82967A Speech Synthesis Module (partial list shown here). All the words must be entered exactly as shown. Sound effects are denoted by words ending in a closing parenthesis); prefixes and suffixes end in an apostrophe '. Three of the words generate pauses: PAUSE1 (.1 second), PAUSE2 (.2 second) and PAUSE10 (1 second).

A
A1
A2
A3
A.M.
A.M.1.
ABEAM
ABORT
ABORT1
ABOUT
ABOVE
AC
ACCELERATED
ACCEPT
ACCEPT1
ACCEPTED
ACCEPTED1
ACCESS
ACCESSING
ACCOUNT
ACCOUNT1
ACE
ACES
ACKNOWLEDGE
ACTION
ADD
ADDRESS
ADJUST
ADSORBER
ADVISE
AERIAL
AFFIRMATIVE
AFTER
AFTERNOON
AGAIN
AGAIN1
AHEAD
AIR
AIR1
AIR BRAKES
AIRCRAFT
AIRPORT
AIRSPEED
ALARM
ASSOCIATED
ASTERISK
AT
AT1
ATIS
ATTACK
ATTENTION
ATTENTION1
AUGUST
AUTHORIZATION
AUTHORIZATION1
AUTO
AUTOMATIC
AUTOMATIC1
AUTOMATICALLY
AUXILIARY
AVAILABLE
AVAILABLE1
B
B1
BACK
BAD
BALANCE
BALL
BALL BEING CAUGHT)
BAND
BANDIT
BANK
BANK1
BASE
BASE1
BASKET
BAT HITTING BALL)
BATTLE
BEEN
BEFORE
ALERT
ALERT1
ALERT2
ALERT3
ALIEN
ALL
ALL1
ALL2
ALOFT
ALPHA
ALTERNATE
ALTIMETER
ALTITUDE
AM
AMERICAN EXPRESS
AMOUNT
AMOUNT1
AMPS
AN
AN1
AND
AND1
AND2
AND3
ANSWER
ANSWERED
ANTE UP
APPOINTMENT
APPROACH
APPROACHES
APPROVED
APRIL
APU
ARE
AREA
AREA1
ARRIVAL
AS
AS1
ASKED
BEGIN
BEHIND
BEING
BELL
BELOW
BET
BETWEEN
BID
BIG
BIRTH
BLIND
BLOCKED
BLOWING
BONUS
BOOKLET
BOOST
BOTH
BOTTOM
BOX
BRAKE
BRAKING
BRAVO
BREAK
BROKEN
BUFFER
BUGLE CALL)
BUILD
BUS
BUSINESS
BUSINESS1
BUSY
BUSY1
BUT
BUTTON
BUTTON1
BUTTON2
BY
C
C1
C2
CABIN
CALENDAR
CALIBRATE
CALL
CALL1
CALL2
CALL3
CALL4
CALL5
CALL6
CALL7
CALL8
CALL9
CALLING
CALM
CAN
CANCEL
CANCEL1
CANCEL2
CANCELLED
CANNOT
CAPACITANCE
CAR
CARD
CARD NUMBER
CARDHOLDER
CARDS
CARTE BLANCHE
CASING
CAUTION
CAUTION1
CEILING
CELSIUS
CENT
CENT1
CENTER
CENTER1
CENTI
CENTS
CENTS1
CHANGE
CHANGED
CHARLIE
CHARLIE1
CHECK
CHECK1
CHECK2
CHIME1)
CHIME2)
CHOICE
CIRCUIT
CLEAN
CLEAR
CLEAR1
CLERANCE
CLEARANCE DELIVERY
CLIMB
CLOCK
CLOCK1
CLOSE
CLOSED
CLOSED1
CODE
CODE1
COIN
COIN1
COLD
COLLECT
COLLECT1
COME
COMMAND
COMPANY
COMPLETE
COMPLETE1
COMPONENT
COMPUTER
CONDUCTIVITY
CONFERENCE
CONGRATULATIONS
CONGRATULATIONS1
CONNECT
CONTACT
CONTAMINATED
CONTINUE
CONTROL
CONTROL 1
CONVERGING
COPY
CORRECT
COUNT
COURSE
COURT
COWL
CRANE
CRANE1
CREASE
CREDIT
CROSS
CROSSWIND
CROWD NOISE)
CRYSTALS
CURRENT
CUSTOMER
CYCLE
CYCLE1
CYLINDER