on the human ear and brain to fill in gaps. For example, the range of frequencies that can be transmitted over a telephone line gives only one-fifth the quality we would expect from a reasonable hi-fi system, yet the speech we hear through the ear-piece is perfectly understandable. This is because our brain fills in the gaps.

The second method of synthesis, called 'digitised speech', uses the same phenomenon. With the reduction in cost of computer memory it is now possible to convert speech into digital information by means of an analogue-to-digital converter. The resulting data is then compressed many hundreds of times and stored in a ROM — thereby creating the gaps which your ear can compensate for.

To cause any of the stored words to be spoken we simply give the computer the address of that word in the ROM, and the digital information is recovered and converted back into sound. Because the original speaker's words are stored, the personal characteristics remain. Acorn's speech chips for the BBC Micro, for example, can be clearly identified as the voice of newscaster Kenneth Kendall.

Some computers, notably the Sirius 1, feature built in hardware and disk-based software to allow the user to digitise his own voice using a microphone. The resulting data is stored on disk — one second of speech occupies about one Kbyte — to be recalled from an applications program as verbal messages and warnings.

The uses for speech synthesisers are so many and varied that it is almost impossible to list them. To start with, speech synthesis can replace taped announcements at railway stations, airports and other terminals. In the USA it is widely used on the telephone system to inform callers of wrongly dialled numbers, engaged numbers or withdrawn services. Many automated ordering systems now feature speech response. An order number is keyed in to a computer, which speaks the description as a double-check. The computer can also inform the customer of the current stock level or the likely waiting period so that the order may be modified at the time it is placed.

Speech synthesis units are now incorporated into cars — the BL Maestro, for example — as part of the standard instrumentation. More than a mere sales ploy, the synthesiser provides warnings that the driver can hear and act on without having to take his eyes off the road.

In the home computer and electronic games market speech synthesis is used to enhance games: scores are called out and warnings of enemy attack can be given verbally, leaving the player free to concentrate on the tactics of the game rather than having to consult messages printed at the bottom of the screen.

Finally, there are educational devices such as Texas Instruments' Speak'n'Spell, which recites a word that must then be spelt correctly, and foreign language dictionaries that speak the words as they are displayed.

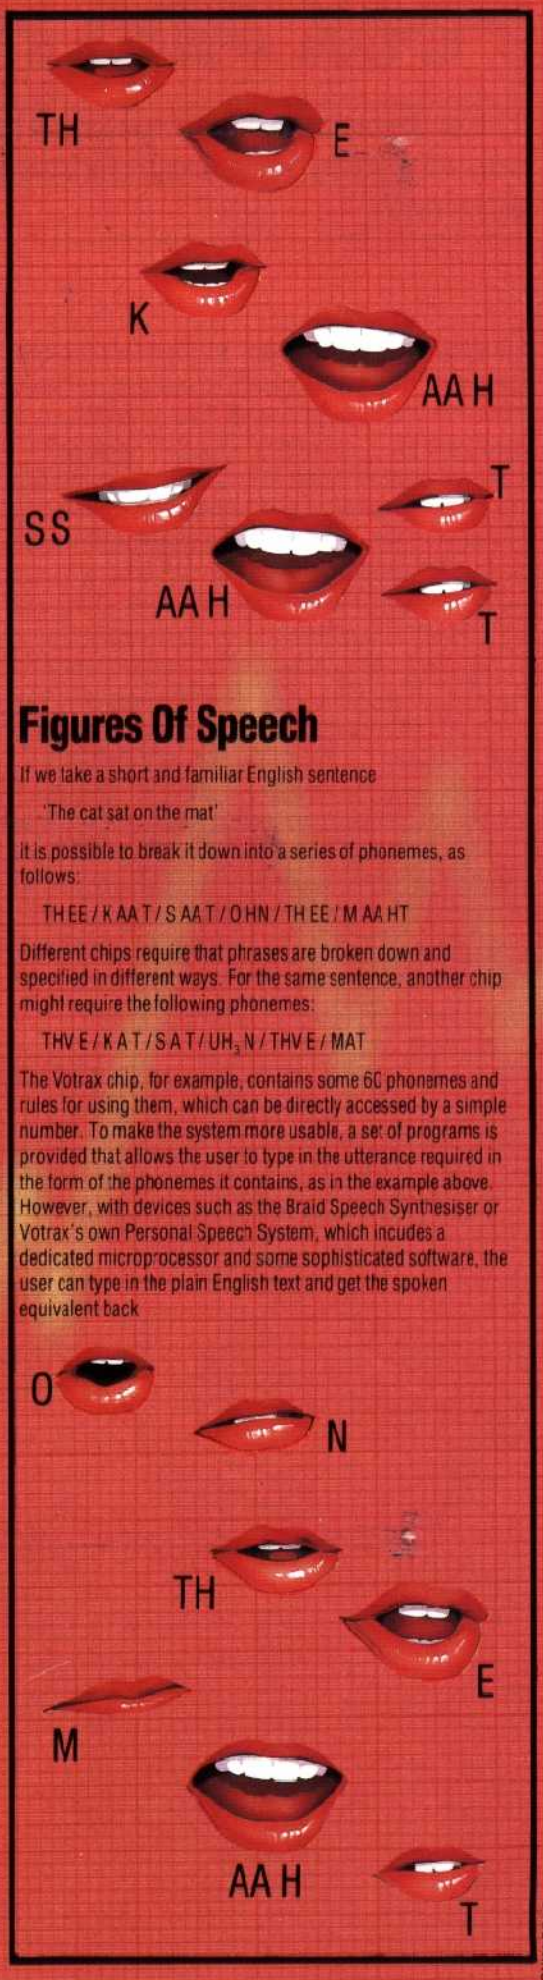

Figures Of Speech

If we take a short and familiar English sentence

'The cat sat on the mat'

it is possible to break it down into a series of phonemes, as follows:

THEE / K AA T / S AA T / O HN / THEE / M AA HT

Different chips require that phrases are broken down and specified in different ways. For the same sentence, another chip might require the following phonemes:

THV E / K A T / S A T / UH$_3$ N / THV E / MAT

The Votrax chip, for example, contains some 60 phonemes and rules for using them, which can be directly accessed by a simple number. To make the system more usable, a set of programs is provided that allows the user to type in the utterance required in the form of the phonemes it contains, as in the example above. However, with devices such as the Braid Speech Synthesiser or Votrax's own Personal Speech System, which incudes a dedicated microprocessor and some sophisticated software, the user can type in the plain English text and get the spoken equivalent back