Versatile speech synthesiser

This speech synthesiser, originally designed for the Tasman Turtle robot (see ETI, April/May/June '82), is an extremely versatile unit with functions and features offered in no other speech synthesiser system. Applications range from a talking doorbell, to a communications aid for the handicapped, from a self-learning tool to an industrial equipment voice indicator. Or whatever you dream up.

Allan Branch
Flexible Systems, Hobart, Tasmania

THERE ARE MANY speech synthesisers about now. In the last few years a number of companies have produced specialised integrated circuits or IC sets designed to produce speech, electronically, in various ways. The result is a maze of talking electronics — clocks, cars, ovens, calculators, alarms and learning aids to name a few. Just what led up to this modern electronic wonder?

For some time, psychologists and artificial intelligence researchers have been developing a careful understanding of the human speech communication mechanism. The ultimate aim, as could be expected, is to develop machines that can listen, understand, answer back or ask questions and then carry out an appropriate task.

The problem has always been the *enormous* amount of data or information stored in even very short utterances. Electronic memory has not yet reached the stage where unlimited 'words' can be stored and recalled at will.

That is not to say that the data cannot be recorded at all. Mankind has been coding verbal speech in various ways and storing the coded information, so that it is available for decoding relatively easily, for thousands of years! Some of these techniques are not usually realised for what they are since they are a part of every day life and a major activity in communication.

The types of encoding and decoding used so that the speech data can be stored varies in these old techniques. The most common is the use of written symbols such as hieroglyphics, alphabet characters, shorthand symbols or pictures. Speech synthesis in this case goes under the common name of 'reading'. The trouble is that man is still the instrument needed to do the coding (writing) and then the decoding (reading).

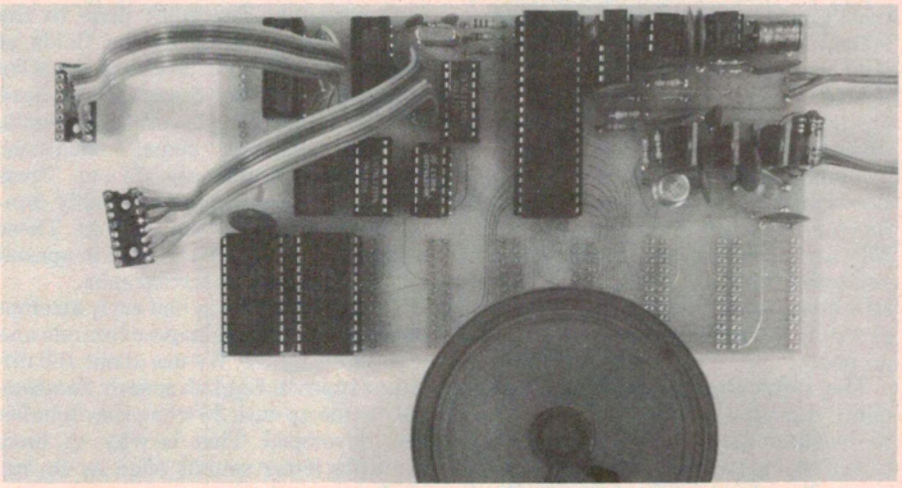

The speech synthesiser uses National Semiconductor's Digitalker chip set on a board just 155 x 90 mm.

Recording or storage

The first cylinder recording and playback machines were the start of the electronic version of speech coding and eventual synthesis. The speech data were stored, and still are, as a code of displacement values from a reference line on a physical medium such as plastic or metal. A machine does the coding (disc-cutting lathe) and a machine (record player) does the synthesis on decoding. The playing back can be done when and in what order the user requires. It is normally done by a human who selects the record, the track and the time of playback (synthesis) although it can be done by machine (a jukebox). Tape recordings are another example.

The ease with which these types of equipment work and their abundance in western society obscure a very important fact. The information content or number of bits of information stored is *colossal.* A lay person would tend to think of a unit of information as being a sentence or a word or a letter. This is not so.

Sampling

Sampling theory dictates that a sampling rate of at least twice the maximum frequency to be reproduced is necessary to record and playback an analogue signal by taking 'samples' of that signal. Imagine a voice with a highest frequency of 1000 Hz (normal female speech). A sample rate of at least 2000 Hz is required. For an analogue device such as a microphone or record stylus, this means the physical equipment has to have a response of *at least* 2000 Hz (i.e: be able to vibrate that quickly).

Now each single sample bit (analogue) has a particular value, say between 1 and 256. A binary number having eight bits would be necessary to represent each value. Therefore *one second of speech* would require the equivalent of approximately *two kilobytes* of memory or storage information. This is about half a track at the start of an LP record or a couple of 2114 memory IC chips.

A total of *20 minutes of speech,* therefore, occupies some 20 x 60 x 2000 bytes ▶

or *2.4 megabytes* of memory, the equivalent of some 36 microcomputers having 64K of RAM!

Alternatively, this occupies one side of a normal LP record, and LP records have frequencies greater than 1000 Hz.

Not yet convinced? Twenty minutes of speech is recorded in about 20 pages of a paperback book, so a 200 page book, which is not a large book, holds ten times as much again — 24 megabytes!

Redundancy

Language, written or spoken, is a remarkable thing. Despite its apparent efficiency in storage capacity, it contains a considerable amount of redundant information. Try this as an example: get a friend to write out a sentence in pencil on a single line. Then have them rub out the bottom half of the sentence. Now you try and read it. Providing your 'friend' hasn't used some strange Gaelic expressions or something similar, you'll find that you can decipher the sentence — with a little difficulty perhaps. This is a crude example because both essential and redundant information is removed but you can see that you *can* synthesise the original sentence *even though 50% has been removed.* If the right redundant bits are removed, the sentence can be reduced even further than that and yet still be restored through synthesis, and probably with greater speed and accuracy.

The same can be done with speech using digital electronics. By removing the redundancies and with appropriate sampling, speech patterns can be 'compressed' and stored for later synthesis into what we recognise as speech.

There are three techniques currently used to achieve this extraction or compression of data and they are based on different ways of looking at the speech waveform.

If one looks at a speech waveform on an instrument such as an oscilloscope then a waveform like the one below is seen:

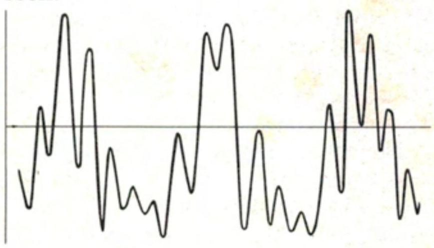

Figure 1. **Typical waveform.**

If the amplitude of the waveform is recorded for consecutive periods of time then the speech data has been *coded* and if the code is read back so that the original amplitude variation can be regenerated, then the speech is *synthesised.* The technique is called *waveform digitisation.*

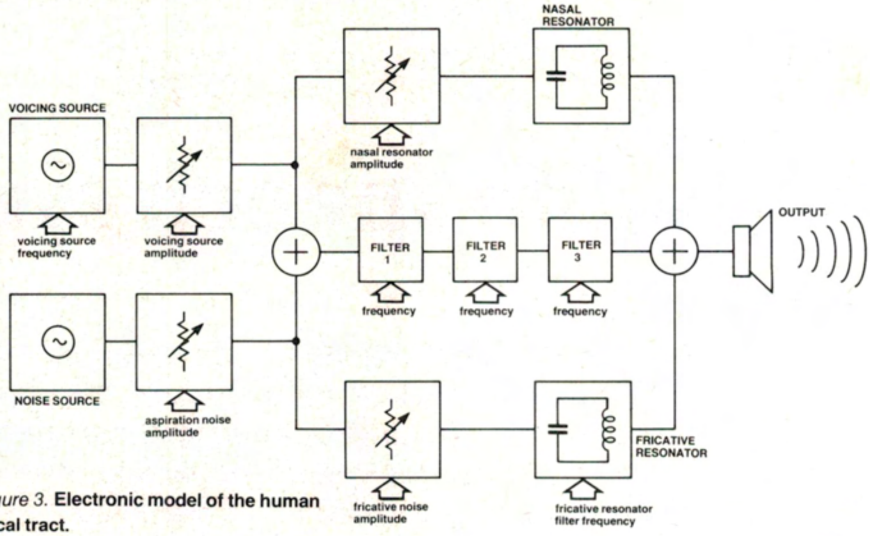

Figure 3. **Electronic model of the human vocal tract.**

A second technique is called *phoneme reconstruction.* The work done in the last 50 years or so in the fields of linguistics and acoustics have revealed 'subunits' of human speech. These have various names like phonemes, formants, glottals, nasals, stops, fricatives, morphemes, trills, laterals etc. These can be put together in reasonably complicated ways to form sounds. These parts or subunits can be seen in special recordings called spectrographs.

The alphabet was a bad early attempt to distinguish these parts of sounds and give them names. We use about 160 different types in English speech. Someone only came up with 26 when the alphabet was developed. That is why we have multiple letter sounds such as ch, ng, ough, ea, oo, sh, etc.

It is possible to get by with only 40 if you don't mind 'mechanical' speech. A piece of speech written with this sort of data should look something like:

" ðə 'lıŋgwıst 'sɔ: ðə pèlıkən# "

Figure 2.

It becomes relatively easy to synthesise speech from these subunits since only a small number of sound subunits have to be stored and each word can be generated by recalling which phonemes to use and what order to use them in. This uses less memory than the waveform digitising mentioned earlier. A data rate of about 400 bits per second will suffice. The problem is what symbols to use when specifying various phonemes. A conventional typewriter or computer keyboard runs out after 26.

A more sophisticated (but not yet more satisfying) method approaches the problem from a completely different viewpoint. In fact, this method was one of the earliest attempted in the young days of acoustics. It models the human vocal tract and is called (by some people) *linear predictive coding.* A 'noise' generator, either mechanical (tubes and wind pump) or electronic (an oscillator or oscillators), is used and various things are done to the noise to copy the things that we do when we cause speech to occur.

An electronic model of the vocal tract is shown in Figure 3. When we speak we do various things at different times with different parts of the vocal tract. These variables include things like alter the diameter of the pharynx at set places, block the nasal cavity, raise or lower the tongue, oppose lips and teeth, stretch vocal cords, close the glottis, etc.

With synthesis, the variables are identified for each sound to be coded and only information about these actions are stored in memory. When recalled, these variables tell the mechanical or electronic equipment what to do as it generates noise. With luck, the noise sounds like speech. The memory requirements are about five times that of the phoneme reconstruction technique.

The three major techniques of speech synthesis — waveform digitisation, phoneme reconstruction and linear predictive coding, are employed in digital electronic speech synthesisers produced by specialist IC manufacturers. Each produces a chip or chip set capable of producing clearly recognisable words, sounds and phrases. There are limitations in the 'vocabulary' each produces and the quality of the sound generated, but having heard the result one cannot help being amazed. The companies involved are National Semiconductor, Texas Instruments and Votrax. The technique employed and chips or chip set are as follows:

Waveform digitisation
National's DT1050 Digitalker
Phoneme reconstruction
Votrax's SC-01
Linear predictive coding
TI's TMS-5100 'speak 'n spell'

ETI-647 'TURTLE TALK' SPEECH SYNTHESISER BOARD

- Mounts in 'Tasman Turtle Robot' (ETI-645 — April/May/June '82)
- Can be used as 'stand-alone' speech synthesiser
- Can interface to almost any computer
- Produces fixed words using simple program
- Can say alphabet and numbers
- 109-word minimum vocabulary, plus prefixes, suffixes, tones
- Vocabulary expandable to about 600 words (up to 8 ROMs)
- Words can be spelled
- Different languages available (mixture permitted)
- Different voices (male, female, child) possible
- Sounds other than speech can be used
- Phrases, sentences easily programmed
- World exclusive 'mute' facility to create other words, etc
- Clock oscillator on-board
- Audio amp. on-board
- Voltage regulators on-board
- ROM power saver circuit

An excellent discussion on the three techniques, well worth reading, was published in the November 1981 issue of *Practical Computing* (p. 112-114).

Turtle Talk

Inevitably, a speech synthesiser system must be a conglomeration of compromises between storage capacity, data rate, speech quality and cost. The higher the data rate, the better the speech quality. When setting out to select a system to add speech to the Tasman Turtle Robot (see ETI, April/May/June '82) we looked for a system that gave compromise between data rate and speech quality. The National Digitalker system was selected. It applies a series of data compression techniques to remove as much as possible of the data not absolutely necessary. Word and sound reconstruction is done by a complex algorithm, but the Digitalker, we thought, provided the best quality available and still required only modest memory space.

A number of extra functions, plus interfacing, were designed into the Turtle Talk Board using the Digitalker so that an extremely powerful speech synthesis board results. The extra functions include a 'ROM power saving' circuit and a 'mute' facility that permits other words to be constructed from the standard vocabulary words. This feature is a *world first* and the subject of a patent application. The features of the ETI-647 'Turtle Talk' Speech Synthesiser are summarised in the accompanying panel.

The design

An overall block diagram of the speech synthesiser board is shown in Figure 4. From this you can see that the system is virtually self-contained. The only necessary external inputs required to produce speech are 13.8 Vdc power (or 12 Vdc) and data in. The latter can be provided from a set of switches, from discrete logic or from a microprocessor or microcomputer. All necessary interfacing and decoding is done on-board as well as power regulation and audio filtering and processing.

In all, 12 ICs are incorporated, although not all of them may be required, depending on the way you interface the synthesiser to the 'outside world'. Up to eight 'speech' ROMs can be included on-board, although only the 'basic' two will be provided with kits. These provide a total of 144 'words'. This includes 109 'real' words, the 26 letters of the alphabet (ay, bee, see, dee . . .), the numbers one to twenty plus 'thirty', 'forty', 'fifty', 'sixty', 'seventy', 'eighty', 'ninety', 'hundred', 'thousand', 'million' and 'zero'. Also included are the word prefixes 'centi', 'milli' and 're', as well as the suffix 'ss' for making plurals. Two tones and five 'silence' periods make up the remaining 'words'.

More speech ROMs can be added if a larger vocabulary or a foreign language is required.

A block diagram of the National Semiconductor speech processor chip used in this project is shown in Figure 5. This IC, designated MM54104, uses the waveform digitisation technique discussed earlier.

Word selection input (SW1 — 8) is held in a register. This is multiplexed with data from the speech ROMs (DATA 1 — 8). From this, two digital control signals emerge — a 'control word address' and a 'phoneme address'. These are held in two registers and multiplexed onto the speech ROM address buss (ADR 0 - 13). 'Control words' from the ROM (via the ROM data input) are held in a 'control word register', passing to the 'control logic' which accepts 'chip select' (CS), 'com-

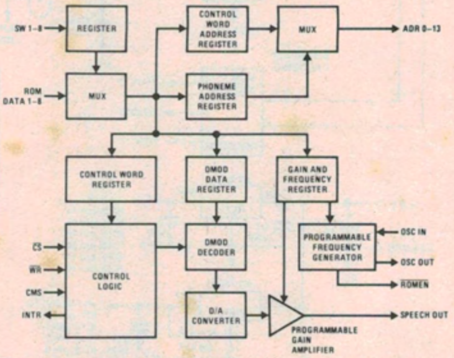

Figure 5. **Internal block diagram of the Digitalker speech processor chip.**

mand select' (CMS) and 'write' (WR) inputs for overall control of the IC, and also puts out an 'interrupt' (INTR) signal for communication back to the equipment sending input to the chip.

The frequency and amplitude information of speech phonemes (word subunits) are reassembled by the speech processor chip from binary data fetched from the ROMs at the appropriate time. The encoded amplitude information is passed to 'deltamodulation decoder' (the DMOD DECODER block) which drives a digital-to-analogue converter (the D/A CONVERTER block). This changes the digital amplitude information to partial amplitude information which passes through a programmable gain amplifer. The encoded gain and frequency information is held in a 'gain and frequency register'. This selects frequency information from the programmable frequency generator and sets the gain of the programmable amplifier. The resultant is speech or whatever sound is programmed. ▶

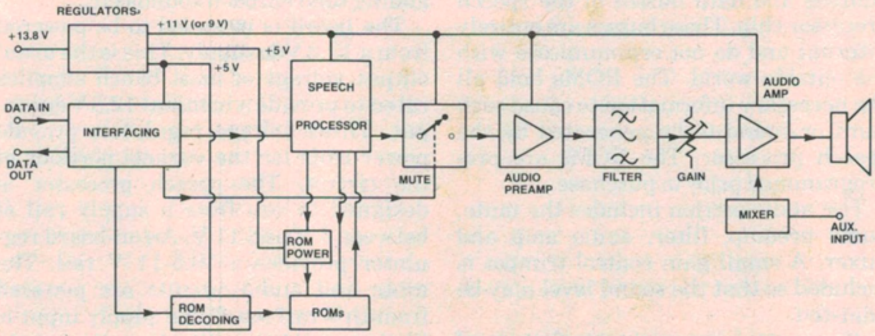

Figure 4. **Block diagram of the 'Turtle Talk' speech synthesiser project.**

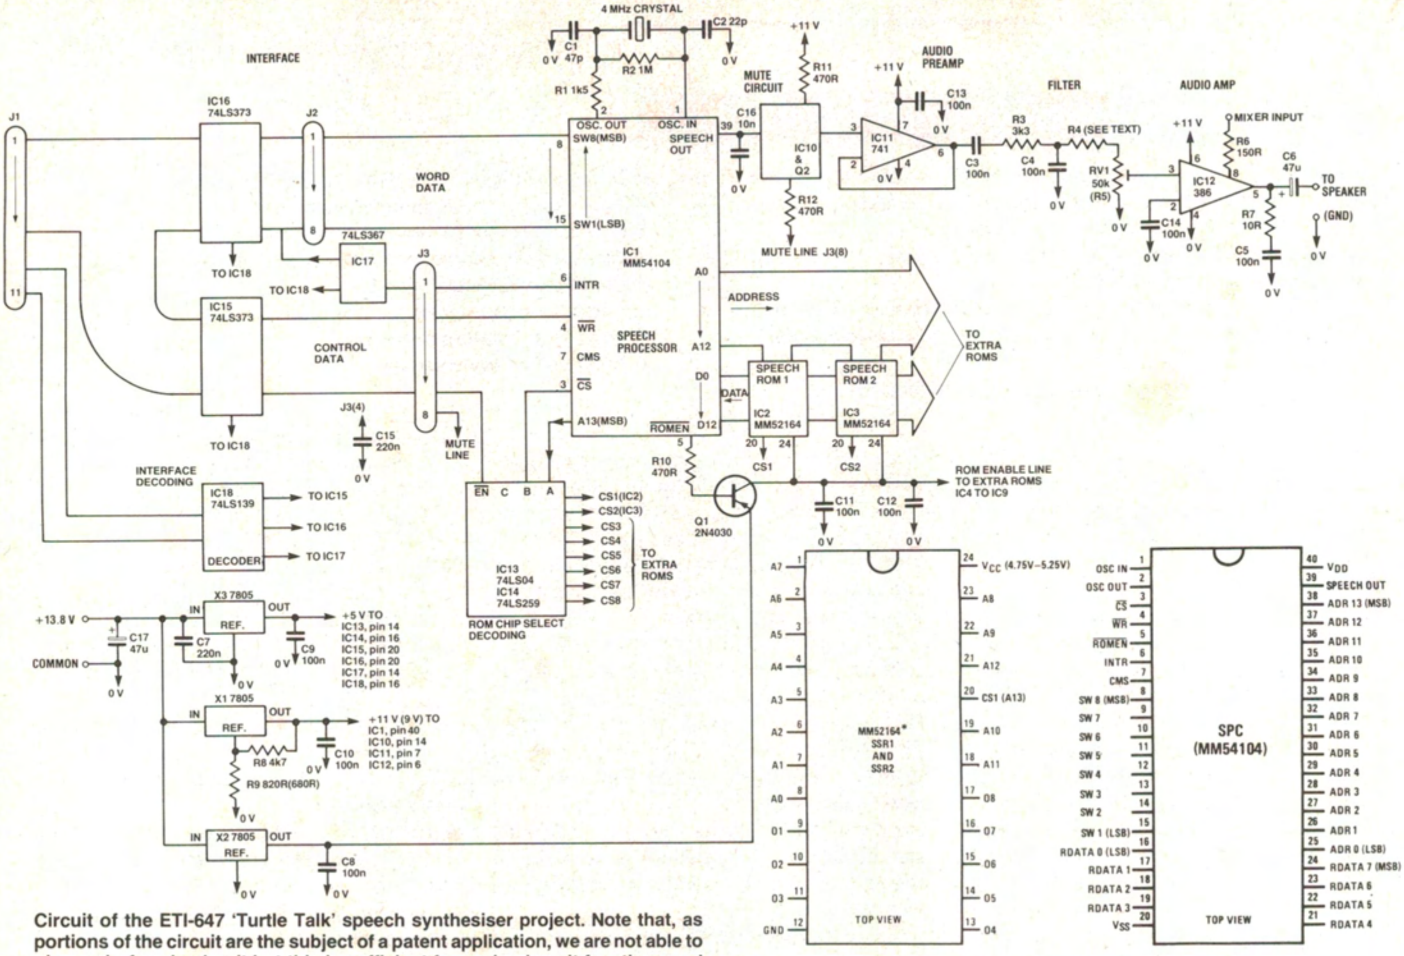

Circuit of the ETI-647 'Turtle Talk' speech synthesiser project. Note that, as portions of the circuit are the subject of a patent application, we are not able to give a pin-for-pin circuit but this is sufficient for seeing how it functions and fault-finding.

Note: 'Turtle Talk' is a trademark of Flexible Systems.

The MM54104 IC has facility for a crystal oscillator or can accept an external 'clock' input. This facility provides all the required clock facilities for the chip and provides a frequency reference from which speech frequency information is derived. A 4.000 MHz crystal is employed in this project, the 'osc. in' and 'osc. out' pins of the IC being used to form an oscillator.

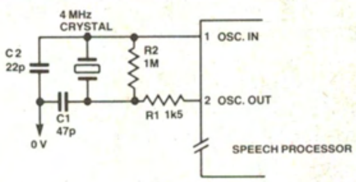

Each speech ROM is connected to the address and data busses of the speech processor chip. These busses are entirely internal and do not communicate with the outside world. The ROMs hold all the necessary information to cause each word or sound to be generated by the speech processor. The ROMs are pre-programmed prior to purchase.

The audio section includes the mute, audio preamp, filter, audio amp and mixer. A small gain control trimpot is included so that the sound level may be adjusted.

The mute is an exclusive feature of this design, adding extensive sound handling capabilities to the conventional Digitalker. It nulls the audio output to the speaker without resetting the processor.

The filter is designed to give an optimal frequency response to the stepped analogue data emenating from the speech processor chip. The filter provides a lower cutoff at 200 Hz and a roll off at the high end at around 8 kHz. The speech signal is smoothed and residual sampling frequency noise is removed. An LM386 low voltage audio power amplifier chip provides drive to an external loudspeaker.

The mixer is primarily designed to take the horn input from the Tasman Turtle but could be adopted for other uses, perhaps microphone over-talk, adding tone encoded sounds etc.

The board is intended to be powered from a 13.8 Vdc supply. This is the usual output voltage of most bench supplies rated to provide a nominal 12.5 Vdc output. Three voltage regulators provide power rails for the various portions of the circuit. The speech processor is designed to run from a supply rail of between 7 V and 11 V. An on-board regulator provides a 10.5-11 V rail. The mute and audio circuits are powered from this rail too. If the supply input is likely to fall below 13.2 V, or the supply input is 12-12.6 V, then this rail should be set at around 9 V. This is simply done by changing the value of one resistor.

There are two other regulators, each providing a +5 V rail. One supplies the interface and decoding circuitry, the other is used to supply the speech ROMs. As these only need to have the supply voltage present when being accessed, power consumption can be reduced by only turning on the supply to them as required. Accordingly, a 'power saver' circuit is incorporated to effect this.

If a separate source of power is available, the on-board regulators can be dispensed with.

Construction

A double-sided circuit board with plated-through holes is employed. Sockets are used to mount all the ICs. All the other components are soldered directly to the board. We recommend you use a temperature-controlled soldering iron with a small 'wedge' or 'chisel' point as the tracks and pads on the pc board are quite fine. All soldering is done from the *rear side* (i.e: non-component side) of the pc board. Use 20g or 22g resin-cored solder. The pc board tracks are solder plated to aid assembly and help prevent poor joints.

Examine the overlay diagram and identify which way up and which way

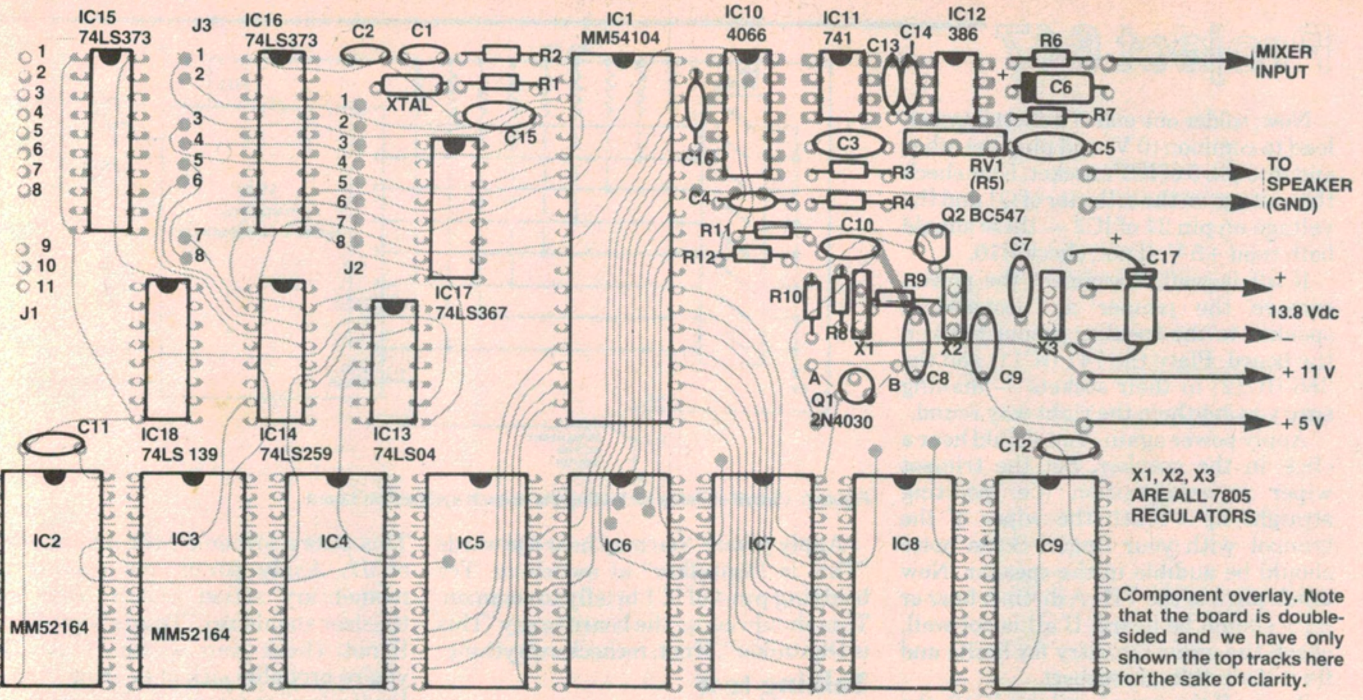

Component overlay. Note that the board is double-sided and we have only shown the top tracks here for the sake of clarity.

around the board goes. Commence construction by soldering all the IC sockets in place. Note that they are all orientated in the same direction. If you have your board facing the same way as our overlay shows, pin 1 of all IC sockets is at the *top left hand corner.*

Solder all the resistors in place next. Note that a link is used in place of R4. A resistor is only used here if the synthesiser is to drive a subsequent audio amplifier. The function of R4 is to reduce the level out of the audio amp, IC12 (the 386), while still permitting a reasonable setting range for the gain control trimpot. If contemplating this, a good value to start at would be 10k. Don't solder in the trimpot yet.

If you intend powering the unit from a supply having an output less than 13.2 V, then use a 680 ohm resistor for R9 instead of the 820 ohm value specified, otherwise regulator X1 will not function as the V_{in} — V_{out} voltage may fall below the regulator dropout voltage of 2 V, causing a malfunction of your synthesiser.

Solder all the capacitors in place next. Start with the ceramic capacitors, then solder the greencap (C16) in place and finish with C6. Leave C17 off for the present. Ensure the polarity of C6 corresponds to the overlay.

Now solder in the gain control trimpot — RV1(R5). The trimpot pins may be a tight fit in the holes, but a little gentle pressure will seat them home. Make sure the pads are well heated before applying solder. You might lightly tin the pins of the trimpot before mounting it, also.

The three voltage regulators can now be soldered in place, unless you have provision for off-board power supply rails and don't intend running the board from a 12 V or 13.8 V supply. Make sure you orient them correctly. The overlay shows a thick line on that part of the regulators which have the metal mounting tab, i.e: the tabs of all the regulators face toward the right side of the board, where the dc supply input connects.

Solder the two transistors in next. Make sure these, too, are the right way round. Don't seat Q1 right down on the board. Insert it such that its bottom flange is about 5 — 6 mm off the board. Solder the 4 MHz crystal in place next. Do it quickly so that excessive heat cannot damage the quartz crystal inside the can.

Now solder in place a length of figure-8 flex for the dc supply input. Use a type that has one lead marked (i.e: 'speaker' flex), so that you can identify which is the positive and which is the common lead. Having done that, C17 can be soldered in place. Make sure you get it the right way round.

Now you're ready for an initial test. DON'T PUT ANY ICs IN THEIR SOCKETS YET.

First checks

Make a complete physical check of your work *first*. See that there are no poor joints, no unsoldered leads and no solder 'dags' lying about or solder 'bridges' between tracks or pads. Check that all components are correctly placed. When you're satisfied, you can apply power.

Check voltages at the following points (all with respect to common):

- Input pin of X3 — +13.8 V (or +12 V)
- Output pin of X3 — +5 V
- Output pin of X2 — +5 V
- Output pin of X1 — +11 V (or +9 V)
- Collector of Q1 (case) — 0 V
- Pin 40 of IC1 — +11 V (or +9 V)
- Pin 14 of IC10 — +11 V (or +9 V)
- Pin 20 of IC15 — +5 V
- Pin 24 of IC2 — 0 V

PARTS LIST — ETI-647

Resistors	all ½W, 5% unless noted
R1	1k5
R2	1M
R3	3k3
R4	(link)
R5	(see RV1)
R6	150R
R7	10R
R8	4k7
R9	820R
R10, 11	470R
R12	15k
RV1	10k min. vert. trimpot
Capacitors	
C1	47p ceramic
C2	22p ceramic
C3, 4, 5, 8, 9, 10, 11 12, 13, 14	100n disc ceramic
C7, C15	220n ceramic
C6, C17	47u/16 V electro.
C16	10n greencap
Semiconductors	
IC1	MM54104 SPC
IC2, 3, 4, 5, 6, 7, 8, 9	MM52164 SSR
IC10	4066
IC11	741
IC12	386
IC13	74LS04
IC14	74LS259
IC15, 16	74LS373
IC17	74LS367
IC18	74LS139
X1, 2, 3	7805, LM340T5
Q1	2N4030
Q2	BC547

Miscellaneous

ETI-647 pc board (TTB-A) — this board is copyright to Flexible Systems who are the sole supplier; 4 MHz HC18/u crystal; IC sockets — 1 x 40 pin, 2 x 24 pin (or 8 x), 2 x 20 pin, 3 x 16 pin, 2 x 14 pin, 2 x 8 pin; ribbon cable and DIP headers or plugs to suit connection to computer; loudspeaker to suit; figure-8 cable, etc.

Price estimate $240 — $250

Project 647

Now, solder one end of a short 'jumper' lead to common (0 V) and plug the other end into pin 5 of IC1's socket. Now check the voltage on the collector of Q1 and the voltage on pin 24 of IC2 — these should both read +5 V. If not, check R10.

If all is well, disconnect the power, remove the jumper and connect a speaker to the speaker output pads on the board. Place the 741 (IC11) and the 386 (IC12) in their sockets — making sure you get them the right way round.

Apply power again. You should hear a click in the speaker. Set the trimpot wiper to mid-position (i.e: pointing straight up). Touch the wiper of the trimpot with your finger. Some noise should be audible in the speaker. Now touch pin 3 of the 741. A distinct buzz or hum should be heard. If all is not well, check the audio circuitry for faults and fix it if anything is suspect.

Turn off the power and put IC10, the 4066, in its socket. Apply power again. Touching pin 1 should cause a slight noise to be heard in the speaker, touching pin 2 should do the same. Temporarily jumper pin 8 of J3 to the +5 V aux. supply input. Now touch pin 2 of the 4066. No noise should be heard in the speaker. If not, check that R12 is correct and the Q2 is OK and correctly oriented. Remove power. Put all the other (LS-TTL) ICs in their sockets. Apply power and check that +5 V appears on the upper right hand pin of each. Remove power.

Now you can insert IC1 in its socket. Note that it is a MOS chip and should be supplied in conductive foam or other conductive package. Use the usual MOS-handling precautions when inserting it in its socket. Make sure you don't bend any leads under the chip package. Apply power once again and check that you get +11 V (or +9 V) on pin 40. If you have a CRO, check pins 1 and 2 for signs of oscillation — you should get a 4 MHz sinewave here. The amplitudes on the two pins will differ. (Pin 1 will have the greater amplitude signal — you may need a x10 probe to avoid 'loading' of the circuit causing it to cease oscillation). Alternatively a receiver covering this frequency could be used. You should find a strong signal on 4 MHz. If no sign of oscillation can be found, remove power and check the circuit and components around the crystal.

If all is well, remove power and insert the two speech ROMs, IC2 and IC3. These, like IC1, are also MOS devices so take care of them. One of these is marked SSR1 — this is IC2. The one marked SSR2 is IC3. Above all, make sure you insert them the right way round. See that no pins get bent under the package when inserting them into the sockets.

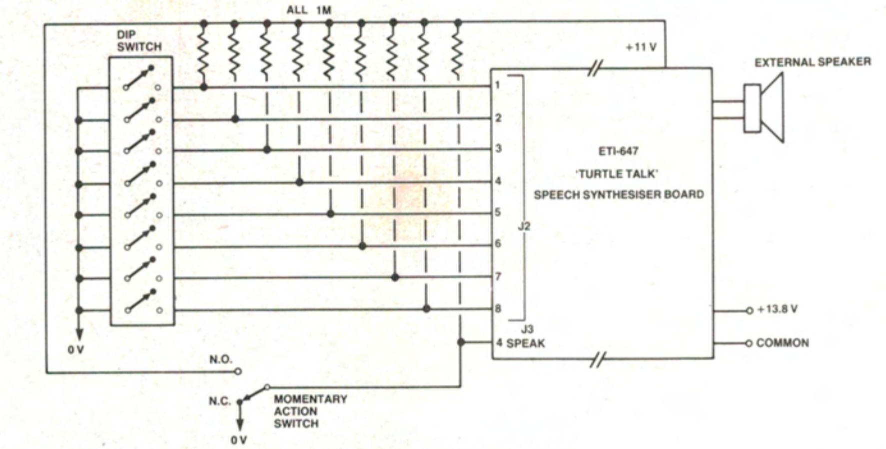

Figure 6. **Circuit suitable for testing the speech synthesiser board.**

Apply power. You may hear the words "This is Digitalker" at switch-on. Try bridging pin 4 of IC1 briefly to common. This should cause the board to say "This is Digitalker". If not, recheck everything.

Talking test

For this you will need an 8-way DIP switch, nine 1M resistors, a momentary action SPDT switch and either some ribbon cable or a host of hookup wire. Wire up the circuit shown in Figure 6. (Note that 0 V goes to COMMON on the pc board.) Set all the DIP switches ON. This puts a low on all the address lines of J2. Apply power and operate the momentary action switch. Your synthesiser should say "This is Digitalker". If not, check your wiring. If it's OK, you're probably sick of hearing "This is Digitalker" so set up another address on the DIP switch by consulting the accompanying table. Upon operating the momentary action switch you should hear the word set-up. Try a few others.

All OK? Well, you're ready to 'Turtle Talk'! ●

(. . . to be continued).

TABLE I. DT1050 MASTER WORD LIST

Word	8-Bit Binary Address SW8 SW1		8-Bit Binary Address SW8 SW1		8-Bit Binary Address SW8 SW1
THIS IS DIGITALKER	00000000	Q	00110000	IS	01100000
ONE	00000001	R	00110001	IT	01100001
TWO	00000010	S	00110010	KILO	01100010
THREE	00000011	T	00110011	LEFT	01100011
FOUR	00000100	U	W0110100	LESS	01100100
FIVE	00000101	V	00110101	LESSER	01100101
SIX	00000110	W	00110110	LIMIT	01100110
SEVEN	00000111	X	00110111	LOW	01100111
EIGHT	00001000	Y	00111000	LOWER	01101000
NINE	00001001	Z	00111001	MARK	01101001
TEN	00001010	AGAIN	00111010	METER	01101010
ELEVEN	00001011	AMPERE	00111011	MILE	01101011
TWELVE	00001100	AND	00111100	MILLI	01101100
THIRTEEN	00001101	AT	00111101	MINUS	01101101
FOURTEEN	00001110	CANCEL	00111110	MINUTE	01101110
FIFTEEN	00001111	CASE	00111111	NEAR	01101111
SIXTEEN	00010000	CENT	01000000	NUMBER	01110000
SEVENTEEN	00010001	400HERTZ TONE	01000001	OF	01110001
EIGHTEEN	00010010	80HERTZ TONE	01000010	OFF	01110010
NINETEEN	00010011	20MS SILENCE	01000011	ON	01110011
TWENTY	00010100	40MS SILENCE	01000100	OUT	01110100
THIRTY	00010101	80MS SILENCE	01000101	OVER	01110101
FORTY	00010110	160MS SILENCE	01000110	PARENTHESIS	01110110
FIFTY	00010111	320MS SILENCE	01000111	PERCENT	01110111
SIXTY	00011000	CENTI	01001000	PLEASE	01111000
SEVENTY	00011001	CHECK	01001001	PLUS	01111001
EIGHTY	00011010	COMMA	01001010	POINT	01111010
NINETY	00011011	CONTROL	01001011	POUND	01111011
HUNDRED	00011100	DANGER	01001100	PULSES	01111100
THOUSAND	00011101	DEGREE	01001101	RATE	01111101
MILLION	00011110	DOLLAR	01001110	RE	01111110
ZERO	00011111	DOWN	01001111	READY	01111111
A	00100000	EQUAL	01010000	RIGHT	10000000
B	00100001	ERROR	01010001	SS (Note 1)	10000001
C	00100010	FEET	01010010	SECOND	10000010
D	00100011	FLOW	01010011	SET	10000011
E	00100100	FUEL	01010100	SPACE	10000100
F	00100101	GALLON	01010101	SPEED	10000101
G	00100110	GO	01010110	STAR	10000110
H	00100111	GRAM	01010111	START	10000111
I	00101000	GREAT	01011000	STOP	10001000
J	00101001	GREATER	01011001	THAN	10001001
K	00101010	HAVE	01011010	THE	10001010
L	00101011	HIGH	01011011	TIME	10001011
M	00101100	HIGHER	01011100	TRY	10001100
N	00101101	HOUR	01011101	UP	10001101
O	00101110	IN	01011110	VOLT	10001110
P	00101111	INCHES	01011111	WEIGHT (Note 2)	10001111

Note 1: "SS" makes any singular word plural

Note 2: Address 143 is the last legal address in this particular word list. Exceeding address 143 will produce pieces of unintelligible invalid speech data.