



Stereo sound generator for sound effects and music synthesis

All the sophisticated application software in the world can't compensate for a home computer that communicates with its user in an unattractive or confusing way. To enhance present levels of visual and audio communication from computer to user, a stereo sound generator, SAA1099, and a 64-colour encoder, TEA2000 (Ref.1), have been developed, each giving software and hardware designers the scope to add more realism to their products.

The sound generator SAA1099 is described in this article. It can produce a wide variety of sound effects including simulations of musical instruments, and the sounds required for arcade/home-computer games. Whereas most integrated sound generators have only three frequency generators, the SAA1099 has six, and the tones from each can be mixed with several kinds of noise. Because there are six frequency generators, full musical chords (including the tonic) can be produced, and two chords (excluding tonics) can overlap. All musical notes in an 8-octave range can be produced. A stereo effect that can give width to scenes in video games is produced by duplicating the six sound components to form identical left and right-channel signals, weighting the signals of each channel, and combining them to form a stereo signal.

Besides the stereo sound facility, effects such as Doppler shifts can be imitated with a minimum of software control. Outside the home computer market, the SAA1099 can be used to good effect in model railways and cars, electronic musical instruments and audible alarms, to name just a few applications.

The SAA1099 readily interfaces with most 8-bit micro-controllers and requires only a simple filter to suppress any high-frequency components in the audio output. It has been designed in such a way that a minimum of external components are required. Table 1 gives additional data on the SAA1099; see also Refs 2 and 3.

TABLE 1
Brief data on the SAA1099 (all values are typical)

supply voltage	V_{DD}	5 V
supply current	I_{DD}	70 mA
reference current (pin 6)	I_{ref}	250 μ A
total power dissipation	P_{tot}	500 mW
external clock		8 MHz
data input		8-bit parallel (TTL-compatible)
output frequency range		31 Hz to 7.81 kHz (8 octaves)
output		pulse width modulated
operating ambient temperature range	T_{amb}	0 to -70° C
package		18-pin plastic DIL

COMPLETE SOUND GENERATOR

Figure 1 shows a complete sound generating system. From an 8-bit wide data input from a microprocessor, the SAA1099 generates a variable-amplitude stereo analogue signal chopped at a rate of 62,5 kHz. A simple external low-pass filter suppresses the high frequency components of the output signal. The incoming data which shapes the spectrum of the audio output is multiplexed to simplify interfacing, the signal A0 being used to indicate whether the data is a register address, or data for the register. The A0 signal is used with the \overline{CS} and \overline{WR} signals to control the data transfers from the microprocessor to the SAA1099. These control signals are compatible with a wide range of microprocessors. In addition, for optimum interfacing with an SCN68000 series microprocessor, the SAA1099 has a DTACK output. All internal timing is derived from an external 8 MHz clock.

TABLE 2
Function of the A0 input

A0	data bus input								function
	D7	D6	D5	D4	D3	D2	D1	D0	
0	D7	D6	D5	D4	D3	D2	D1	D0	data for internal registers, see Table 3
1	X	X	X	X	A3	A2	A1	A0	internal register address, A3 is the MSB

X = don't care

Table 2 shows the function of the A0 input. When A0 = 1, the bus data indicates the address of the control register in the SAA1099 to be written and this address is loaded into the command register. The next data byte on the bus, which contains the control information for the register that has been addressed, is written to the register when A0 = 0. Once addressed, a control register can be updated without further addressing.

CIRCUIT DESCRIPTION

Frequency generators

The SAA1099, see Fig. 2, has six frequency generators each of which can generate 256 tones in each of eight octaves from 31 Hz to 7,81 kHz. Each generator can be switched on and off individually, making it possible to preselect a tone and to make it audible when required. To simplify the software, the frequency generators can be synchronized at start-up

and when changing frequencies, octaves and envelopes using the frequency reset bit, see 'Synchronization'. The outputs of frequency generators 0 and 3 can each control a noise generator while those of generators 1 and 4 can each control an envelope generator for creating special effects.

Table 3 gives the addresses and the bit allocation of the SAA1099's internal registers used to control the frequency generators and its other sound-generating circuitry.

Noise generators

Two noise generators each have a programmable output controlled by the contents of register 16 which determines whether the output is:

- software-controlled via frequency generator 0 or 3 (which then produce no tone). The 'colour' of the noise generated is derived from twice the frequency of the frequency generator output, i.e. from 61 Hz to 15,6 kHz.
- one of three pre-defined noises based on clock frequencies of 7,8 kHz, 15,6 kHz or 31,25 kHz. In this case, the output of noise generator 0 can be mixed with the outputs of frequency generators 0, 1 and 2, and the output of noise generator 1 can be mixed with the outputs of frequency generators 3, 4 and 5, see Fig. 2. For mixing, the amplitude of the tone is increased relative to that of the noise.

Noise/frequency mixers

The SAA1099 has six mixers, one per frequency generator, for mixing tones with noise. Dependent on the status of bits D0 to D5 of registers 14 and 15, each mixer can be set:

- to mix the noise and the tone
- to pass the tone only
- to pass the noise only
- to pass neither tone nor noise.

Amplitude controllers

The SAA1099 has six amplitude controllers used, for example, to create a stereo effect. Each controller duplicates the signal from one of the noise/frequency mixers to form left and right-channel components and assigns one of sixteen amplitudes set by the contents of the control registers 00 to 05 to each component. A stereo effect can be produced simply by varying the amplitude of each component. To move a sound from one channel to another requires, per tone, only one update of the contents of the appropriate amplitude register.

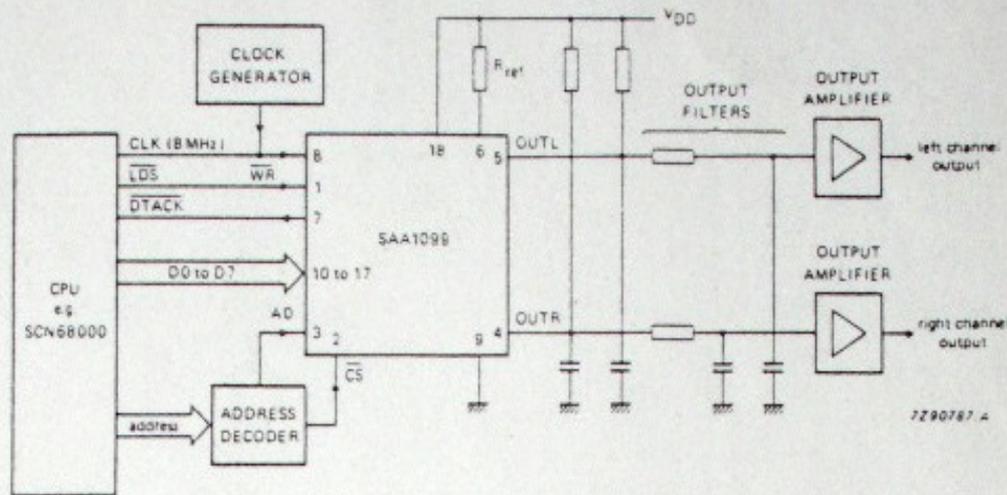


Fig. 1 Sound generation system for home computers and video games equipment.

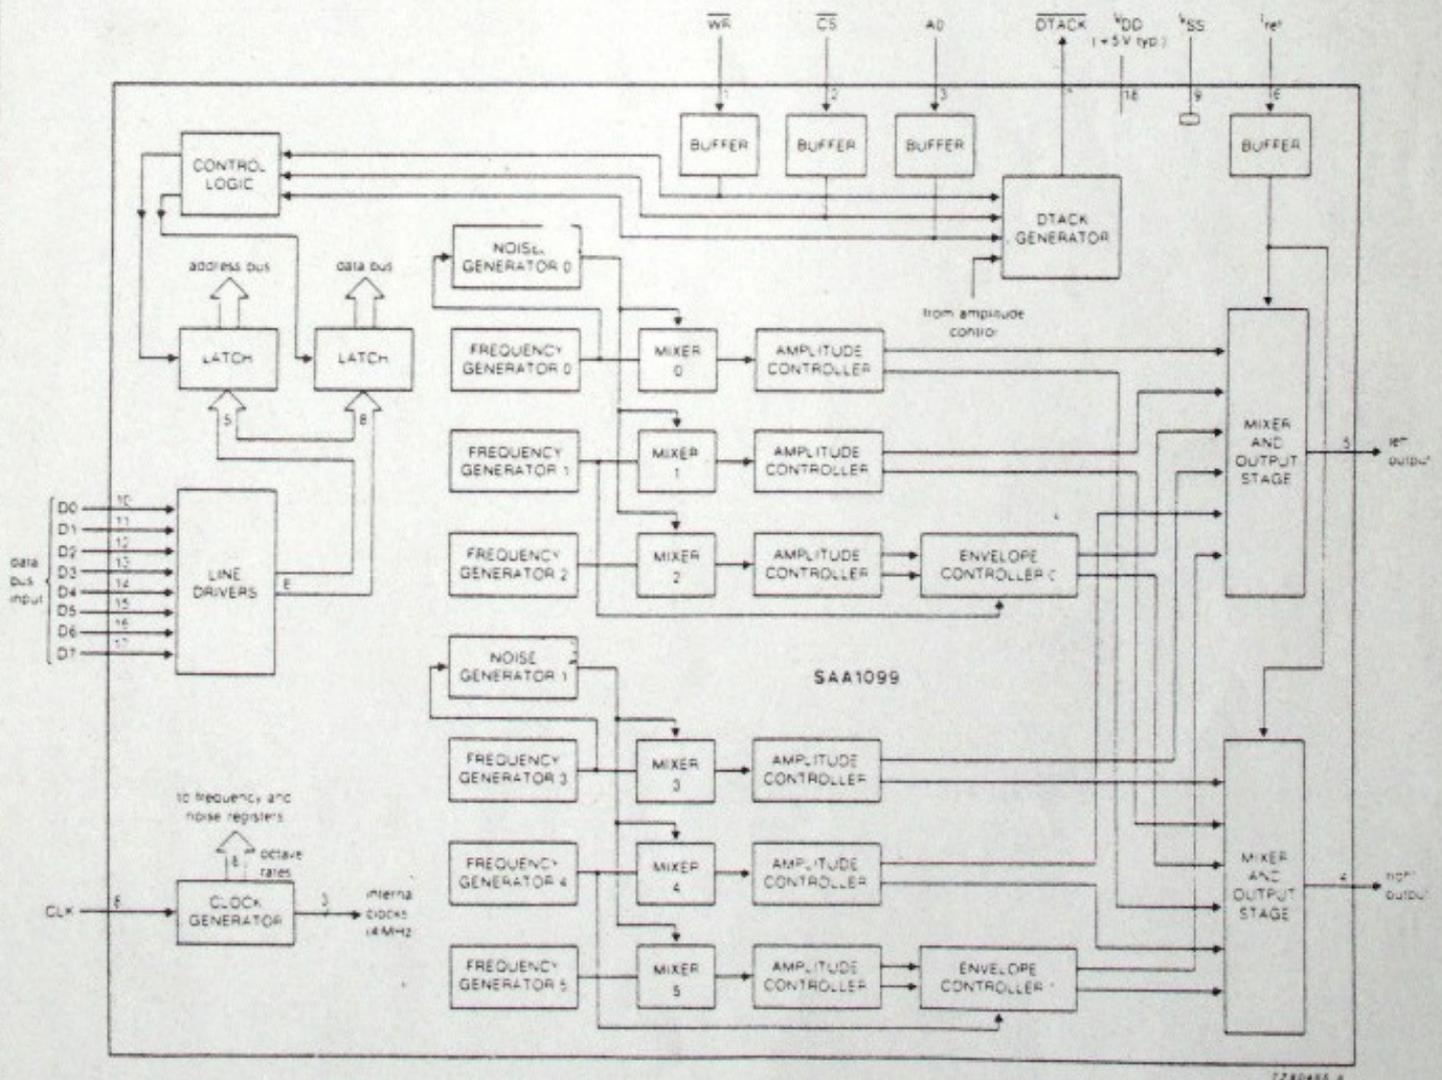


Fig. 2 SAA1099 stereo sound generator.

TABLE 3
Bit allocation of the internal registers of the SAA1099; A0 = 1

register address	data bus input								description/remarks
	D7	D6	D5	D4	D3	D2	D1	D0	
	MSB				LSB				
00	amp. right channel				amp. left channel				controller 0
01	amp. right channel				amp. left channel				controller 1
02	amp. right channel				amp. left channel				controller 2
03	amp. right channel				amp. left channel				controller 3
04	amp. right channel				amp. left channel				controller 4
05	amp. right channel				amp. left channel				controller 5
06	X	X	X	X	X	X	X	X	reserved for possible expansion
07	X	X	X	X	X	X	X	X	reserved for possible expansion
08	tone number for frequency generator 0								$\text{frequency of tone (Hz)} = \frac{15625 \times 2^{\text{octave number}}}{(511 - \text{tone number})}$, see notes
09	tone number for frequency generator 1								
0A	tone number for frequency generator 2								
0B	tone number for frequency generator 3								
0C	tone number for frequency generator 4								
0D	tone number for frequency generator 5								
0E	X	X	X	X	X	X	X	X	reserved for possible expansion
0F	X	X	X	X	X	X	X	X	reserved for possible expansion
10	X	octave no. of freq. gen. 1			X	octave no. of freq. gen. 0			octave 0 (0 0 0): 31 Hz to 61 Hz; octave 1 (0 0 1): 61 Hz to 122 Hz octave 2 (0 1 0): 122 Hz to 244 Hz; octave 3 (0 1 1): 245 Hz to 488 Hz octave 4 (1 0 0): 489 Hz to 977 Hz; octave 5 (1 0 1): 978 Hz to 1.95 kHz octave 6 (1 1 0): 1.96 kHz to 3.91 kHz; octave 7 (1 1 1): 3.91 kHz to 7.81 kHz
11	X	octave no. of freq. gen. 3			X	octave no. of freq. gen. 2			
12	X	octave no. of freq. gen. 5			X	octave no. of freq. gen. 4			
13	X	X	X	X	X	X	X	X	
14	X	X	5	4	3	2	1	0	frequency enable (active-HIGH); 0 to 5 refer to the noise/frequency mixers
15	X	X	5	4	3	2	1	0	noise enable (active-HIGH); 0 to 5 refer to the noise/frequency mixers
16	X	X	generator 1		X	generator 0			noise generator clock frequency:
			0	0		0	0		31.25 kHz
			0	1		0	1		15.6 kHz
			1	0		1	0		7.8 kHz
			1	1		1	1		61 Hz to 15.6 kHz (freq. generator 0 or 3 controlling noise generator 0 or 1 respectively)
17	X	X	X	X	X	X	X	X	
18	envelope generator 0								see Table 4 and Fig. 3
19	envelope generator 1								see Table 4 and Fig. 3
1A	X	X	X	X	X	X	X	X	
1B	X	X	X	X	X	X	X	X	
1C	X	X	X	X	X	X	RST SE		RST: reset for all freq. generators (active-HIGH); SE: sound enable for all channels (active-HIGH), see notes
1D	X	X	X	X	X	X	X	X	
1E	X	X	X	X	X	X	X	X	
1F	X	X	X	X	X	X	X	X	
	D7	D6	D5	D4	D3	D2	D1	D0	

This block of 32 registers is repeated eight times between addresses 00 and FF in the full internal memory map. All don't cares (X) should be written as zeroes.

Tone numbers of 1 to 256 are valid.

At power-on, the sound enable bit is set to 0 (all channels disabled). When the frequency reset bit is set, all frequency generators are reset and synchronized.

Envelope controllers

Two envelope controllers enable the left and the right components of two stereo channels to be modified for:

- single attack
- single decay
- single attack and decay (triangular)
- maximum amplitude
- repetitive attack
- repetitive decay
- repetitive attack and decay (triangular)
- zero amplitude.

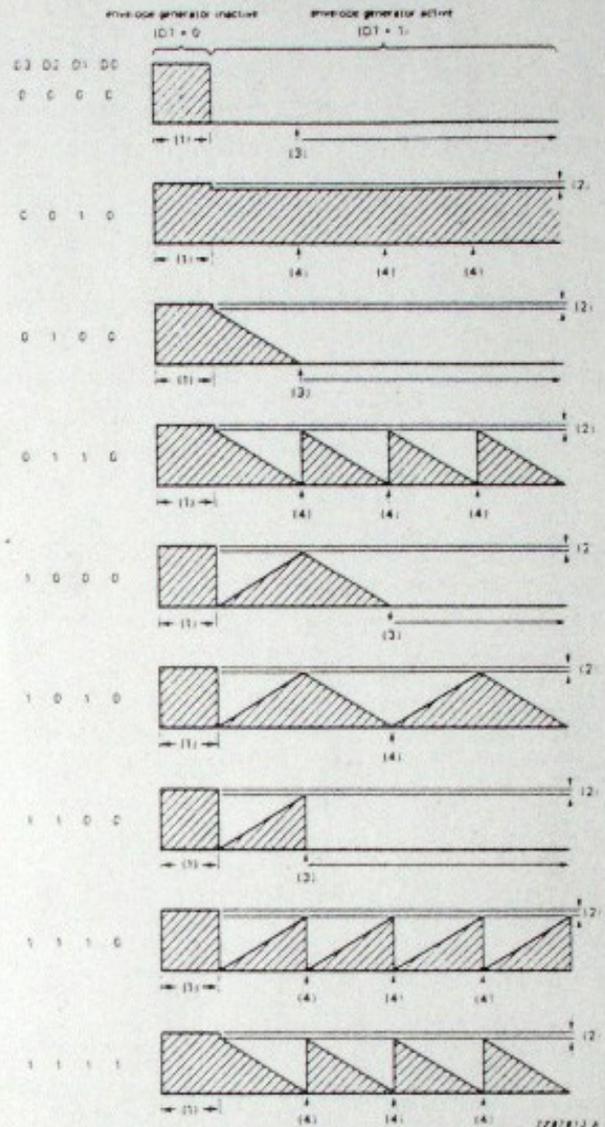
Figure 3 shows the envelopes which are selected by bits D1 to D3 of the envelope registers 18 and 19, see Table 4. The repetition frequency of the envelopes can be software-controlled by writing to the envelope address register (the data written is irrelevant), or clocked internally at the output frequency of the frequency generator (1 or 4). An envelope will always be completed before a new envelope instruction is implemented. In the case of repetitive envelopes, the envelope returns to its starting level which may not necessarily be zero.

When the envelope facilities are used, and the maximum amplitude set by the contents of registers 00 to 05 is chosen, the amplitude is 7/8ths that normally available.

TABLE 4

Bit allocation of the envelope generator registers (addr. 18 and 19)

bit	function
D0	
0	identical envelopes for the left and right channel components
1	inverse envelopes for the left and right channel components
D3 D2 D1	
0 0 0	zero amplitude
0 0 1	maximum amplitude
0 1 0	single decay
0 1 1	repetitive decay
1 0 0	single triangular
1 0 1	repetitive triangular
1 1 0	single attack
1 1 1	repetitive attack
D4	
0	4-bit envelope resolution
1	3-bit envelope resolution
D5	
0	internal envelope clock (frequency generator 1 or 4)
1	external envelope clock (address write pulse A0)
D6	don't care
D7	
0	reset (no envelope control)
1	envelope control enabled



Notes:

- (1) normal amplitude control; envelope generator inactive.
- (2) when the envelope generator is active, the maximum amplitude is 7/8ths of the maximum set by registers 00 to 05.
- (3) after position (3), a new envelope will be implemented as soon as the new buffered controls are received.
- (4) when an envelope is interrupted, the new envelope starts at the positions indicated (4).

Fig.3 Envelope waveforms. (a) to (h) the waveforms of one channel, the left say, of an envelope generator; (a) zero amplitude, (b) maximum amplitude, (c) single decay, (d) repetitive decay, (e) single attack and decay, (f) repetitive attack and decay, (g) single attack, (h) repetitive attack, (i), as (h), but for the right channel showing the operation of the inverse envelope bit D0, see also Table 4

Two types of envelope control are stored in the envelope registers – direct-acting controls and buffered controls. The direct-acting controls always take immediate effect and are:

- the envelope enable/reset (bit D7)
- the envelope resolution: 16 levels up to an envelope repetition frequency of 977 Hz, 8 levels above 977 Hz (bit D4).

The buffered controls are acted upon only at the times shown in Fig.3 and determine:

- the envelope waveform (bits D1 to D3)
- the type of envelope clock (bit D5)
- whether the left and right channels are inverted (bit D0).

When an external envelope clock is selected, an envelope is only created when address 18 or 19 is written to (that is when A0 is set to 1 and there is a 'write 18 or 19' command).

ix-channel mixers/output stages

The six components of the left channel are combined in a mixer. The output stage of the mixer contains six equally-weighted current sinks which provide a PWM output from which an analogue output is derived by low-pass filtering. An identical mixer is used to combine the components of the right channel.

SYNCHRONIZATION

To simplify the software writer's work, several synchronization functions are incorporated in the SAA1099. They affect:

- the starting of frequency generators
- the changing of frequencies and octaves
- the changing of envelopes.

Synchronizing frequency changes

The internal architecture of the SAA1099 is such that when a new tone in an octave different from that currently selected is required, the frequency register should be written to before the octave register. Failure to write to the registers in this order with the sound enabled, may produce a click in the audio output. The frequency and octave registers can, however, be written at any time, but data can only be acted upon by the SAA1099 on a transition of the associated frequency generator, that is, data won't be acted upon until half the period of the current frequency has elapsed. This means that:

- when the frequency and octave registers are set for the lowest frequency of 31 Hz, the new frequency data or octave data may not be acted upon for up to about 17 ms (half a period). Therefore, to ensure this interval has elapsed, a delay corresponding to half a period of the existing tone should be written in the software between writing the new frequency or octave and enabling the output.

- at the higher end of the spectrum, if it is required to change the frequency and octave registers simultaneously, both new values must be written (frequency first; octave second) within half the period of the current frequency (i.e. within 64 μ s for a 7.74 kHz tone).

Synchronization on reset

All frequency generators can be reset by setting bit D1 (RST) of the register at address 1C. In this state, frequency and octave data can still be written to the SAA1099, but will not be acted upon. Therefore, as long as RST is active, a register value can be overwritten with new data. However, any new data in the register will not be acted on until half the period of the frequency whose value is held in the register when RST was set has elapsed. This is because RST not only sets all generators to a known state, it synchronizes their start-up.

APPLICATIONS

Probably the most obvious application for the SAA1099 is in video games where the wide range of sound effects available can be used to make games more appealing. The stereo effect, for example, can be used to give width to scenes and to create the impression of movement of objects in the scene. More interesting is the possibility of relating both channel amplitudes and Doppler shift in frequency to the position of an object relative to the user. For example, the sound of passing vehicles or swooping spaceships can be realistically produced.

Many of the sounds in computer games are based on 'coloured' noise (e.g. aircraft, gunfire and car engines). The two noise generators of the SAA1099 with full software control of the noise colour, and the ability to mix the noise with tones, enable two separate 'coloured' noises to be produced in stereo for increased realism.

As mentioned earlier, the SAA1099 can produce all musical notes across eight octaves from 31 Hz to 7.81 kHz. The availability of six frequency generators enables full musical chords (including the tonic) to be produced and allows two chords (excluding tonics) to overlap.

The advanced envelope generation facilities and software control of amplitude and frequency enable musical instruments to be mimicked including vibrato and tremolo effects.

Software modules that generate specific sounds e.g. piano and trumpet, as well as modules that generate sounds that can be altered to suit the situation in a video game can be created, e.g. laser gun, sirens and error warnings.

The following sections outline how several sound effects can be produced. Programming details aren't given (these are published in Ref.3), but some salient points and data are.

Notes and chords

A scale of middle C to high C can be produced using just one frequency generator set for maximum amplitude on both outputs, with the signal from the appropriate noise generator disabled. Table 5 shows the tone numbers and octave numbers required. Note that the same note in different octaves has the same tone number, e.g. C is tone H21 in octave 3 and octave 4, with an 8 MHz crystal-derived clock, all the notes are produced to an accuracy better than 0.1%.

Chords, for example C major and A minor, can be produced using all six generators by writing the appropriate values to the tone and octave registers, then enabling the notes in sequence, without disabling those already active.

Musical arrangements are produced in a similar fashion by an appropriate combination of writing/enabling of tone and octave registers. A stereo effect is produced by giving the treble clef predominance on one channel and the bass clef predominance on the other.

TABLE 5
Chromatic scale

note	tone number (hex)	octave number	required frequency (Hz)	actual frequency* (Hz)
middle C	21		261,626	261,506
C #	3C		277,183	277,162
D	55		293,665	293,427
D #	6D		311,127	310,945
E	84		329,628	329,815
F	99	03	349,228	349,162
F #	AD		369,994	369,822
G	C0		391,995	391,850
G #	D2		415,305	415,282
A	E3		440,000	440,141
A #	F3		466,164	466,418
B	05		493,883	494,071
C	21	04	523,251	523,013

* with an 8 MHz crystal-derived clock

Siren

The sound of a siren can be simulated using only one frequency generator. All possible values (from high to low) are written to a tone register, changing the octave where necessary, and repeated. A repetitive triangular envelope with the left channel set to be the inverse of that of the right gives the impression of side-to-side movement. When the envelope generator is externally clocked at 85 clocks/second using the address write pulse A0, the sound repeats every 2.7 s.

Aircraft

A dogfight between two aircraft can be simulated using both envelope controllers. One envelope should be set for repetitive attack to simulate gunfire, the other set for repetitive decay to simulate the noise of a propeller. These can be combined with two different types of noise and can be internally clocked. Writing different values to the amplitude registers to vary the volume in the two outputs creates the impression of movement.

A combination of noise and tone can be used to simulate a jet plane preparing for take-off. The colour of the noise should be held constant while the frequency of the tone should be increased smoothly by successive writing to the frequency register. A sawtooth envelope with the invert bit set suggests taxiing movement. If the envelope is clocked to its maximum value on the left output before the sound is enabled, then not clocked during tone changes, the right channel will not be heard, and the plane will appear to be located on the left of the scene.

To give an impression of movement from left to right, the envelope should be clocked eight times when the tone reaches some desired amplitude, reducing the level on the left and increasing that on the right. The left output can then be disabled and the envelope clocked another eight times to reduce the level on the right to zero, giving the impression of movement to the right.

Gun, cannon-fire and laser guns

The sound of a machine gun can be produced using an internally-clocked sawtooth envelope to shape continuous coloured noise. Another gun can be simulated using noise of a different colour and a repetitive decay envelope. Lulls in firing are produced by enabling and disabling the sound enable bit.

The sound of a cannon can be simulated in the same way as the first machine gun just described with the envelope clocked externally and much more slowly. The direction and distance of the cannon is determined by the contents of the amplitude register.

A space gun or 'laser sound' can be produced using a 500 Hz tone (tone H0F in octave 4) and a sawtooth envelope with external envelope clock.

Space-ships

To suggest two swooping space-ships, two identical envelopes clocked at different rates and each with the invert bit set are used with two tones whose frequencies are varied between preset points at different rates.

Steam locomotive

To simulate the sound of a steam locomotive, coloured noise (15,6 kHz clock) and a repetitive triangular envelope are used. With the sound disabled, the envelope is clocked through to its peak, then with the sound enabled clocked slowly to its trough to give the sound of a gentle hiss of steam. After a pause, the envelope should be clocked continuously (beginning for example at 3,5 clocks per second and increasing to 28 clocks per second to suggest a locomotive gathering speed).

The sound of a locomotive's whistle can be produced using a 480 Hz tone (tone HFF in octave 3) mixed with noise.

Telephone ringer

The sound of a telephone ringer such as that fitted to a Trimphone® is simulated using an internally-clocked sawtooth envelope mixed with an appropriate tone (3,3 kHz, tone H6F in octave 6). The cadence is produced by enabling and disabling the sound (0,46 s enabled, 0,23 s disabled, 0,46 s enabled and 2,3 s disabled).

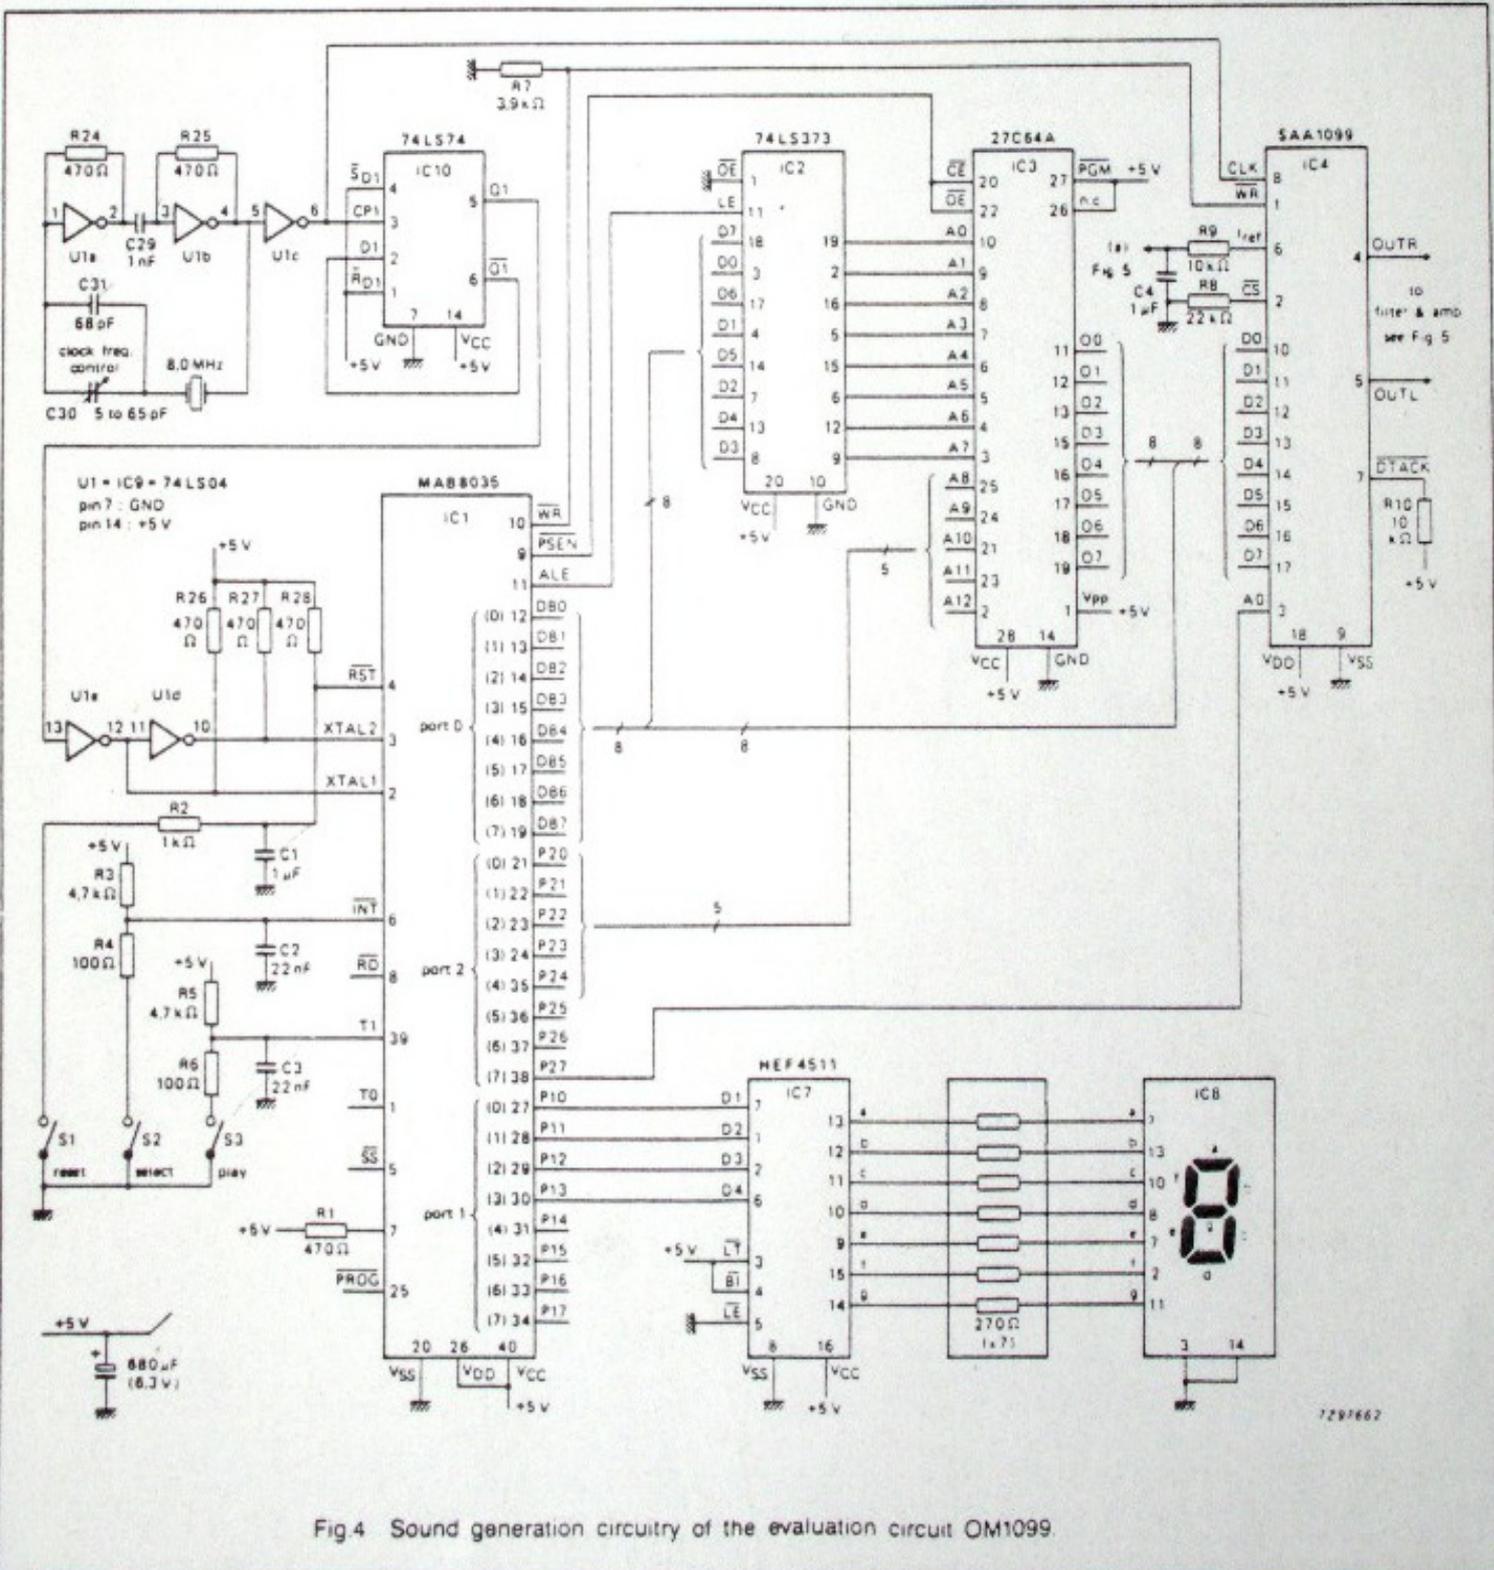


Fig.4 Sound generation circuitry of the evaluation circuit OM1099.

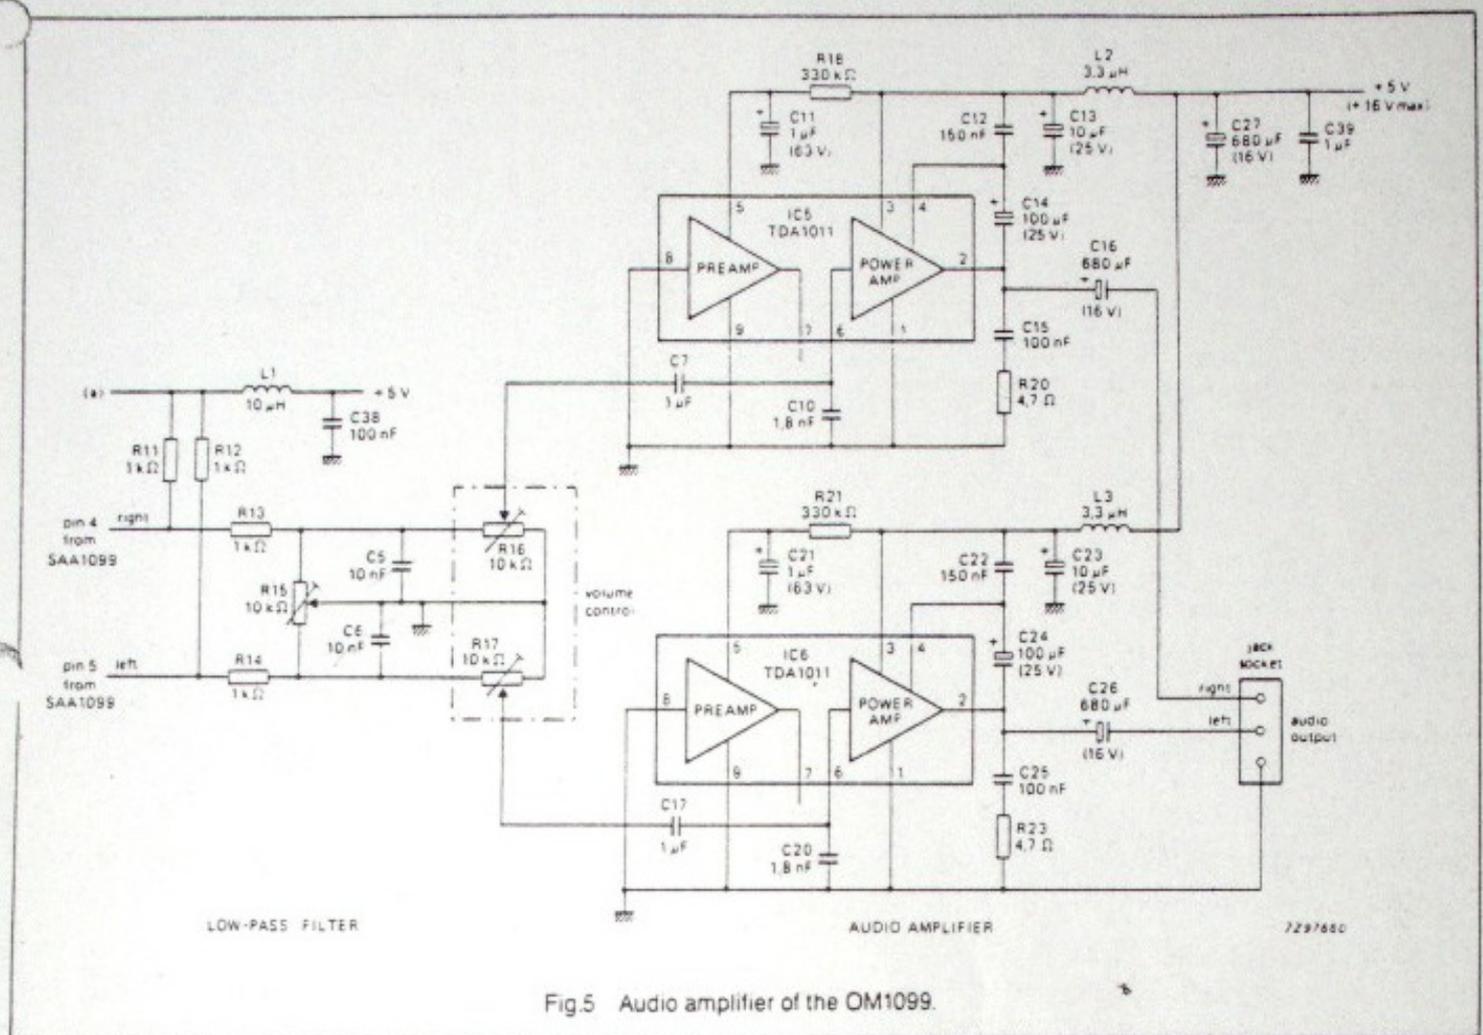


Fig.5 Audio amplifier of the OM1099.

EVALUATION CIRCUIT

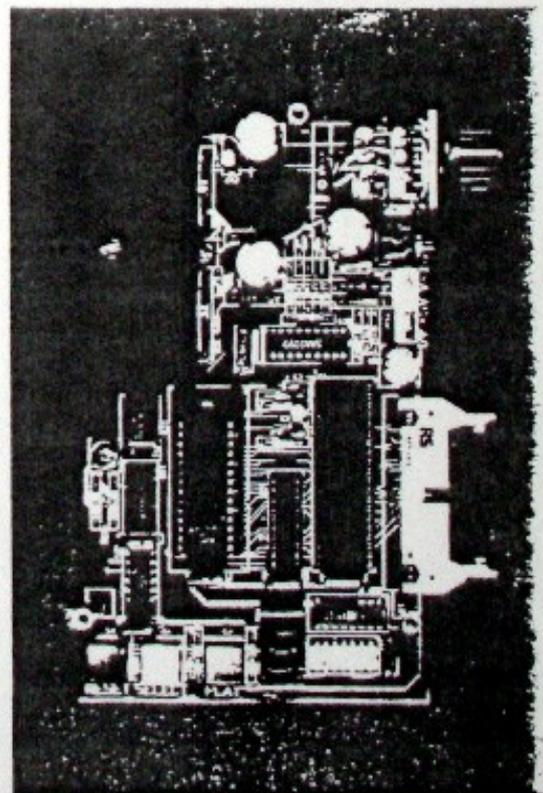
Figure 4 shows a circuit that is available for you to evaluate the SAA1099 in your application. The circuit which includes an integrated stereo power amplifier can be ordered using part number OM1099. A description of the circuit's hardware and software appears in Refs 4 and 5 respectively.

Data bytes held in an EPROM (the MAB8035HL microprocessor has no on-chip ROM) are sent via the microprocessor to the data inputs of the SAA1099 together with the A0 and \overline{WR} signals.

Pre-programmed sound modules can be selected using the 'select program' key which interrupts the current sound module being played. A display indicates the program selected. The 'play' key is used to play the selected program. Software support is available on request. For timing details, the reader is referred to the SAA1099 data sheet (Ref.2).

The standard MAB8035 requires a clock in the range 1 to 6 MHz. It is convenient to use the 8 MHz crystal oscillator for the SAA1099 with a divider to provide a 4 MHz clock for the microprocessor.

Since the SAA1099 can always accept data faster than the rate at which the MAB8035 can produce it in the circuit of Fig.4, the \overline{DTACK} output has not been used. However, when a faster processor such as the SCN68000 is used, \overline{DTACK} must be used to prevent data being missed by the SAA1099.



Laboratory model of the sound generation circuit shown in Fig.4.

After low-pass filtering (17 kHz, -3 dB), the output of the SAA1099 is amplified by two TDA1011 integrated amplifier circuits. Because the output of the SAA1099 is about 1.3 V (peak-to-peak), only the power amplifier sections of the TDA1011s are used, providing about 23 dB gain which is equivalent to an audio output of about 6 W per channel in a 4 Ω system with a supply voltage of 16 V. Even though the pre-amp sections are not used, some components (R18, R21, C11 and C21) are necessary to ensure correct biasing of the output stage. Some additional filtering has been added to reduce any high frequency coupling between the amplifiers and the SAA1099 via the supply lines.

SAA1099 AS A HOME COMPUTER PERIPHERAL

Figure 6 shows how the SAA1099 can be connected to a typical home computer - the BBC Model B. Although the following description is based on this arrangement, it is applicable for many home computers.

The audio amplifier is a TDA7050 stereo integrated amplifier providing about 140 mW output power per channel in a 16 Ω system with a 4.5 V supply. The TDA7050 requires few external components and is therefore suitable for low-cost applications.

The operation of the inputs to the SAA1099 is entirely software-controlled. Therefore, the software must provide suitable timing, for example, for the length of the write pulse and its timing relative to the data and changes of A0.

DTACK

Because the $\overline{\text{DTACK}}$ signal is too fast to be acknowledged by the BBC home computer, the $\overline{\text{DTACK}}$ pin is connected to one of the user ports (CB1), which is configured in a non-driving state (HIGH) so that when $\overline{\text{DTACK}}$ is LOW, the sink current is limited to about 3 mA.

Chip select

The chip select ($\overline{\text{CS}}$) line is connected to earth via a 22 k Ω resistor and therefore the SAA1099 is normally enabled. The $\overline{\text{CS}}$ input is also connected to CB2 on the microcomputer and this port or equivalent should be held LOW for normal operation and HIGH to inhibit the SAA1099.

