EXPERIMENTS IN COMPUTER CONTROLLED ACOUSTIC MODELLING (A STEP BACKWARDS ??) 8-AUG-1988

R.W.Berry

School of Music, University of Durham, Palace Green, Durham City, County Durham, DHI 3RL, ENGLAND Tel, 091-384 3224

ABSTRACT.

This paper describes a series of experiments in the modelling electronically of the acoustics of a variety of real instruments and a truly affordable computer system used to construct these. This work arose out of the desire to create electronic instruments very cheaply and in a reasonable amount of time but which nevertheless sounded like REAL instruments. As a musician working with and tiring of traditional synthesis methods, obtaining that peculiar quality of sound that only real instruments have was the goal, within limits set by time and expense.
The electronic modelling of the acoustics of real instruments seemed therefore to be a new research area of great promise, providing these limitations imposed by only a modest outlay could be tolerated.

TAPE OF SOUND EXAMPLES

The success or otherwise of this work can only really be judged on the quality of the sounds produced by the various models discussed here. An important part of the presentation of this paper therefore will be results demonstrated in the form of several taped audio examples.

MODEL BUILDING BLOCKS

In the following, block diagrams are used to illustrate model construction. The blocks are in the main familiar: multipliers, filters, envelope shapers etc. all requiring no comment. There is a however a new ingredient; the special delay line module shown below.

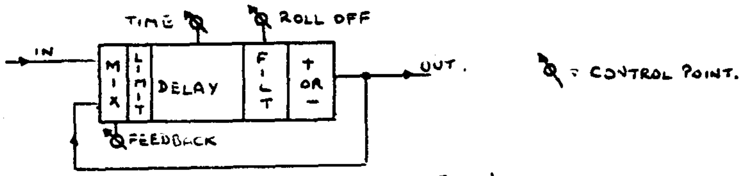

FIG. 1

It consists of a real time variable delay line, to which has been

added a two input mixer on the front, a built in soft limiting device (to be described in detail later),a variable first order low pass filter, and an optional output signal inverter. One input of the front end mixer is fed directly from the output. The overall feedback can be varied by the mixer.

SCOPE OF MODELLING EXPERIMENTS.

1) A simple assessment of the acoustics of several instruments. 2) A reasoned description of the models arrived at.
3) A recorded example of the sound produced by the models.

4) Comments on performance and experimentation using models and model variations together with any general conclusions.
The following instruments are covered : Plucked strings ,Gongs and bells, Trumpet, Saxophone, Clarinet, Flute.

PLUCKED STRINGS

There has been much written about the modelling of these. It seemed therefore to be a good testing ground for building block ideas, and general system proving.

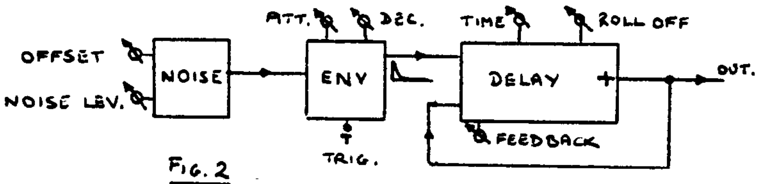

FIG. 2

The basic string model used is shown above. It has many similarities to the KARPLUS STRONG algorithm(l). Here though for reasons that will become apparent later, a noise source is used fed through an envelope shaper (adsr type) in order to produce a variety of transient stimulii for the delay line. Also JAFFE,SMITH (2). The noise source also has a d.c. offset associated with it.
This worked well producing some very pleasing acoustic sounding string tones. Varying the overall feedback level altered the string decay time. Reducing the filter roll off frequency had the effect of 'softening' the string material due to more rapid hf decay.

Varying the envelope parameters and d.c. offset produced different pick effects.

NEW EXPERIMENTS and DEVELOPMENT.

To this simple arrangement the addition of a few extra processes produced some interesting results.

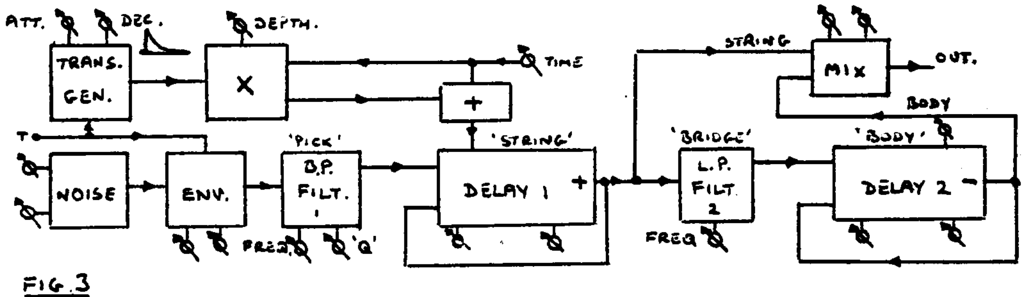

FIG.3

1) In the basic model (fig.2), the pitch remained absolutely constant. This sounded OK but too perfect. In the real world pulling on a string increases the tension. On release this quickly reduces to its original value. The quick pitch slur affects the initial transient conditions in an interesting way as the string tries to organise vibrations into a rapidly changing an then more constant harmonic series. Modelling this with the added transient generator on the delay line pitch control made quite a difference and added an extra degree of realism to the string sound. In effect it is a 'string tension' control.

2) To increase the variety of 'pick' effects a band pass filter was added to the noise source. This had the effect of plucking at differing distances from the 'bridge when the centre freq. was changed.

3) In fig.3 an attempt at adding a 'body' has been made. A second delay line with heavily filtered inverted feedback was used for this, in order to model a short stopped pipe or hollow box. Coupling between string and body was via a low pass filter to effect 'different bridge materials'. Some success was achieved with this and a 'Koto' type of instrument was produced. However it was found that really convincing body models could not be produced so simply.

(TAPE SOUND EXAMPLE 1 - Continuous string sound development
Short noise impulse. Delay feedback gradually added. Feedback loop low pass filter roll off frequency reduced, 'softening' the string. Initial quick pitch slur added. Finally simple body resonance added to end up with a hollow wooden bodied soft string instrument.]
[TAPE SOUND EXAMPLE 2 -'Koto type instrument'.]

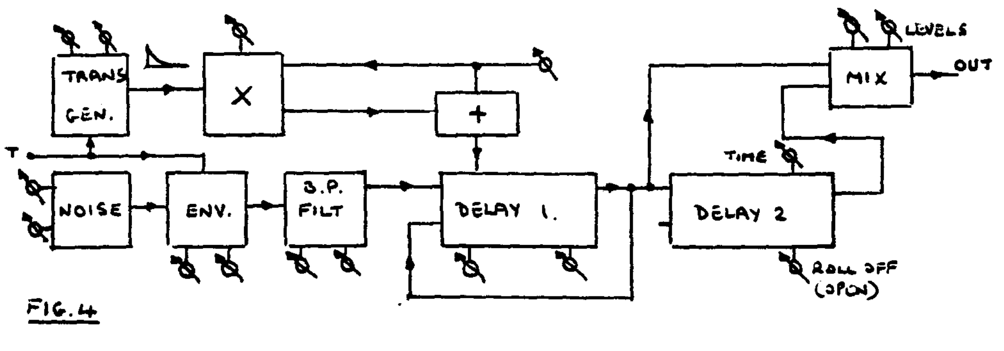

FIG.4

4)The above (fig.4) shows the body resonator removed and being replaced by a delay line in a comb filter configuration.
The experiment was to try to produce the same effect as that of an electronic pick up device such as can be found on an electric guitar. Here string harmonic nodes coincident with the pickup position are not reproduced, hence the comb effect.
Varying the delay time should have a similar effect to moving the pickup position in relation to the bridge, the shorter the delay the 'nearer the bridge'.

The results from this model produced some quite remarkable 'electronic string instrument' sounds. The most successful to date being a thoroughly convincing set of bass guitar sound models complete with the above software-movable 'pickup positions'.

[TAPE SOUND EXAMPLE 3 - Bass guitar. Various pick and string changes .Finally pickup position change.]

Having obtained some very encouraging acoustic results with plucked strings and in the process sorted out bugs in the system (to be described later) it was decided to explore some more new territory.

TWO DIMENSIONAL MODELS

The simple model of a string is essentially one dimensional, but what of more than one mode of resonance. A rectangular plate must have two at least, an interacting standing wave series down the length with one along the width. A large bell probably has many more.

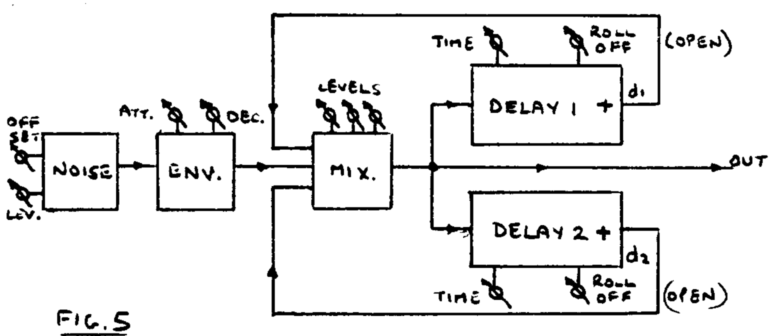

FIG.5

The above arrangement was tried as a simple experiment in two dimensions with spectacular results! When the delay lines' filters were fully open, the two delay lines set to unrelated delay times in the order of a few milliseconds and the overall loop feedback adjusted to just below the point of sustained oscillation, a noise impulse from the envelope generator produced the distinctive sound of a struck bright metallic plate or gong. The sound was rich real and extremely lifelike.

[TAPE SOUND EXAMPLE 4 - Sounds from simple two dimensional model gongs, bells etc.]

When the two differing delays are connected in parallel then resonant loops of say delay time d 1 and d2 are set up, plus a 'free' extra two loop delay (d 1 + d2); the sum of once round each loop. Each time round the signal is being split, delayed by different but fixed amounts and added together again. The original signal looped becomes two signals separated by (dl - d2) and so on. It seems therefore that the realistic bell sound is the result of a myriad of phase additions and subtractions as the initial noise impulse becomes PROGRESSIVELY more averaged into a number of

ordered harmonic series, some which of which have an additive and/or subtractive relationship to d 1 and d2.

FURTHER EXPERIMENTS.

Further experimentation resulted in the following:

1) Trying different relationships of the two delay times and overall feedback a wide variety of struck plates, bells, gongs, tin lids, small, large, resonant and dead were produced. In the extreme case of one large delay and one small, the sounds were more like that of struck metallic bars.
2) Reducing the filter roll off frequency softened the sound considerably into a more wooden quality. In the extreme, this became more drum like.

3)Using the inverting output of one or more of the delay lines, caused little change to the basic bell quality but there was a definite change into a more hollow sound presumably due to the lack of even harmonics in the sustained part of the sound.
4)Three delay lines in parallel were tried with surprisingly little difference.
5) Larger arrays of delay lines.

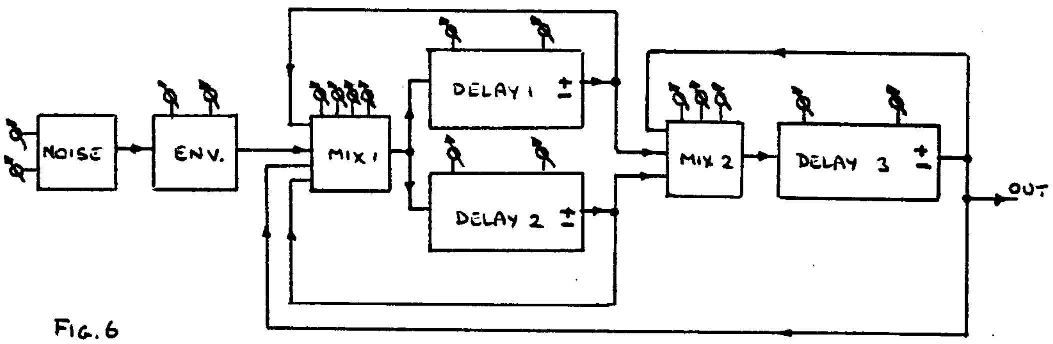

FIG. 6

Here an extra delay is added in series to the simple plate model in a 3 delay array. This was an attempt to model a larger more complex bell or gong. It was thought that perhaps a bell could have an interacting standing wave series round the circumference, down the length plus a lower frequency series caused by the deformation of the cylindrical shape (flan tin effect?) also. Delay 1,2 form the time delay system over the surface of the bell, and delay 3 the deformation series at a much lower frequency.

This view may be wide of the mark, but the model did indeed produce an extra depth to the sound, a more complex and rich sounding decay. With 3 delays, many feedback balance points, 3 filter settings, plus a large number of different 'strikes' to play with, an immense variety of gong and bell type instruments have been constructed with this patch. The computer is of course very adept at changing large numbers of parameters quickly so rapid transformation from one sound to another is possible here with good effect.
6) Stereo width.
For large instruments, the effect of listening to a tiny point of mono sound somewhere between two loudspeakers is not convincing. The addition of reverb can sometimes help, but often is not enough. The addition of a stereo multi-path chorus type of arrangement has been tried therefore as an alternative. 6 more delays (3 per side) slowly moving in time were used to form this. Stereo width was improved and the quality of the different bell sounds etc. was given some inner movement, which may be anyway when bells are rung.
An alternative approach tried was to take left and right mixes from different points in the patch. This was quite successful in many cases and useful for positioning certain parts of the model.
Finally of course it is possible to model a variety of bell 'strikers', but what if the striker part of the model was removed to be replaced by something else, like a human voice for example. Then a strange physically impossible 'instrument' will have been created. This was tried producing the interesting effect of a human voice actually exciting the resonant modes of the bell model.

[TAPE SOUND EXAMPLE 6 - Various bells and gongs, Human bell.]

------ PIPE INSTRUMENTS -----TRUMPET MODEL. Refs.(4),(5),(6)

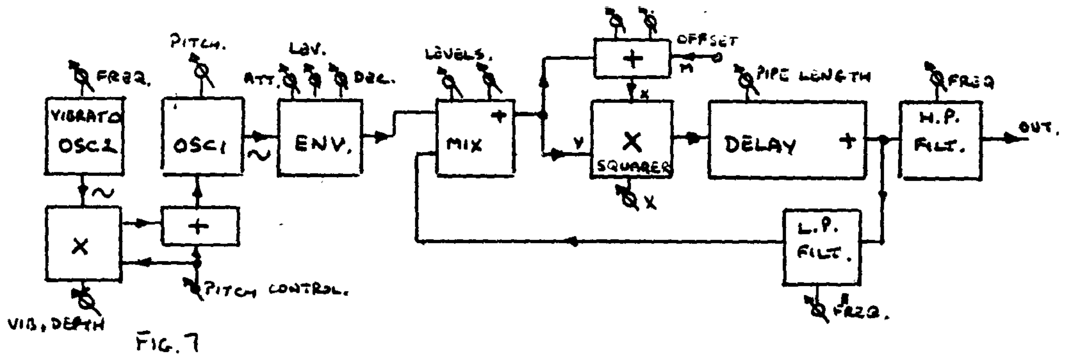

Fig.7

(This is a simplified arrangement from the original attempt, which had two delays one for each direction of sound down the pipe. The simplification seemed to make no difference to the sound).

This model was arrived at after much thought and difficulty. The main difficulty was in trying to visualise in the form of the connection of electronic modules, of the workings of the trumpet, in particular the mouthpiece mechanism, which seemed to be of vital importance. In trumpet playing, the player tenses the lips and then forces air trough them. When the lips are closed there is infinite impedance to the air flow. As the lips open the AREA of the aperture formed increases and the airflow impedance falls proportionally. There is a suggestion of a square law relationship here. There is soon reached a point when the lip tension overcomes the force of the outward rushing air and the whole thing collapses again, forming an oscillating system. Because of lip tension and the geometric relationship to the trumpet pipe the airflow impedance at the lips must be much greater than in the pipe. We seem therefore to have a resonant system driving hard another resonant system (the pipe), one a much higher impedance than the other to the airflow.

Considering the pipe, the air from the lips travels down the length where it meets a sudden pressure change on discharging from the end. This reflects back up the pipe where it meets the mouthpiece mechanism again. Conditions here are very changeable however depending on the phase of the incoming airflow. A further reflection will therefore be grossly affected by this and so on. In fig.7 these elements are simply represented electronically. As lips are basically rather soft, it was thought that the driving resonant system would produce a sineish sort of waveform, so a simple sine generator was used. This is fed via a mixer into a multiplying element. The multiplier has the same signal to both inputs producing

the squaring effect of the dynamic aperture. One input, however, has a d.c. offset added to make a central bias point for the 'lips' reflecting action.
Thus if X & Y are the multiplier's inputs, Z the output and M the manual offset, then

$$Z - Y(X + M)$$

Because the squarer is also in the delay loop, each time round the loop the signal gets repetitively squared. X**2,X**4,X**6 etc. The trumpet bell acts as a filter represented in the model by the delay's low pass filter and high pass on the output.
The model though simple worked well.

[TAPE SOUND EXAMPLE 7 - Trumpet sound and glissandi)

It was soon realised that what makes a real trumpet sound so wonderful is this repetitive non-linear action on successive delays down the tube. Layer on layer of these a few milliseconds apart. If the pitch of the driver is not absolutely constant, which in never is with a real trumpet player, then these successive layers arrive in all manner of levels and phases. Thus the repetitive non-linearity and layering produces the rich and extremely complex time varying sound that is so pleasing to the ears. With this in mind it is hard to imagine any other equally good method of synthesising the sound this effect produces.

Using the model some further experiments were carried out to test this viewpoint.
I) In fig.7 a second sinusoid is employed to provide crude vibrato in order to keep the trumpet sound alive with a continual variation. Removing this caused the liveness to collapse resulting in a completely static and electronic sound.
2) The driving sinusoid was replaced with other waveforms. This surprisingly did not affect the basic sound greatly. Whatever waveform was used a strong characteristic trumpet type sound appeared. In the extreme case a human voice was easily given trumpet characteristics.

[TAPE SOUND EXAMPLE 8 - The human trumpet!]

Any simple model must have shortcomings. The main one here is reproducing the actions of the player. A keyboard and a bit of vibrato is not really good enough. The keyboard alters the 'pipe length' together with the driving sine tone pitch which is a bit of a cheat. Also, driving tone to pipe is in the model a one way communication. This is not so with a real trumpet.

Some attempts have been made to try to improve on these. Firstly a joystick was used to vary the sine tone pitch, and secondly the keyboard was limited by the interface program to only respond to chromatic 'tube' extensions. Trying to pick up pipe harmonics like this was an interesting experience! Acoustically the trumpet effect was much more convincing, but playing presented some difficulties to say the least. I feel Stolzel and Bluhmel would have been amused)

The trumpet tone could be varied in several ways. Changing the amount of offset on the squarer, softened or hardened the tone. Changing the sinusoid amplitude, had a similar type of effect. Using the low pass filter in the delay line, had the effect of 'widening' the bore, again softening the sound.
By altering this model and its parameters, it should be possible to make simple models of other members of the brass family. To date only a tuba (a bass trumpet model) has been tried.

[TAPE SOUND EXAMPLE 9 - Tuba.]

SAXOPHONE

After obtaining some success with trumpet sounds, it was thought that the modelling of the non-linear mouthpiece action on layers of repetitive delays down the instrument tube could be applied to other pipe instruments just as effectively. It was therefore decided to try a model of a reed instrument.
Again much difficulty was encountered in solving the basic problem of the mouthpiece action. Many failed attempts followed until the following model was arrived at.

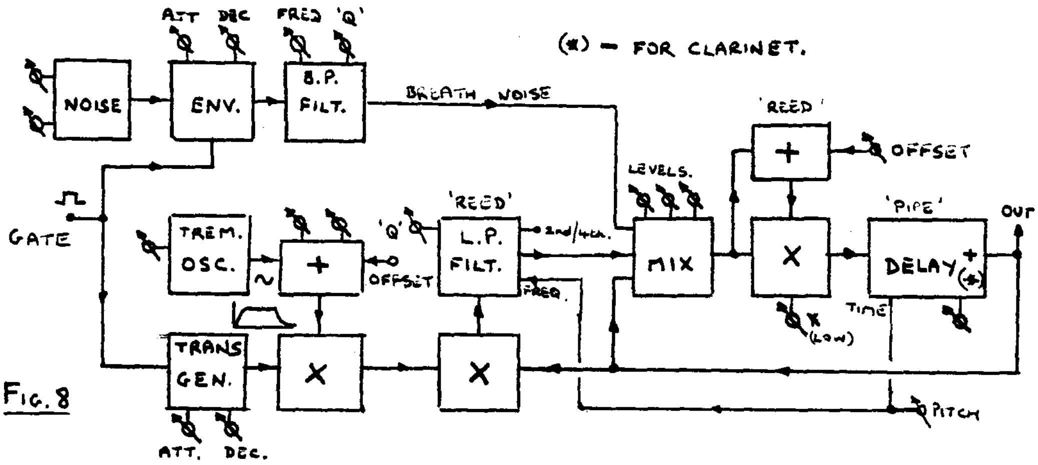

FIG. 8

Here the hard driving action of the trumpet 'mouthpiece' has been replaced by a reed model closely coupled to an overall resonant system. For the valve action, it was thought that the reed would still vary the area of the mouthpiece aperture as before, perhaps producing the same kind of squaring action but this time mainly under the influence of the pipe resonance. It was felt that because of reed geometry and the lack of any equivalent lip tension effect the reed airflow impedance would be much closer to the pipe than in the trumpet. It was also though that the reed being a flexible shape and made of softish material it would tend to be sluggish to respond to high frequencies. This would produce its own additional 'fixed' frequency dependent phase delay to the pipe’s resonant loop. A variable low pass filter 2nd./4th. order with resonance (Q) control was all that was used to model this and any damped resonance the 'reed-lip' arrangement might posses. If the system is to oscillate, when producing a sustained note, then like all oscillators, there must be some form of amplitude limiting. In reed instruments it seems this is a function of the limiting spring action of the reed in the mouthpiece. It will sooner or later reach its limit of elasticity or even collide with the body. Electrically in the model a soft limiting function was built in to the delay line. We can see that saxophones have conical bores while clarinets for example have straight cylindrical bores. This is all very well but with electronics unfortunately it is very difficult to construct a conical delay line! All we have is either positive feedback or negative. Strangely though both forms at unity gain will readily oscillate: when positive, at all harmonics of the fundamental: when negative, at odd harmonics but an octave lower. In the saxophone the conical bore apparently acts as an open pipe resonant at all harmonics so in the model a positive loop was used as the nearest practical approach to this.

To sum up the simple model has 4 basic elements:

A vibrating 'reed' having two effects: 1)a non-linear distorting action on the incoming airflow and subsequent pipe reflections, and 2) loop phase modifications caused by its own damped resonant inertia plus: 3) a symmetrical limiter capable of producing odd harmonics, 4) a delay line resonant system repetitively filtering into an all harmonic series.

In the model, a second multiplying element in series with and before the (reed) low pass filter is used to control the overall gain of the system. Just below unity, results in a breathy pipe sound, just above, results in the normal saxophone sound. The remaining multiplier introduces a small amount of tremolo to keep the sound alive. The filter freq. ('lip-reed position') is varied also with pitch. The simple model in fig.8 produced an very realistic sound with all the lifelike qualities experienced with the trumpet model.
A small quantity of band limited noise has been added however just to add a touch of breathyness to finish the model off with

[TAPE EXAMPLE 10- Saxophone.

When performing on the saxophone model, several properties were noted:
1)Like the trumpet model, if the system was allowed to become absolutely static, the sound became lifeless.
2)Much less squaring effect was required. $Z = Y(X + M)$ for the multiplier had $X << M$.
3)The model tended to 'OVERBLOW' easily [Ref tape ex.91.
4) Altering the various parameters, produced wide ranging effects on the sound. The reed filter settings in particular had interesting effects. Setting to a low roll off frequency the saxophone sound was wooly, with higher notes sounding strained while tuning was seriously affected because of the frequency dependent added phase delay caused by this device. Set much higher, the 'sax.' sounded bright and crisp but unintentional 'overblowing' became more of a problem. The

very extreme high notes still sounded strained as the sound quality gradually changed with increasing pitch. With a large number of parameters to play with, many varieties of saxophone have been produced, some with a distinctly oriental flavour.

(TAPE EXAMPLE 1 1 - Oriental saxophone.)

FURTHER EXPERIMENTS

Some brief experiments were carried out with the basic elements of the saxophone model to explore the possibilities of modelling other members of the woodwind family.

CLARINET

A major difference in the sound of a clarinet is that this instrument behaves as a stopped pipe, and therefore produces an overtone series of predominantly odd harmonics. In the basic saxophone model, it follows that reversing the overall feedback from positive to negative should have the same effect.
This was tried with promising results. The sound quality once again very real and lifelike, being very much of the clarinet family. Accidental 'overblowing ' was something of a problem here also with certain bright sounding parameters. The interesting phenomena here is the contrasting and conflicting effects of a 'reed' producing even harmonics as it squares repetitively round a delay loop trying to filter this into a series of odd harmonics.
Much work has already been done on clarinets by SMITH and others (3).

[TAPE EXAMPLE 12 - Clarinet family.]

FLUTE

By reducing one of the multipliers inputs the non-linear overtone series was gradually reduced until a flute like sound was produced. Once again the sound had that real quality about it. It was quite easy for the computer to control the model to perform a gradual transformation from saxophone sound to flute sound.

[TAPE SOUND EXAMPLE 13 -Transformation from saxophone to flute on one note of a musical phrase.

----- SYSTEM HARDWARE AND SOFTWARE ----

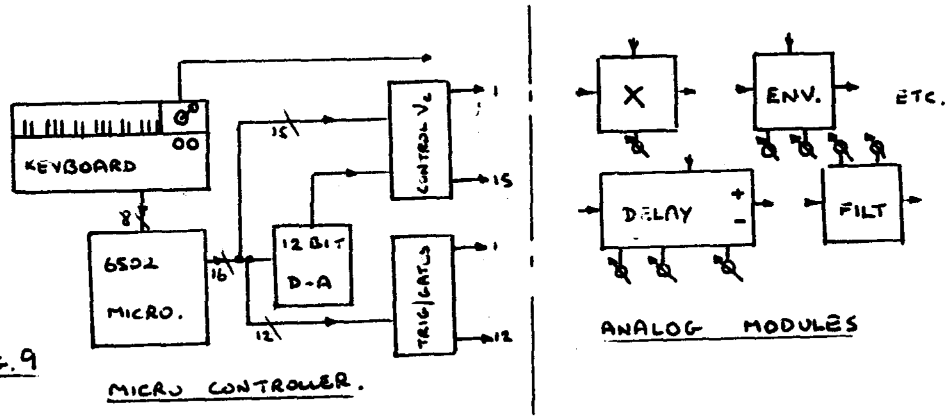

Fig. 9

SYSTEM HARDWARE.

The hardware used to build and control the above models is by factors of time and finance extremely modest. A 6502 based one board microprocessor (SYM 1.1 by Synertec) is interfaced to; on the input side, a keyboard with 8 bit digital output and on the output, side to a 12 bit d-a converter, subsequently multiplexed to form 15 voltage controlled outputs plus 12 trigger/gate outputs. The rest of the *hardware* is analog.

In addition to a range of conventional analog synthesis modules such as envelope shapers, voltage controlled filters etc., there are provided, adders in the form of mixers with normal and inverted outputs, multipliers in the form of 4 quadrant voltage controlled amplifiers but with an extra control input; a positive manual offset which forms a gain bias point, and to date, 6 voltage controlled analog delay line modules. These are all used to build a specialised form of the old fashioned ANALOG COMPUTER. So here the system basically consists of a modest digital computer controlling an analog computer. A not ideal but nevertheless a very cost/time effective solution capable of producing some excellent results. This is partly due to the analog side having the advantage of very rapid and easy user interfacing in real time, which in this case makes up for any lack of precision.

DELAY LINE MODULES

The delay lines are based on the ccd TDA 1022, TDA 1097 ccd (bucket brigade) devices. [The TDA 1022 has 512 stages, the TDA1097 has 1536 stages].

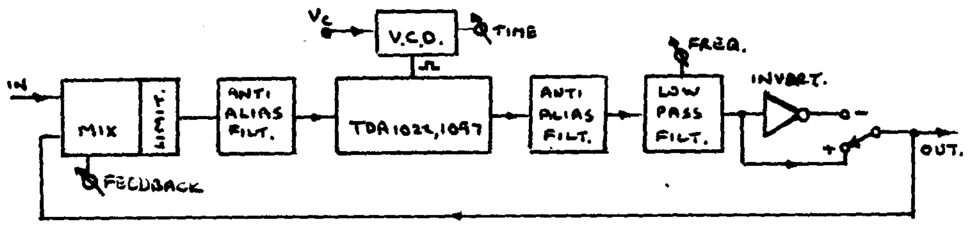

FIG.10

The block diagram shows this in more detail. The problems encountered in getting these devices to function as required were considerable.
Firstly steep cut anti-aliasing were needed to prevent the usual inter-modulation effects.
Secondly a reliable and virtually drift free clock VCO capable of good linearity up to the MHz region had to be devised.
Thirdly the gain of the above ccd devices is unfortunately not quite constant with clock frequency or temperature variations. When loops of very near unity gain are being employed this is important. In near unity gain feedback situations they can be noisy, especially at lower clock frequencies.
Fourthly, delay range is limited. About 10 times (50KHz-500KHz clock) is available with good fidelity. About 30 is generally usable (100Hz – 3Kz resonance fundamental for a TDA1022).
Fifthly, under conditions of sustained oscillation as in the saxophone model, an amplitude limiting device is needed. A simple zener diode/resistor was used to give symmetrical limiting quite effectively.

A good deal of development time was spent overcoming these problems. Most were overcome leaving the noise and limited clock range unsolved. Noise is fortunately not always a serious problem however. When the inbuilt low pass filter is being used this tends to filter off the reiterative noise build up. When noise is a nuisance, a noise gate is employed on the model's output.

SYSTEM SOFTWARE

Software is almost all in real time. Only minor editing functions are not real time. It consists of a modest package of interface and sequencing routines. Data from the digital output of the keyboard can be intercepted and used in a wide variety of ways to control the system or edit in real time the sequencing routines. Sequencing routines consist of simultaneous sequencing of all 15 voltage control outputs, plus trigger/gate outputs, in a very simple but open ended arrangement. Any sequence output can access the stack of temporary variables if desired, thereby gaining access to the current parameters of any other sequence. Thus for example one sequence of notes can be played using the rhythm pattern of another of differing length and maybe tempo also.
Keyboard interface routines can for example give different musical scales or convert to an odd scale such as for simulating valve operations on 'brass ' instruments, or turn the keyboard into a typewriter keyboard with 'cursor' control when accessing sequencing arrays of data etc.

FUTURE WORK

Having achieved some success in obtaining that 'real instrument sound' by very modest means and in the process finding some pointers as to the reason why some instruments sound the way they do, future work will therefore involve improving on these crude and often oversimplified models, as well as exploring new ones. Also the building of strange hybrid models or totally imaginary ones has now become an exciting possibility through this work. Much work has to be done on the hardware and software too as there is obviously a great deal of room for improvement here.

REFERENCES

(1) Kevin Karplus and Alex Strong - Digital synthesis of plucked string and drum timbres. Computer Music Journal Vol.7 No.2. (2) David A Jaffe, Julius 0 Smith - Extention of Karplus-Strong algorithm. Computer Music Journal Vol.2 No.2.
(3) Julius 0 Smith - Efficient simulation of Reed-bore and bow string mechanisms. ICMC Proceedings 1986.
(4) John Backus -Input impedance curves for the brass instruments. J. Acoust. Soc. Am., Vol 60, No.2, August 1976.
(5) John Backus and T.C. Hundley - Harmonic generation in the trumpet. J. Acoust. Soc. Am., Vol. 49,2ii 1971.
(6) Stephen Elliot and `John Bowsher - Input transfer response of brass wind instruments. J. Acoust. Soc. Am., Vol. 72,6 Dec. 1982.

R.W.BERRY 9-AUG-1988