A Toolkit For Chaos

[Original March 1991. This version reconstituted January 2019.]

Electronic engineers are used to dealing with devices which display various forms of electronic non-linearity. It is virtually impossible to find a radio or microwave system which does not contain a non-linear resistance (i.e. an ordinary diode). Varactor diodes are also widely used as capacitors whose value may be controlled by an applied voltage. Yet outside of a few well-defined applications, non-linearity has generally been an unwelcome guest. One of the main tasks of engineers is often to suppress unwanted non-linear behaviour in electronic systems. Circuits which display non-linearity or complex, unpredictable oscillations are likely to be thrown in the bin as 'useless'.

Chaos has become a major growth area in science, and the behaviour of non-linear systems is a hot topic for study. For many academic scientists, the fascination of experimenting with non-linear systems and discovering amazing complexity can be an end in itself. A practical engineer is likely to be more sceptical and ask, "What is it good for?"

It will probably be some time before we can really know what the full range of benefits might be from understanding chaotic systems. Until now, one of the most obvious advantages is negative – "Know your enemy". If you can define just what produces chaotic behaviour in a given system then you are half-way to knowing how to avoid it! This can be very valuable in preventing system failures. In the future, more positive results may be expected, although it is a bit soon to predict just what they will be.

At present, the most immediate positive results look like being in the direction of improving the performance of existing non-linear applications – e.g. frequency multipliers, mixers, etc. Beyond that, as yet, there are few obviously new applications, although work is beginning on some ideas. For example, the use of non-linear systems to produce 'chaotic oscillators'. These may prove useful as broadband sources, replacing noise sources for some applications, providing spread spectrum signals for communications or radar, and complex deterministic signals for signal encryption. Given these developments it is probably time for engineers to begin familiarising themselves with the new toolkit of chaos.

A diode is the simplest form of non-linear device. In general, diodes have a conductivity and reactance that are voltage dependent. In principle we can imagine building any sort of non-linear system from a suitable set of diodes along with some passive linear components, and amplifiers to provide any gain we require.

The main properties of an ordinary diode can be defined in terms of how the current passing through the device depends upon the applied voltage. The non-linear properties of a varactor are less obvious, but can be explained by comparison with an ordinary capacitor. The voltage, V, of a normal capacitor is proportional to the charge, Q, it holds. If we start with a discharged capacitor and add an amount of charge, Q, the capacitor acquires a voltage

$$V = Q / C \qquad \text{... (1)}$$

To change the voltage by a small amount, dV, we must add an extra amount of charge

$$dQ = C . dV \qquad \text{... (2)}$$

We may now define the capacitance using expression 1 or 2. Either way we will get the same value, irrespective of our choice of V.

If we repeat the same process with a Varactor we find that the voltage produced is not simply proportional to the charge. To change the diode voltage by dV we must now add a charge

$$dQ = C\{V\}.dV \qquad \text{... (3)}$$

where $C(V)$ depends upon the device voltage, V. To take a discharged varactor and charge it to a voltage, V, we need a charge

$$Q = C'\{V\}.V \qquad \text{... (4)}$$

where

$$C'\{V\} = \left(\frac{1}{V}\right)\int_0^V C\{V\}.dV \qquad \text{... (5)}$$

Unlike for an ordinary capacitor, expressions 3 and 4 *won't* now usually provide the same capacitance value for any randomly chosen value of V.

This situation is analogous to the resistance of a normal diode which may be given either as a static value $R = V/I$ or as a dynamic or slope value $r = dV/dI$. Here we can define a static capacitance, $C\{V\}$, or a dynamic one, $C'\{V\}$.

In general, these devices are used to carry out fairly basic functions like frequency conversion and tuning. Their analysis as presented in most textbooks seems straightforward. Yet, surprisingly, even quite simple arrangements containing these devices can behave in ways which are very complex - even chaotic.

Ideally, we can explore the behaviour of these systems by building circuits and watching what they do as we vary circuit element values or applied signals. Unfortunately, this approach isn't always practical. For example, millimetre-wave engineers have to deal with signal frequencies of 70 - 100 GHz or more. At these frequencies it can be quite difficult to define the electrical properties of any circuit you have built. Measuring the details of any output signals can - unless they are fairly simple - also prove rather difficult and expensive.

In some cases we are also faced with the question, *"What would this new sort of device do if we were to make one"*. Solid state engineers can now manufacture a wide range of Low Dimensional Structure Devices which - although they only have two terminals like the humble 1N4148 - can have a wide range of complicated electrical properties. Naturally, the people who have to grow these devices would like some advance idea of which types will be most useful.

When asked 'what if' questions of this type, most engineers and scientists resort to some kind of computer model, based upon a algorithm running on a digital computer. This technique is a very powerful one, but is does present us with some problems of its own. These can be roughly divided into three categories:- i) is the program right? ii) is the computer big and fast enough? iii) is there a problem with computational errors?

All being well, we can hope to produce a program which represents the system we wish to model. The speed and size of the computer is likely to be determined by the available funds. Computational errors are an unavoidable risk, and it can be very difficult to decide whether the results we see are being affected by rounding errors or by choosing finite time steps between calculations which are too large.

The problem of finite computational precision is particularly important when examining the behaviour of circuits which may become chaotic. Imagine a non-linear circuit or system whose state at any moment can be defined by the values of, say, just four quantities (perhaps two voltages & two currents). The properties of the circuit will determine how these quantities change with time in response to some given input signal. If the process is periodic, after some appropriate time the quantities will return to a set of values which they have already had. The process then repeats its previous behaviour over and over again. A chaotic process cannot, by definition, be periodic. Hence if the process is chaotic the quantities which describe the state of the circuit will never repeat any set of values.

A digital computer model of this circuit in operation will store information about the state of the process in terms of four values. Each of these will be recorded in the computer as a finite number of bits. For standard IEEE floating point precision a number is recorded as four 8-bit bytes - i.e. as 32 bits. Hence a set of four variable values will be stored as 4x32=128 bits.

Now a finite number of bits, N, can only record $M = 2^N$ possible arrangements. So far as the computer is concerned, the simple digital representation we have described above must always be in one of the 2128 states which the computer model can distinguish. The computer can only calculate and recognise this finite number of states. As a result any process whose state we represent in terms of N bits must repeat itself after - at most - 2^N recalculations since the computer will have run out of 'fresh' patterns of bits.

We can, of course, use extended precision values to increase N. More complex circuits and systems will also require more than four quantities to describe their state. Hence we can produce computer models where M is very large indeed - but in practice it can never be infinite. This problem is made worse by the possibility that rounding errors during a calculation may force an 'accidental' repeat of a set of values after far less than M updates of the state of the model. As processes approach becoming chaotic it becomes increasingly difficult to assess their behaviour reliably using a computer model.

In general, the risk of computational error can be reduced by buying a bigger, faster computer and performing the digital calculations with extended precision and smaller time steps. Yet this does not banish the risk, although it does increase the cost. However, in order to avoid time step (sampling) and precision (quantisation) errors we can adopt a very different technique and use an *analog* model of a non-linear device. This can then be used to cross-check the predictions of a digital computer model, or used by itself for the sake of cheapness and speed.

In general, we can define two basic types of passive non-linear device – a non-linear resistance and a non-linear reactance. These can then be treated as building blocks to assemble other non-linear systems. To see how they can be used to explore the behaviour of complex or chaotic processes we can concentrate here on just one type in the form of an artificial varactor. Figure 1a shows a circuit which behaves as a voltage dependent capacitance.

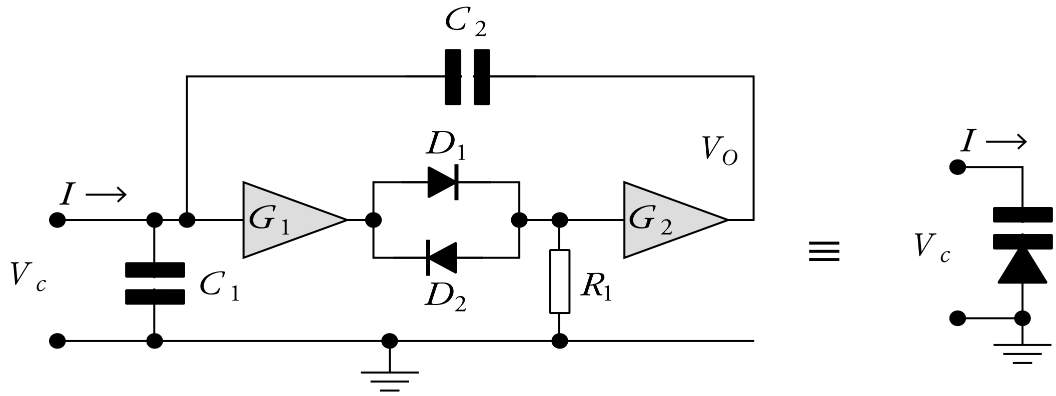

Fig1a. Circuit used to create a capacitance that varies with the applied voltage – i.e. equivalent to a varactor.

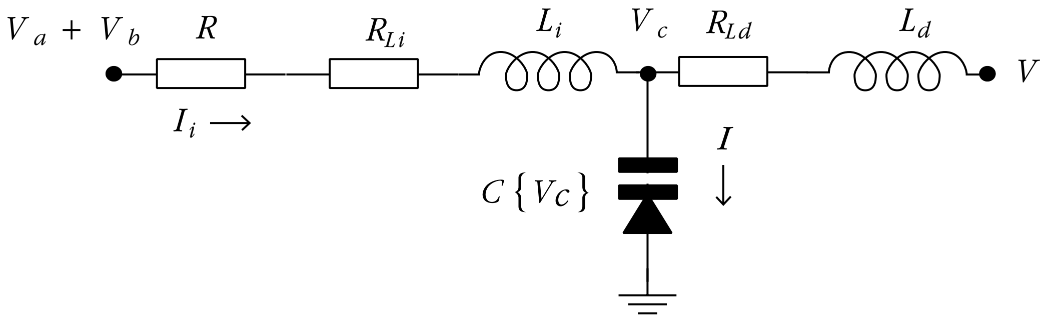

Fig1b. Embedding ‘T’ network used with artificial varactor.

Fig 1. Analog non-linear capacitance and experimental network.

For simplicity, we can assume that the gains, G_1 & G_2, of the two amplifiers are both unity – i.e. they act as buffers. In practice these gains may be chosen to have some other value (they may also, of course, be voltage dependent). In this way the voltage dependence of the capacitance presented at its terminals can be chosen.

The diodes, D_1 & D_2, may be assumed to be ordinary Silicon signal diodes (e.g. 1N4148's or similar). When $|V| < 0{\cdot}5$ Volts the diodes will have a high effective resistance and V_0 will be held at about 0 Volts. In this situation any input current, I, must charge both C_1 & C_2 in order to alter the voltage, V , Hence for low input voltages the circuit behaves as if it were a capacitance,

$$C\{V \approx 0\} \approx C_1 + C_2 \qquad \text{...(6)}$$

When $|V| \gg 0{\cdot}5$ Volts any change in V will produce a corresponding alteration in V_0. i.e. no charge need flow into or out of C_2. Hence in this situation the circuit behaves as a dynamic capacitance

$$C\{|V| \gg 0{\cdot}5\} = C_1 \qquad \text{...(7)}$$

At intermediate voltages, the effective capacitance depends upon the behaviour of the diodes and the chosen value of R_1. As $|V|$ increases from zero, through 0·5 Volts, the effective capacitance tends to fall smoothly from $C_1 + C_2$ to C_1.

To compare the digital and analog approaches a circuit of the form of 1a was built using a pair of TL081 operational amplifiers. D_1 & D_2 were a pair of 1N4148's, R_1 = 5kOhms (two 10k's in parallel). The nominal capacitor values used were $C_1 = 0{\cdot}1\mu F$, and $C_2 = 1\mu F$.

The $C\{V\}$ behaviour of the circuit was determined by connecting it in series with a known inductor and then noting how the series resonance frequency varied with an applied d.c. voltage. From these measurements it was found that the dynamic capacitance behaviour was essentially of the form

$$C\{V\} = C_1 + \frac{C_2}{1 + (\alpha V)^4} \quad \text{...(8)}$$

where $C_1 = 0{\cdot}14\mu F$, $C_2 = 0{\cdot}94\mu F$, and the coefficient, $\alpha = 2{\cdot}2$ per Volt.

This artificial varactor was then connected to a 'T-network' consisting of a pair of inductors and a capacitor as illustrated in Fig 1b. The properties of this fairly simple system could then be examined as a function of the applied signal voltage.

The actual network used consisted of two inductors whose values were L_1 = 3·24 mH, and L_d = 3·38 mH. The d.c. resistances of these inductors were R_{Li} = 6 Ohms and R_{Ld} = 4 Ohms. These particular components were chosen purely because they were to hand - i.e. their choice was 'at random'. Their order of magnitude was selected to ensure that any affects of interest would be likely to arise in the 100 Hz - 10 kHz frequency range. This would make them easy to observe on a normal oscilloscope and be well within the bandwidth of the op-amps. An input sinewave, V_a, and nominal d.c. bias level, V_b were coupled into the circuit via a series resistor, R = 100 Ohms.

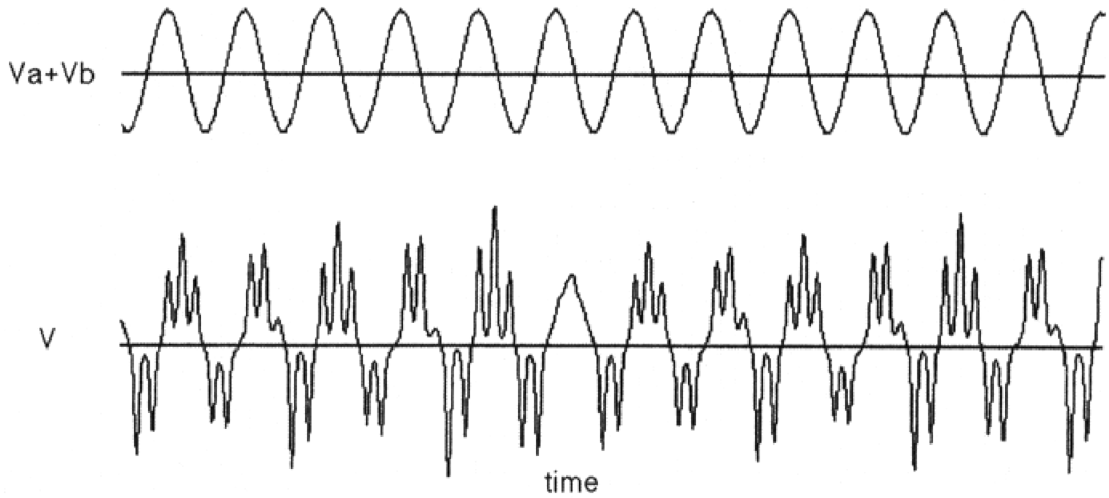

Fig 2. Typical input and output voltages vs time.

A computer program was also produced to act as a comparative digital model. This was written in **BBC Basic V** and run on an *Acorn Achimedes 410* machine fitted with an Arm3 processor. Although not, by any means, a 'super-mini', the Arm-3 is a 32-bit, 30MHz, RISC processor with an on-chip cache memory. It is therefore well suited to this sort of algorithm involving rapid iterations of a simple set of difference equations. In this case the set of equations used were

```
REPEAT
Vi=Va*SIN(W*t)+Vb
V+=I*dt/FNcap
Vc+=(Ii-I)*dt/C
I+=(Vc-RLd*I-V)*dt/Ld
Ii+=(Vi-Ii*(R+RLi)-Vc)*dt/Li
t+=dt
UNTIL FALSE
:
DEFFNcap    =C1+C2/(1+(2.2*V)^4)
```

where V_a is the half-peak magnitude of a nominal input sinewave at the (angular) frequency, W. V_b is the d.c. bias level. dt is the time-step increment (here normally preset to a 1/256'th of a cycle at the chosen input driving sinewave frequency). **BBC Basic V** is a convenient language for this type of computation for two reasons; firstly, floating point values are stored as four bytes of mantissa and one of exponent - i.e. more accurately than standard 32bit IEEE floating point; secondly, the provision of the unary addition operator, "+=", which speeds up differential calculations.

The possible effects of finite precision and time-step value were tested by compiling a double precision (to IEEE standard) version using the **ABC Basic** compiler and running this with smaller *dt* values. This gave results visibly identical to the normal interpreted basic program, as viewed by plotting V_d against *t* on the computer display. When plotting a useful output display, the computer program ran at an actual drive waveform rate of just a few cycles per second. The analog system operated in real time and provided essentially identical behaviour with drive frequencies in the kHz range. Hence it was found that the analog system operated around 1000 times faster than the computer program.

The correspondence between analog and digital models allowed a rapid exploration of the system behaviour using the analog system. Detailed analysis of interesting regions of operation could then be carried out using the computer model without having to spend considerable time checking a wide range of conditions which produce less interesting results.

An example of interesting complexity was found, using the analog system, to occur for a drive frequency of around 1300 Hz at V_a =3·5 Volts (half-peak) combined with a d.c. bias level of around 0·25 Volts. Figure 2 illustrates some typical results for the time variation of *V* in this region. This figure was obtained from the computer model, but identical results could be obtained using an oscilloscope connected to the analog circuit.

We are accustomed to examining signal versus time variations on the screen of an oscilloscope, but in this case such a display does not tell us very much about what is happening. Fortunately, various methods have been developed for examining complex and chaotic processes. We can adopt one of these – a ‘phase’ plot – to explore the behaviour of our non-linear system.

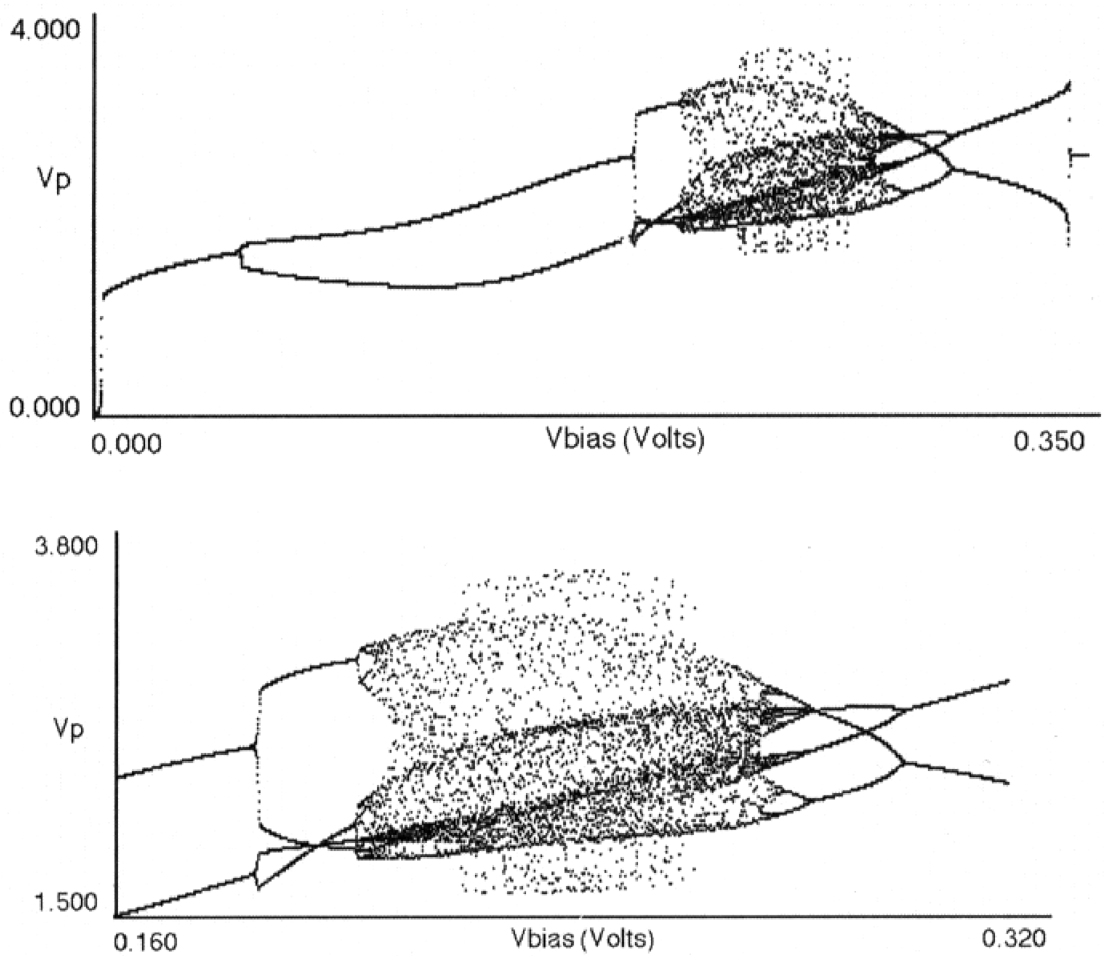

Figs 3a & 3b. Examples of phase plots showing pattern of bifurcations.

The concept of a phase plot is fairly simple. We drive a non-linear arrangement with a periodic signal and note the value of its output at a particular phase of each drive cycle. We then alter one of the system parameters and repeat the measurement. In this case, we drive the system with a sine wave

$$V = V_a \sin\{\alpha t\} \qquad \text{... (9)}$$

and note the output voltage (i.e. the voltage on the non-linear capacitor), V_p, at some chosen phase angle, ψ, of each cycle. We can now plot how V_p varies with, say, the d.c. bias voltage.

[N.B. See the appendix for a more detailed explanation of the use of a phase plot.]

Figures 3a & 3b show some typical phase plot results for the circuit we are examining. Figure 3b is a ‘close up’ view of the most interesting section of the phase plot results. At low d.c. bias levels, the output appears simply to be a distorted version of the input. Each cycle is identical to the next. If we were to just look at the output we would conclude that its period was identical with that of the input drive signal.

When the d.c. voltage reaches about 0·05 Volts the situation suddenly changes. Now the output level alternates from input cycle to cycle, back and forth between two values. The output signal now appears to have a period which is double that of the input. The phase plot shows a *bifurcation*, a place where one line splits into two. Another bifurcation occurs when the bias voltages reaches about 0.18 Volts. V_p then cycles around four values and the basic period of the output appears to be four times that of the input. These bifurcations are also referred to as 'period doublings' for obvious reasons.

As we continue, the change in bias required to produce extra bifurcations rapidly decreases and we see a 'cascade' of bifurcations. The basic period rises rapidly and, in principle, may go to infinity. If that occurs the system becomes 'chaotic' and the output waveform *never* produces any two cycles which are identical. In this case the number of bifurcations appears to reach a maximum value at a d.c. bias level around 0·25 Volts. Any further increase in bias causes a reduction in the periodicity. A system which has undergone a finite number of bifurcations can be said to show 'semi-chaotic' behaviour. Its output period may be much longer than that of the 'driving' input sinewave, but it eventually will repeat.

We could continue to examine the behaviour of this particular system in more detail, exploring the effects of altering other variables, and applying other analysis methods. However, the purpose of looking at the complex behaviour of this particular circuit was to demonstrate the principle that we may easily and cheaply construct an analog system which may replace a digital computer model. To understand how this approach may be extended we can use the example of an artificial non-linear resistance as shown in Figure 4. This represents a rough approximation to a Gunn diode.

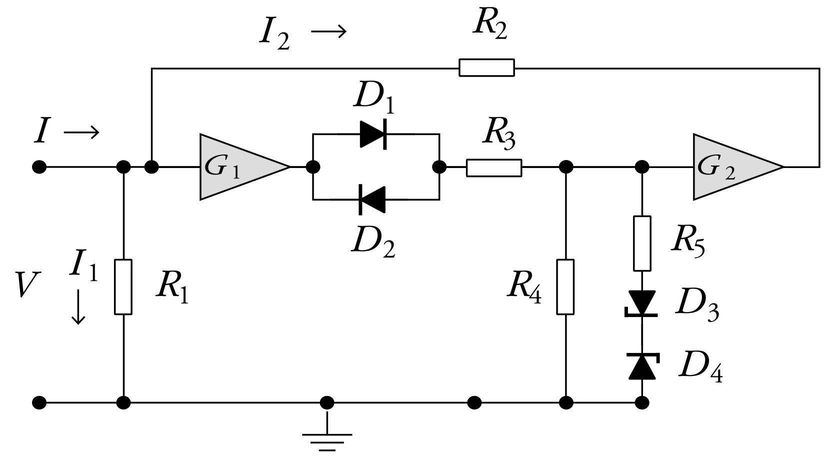

Fig4a. Circuit that mimics a Gunn Diode

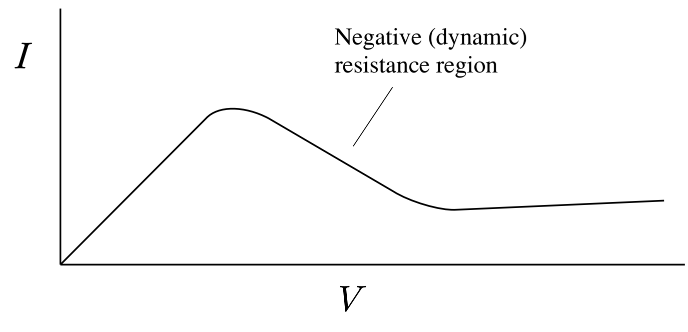

Fig4b. Resulting Voltage/Current relationship

Gunn diodes are widely used at microwave and millimetre-wave frequencies as oscillators and amplifiers. i.e. although they are two-lead devices, they can provide gain. This is because they exhibit the property of negative resistance. (Strictly speaking, this is more correctly called negative differential conductance, but the term negative resistance is commonly used.)

Consider a potential divider made of two resistors, R_1 & R_2 connected in series. When we apply an overall voltage, V, the voltage across, say, R_2, will be

$$V_2 = \frac{R_2}{R_1 + R_2} \qquad \text{... (10)}$$

Normally we will expect $R_1 + R_2 \geqslant R_2$, and hence that $V_2 \leqslant V$. If, however, R_1 or R_2 is *negative* we may find that $R_1 + R_2 < R_2$, which means that $V_2 \geqslant V$! The usual potential 'divider' will then act as an amplifier.

Clearly we can't get 'owt for nowt. Any power which is added to an amplified signal must come from somewhere. As a result, for a two-terminal device like a Gunn diode we find that the static resistance must always be positive, but the dynamic (slope) resistance can be negative over some range of applied voltages. This means that an input bias signal must be applied and this acts as the source of power which can be added, making it possible to amplify signal variations.

The circuit in Fig4a produces a system that shows a region of negative dynamic resistance, acting to mimic a Gunn diode. The behaviour of the circuit may be divided into three sections:-

i) Low voltage: $|V| < V_D / G_1$ where V_D is the turn-on voltage of the diodes, D_1 & D_2;

ii) Medium voltage: from $V_Z / (G_1 A) - V_D / G_1 \leqslant |V| < V_D / G_1$

iii) High voltage: where $|V| \geqslant V_Z / (G_1 A) - V_D / G_1$.

where V_Z is the turn-on voltage of the zener diodes, and $A = R_4 / (R_3 + R_4)$;

In the low voltage range the apparent dynamic input resistance, $r\{V\}$, will be

$$r\{\text{low } V\} = \frac{R_1 R_2}{R_1 + R_2} \qquad \text{... (11)}$$

In the medium range, a change in input voltage, dV, will alter the current flowing through R_1 by an amount

$$dI_1 = dV / R_1 \qquad \text{... (12)}$$

and that through R_2 by an amount

$$dI_2 = \frac{dV - G.dV}{R_2} \qquad \text{... (13)}$$

where $G = G_1 A G_2$.

The dynamic resistance in this range will therefore be

$$r\{\text{mid V}\} = \frac{R_1 R_2}{R_2 + (1 - G) R_1} \qquad \text{... (14)}$$

Both the magnitude and sign of r depend upon the overall gain, G. If we choose a value for this gain which satisfies the inequality

$$G > 1 + \frac{R_2}{R_1} \qquad \text{... (15)}$$

then r will be negative. Any increase in the input voltage will then produce a fall in the input current.

In fact, in this situation we can synthesise a system which has a negative static resistance since the operational amplifiers are powered separately from the input, V. For a real two terminal device such as a Gunn diode, the fall in I must stop before it reaches zero. Here, this effect may be produced by the inclusion of the zener diodes and the extra resistors. In the high voltage range the dynamic resistance will become

$$r\left\{\text{high}\,V\right\} = \frac{R_1 R_2}{R_2 + (1 - G')R_1} \qquad \text{...(16)}$$

where G' is a reduced gain, $G' = G_1 A' G_2$, where

$$A' = \frac{R_4'}{R_4' + R_3} \qquad \text{...(17)}$$

and $R_4' = R_4 R_5 / (R_4 + R_5)$. In this way we can synthesise a resistance non-linearity which is broadly similar to that of a Gunn device.

Given the ability to construct non-linear resistances and reactances we can employ circuits of the types described as building blocks to assemble systems of arbitrary complexity. For example, we could build a set of artificial varactors and connect them with a series of inductors to form a non-linear transmission line. This could then - at a frequency range of our choice - be used as a form of analog parallel processor to examine harmonic/anharmonic and chaotic processes and the propagation of solitons along a non-linear medium.

Of course, analog systems are no more perfect than digital ones. Any analog model we build will suffer from some unavoidable defects. For example, the actual signal levels in the system will be affected by random noise fluctuations. This noise has a similar effect to the quantisation of digital values in limiting the precision with which we can specify any individual value. When considering a time-series of values, however, the effects of noise and quantisation are different. We can recover low-level signals from noise by a suitable averaging or signal integrating technique. Quantisation is a form of distortion, and is not inherently random in its effects.

Quantisation distortions can, in some systems, be transformed into random errors by the addition of an extra noise signal (usually referred to as 'dither'). This method is very useful in systems like CD players for producing low distortion and recovering sub-quantisation signals. However, this approach is less easy to apply when the entire system consists of a digital computational process because dither must be added *before* quantisation.

One of the distinguishing characteristics of chaotic behaviour is the way any small perturbations are rapidly magnified. Non-chaotic processes tend to maintain a periodic behaviour even in the presence of noise. In many ways, analog systems give a rather more immediate picture of the behaviour of a non-linear system. Perhaps it is time to start building analog computers again...

[N.B. The following address/status details were correct in 1991, but not now!]

Dr J.C.G.Lesurf
Lecturer in Electronics & Physics
Physics & Astronomy Department
University of St. Andrews St. Andrews
Fife KY16 9SS Scotland

Appendix: System phase plots

The output waveforms generated by non-linear systems can be very complex. This gives us the problem of trying to devise some way of 'tidying up' some of this complexity so we can see more clearly just what is going on. Mathematicians have devised various ways of tackling chaos and complexity. One of the simplest of these leads to the phase plots shown in figures 3a & 3b.

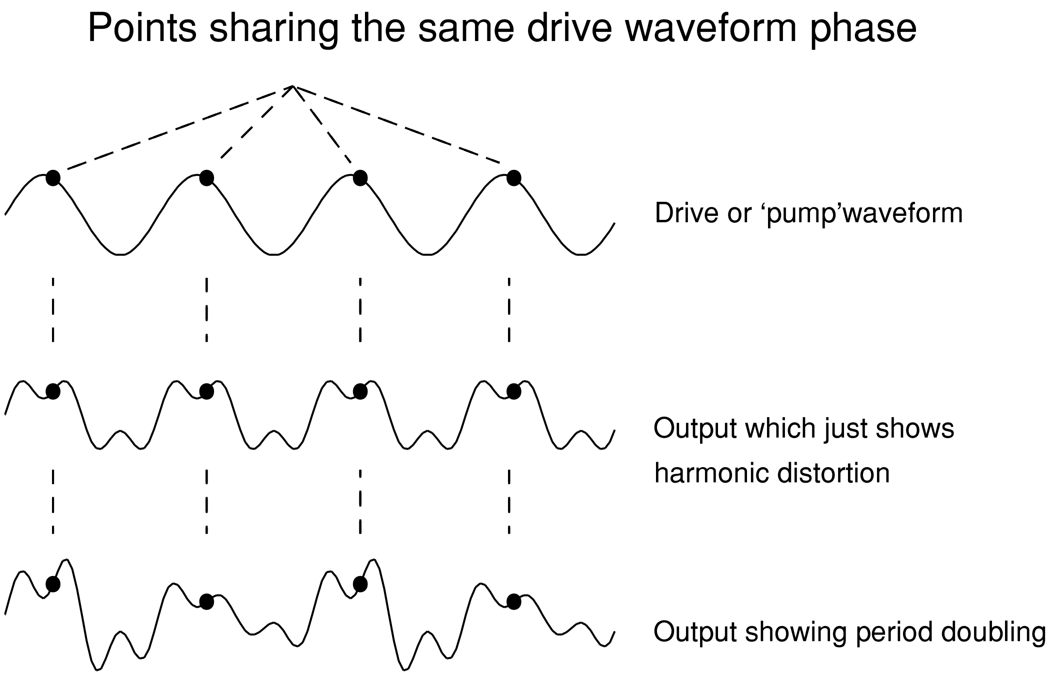

The above Figure illustrates how the values shown on a phase plot are produced. The top waveform represents the time-variations of an input signal which drives or 'pumps' the non-linear system at a particular frequency, *f*. The output this produces will depend upon both the circuit and the details of the input. The middle and bottom waveforms represent two possible kinds of output.

In many cases, the output is simply a distorted version of the input – i.e. it only contains frequency components which are harmonics of those present in the input. Although its shape is different, it repeats itself with the same basic frequency, *f*. This kind of output is represented by the middle waveform.

The bottom waveform is typical of more complex behaviour where the output does *not* repeat itself with the same frequency as the input. Under these circumstances the output contains frequency components which are *not* harmonics of those present in the input.

When dealing with simple waveforms we would normally collect information about their time variations (or frequency spectra) and hope to see recognisable patterns. This isn't much help with very complex waveforms as the patterns tend to be so complicated as to be meaningless. To avoid becoming bogged down in all the details we can decide to ignore most of them and only note the output signal level at a single moment during each input cycle. The times at which we sample the output can be chosen on the basis that they all correspond to the same phase of the periodic input signal.

When the output is just a distorted version of the input we find that every output value we record is the same as every other. The ‘complexity’ of such an output waveform can then represented by just one number, no matter how different it may look from the input.

Figure 3 uses this approach to explore how the complexity of behaviour of a particular non-linear system varies when we drive it with a steady 1300 Hz sinewave and slowly vary and added applied d.c. voltage. For input d.c. voltages between 0 and about 0·05 Volts the output is just a distorted version of the input. The details of the distortion (and hence the sampled output level) vary with the applied d.c., but the output waveform does not change its basic nature. As a result we can plot a single line showing how the sampled level varies with the applied d.c.

Above about 0·05 Volts, things abruptly change. Now the sampled output level alternates back and forth between two different levels. This is the kind of behaviour represented by the lower waveform of the above illustration. Now we must record *two* values to indicate the complexity of the output waveform. When we plot the sampled output levels against various choices for the d.c. input level we find that the line has split into two distinct branches. The waveform is said to have undergone a bifurcation at about 0·05 Volts d.c.

The precise output level we record during each cycle isn't particularly important. What matters is how many levels or lines we see, one after another, under a given set of conditions. The more there are, the more complex the waveform has become, and the longer the time taken before the output pattern will repeat itself. The phase plot therefore gives us a way of assessing how the complexity of the circuit's behaviour alters without having to record all the details of every possible waveform.

4500 Words
This version Jan 2019