

Thermal Compensation of Analog Exponential Converters

openmusiclabs – March 7, 2015

I. INTRODUCTION

The human sense of hearing is amazingly broad and accurate. It can hear leaves rustling in a forest, or a blaring guitar riff at a rock concert – sounds that are 1,000,000 times different in volume. It can detect pitches as low as 20Hz, and as high as 20kHz – again, a wide range of 1,000 times. The amazing part is that the relative accuracy of each of these perceptions does not change much as we move through these ranges. For this reason, it is often useful to express these quantities in terms of exponents (or logarithms). For example, we talk of pitch in terms of octaves, or a doubling of frequency, which follows an exponential: $F = 2^{(n/12)}$ where n is the note of interest, and F is the frequency. We also talk of volume in decibels, which is the logarithm of the air pressure we experience: $dB = 20\log_{10}(P_1/P_0)$. So, for electronic music systems, an exponential converter or logarithmic compressor is a useful tool. It allows us to move smoothly over a large range of values without losing resolution at any part of the range. The inability to have this level of control can be quite frustrating, as anyone who has used a linear potentiometer on a volume or attack/decay knob can attest to. The parameter being controlled seems to move too quickly through the values, and can never be set quite right.

Analog exponential converters have been used for quite some time, and to great effect. The volts per octave control of Voltage Controlled Oscillators (VCOs) was a great innovation, and allowed them to be far more stable as control voltages (CVs) were passed between instruments. But, the limiting factor in their use comes down to this stability, as they often drift with temperature. This paper will discuss the causes of this drift, and explore some of the solutions used to compensate for it. The main goal of this work is to see what the limits are for creating a very accurate, and temperature stable exponential converter for VCOs. VCOs are chosen as they are a very demanding application, with stricter tolerances than other applications (e.g. volume control). A well trained ear can hear a change in pitch down to a few cents (a cent is $1/100^{th}$ the distance to an adjacent note). As the change in pitch required to move up to the next note is only 6%, this is a very small band that the exponential converter must be kept within (1 cent is therefore 0.06%). And, this tolerance needs to be held over 10 octaves! A difficult task indeed.

II. BACKGROUND

It will be useful to have a mathematical representation of the exact exponential characteristic we want, so that all converter topologies can be compared to a standard. Since we

are looking at VCOs, the 1V per octave CV “standard” will be a good reference. Assuming an oscillator that produces a frequency proportional to an input current, we are looking for a relationship as follows:

$$I_{out} = I_{ref} \cdot 2^{CV/1V} = I_{ref} \cdot e^{\ln(2) \cdot CV/1V}, \quad (1)$$

where I_{out} is the current to the oscillator, I_{ref} is a fixed reference current, CV is the input control voltage in volts, and $\ln()$ and e are the natural log and its base. If you are not familiar with the natural log, do not fret, as it’s just another way of expressing a logarithm, but with a base of ~ 2.71828 (which is a fixed number abbreviated as ‘ e ’) instead of the usual base of 10. One of the useful properties of logarithms is that they can be used to easily do multiplication by simply adding them together, and then exponentiating them (its how slide rules work). If you don’t know much about them, there are plenty of resources to learn more, but it is not critical to understanding the work done here.

Most exponential converters are implemented with Bipolar Junction Transistors (BJTs). This is because they are inherently exponential; the current they produce is exponentially related to the base voltage applied. A simple model of this relationship is as follows:

$$I_c = I_s \cdot e^{V_{be}/V_t}, \quad (2)$$

where I_c is the collector current, I_s is a device specific parameter, V_{be} is the voltage between the base and emitter, and V_t is the “thermal voltage” ($\sim 26\text{mV}$). There are more accurate models, but as it turns out, we can not compensate for the drifts that this model predicts, so there is no need to go to lower levels just yet.

Both the I_s and V_t terms are temperature dependent. I_s can be mostly canceled out by using a transistor with a very similar I_s parameter as a reference transistor. This is the ‘matched pair’ that is often talked about. In this case, the base voltage applied is actually a differential voltage between the bases of the two transistors (the reference transistor, and the exponentiating transistor). A typical implementation is shown in Figure 1.

The current through the reference transistor is I_{ref} , and is held constant by the action of the op-amp. The current through the exponentiating transistor is our output, and controlled by the differential base voltage dV . Since the reference transistor’s base is held at ground, this difference is merely the voltage at the base of the exponentiating transistor. Using math we can show that this cancels out the I_s term (assuming I_s is identical between the two transistors):




Fig. 1. The basic exponential core

$$dV = V_{be2} - V_{be1} \quad (3)$$

$$I_{ref} = I_s \cdot e^{V_{be1}/V_t} \rightarrow I_s = I_{ref} \cdot e^{-V_{be1}/V_t} \quad (4)$$

$$I_{out} = I_s \cdot e^{V_{be2}/V_t} \rightarrow I_s = I_{out} \cdot e^{-V_{be2}/V_t} \quad (5)$$

$$\Rightarrow I_{out} = I_{ref} \cdot e^{(V_{be2} - V_{be1})/V_t} = I_{ref} \cdot e^{dV/V_t}. \quad (6)$$

But, you can also intuitively think of the reference transistor as canceling out any base voltage variations that might occur over temperature, as both transistors will increase by the same amount, and the applied voltage (the difference between the two transistors) stays the same.

This just leaves us with the pesky V_t term. V_t is called the “thermal voltage” because it is proportional to absolute temperature:

$$V_t = \frac{KT}{q}, \quad (7)$$

where K is Boltzman’s constant, T is absolute temperature (in Kelvin (K)), and q is the fundamental charge of an electron. So, at room temperature (300K) it’s around 26mV, but it increases with temperature. How bad this drift is depends upon two things: the difference in temperature, and the difference in base voltages between the reference and exponentiating transistors. To illustrate this effect, Figure 2 shows the error produced as a function of temperature for various applied CVs.


Two very important things to note are that the error is out of any usable range for almost any temperature drift, and that this error increase linearly with the CV. This means that it is easier to maintain accuracy if you limit the applied CV. Using equations 1 and 6, we can figure out the maximum dV required at a 10V CV:

$$I_{out} = I_{ref} \cdot e^{(\ln(2) \cdot CV/1V)} = I_{ref} \cdot e^{dV/V_t} \quad (8)$$

$$\rightarrow \ln(2) \cdot CV/1V = dV/V_t \quad (9)$$

$$\Rightarrow dV = \ln(2) \cdot CV \cdot V_t/1V. \quad (10)$$

So, for a 0V to 10V CV, 180mV is required. But, for a $\pm 5V$ CV, which would cover the exact same range, only half of that ($\pm 90mV$) would be needed. This has the double benefit of halving any V_t related errors, and minimizing errors in the most commonly used portion of the musical range (the middle). For example, if a 100nA reference current was used, and the full output range was 100nA to 100 μ A, then the

Fig. 2. V_t induced error [%] versus temperature [K] and CV [V].

middle of the keyboard would be at 3 μ A (+5V CV), and would have all the drift associated with that large CV value. But, if the reference current was 3 μ A, and the output swung from 100nA to 100 μ A via a $\pm 5V$ CV, that same middle of the range would have negligible error, and the extremes of the keyboard would only have half the maximum error of the previous example. So a system that allows for bipolar CVs will inherently have less drift.

There are a few limitations which are common to all exponential converters. for example, the Early effect, which causes a small variation in output current with changes in V_{ce} (the collector to emitter voltage). A more accurate model of transistor behaviour would look like:

$$I_c = I_s \left(1 + \frac{V_{ce}}{V_a}\right) e^{V_{be}/V_t}, \quad (11)$$

where V_a is the Early voltage (usually 50V to 100V). So the op-amp in Figure 1 is actually accomplishing two goals, it is keeping both I_c and V_{ce} constant in the reference transistor. It is very important, therefore, to use a similar regulating op-amp on the exponentiating transistor, so the $I_s(1 + V_{ce}/V_a)$ terms between the reference and exponentiator are identical, and will cancel. Again, we need accuracy on the order of 0.06%, and if V_{ce} is allowed to vary by just a few volts, there will be a discrepancy of 1.4%! A typical V_{ce} on the reference transistor is 0.6V, so if the exponentiating transistor has a value of 2V this would give a difference of $(1 + 2V/100V) - (1 + 0.6V/100V) = 1.4V/100V = 1.4\%$.

The Early effect is important to keep in mind for other parts of an exponential converter as well. Many of the active gain compensation techniques use discrete multipliers comprised of differential pairs, who’s V_{ce} ’s must also be kept constant for accurate results. It is not only good to keep V_{ce} constant, but also as small as possible. Ideally, V_{cb} (the collector to base voltage) would be 0V. This helps reduce collector to base leakage, which increases exponentially with temperature. Topologies which hold V_{cb} at 0V on all transistors will perform better than those that do not.

Another common source of error is fluctuations in op-amp parameters with temperature. For example, a more accurate version of the exponential core showing error sources is shown in Figure 3. The current through the reference transistor is set




Fig. 3. op-amp error sources in the basic exponential core.

by the value of the supply voltage, the value of the resistor, and the values of the offset voltage and bias current in the op-amp:

$$I_{ref} = \frac{V_{ref} - V_{os}}{R_{ref}} - I_b. \quad (12)$$

So R_{ref} must move less than 0.06% over the temperature range of interest; this means $<10\text{ppm}/^\circ\text{C}$ metal film resistors ($10\text{ppm}/^\circ\text{C} \times 50^\circ\text{C} = 500\text{ppm} = 0.05\%$). Likewise, V_{ref} must be a constant source with less than $10\text{ppm}/^\circ\text{C}$ drift. But, V_{os} and I_b do not need to be as carefully controlled, as long as they are kept significantly smaller than V_{ref} and I_{ref} . For example, if V_{ref} is 5V and V_{os} is 1mV, the voltage across the resistor is 4.999V. A 0.05% decrease in this voltage would give 4.9965V, which is the equivalent of a 2.5mV increase in the offset voltage. For a 50°C temperature change, that would be a $+50\mu\text{V}/^\circ\text{C}$ drift. The case for I_b is similar, although for JFET input op-amps the bias current increases exponentially with temperature, making it a bit more difficult to control. It is often easier to just ensure that for all operating temperatures $I_b < 0.0005 \times I_{ref}$, rather than make sure its change is within specification. This is also an argument for keeping I_{ref} as large as possible (within reason) so that op-amp bias currents are not as detrimental.

It is important to keep in mind that the above offset voltage and bias current calculations are for just one portion of the exponential converter, and that other parts might have tighter constraints. For example, any parameter that affects the exponent will have a much a larger effect than one that effects the reference current, as the reference current is linear to the output, whereas the exponent is, as it sounds, exponential to the output. Also, it is good to keep variations well below the 0.06% mark, as all of the errors will add together, and could easily push things ten times past that limit.

Finally, finite base currents need to be considered in the converter. Ideally, transistors with very high forward current gain (h_{fe} or β) should be used. Although β tends to increase with temperature, reducing the base current, leakage currents will increase with temperature at a much faster rate. And leakage can occur between any parts of multiple transistors on a die: collector to base, collector to collector, collector to substrate, etc. This creates a tricky problem of picking an

operating point, as a lower I_{ref} will have a lower base current due to finite beta, but leakage currents will be a much larger percentage of the reference current at higher temperatures.

There are a couple of tricks that can be used to minimize the errors associated with base currents. First, applying the exponentiating voltage to the base of the reference transistor will keep any source resistance induced errors relatively constant, and small. This is because the base current is proportional to the collector current, and the collector current in the reference transistor is constant, whereas the output current moves over three decades of range. Second, keep the source resistance as small as possible. Ultimately, this is limited by the amount of current you are willing to source from your op-amps, as drawing more current limits op-amp accuracy and creates extra heat. Usually 1mA is a reasonable amount, which would require a base resistor of 90Ω to generate the $\pm 90\text{mV}$ needed to cover the full 10 octave range. Finally, matching the resistance seen by both bases can help improve linearity in some cases. Usually this hurts an exponential converter, as any voltage generated by a base current will not be temperature compensated, but in the case of leakage currents, it helps cancel them out.

Ultimately, to compensate for the variation in V_t , either the temperature must be held constant, or the voltage applied to the base must be scaled in proportion to V_t . If it is not, the exponential current will drift. These two techniques are applied in a number of different ways, each with their relative benefits. The following topologies will be discussed, each in their own section:

- Thermal oven compensation (constant V_t technique)
- Thermistor compensation (passive gain compensation)
- Exponential V_t approximation (active gain compensation)
- OTA V_t multiplier (active gain compensation)
- Inverted OTA V_t multiplier (active gain compensation)
- Gilbert cell V_t multiplier (active gain compensation)
- Feedback V_t multiplier (active gain compensation)

They will be compared for cost, frequency response, linearity, and temperature drift. Some of them were built and tested, and for these experimental data are given, and partial schematics of the circuit are shown. Full schematics are not shown to improve clarity of their basic functionality, and in some cases because of poor documentation (i.e. i don't remember!). Not all were tested, as the theoretical performance did not seem good enough to warrant testing, or adequate parts were not available.

III. THERMAL OVEN COMPENSATION

In a thermal oven compensation scheme, two or three transistors and a heating element (often another transistor) are required to do the exponentiation, measure the temperature of the transistors, and apply heat through a feedback loop which keeps this temperature constant. This is a very old technique, often employed in Buchla and Moog synths with the $\mu\text{A}726$. The $\mu\text{A}726$ was a matched pair on a single die, with a temperature sensor, heater and feedback loop to keep the whole part temperature stable. It worked very well, but has not been in production for a long time. Parts are hard to come by,




Fig. 4. SSM2033 functional diagram (from datasheet).

and very expensive (100\$) when they are available. Another IC which used this method was the SSM2033, which also can not be found these days. The schematic of the SSM2033 from its datasheet is shown in Figure 4. Luckily you can build your own with a matched pair and an external heater, or a matched quad like the LM3046, MAT14, or THAT300. Using an external heater is not recommended, as the time delay from applying heat to raising the temperature of the die will be much longer than on-die heaters (like a third transistor) and it will be more difficult to maintain a stable feedback control loop.

It is important to point out that the heating element can only heat, it can not take heat away (cool). so the die must be kept at the maximum temperature you want it to withstand at all times. For example, if you need your exponential converter to work up to 85°C, then the die must be kept above 85°C at all times. This has three serious side effects. The first is that the die will take time to warm up to this temperature when the circuit is first turned on. This warm up process may take as long as a half hour for everything to settle out. Secondly, leakage currents are much, much worse at higher temperatures, so the exponential converter will always be operating at a reduced accuracy. With discrete implementations, the op-amps and feedback loop are off-die, so they are not affected by the raised temperature. Finally, it creates thermal gradients across the die between the heater element and the outside of the case. If the sensing, reference, and exponentiating transistors are not equidistant to the heater, they will be at different temperatures, and this difference will vary with the ambient temperature.

I have not built and tested one of these circuits, but if warm-up time and power consumption (it may draw 15mA to run the heater) are not concerns, it is probably a good choice. I would expect its linearity to be worse at low currents, due to leakage, and there might be some thermal oscillation if the control loop is not completely stable. The gain on the control loop is necessarily a function of temperature, as the amount of heat being removed reduces with increased ambient air temperature, making it harder to regulate. But, it is generally

low cost and in theory has no thermal drift. An LM3046 and a few precision op-amps could be as low as 5\$ total. Using a MAT14 and some really nice op-amps and resistors could push that to 20\$. The frequency response of the converter will also be very fast, only limited by the bandwidth of the transistors, which is usually in the MHz range.

IV. THERMISTOR COMPENSATION

With thermistor compensation, the value of the applied base voltage is varied with temperature via an attenuator built with a temperature dependent resistor. There are 3 kinds of temperature dependent resistors: Negative Temperature Coefficient (NTC) thermistors, Positive Temperature Coefficient (PTC) thermistors, and Resistance Temperature Detectors (RTDs). Thermistors are generally less expensive, have higher temperature coefficients, and are less accurate than RTDs. RTDs are made of pure metals (platinum is very common), and have a very linear response. NTC thermistors have an inverse exponential response ($R = R_0 \cdot e^{1/T}$), and can be used, but do not do a good job of compensating for temperature changes. A more in-depth explanation of why they make poor temperature compensators can be found in Section V, where the same technique using transistors is evaluated. But, they are cheaper, often more well defined, and come in a wider selection of values.

For the above reasons we will focus on PTC thermistors and RTDs, as they have a response that is more linear with absolute temperature. Unfortunately, PTC Thermistors and RTDs tend to be more expensive and harder to come by, especially in a particular value. DIY audio suppliers will often have through-hole mount thermistors in common values for anywhere from 1\$ to 10\$ a resistor. Suppliers like DigiKey have SMT thermistors (Panasonic ERA series) in a range of values for under a dollar, and recently started carrying platinum RTDs (Vishay PTS series) for similar prices. I have not tested any of the RTDs yet, but they seem to be about as good as it gets. They are extremely linear, and the temperature coefficient is specified to 0.3% (the PTC thermistors are only specified to 10%). The two ways a thermistor can be used are shown in Figures 5 and 6.

In Figure 5, The thermistor sets the gain of an op-amp stage: R_1/R_2 . So, as the temperature increases and the resistance increases, so does the output voltage, scaling the CV directly with temperature. In Figure 6 the output is $R_1/(R_1 + R_2)$ times the input, and as such is not directly scaled with R_1 , but also has an $R_1 + R_2$ quantity on the bottom. This lower quantity “waters down” the effective temperature coefficient, changing the scaling factor. This is why you will often see two different temperature coefficients specified for compensation circuits: 3300ppm/°C and 3500ppm/°C. The former is what you want for op-amp scaling, and the latter would be used in resistor divider circuits, as it gets “watered down” to the appropriate value.

But what is the appropriate value? This depends heavily on how the thermistor is specified. Most thermistors will give a nominal value at room temperature (25°C) and then take another measurement at some higher temperature (75°C).




Fig. 5. Thermistor compensation with an op-amp.




Fig. 6. Thermistor compensation with a voltage divider.

The change in resistance is then divided by the change in temperature and the nominal resistance to find the percent change per degree, which is the Temperature Coefficient of Resistance (TCR).

$$TCR \text{ (ppm/}^\circ\text{C)} = \frac{R_{75} - R_{25}}{R_{25} \cdot (75^\circ\text{C} - 25^\circ\text{C})} \times 10^6. \quad (13)$$

Since V_t varies linearly with absolute temperature (1K is equivalent to 1°C , and $0\text{K} = -273^\circ\text{C}$), we need a resistance that also varies linearly with absolute temperature. This means we need a line that goes through 0K at 0Ω , $(25 + 273)\text{K}$ at R_{25} and $(75 + 273)\text{K}$ at R_{75} . Taking the slope of these line segments and setting them equal to each other (so they actually form a line) gives:

$$\frac{R_{25} - 0\Omega}{(25 + 273)\text{K} - 0\text{K}} = \frac{R_{75} - R_{25}}{(75 + 273)\text{K} - (25 + 273)\text{K}} \quad (14)$$

$$\rightarrow \frac{R_{25}}{298\text{K}} = \frac{R_{75} - R_{25}}{(75 - 25)\text{K}} \quad (15)$$

$$\Rightarrow \frac{1}{298\text{K}} = \frac{R_{75} - R_{25}}{R_{25}(75 - 25)\text{K}} = \frac{TCR}{10^6} \quad (16)$$

$$\Rightarrow TCR = (1/298\text{K}) \times 10^6 = 3356\text{ppm/}^\circ\text{C}. \quad (17)$$

So the temperature coefficient needed is just one over the temperature (in Kelvin) at which it is specified. But, this is only for the op-amp scaled version. to calculate the amount of “watering down” that occurs in the resistor divider version we need to know the relative value of the other resistor. To find the effective temperature coefficient, we find the change in output voltage that occurs for a given temperature change.

$$\text{Effective TCR} = \frac{\frac{R_{75}}{R_{75} + R_2} - \frac{R_{25}}{R_{25} + R_2}}{\frac{R_{25}}{R_{25} + R_2} \cdot (75^\circ\text{C} - 25^\circ\text{C})} \quad (18)$$

$$= \frac{R_{75}(R_{25} + R_2) - R_{25}(R_{75} + R_2)}{(R_{25})(R_{75} + R_2)(75^\circ\text{C} - 25^\circ\text{C})} \quad (19)$$

$$= \frac{R_2(R_{75} - R_{25})}{R_{25}(R_{75} + R_2)(75^\circ\text{C} - 25^\circ\text{C})} \quad (20)$$

$$= \frac{R_2}{R_{75} + R_2} \times TCR \quad (21)$$

$$= \frac{R_2}{R_2 + (R_{25} + R_{25} \times TCR \times 50^\circ\text{C})} \times TCR. \quad (22)$$

The final substitutions are based on the definition of the TCR from Equation 13. As you can see, the effective TCR now slightly changes with temperature, but for $R_2 \gg R_1$, this is a negligible factor, and since $R_1 = R_{25}$ the effective TCR can be approximated as:

$$\text{Effective TCR} \approx \frac{R_2}{R_1 + R_2} \times TCR. \quad (23)$$

So, for an initial TCR of $3500\text{ppm/}^\circ\text{C}$, R_2 should be 23.3 times R_1 to give an effective TCR of $3356\text{ppm/}^\circ\text{C}$. For TCR values that are significantly different from the ideal TCR (more than 10%), the R_1/R_2 ratio is no longer negligible and the TCR is no longer linear. An example of this is shown in Figure 7, where an initial TCR of $3850\text{ppm/}^\circ\text{C}$ is resistor divided down to $3356\text{ppm/}^\circ\text{C}$. A strong curve is present, which reduces the useful range. This curve is for a platinum RTD, and although the curve is bad, the repeatability is much better than the PTC thermistors. In other words, a perfect thermistor will be better than a perfect RTD, but the average thermistor you can obtain will be much, much worse than the average RTD you can obtain.

Thermistor compensation is a decent technique that has been used on a lot of synths. It works very well if the ambient temperature does not vary much from the die temperature, and if the TCR is accurate. Unfortunately, both of these might not be the case. The common Panasonic PTC thermistors (as used in old Roland gear) are only specified to $\pm 10\%$, which is not accurate enough for large temperature swings, as shown by the error curves in Figure 8. Also, if there is any airflow near the circuit, the die and thermistor temperatures will vary from one another. On tests with a thermistor glued to the face of 5-pin SIP matched pair, i found the current to vary quite a bit with airflow, and even without airflow the current output had a lot of low frequency noise. This is because the compensation sensing element (thermistor) is not directly coupled to the reference transistor, so the temperature of one is not directly related to the other and it takes time for the heat to transfer from one to the other. Another issue is self heating of the thermistor. Small values of resistance are good for reducing base current errors (in voltage divider mode) but will also increase the power dissipated in the thermistor, changing its temperature at higher CVs.

One workaround to the thermal coupling problem is to use a special IC that has a matched pair and a thermistor on the same die. This is the technique that Texas Instruments (TI) use




Fig. 11. SSM2164 used as exponential converter.




Fig. 12. Discrete exponential compensation scheme.

transistor, it can be used to compensate for the temperature drift. Two different versions of this are shown Figures 11 and 12, one implemented with the SSM2164 and the other with a matched quad. They both operate on the same principle since the SSM2164 is an exponential amplifier, and they both multiply the incoming signal by 26mV (which is 260mV for the SSM2164 which has an internal divide by 10 circuit).

So why use exactly 26mV? And is it mere coincidence that this matches V_t ? To best answer these questions we will again have to use some math. The basic structure of each circuit is to multiply the incoming signal by an exponential, and then put that value into the exponent of the final output stage. As shown before, the gain of an individual stage is:

$$I_{out} = I_{ref} \cdot e^{dV/V_t}. \quad (24)$$

so, if we have a CV that is converted to a current I_{in} , it will be modified by the first stage to become:

$$I_2 = I_{in} \cdot e^{dV_1/V_t}, \quad (25)$$

where I_2 is the output current and dV_1 the exponentiating voltage of the first stage. I_2 is then converted to a voltage via a resistor R , and applied to dV of the second stage while I_{ref} is held constant. This gives the final output as:

$$I_{out} = I_{ref} \cdot e^{I_2 R/V_t} = I_{ref} \cdot e^{(I_{in} R \cdot e^{dV_1/V_t})/V_t}. \quad (26)$$




Fig. 13. $(e^{dV_1/V_t})/V_t$ [1/V] for $dV_1 = -26\text{mV}$ versus temperature [K].

Now the only decision to make is the value of dV_1 . Ideally we would want the e^{dV_1/V_t} term to track linearly with V_t , so that $(e^{dV_1/V_t})/V_t$ remains constant for all values of V_t . Another way of saying this, is that we would like the rate of change of that quantity with respect to V_t to be zero. The best way to find the rate of change is to take the first derivative. Once we have this rate of change, we can then solve for the value of dV_1 that makes this rate zero.

$$\frac{d}{dV_t} \left[\frac{e^{dV_1/V_t}}{V_t} \right] = \frac{-e^{dV_1/V_t} \cdot dV_1}{V_t^3} + \frac{-e^{dV_1/V_t}}{V_t^2} = 0 \quad (27)$$

$$\Rightarrow dV_1 = -V_t. \quad (28)$$

So if we apply $-V_t$ to our first gain stage it should completely cancel out the temperature effect. A graph of $(e^{dV_1/V_t})/V_t$ versus temperature is shown in Figure 13 for $dV_1 = -V_t = -26\text{mV}$. You can see that there is a flat spot in the curve at 300K (27°C, or room temperature) where the value barely change for large changes in temperature. On either side of this point, the curve bends down at an increasing rate, making these less desirable points to operate at, as a small change in temperature will have a much larger change in the output value.

The total error for this scheme is quite small for small temperature changes. In Figure 14 the exponential compensation error is plotted against the thermistor error graph for a 5V applied CV. For temperature changes less than $\pm 20^\circ\text{C}$ it is far more favorable than a thermistor with moderate TCR accuracy. But, for large temperature changes, the nonlinear nature of the compensation becomes more apparent, and the effectiveness drops considerably. Using an NTC thermistor, because of its exponential response, would have similar effectiveness as this method (except for the errors due to not being on-die with the transistors).

Figure 15 shows a plot of the linearity and temperature drift of an exponential converter implemented with an SSM2164. The SSM2164 is built to accommodate bipolar signals, so it has two complementary exponential converters inside. This is beneficial as it allows bipolar CVs to be used, which

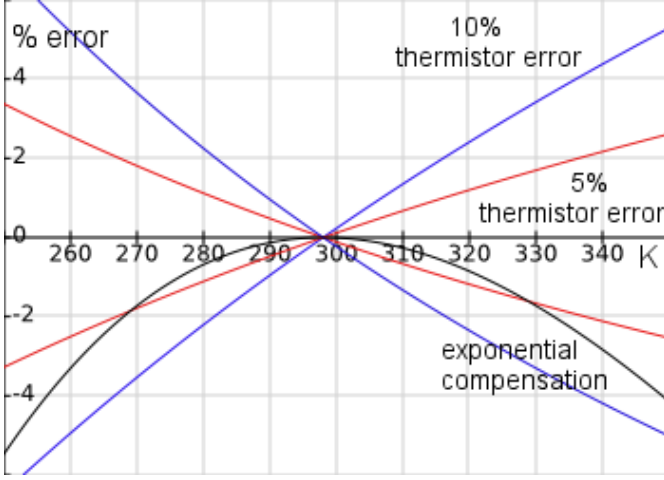


Fig. 14. Exponential compensation and thermistor compensation error [%] versus temperature [K] at 5V CV.




Fig. 15. SSM2164 compensation error [%] versus CV [V].

reduces thermal drift errors. But it also introduces a cross-over distortion as the circuit switches between the two stages. This can be partially trimmed out with the “Mode” input, but the output becomes inaccurate at the extremes of current, so the slight bump actually helps keep the current within a linear range. After trying a number of different operating points, Figure 15 shows $\pm 0.2\%$, which is the best linearity I could achieve from the SSM2164. It is completely usable, but not as good as a single differential pair doing the exponential conversion. The temperature drift was measured at a single point (+3V CV), and it matched the theoretical drift of 1.5%.

The matched quad implementation’s performance is shown in Figure 16, and has much better linearity. It stays within $\pm 0.05\%$ error for the full 10 octave range. Unfortunately, this version only accepts unipolar CV inputs, so the worst case drift is twice as bad as for the SSM2164. Another issue with unipolar CV schemes is that I_{ref} must be either really small (50nA) or really large (200 μ A), as it needs to be outside of the output current range. On the small side, leakage currents and offset voltages have a much larger effect, since the current being controlled is so small. On the larger side, the base current into the reference transistor is correspondingly quite




Fig. 16. Discrete exponential compensation error [%] versus output CV [V].

large, and any base resistance causes errors. Therefore, besides for the worse thermal drift due to double the CV input, unipolar designs also suffer from linearity issues and thermal drift issues due to either a too large or too small I_{ref} .

The temperature drift of both of these schemes matched quite well to the theoretical error. The main exception to this was the SSM2164 implementation, which had large bias currents at its inputs, and bias current and offset voltage drifts with temperature that were unacceptable. The results shown above employed a current source driving the SSM2164 input to help minimize these effects. The SSM2164 also uses a moderately large base resistor (500 Ω) which does not help its linearity. The THAT2162 might give better performance due to not using base resistors or current mirrors internally, which would eliminate those sources of error. But, it was not tested for this work.

Ultimately, the exponential V_t approximation technique is pretty good. It has good frequency response, only limited by the internal compensation of the SSM2164. The temperature compensating element is on-die with the exponentiating transistor, making it less susceptible to external heat sources, and more responsive to internal changes. This makes the output much more stable than thermistor variants. Between the SSM2164 and matched quad implementations, I would highly recommend the SSM2164 version. Any loss in linearity is greatly made up for in better temperature stability and reduced part count. Although the SSM2164 is no longer available, CoolAudio makes a clone which sells for 4\$ in single quantities; and that gets you two temperature compensated exponential converters. A THAT2162 sells for 5\$ in single quantities, and is only a dual device (makes a single converter), but might give better linearity. A matched quad version could run anywhere from 3\$ to 20\$ depending upon which matched quad and op-amps are used.

VI. OTA V_t MULTIPLIER

The ideal way to cancel out the V_t term is to multiply the input CV by V_t before it enters the exponentiator. It would then perfectly cancel out the temperature drift. Of the ways to implement this, the OTA multiplier is the easiest and least

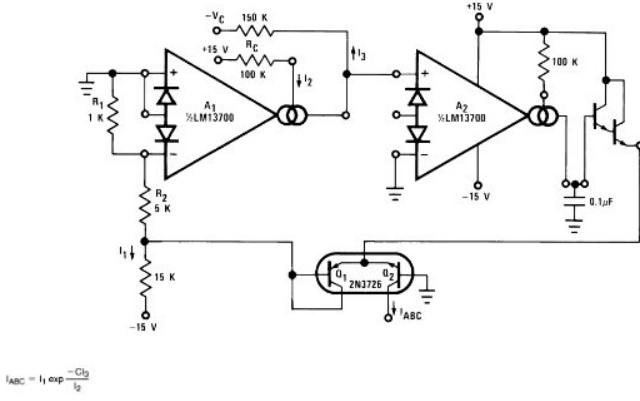


Figure 51. Logarithmic Current Source



Fig. 17. LM13700 as V_t multiplier from datasheet.

Fig. 18. PAIA VCO schematic with OTA compensation.

expensive. This method was first detailed in the LM13700 datasheet and was later used in the PAIA 9700 Series VCO, amongst others.

Despite the PAIA schematic above, which claims the linearizing diode compensates for temperature drift, the basic functionality is actually based on the fact that an Operation Transconductance Amplifier (OTA) has an inherent thermal drift based on V_t . Usually this is a problem to be minimized, but in this case it can be exploited to our advantage. The simplified transfer function of an OTA is $I_{out,ota} = I_{bias} \cdot dV_{ota}/2V_t$, where I_{bias} is the bias current into the OTA and dV_{ota} is the input voltage. Normally the input is applied at dV_{ota} , but that would give us an output divided by V_t . To get an output multiplied by V_t , we apply our signal at $I_{out,ota}$ and take our output from dV_{ota} . We do this by wrapping the OTA around an op-amp. A schematic of this is shown in Figure 19.

In this case the CV creates a current, I_{in} , which matches $I_{out,ota}$. An op-amp is used to modulate dV_{ota} until this is

Fig. 19. Simplified OTA V_t multiplier schematic.

the case. This dV_{ota} , which is now a function of V_t , is then applied to the exponential converter where it is canceled out by the converter's V_t . A voltage divider is used to reduce the linearity error introduced by the OTA, and it has a ratio of z .

$$I_{in} = \frac{CV}{R} = I_{out,ota} = \frac{I_{bias} \cdot dV_{ota}}{2V_t} \quad (29)$$

$$\rightarrow dV_{ota} = \frac{CV \cdot 2V_t}{R \cdot I_{bias}} \quad (30)$$

$$I_{out} = I_{ref} \cdot e^{(z \cdot dV_{ota}/V_t)} \quad (31)$$

$$= I_{ref} \cdot e^{(z \cdot CV \cdot 2V_t)/(R \cdot I_{bias} \cdot V_t)} \quad (32)$$

$$= I_{ref} \cdot e^{(2z \cdot CV)/(R \cdot I_{bias})} \quad (33)$$

Unfortunately the OTA gain equation is an approximation, and even for a mild input voltage range it has linearity issues. A graph of error versus CV for an ideal implementation with $z = 10$ is shown in Figure 20, along with linearly compensated versions to bring the error into a reasonable range. Even with this large value of z , the best case error for a 10 octave range is still $\pm 0.9\%$, which is a ± 15 cent error.

The error itself is an x^3 relationship. This can be seen if we look at the actual OTA gain equation, and the Taylor series expansion (an accurate “approximation”) of the natural log:

$$dV_{ota} = V_t [\ln(1 + \frac{I_{out,ota}}{I_{bias}}) - \ln(1 - \frac{I_{out,ota}}{I_{bias}})] \quad (34)$$

$$\ln(1 + x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \frac{x^5}{5} + \dots \quad (35)$$

$$\ln(1 - x) = -x - \frac{x^2}{2} - \frac{x^3}{3} - \frac{x^4}{4} - \frac{x^5}{5} + \dots \quad (36)$$

$$\Rightarrow dV_{ota} = 2V_t [x + \frac{x^3}{3} + \frac{x^5}{5} + \dots] \approx \frac{2V_t I_{out,ota}}{I_{bias}} \quad (37)$$

For values of x less than 1, x^3 and x^5 become very small very fast, such that they can usually be ignored. But for our case, the errors we need to achieve are also incredibly small, such that the x^3 term is a significant contributor. An x^3 circuit could be built, with an example shown in Figure 21. But any reasonable implementation would only allow for unipolar CVs, and would come with all the associated errors. For the same additional cost a true multiplier could be built which would give better results. Ultimately the advantage of the OTA multiplier is its simplicity, and adding a cube circuit detracts from this.




Fig. 20. Theoretical OTA V_t multiplier error [%] versus CV [V].




Fig. 21. Analog cube generator schematic.

For the OTA multiplier circuit to give good thermal results, The OTA and exponential converter need to be on the same die. This can be done either with a discrete OTA built from a matched quad, or with an OTA that has linearizing diodes. The linearizing diodes are an extra pair of matched transistors wired as diodes and placed at the input of the LM13700, which can be used to minimize distortion and temperature drift in the OTA. Versions of these implementations are shown in Figures 22 - 24.

In Figure 22, a very simple exponential converter is created with a single OTA by having the dV generated by the OTA directly modulate the linearizing diode exponentiating section. Because of the lack of divisor between dV and the exponential input, the distortion is quite bad, and the linear range is confined to a few octaves. But, if you only need a few octaves, this is certainly a cheap and well compensated way to go. Although realistically, keeping a thermistor compensated exponential converter in tune over a few octaves is also very inexpensive and easy to do.

A more complicated version is shown in Figure 23, where the first OTA creates the V_t multiplied CV, and the second section is only used for its linearizing diodes. This has a more reasonable usable range, but is still limited by the x^3 error. Furthermore, the linearizing diodes make poor exponential pairs, since the base current is added with the collector current at the output, giving a scaling error that become problematic




Fig. 22. OTA V_t multiplier using a single LM13700.




Fig. 23. OTA V_t multiplier using LM13700 linearizing diodes.




Fig. 24. Discrete OTA V_t multiplier.




Fig. 25. Error [%] versus CV [V] for OTA V_t multiplier using LM13700 linearizing diodes.

as I_b and I_{os} are modulated with temperature and collector current. A graph of its error is shown in Figure 25. The linearity error for this configuration matches theory quite well, but temperature drift was 0.7% for a 3V CV and 40°C temperature change (single point test).

The matched quad version shown in figure 24 has a number of advantages. First off, its exponentiating pair is not wired as diodes. Secondly, it does not use transistor current mirrors, but rather a differential amplifier to produce the output current. This eliminates the current mirror errors and replaces them with op-amp and resistor errors, which are far easier to control and predict. Unfortunately i did not take any data for this version, but i would suspect it would have lower temperature drift.

All of these implementations are fundamentally limited by the OTA's lack of linearity. This linearity can be improved by increasing the divider ratio, but going much larger than 10 starts to increase the noise in the current at the output. Low frequency noise is the largest component of the OTA noise spectrum, and this can cause a noticeable quaver in the output. The temperature stability can be improved by using a current source into the I_{bias} input on the OTA, eliminating temperature variations in the current mirror at that input.

The OTA V_t multiplier is generally good for small frequency ranges, and allows for bipolar CVs and their associated benefits. It also can be quite cost effective. An LM13700 is 1\$ in single quantities and can be had in DIP format, which is great for DIY. The matched quad version can also be relatively cheap, as the amount of op-amps required is minimal and could be built for 3\$ to 15\$, depending upon the tolerances required. The frequency response is good, and is only limited by the stability requirements of the CV to I_{out} converter. But, this should be kept as low as possible to limit noise into the exponential converter stage. The temperature tracking is in theory perfect, but in practice it has limitations, although is still very good. And, as a side note, it appears that Maxim uses this technology in their log amps, as the error curves are very similar and they have excellent temperature tracking (see Figure 26).




Fig. 26. Error curves for a Maxim log amp, taken from datasheet.

VII. INVERTED OTA V_t MULTIPLIER



An OTA has two inputs, dV_{ota} and I_{bias} , and these get multiplied together and presented at the output. In the previous section we looked at a signal applied to V_t , which theoretically gave perfect temperature tracking, and a pesky x^3 error at the output. But what if we were to apply our signal to I_{bias} instead, and keep dV_{ota} fixed? By rearranging the OTA gain shown in Equation 34 we can gain some insight into what would happen. The gain as a function of the currents, rather than the voltages, is shown below:

$$I_{out,ota} = I_{bias} \cdot \frac{e^{(dV_{ota}/V_t)} - 1}{e^{(dV_{ota}/V_t)} + 1}. \quad (38)$$

Now our output (which is taken at I_{bias}) and our input (which is applied at $I_{out,ota}$) are linearly related to each other, so there should be no distortion. The price we pay for this, is that the multiplying term which contains V_t , is no longer linear. So we get a theoretically linear exponentiator, and a bit of temperature error. This error can then be set by the value chosen for dV_{ota} . Smaller values will give lower temperature error, but more noise and instability; a similar trade-off as before.


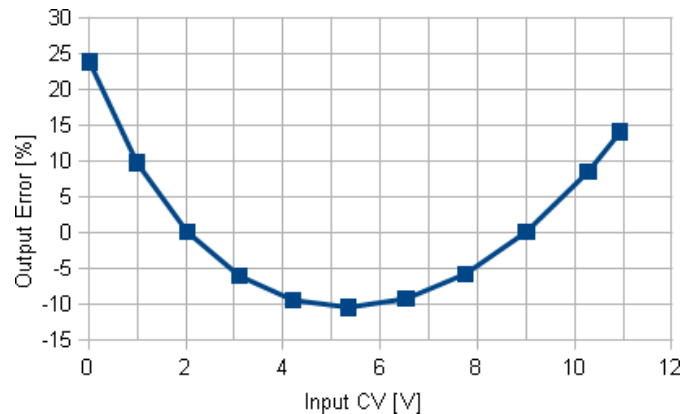
OTA and discrete implementations of this topology are shown in Figures 27 and 28. In both, the input CV is converted to a current and matched to the output current via the op-amp and associated resistor. This op-amp modulates I_{bias} to keep $I_{out,ota}$ equal to CV/R . This I_{bias} is then converted to a voltage via another resistor and fed to the exponential converter.

The main drawback to this configuration is the single sided input. You can only drive current into the I_{bias} pin of an OTA one way. This creates a number of errors in the exponential converter, as discussed earlier. This could be fixed with an

Fig. 27. LM13700 inverted OTA V_t multiplier.Fig. 28. Discrete inverted OTA V_t multiplier.

uncompensated, four quadrant multiplier, like the LM1496, but then you would not have the multiplier on the same die as the exponentiating pair, and would lose the main advantage of this method. Also, because I_{bias} is no longer being held constant, the non-idealities of the transistors must be more carefully considered. For example, Equation 38 is based upon the assumption that the collector currents of the OTA sum perfectly at the emitters. But, due to finite base currents, which vary with temperature, this is not the case. Also, due to the Early effect, the collector currents vary with the collector voltage, which severely hinders the discrete implementation shown in Figure 28. Both of these factors, combined with basic limitations of the simple model of transistor behaviour used to derive Equation 38, give the poor linearity curves shown in Figures 29 and 30.

Further work could be done to improve this linearity, but the level of complexity does not seem worthwhile for a circuit that is not theoretically perfect in temperature tracking. The fact that it is a unipolar design with noise issues does not give a lot of incentive to explore this path much further. Again, this is intended as a simple technique, with minimal parts, that does an acceptable job. Once the topology becomes too complicated, other techniques are far better for the price.

Fig. 29. Linearity error of LM13700 inverted OTA V_t multiplier.Fig. 30. Linearity error of discrete inverted OTA V_t multiplier.

VIII. GILBERT CELL V_t MULTIPLIER

After exhausting all of the “simple” techniques, i.e. ones that can use a single matched pair, a single matched quad, or an OTA, it’s time to try the brute force approach – just multiply by V_t directly. This method was used in the CEM3340, as shown in Figure 31. Since the CEM3340 is also obsolete, a discrete version must be built; in this case using the AD633, a true Gilbert Cell multiplier with high linearity and negligible temperature drift. V_t can be generated a number of different ways, but a common method is shown in Figure 32. In this configuration, two transistors are kept at an equal ratio of currents by an op-amp. The op-amp adjusts the base voltage to keep the voltage constant across 2 different value resistors, creating two different, but constant, currents. The difference in base to emitter voltages is then just $dV = V_t \ln(I_1/I_2)$, which is the inverse of our exponentiating equation.

This V_T base voltage can then be directly multiplied by an AD633 (or other equivalent gain cell) before being applied to the exponential converter. A schematic of this is shown in Figure 34. It is important that the V_t term be applied the the ‘Y’ input of the AD633, and the CV applied to the ‘X’ input. Although the datasheet lists the non-linearity as being higher on the X input, i did not find this to be the case, and it was almost not usable at the Y input.

A graph of the AD633 V_t multiplier linearity and temperature drift is shown in Figure 35. The linearity is not




Fig. 31. V_t multiplier used in CEM3340, taken from datasheet.

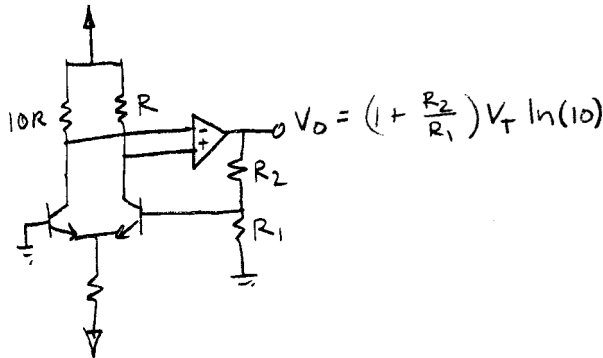


Fig. 32. V_t "bandgap" generator.




Fig. 33. Discrete V_t multiplier.




Fig. 34. AD633 V_t multiplier.




Fig. 35. AD633 V_t multiplier error [%] versus CV [V].

bad, with the characteristic hump still present, but greatly reduced over other methods tried. The temperature tracking is also quite good, and compares favorably to the LOG112 implementation shown earlier. Ultimately there are limitations with how linear the V_t multiplication is, and how accurately the V_t generator tracks the V_t of the exponentiating transistor. One way this was improved is shown in figure 34, where V_t is created by comparing to the reference transistor, leaving the remaining transistor (in a matched quad) free for HF tracking or another exponential converter. It also keeps the collector voltage constant (in contrast to Figure 32), eliminating errors due to the Early effect.

There are other ways the multiplication can be performed, and a few were tried, as the AD633 is a bit expensive (10\$ in single quantities). A discrete method is shown in Figure 33, and an OTA could be used as well. These were tried, and although the linearizing diodes were used on the LM13700, the nonlinearity was too high, and very similar to that seen when using the Y input on the AD633. After testing the discrete implementation, and finding the distortion to be the same, I have to conclude the Y input controls the base voltage on the internal differential pairs, and the X input controls the current through them. Although they both have distortion, the way the current input distorts causes less issues than the crossover distortion as the base voltage input transitions from positive to negative CVs (it appears as an x^2 term). An example of

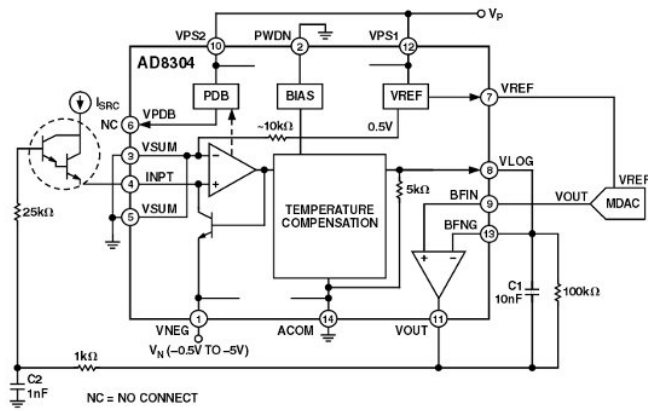


Figure 17. Multidecade Current Source

Fig. 36. AD8304 exponential current source, taken from datasheet.

this error can be seen in Figure 39 in the next section, where a CA3280 was used.



The V_t multiplier works very well, but has a slight issue with nonlinearity. But, this can be kept below 4 cents error, and might be able to be reduced even further by reducing the magnitude of the signal presented to the X input. The temperature tracking is limited by a number of factors, including the simplicity of the temperature model of V_t . Inside of the Analog Device, Inc. (ADI) log amps (Figure 36), they use a more accurate temperature measuring circuit for their V_t multiplier. but, after testing the AD8304, i do not think they are suitable as exponential converters. It was extremely difficult to stabilize the feedback loop around the log amp, the internal op-amps have poor specifications, and the temperature drift was not nearly as good as the LOG112.

The total cost of a V_t multiplier scheme can run as low as 5\$ for a poor linearity LM13700 and LM3046 setup, to 25\$ for a nice AD633 and MAT14 converter. The ADI log amps cost around 20\$, and might be worth revisiting someday. The output current was very stable from the V_t multiplier circuits, and the frequency response was also very good, as the AD633 is faster than audio rate. It has similar performance to the LOG112, but is not dependent upon the accuracy of an internal thermistor, so could be more stable over large production runs.

IX. FEEDBACK V_t MULTIPLIER

Perhaps the most stable technique is the feedback V_t multiplier. This is very similar to the V_t multiplier discussed above, except it nulls out any V_t generator and multiplier errors by having a parallel multiplier track V_t . In this way, the error is reduced to the differences between the two multipliers, which can be kept quite small. A schematic of the implementation tested is shown in Figure 37. It uses a V_t generator balanced against the reference transistor, which is kept at a fixed current by modulating the multiplier connected to a fixed CV. Also controlled by this same multiplier voltage is the exponential CV, so that the two will always be tracking perfectly to hold the V_t generator current fixed.

The linearity and temperature error for this configuration are shown in Figure 38. The linearity is much improved over

Fig. 37. AD633 feedback V_t multiplier.Fig. 38. AD633 feedback V_t multiplier error [%] over temperature and CV [V].

the single multiplier version, and held within $\pm 0.02\%$ for the majority of the range, although it's unclear why this is the case. AD633s were used in both cases, and should have similar distortion. The exact same AD633s were not used, and an extra base resistor was added to the reference transistor, so perhaps one of these affected the results. The temperature performance is also much improved, drifting only 0.2%, which makes sense, as the actual output current of a parallel converter is being monitored and kept constant with temperature.

There are a number of improvements that could be made to this topology if one was to relax the requirement on "analog" exponential converter. The main change would be to use a digital potentiometer, multiplying DAC, or PWM'd resistor/switch combo in place of the multiplier. The ADI X-amp series of VCAs might also be good. Perhaps more true analog methods like dual vactrols or FET opto-isolators could work, but i am not hopeful. A quick test with non-matched vactrols was a complete disaster, and the linearity of a CA3280 was not very good either (see Figure 39). The linearity and matched qualities of an MDAC or digital potentiometer would probably be the best, and would eliminate the 20\$ worth of AD633s and replace them with 4\$ worth




Fig. 39. CA3280 feedback V_t multiplier error [%] versus CV [V].

of parts. The feedback control for these digital techniques might be more difficult, but the functionality could then be expanded to include options like auto calibration. The CV to audio pathway would still be full analog, with only the gain control circuit being digital.

Cost is the main detractor for this topology. Again, a poor LM13700 and LM3046 version could be made for 5\$, but the linearity would suffer. It is probably closer to 35\$ for AD633s and a MAT14. But, for that price, you get a VCO that only drifts 3 cents over a 40°C range, has fast frequency response, and has an output that is rock solid. After getting this circuit up and running, i was completely amazed at the stability of the output current. On my 6 digit multimeter, almost all of the digits stayed completely still after i set the CV. For the other topologies tested, there was always some wavering on the last few digits, but this one held completely still.

X. CONCLUSIONS

There is no perfect way of canceling the temperature drift of an analog exponential converter, and in some ways that imperfection adds to the character of an analog VCO. There are a number of digital solutions that can give you excellent pitch tracking, but for modular synthesis where there are no clear note-off events (giving a pause for recalibration), and modulation frequencies can be in the high kHz range, an analog solution can be advantageous.

For cost and simplicity, not much can beat the PTC thermistor method of compensation. If the exponentiating circuit is closed off from any external air currents or random heat generating parts such as power supply heat sinks, then it can have good tracking, especially if it is only needed for 4 or 5 octaves. It is by no means a perfect solution, but is definitely good enough for monophonic operation, as witnessed by the number of popular early 80's synths that used it. The platinum RTDs show a lot of promise for greatly improving the accuracy and repeatability of this topology, but they have yet to be tested.

Thermal ovens are also well liked, and do a very good job of temperature tracking. Although i have not personally tested this claim, a number of well respected designers swear by

them, so i assume they work very well. If you don't mind a little extra current draw ($\sim 10\text{mA}$) and a bit of wait time for the VCO to warm up, this is probably the way to go. It is slightly more expensive than thermistor compensation, but can be done pretty cheaply if the imperfections of the LM3046 are tolerable.

For relatively small temperature drifts, the exponential compensation technique is very good, and much better than a thermistor. But, for large drifts it is far worse. Discrete implementations are not worth pursuing due to their requirement of unipolar CVs. The SSM2164 has linearity issues, although is completely usable, and the THAT2162 might give better performance.

There are a few conditions under which OTA compensation schemes are worthwhile, due to their low cost. If you only need two octaves, the single OTA scheme is quite cute, but not really worth it, as the temperature drift of a thermistor compensation technique over that small range is negligible as well. The dual OTA scheme is acceptable for 6 octaves, and can be pushed to 8 or 10 if you don't mind 15 cents of error. This will have better thermal tracking and be more repeatable than the thermistor method, but at the expense of nonlinearity. The discrete OTA method adds a bit of complexity to what is intended as a simple solution, but does remove the need to use the linearizing diodes as an exponentiator, which should improve temperature stability. The inverted OTA method is probably not worth pursuing as it has neither good tracking nor linearity.

The V_t multiplier schemes are very good, but also very expensive and complicated. If cost is not an issue, either will do a good job, although the feedback multiplier is much better and has shockingly good output stability. The need to use the AD633 or equivalent multiplier really limits the applicability of this approach. Using some form of digital attenuator would both improve the linearity and decrease the cost. Digital potentiometers can be had for a dollar or two, and with temperature coefficients of a few ppm/°C.

After spending quite some time studying the various topologies, my personal preferences are as follows: If i were to build a high-end VCO, i would use the feedback multiplier with a digital attenuation stage, or AD633s if there was a need to be purely analog. If i wanted to make a cheap monophonic synth, i would use thermistor compensation, most likely with the platinum RTDs. The thermal oven and LOG112 implementations might fit somewhere between those extremes. The LOG2112 is actually not too expensive, as it breaks down to 12\$ per exponentiator. I am not sure i can envision a use for the OTA schemes (perhaps a 6 octave VCO), and the SSM2164 only seems useful if you are using the SSM2164 for filters and VCAs as well, so you have it handy.

For future work, i would like to do an in-depth analysis of the thermal oven technique, and comparisons of the THAT300, MAT14, and LM3046 in the same circuit. From my tests, the MAT14 is superior to the THAT300, and i have not yet tested the LM3046. The digital attenuators and platinum RTDs also need to be evaluated.

XI. ADDENDUM – SEPTEMBER 15, 2016

There were a number of characteristics of the initial results shown in the previous sections of this paper that did not sit well with me, so i went back and re-evaluated them. As it turns out, the contacts on my protoboard and some of the jumper wires were slightly corroded, adding a half of an ohm to a number of important connections, giving slight errors. Therefore, to give a better idea of the accuracy of the various topologies presented, this addendum is added with new results.

A. Platinum RTDs

The platinum RTD error graph shown in Figure 7 is slightly incorrect, due to a misinterpretation of the datasheet. An accurate representation is shown in Figure 40. This is at 5V CV for a 500 Ω , Vishay, PTS series RTD in both op-amp and resistor divider configurations. These configurations are shown in Figures 5 and 6, with R_1 being replaced with the RTD plus a series resistor. The series resistor waters down the platinum RTDs 3850ppm/ $^{\circ}$ C to the required 3356ppm/ $^{\circ}$ C. The op-amp version uses a 29 Ω series resistor, and the resistor divider version uses a 17 Ω series resistor with a 27k Ω divider resistor (R_2 in Figure 6). The accuracy bounds and temperature dependencies are taken from the Vishay PTS datasheet.

The op-amp version has slightly better accuracy due to not having the divider resistor effectively in parallel with it. The larger the ratio of the divider resistor to the RTD, the lower this difference becomes. But, both versions have good temperature dependency, and very low variance. Unfortunately, platinum RTDs only come in values less than 1k Ω , making the op-amp configuration difficult to implement. Due to self-heating, the total voltage across a 1k Ω RTD must be limited to less than 250mV, so op-amp offset voltage drift quickly becomes a problem as it's on the same scale as the total output voltage. A chopper stabilized op-amp is required for this application, which is more expensive than standard op-amps.

As can be seen in Figure 40, the error of the platinum RTD compensation moves down for both increases and decreases in temperature. This gives an opportunity to increase the accuracy if large excursions are expected in one direction, but not the other. For example, it is more likely that a circuit might increase by 20 $^{\circ}$ C (e.g. if put into a rack full of warm gear) than it might decrease by 20 $^{\circ}$ C (e.g. put in a refrigerator). By changing the value of the series resistor, the average error over a given range can be reduced. A plot of these errors for different series resistor values is shown in Figure 41. These are for a 100 Ω RTD with a 10k Ω divider resistor. If large temperature increases are expected, but no decreases are expected, a 3.5 Ω to 4 Ω series resistor would work well. Otherwise, for equal variations, the 4.5 Ω series resistor is best.

The resistor divider RTD configuration was built up and tested using a PMP4201 matched pair and a 100 Ω Vishay, PTS series, 0603 RTD, and two different series resistors (4.5 Ω and 3.6 Ω). The 100 Ω RTD was chosen because it reduces base resistance errors. For a full 10 octave CV input, the self heating on the 100 Ω RTD is close to its limit, so for situations where base errors are not an issue (or higher input impedance is required) the 500 Ω RTD would be better. The 0603 package




Fig. 40. Platinum RTD compensation error [%] versus temperature [C] at 5V CV. The solid lines are for an ideal RTD, and the dashed represent the bounds on RTD variance. The red line is for use in an op-amp, and the blue line is for use in a resistor divider.




Fig. 41. Platinum RTD compensation error [%] versus temperature [C] at 5V CV for resistor divider configuration using different series resistors (10k Ω // (100 Ω RTD + series resistor)).




Fig. 42. Platinum RTD compensation error [%] versus CV [V] with a 4.5 Ω series resistor.




Fig. 46. Linearized LM13700 V_t multiplier schematic.




Fig. 47. Alternative method of driving a controlled current into the bias pin of the LM13700.

issue with FETs at high temperatures. Another option, which eliminates the need for a transistor, is shown in Figure 47. In this configuration, the OTA is driven with an op-amp to keep the current constant into the I_{bias} pin. But, since the negative supply is now varying, the other half of the LM13700 is difficult to use.

The main drawback of the LM13700 is the non-linear output. This non-linearity increases with the input signal as a percentage of the total range. For example, in Figures 48 and 49, the performance of the LM13700 is shown for 50% and 25% of full range, respectively. It can be seen that the total non-linearity is greatly reduced for the lower input range. Unfortunately, at the lower input ranges, both noise




Fig. 48. LM13700 V_t multiplier compensation error [%] versus CV [V] at 50% of full scale input.




Fig. 49. LM13700 V_t multiplier compensation error [%] versus CV [V] at 25% of full scale input.

and temperature compensation begin to suffer. Noise becomes apparent as frequency jitter at the VCO output, which may not be an issue. It's unclear exactly what causes the worsened temperature drift, as lower input ranges should reduce the OTAs dependence upon V_t , but it is most likely due to the OTA offset voltage (and associated drift with temperature) becoming a larger percentage of the input signal. Regardless, both have really good temperature performance, but the non-linearity of the higher range configuration is only good for 7 or 8 octaves.

The larger errors at high currents in comparison to the AD633 are a result of both the distortion and the 560Ω base resistor. An op-amp can be used at the output of the OTA to reduce this error, but this adds op-amp offset voltage and offset current errors, so a precision op-amp must be used. Alternatively, HF tracking can also be applied to reduce this error, although not entirely, as β modulation with temperature will still cause a small change.

To improve the performance of the LM13700, both OTAs in the package can be used to cancel out the inherent distortion.

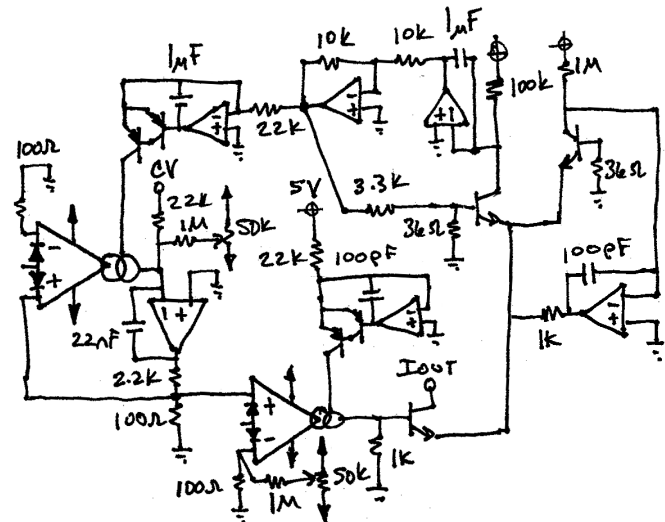


Fig. 50. Dual OTA V_t multiplier schematic.




Fig. 51. Dual OTA V_t multiplier compensation error [%] versus CV [V].

This is shown in Figure 50. The downside is that it requires twice as many components, including two current sources; but, it is far more linear. The dual OTA scheme also allows for both 2nd harmonic and feedthrough trimming, improving linearity even further. The performance of this setup is shown in Figure 51, which is quite good, except for the high frequency error, which is larger due to a 1k Ω base resistor being used. The temperature drift is comparable to the single OTA scheme, but with much better linearity.

C. Conclusions

Ultimately, the V_t multiplier schemes are quite good, but very complicated, both in terms of part count and circuit stability. They necessarily require trimmers, which is costly in terms of calibration time, and are another point of drift. The LM13700 versions are less expensive than the AD633, but come with more op-amps and transistors and their associated errors.

The platinum RTDs are amazingly accurate, and extremely simple to use. Although the RTD itself is expensive (2.50\$US@1 and 1.35\$US@100) they allow the use of a single matched pair, rather than a quad, and require fewer op-amps, making them cheaper than the V_t multipliers. The RTD does need to be in good thermal contact with the matched pair, and also needs to be insulated. I can not stress this last part enough, as it makes an incredible difference in both short term VCO jitter and long term temperature drift. Even a blob of hot glue over the components would be a vast improvement over leaving them exposed to the ambient air.

After this re-evaluation, my personal preferences are as follows. For a super inexpensive VCO, or areas where exact pitch tracking isn't required, compensating with a typical thermistor (e.g Panasonic ERAV series) is fine. For a very high end VCO, a single AD633 V_t multiplier would be good. For almost everything else, a platinum RTD used as shown in Figure 6 (where the RTD is connected to the reference transistor) with a slight shift in the series resistor to favor warmer temperatures is more than adequate. I would most likely use a 10k Ω divider resistor, 100 Ω , 0603 RTD, and a 4.1 Ω series resistor. The 40 $^{\circ}$ C tests performed here are quite

extreme, and typical use will not see such large temperature variations. Over a 20 $^{\circ}$ C range a platinum RTD will give less than 3 cents of frequency drift, which should be good enough for almost all applications.

XII. ADDENDUM II – JANUARY 28, 2018

A. Thermal Oven Compensation

Thermal oven compensation is one of the oldest forms of compensation for exponential converters. It is perhaps the most straight forward technique – keep the transistors at a constant temperature. This is accomplished by a transistor on-die that senses the temperature, and a control loop that increases or decreases heat by modulating the current in another transistor. In a number of early synths this was accomplished with the μ A726 matched pair by Fairchild Semiconductor. Unfortunately, this part is no longer produced and costs around 100\$US per unit for what often ends up being a low quality or fake part. Modern thermal ovens are built around matched transistor arrays like the LM3046, with an external control loop.

An original μ A726 was tested, and its thermal response is shown in Figure 52. An 82k Ω resistor was used for the temperature set input and a 0.25 inch layer of foam was placed around the package to keep thermal gradients down. This latter modification is important, as differing temperatures across the package will force the exponentiating transistors to be at a different temperature than the temperature sensing transistor. This was particularly relevant to the tests performed here, as a heat gun was used to change the ambient temperature; the airflow from the heat gun is directional, and creates thermal gradients.

The temperature tracking of the μ A726 is adequate, but not great. The drift at 5V CV is 0.5%, which is worse than platinum resistor compensation, but better than thermistor compensation. The tracking is actually worse than the data might suggest, as these tests were done at 50 $^{\circ}$ C, whereas the rest of the tests in this paper were done at 65 $^{\circ}$ C. This means that the comparable drift is closer to 1%. There is a large linear offset in the original data taken which was due to resistor drift in the CV attenuator that shifted the base voltage by 200 μ V.




Fig. 52. μ A726 thermal oven compensation error [%] versus CV [V].

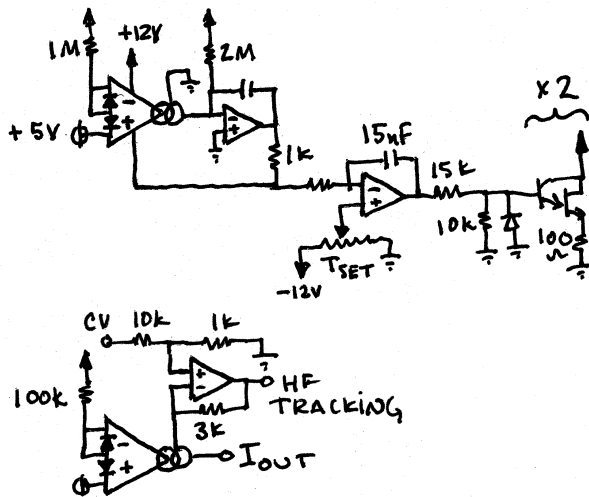


Fig. 53. Implementation of LM13700 heated exponential converter.

Adjusting the graph for these external errors shows only the temperature drift error.

It should also be noted that the exponential tracking of the internal matched pair is inferior to modern devices. Leakage currents begin to become a problem as early as $-4V$ CV, and emitter resistance gives 1% error at $3.5V$ CV. By comparison, the PMP4201 matched pair shows little leakage down to $-7V$ CV, and has 1% emitter resistance error at $4.5V$ CV. For all of the comparisons made in this paper $0V$ CV represents $5\mu A$ of collector current, so CV levels are equivalent between devices.

The temperature drift in the $\mu A726$ is the result of a weak feedback control loop internal to the device. This might be able to be fixed with a current source replacing the temperature set resistor, as this resistor also sets the gain of the feedback loop. Since a current source is a much higher impedance, the gain would go up, and hopefully the temperature regulation would be tighter. But, as the parts are obsolete, it didn't seem worthwhile to test out this hypothesis. Instead, using an external feedback loop on a transistor array was tested. Since the LM3046 has also been discontinued, there are no longer any inexpensive options available. To solve this problem, *elektroewe* on the Muff Wiggler forum proposed misusing the internal transistors of an LM13700, with the Darlington buffers as heaters.

A schematic of the implementation tested is shown in Figure 53. The input current mirrors are used as the exponentiating pair, with the current reference mirror also being used as the temperature reference. Other parts of the LM13700 could serve these functions (see Figure 23), but the current mirror works well as it drives the OTA directly, and the current can be reversed for triangle wave generation. The Darlington pairs are the only parts of the LM13700 that can source enough current for a heater, but are on the opposite side of the die from the input current mirrors. For this reason, the reference mirror needs to be used as the temperature sensor (as opposed to using one of the Darlington buffers) to keep thermal gradients between the exponentiating transistors to a minimum.

The external regulation provides excellent temperature




Fig. 54. LM13700 thermal oven compensation error [%] versus CV [V].

tracking, as is shown in Figure 54. There is, again, an offset due to an unexpected resistor drift in the CV input attenuator, but there is essentially no temperature drift. The high frequency nonlinearity is the worst of all setups tested so far, with 1% error at $3V$ CV. The emitter resistance can be compensated for with feedback of the output current, but there are other problems with this design that make it less than ideal. First off, the requirements on the opamps are quite tight, as offset voltage drifts are directly applied to the exponentiating transistor, with $150\mu V$ giving a 10 cent oscillator drift. Secondly, there is a limit to how hot the internal transistors can get. In the setup used here, the buffers were run at $20mA$ a piece, and the PDIP part insulated with foam, and it still fell out of regulation above $50^\circ C$. Finally, it takes some careful design to ensure that the device does not go into latch-up due to reverse bias of the substrate diode connections.

B. Conclusions

A proper LM3046 or MAT14 heated converter was not tested as the LM13700 performed well enough to assume that these other devices, with better transistors, would perform at least as well. The main drawbacks to the heated converter are the extra current required (LM13700 - $40mA$, $\mu A726$ - $10mA$) in comparison to other techniques, and the thermal settling time. For both devices tested here, it took one minute to settle to within 1% error, three minutes to 0.1% error, and 15 minutes to 0.01% error. This is for both startup, and during ambient temperature changes. The increase in current going from low to high notes heats the device enough to create a change which requires these settling times.

The inexpensive and simple nature of the LM13700 heated core is interesting, but I would prefer the platinum resistor or V_t multiplier schemes for most applications. The warm up time and extra current are large detractors for me, along with the inability to regulate above $50^\circ C$. And going to better transistor arrays might not solve any of these problems. The added cost of low drift opamps starts to make the benefits of the LM13700 less apparent (the TL072 drifts $18\mu V/^\circ C = 1.2\text{cent}/^\circ C$). But, the LM13700 can still be used as an exponential converter with these other compensation schemes for a low cost VCO.