What's the Buzz? – Part II

air·foil (âr'foil')
n.
The shape of a wing or blade (of a propeller, rotor, or turbine) or sail as seen in cross-section.
[http://en.wikipedia.org/wiki/Airfoil]

SJDT170R1400. This "package" may have short legs, but check out those curves! A normally on, 1.4 Ω SemiSouth JFET in a TO-263 surface mount package.

There are some things in engineering and technology where an iconic image of the thing is widely recognized as *the* technology. In other words, where the *shape* of the thing is inseparable from the thing itself. There can be little doubt that an *airfoil* is one such concept. The box "*Curved surfaces with a lift*" shows that when you have the pattern of an airfoil framed in your mind, it is easy to see it repeated throughout nature. (Negative feedback, while less tangible than an airfoil, is another thing widely repeated in nature–but that is a subject for Part III.) An airfoil is a shape that generates lift in the atmosphere (airfoils are useless in space, of course). Movement of an airfoil through a relatively high-pressure gaseous medium at terrestrial speeds is all that is necessary to get more lift than drag. I learned that as a young boy growing up in a household in which our livelihood was provided by my father, a professional aviator. The physical principle underlying the shape of an airfoil was explained to me early, and I have held it in awe ever since. Food on the table creates that kind of respect in a child!

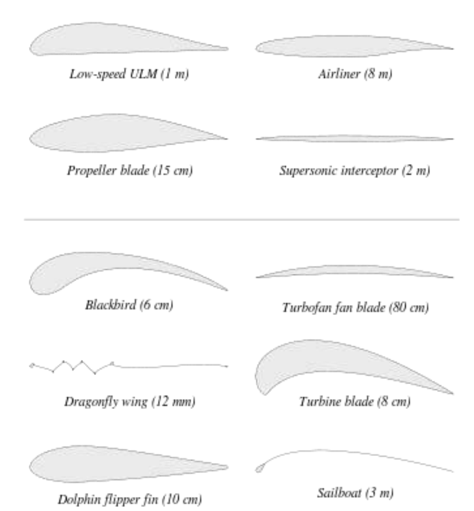

Curved surfaces with a lift. Airfoils are diverse in shape and in origin. Once a physical understanding of how an airfoil works has been gained, it becomes easy to explain their many manifestations. [Image from Wikipedia article on airfoils.]

In this second part of a three part series of articles, I tackle the job of understanding why a SemiSouth JFET makes an audio amplifier deliver the sound a DIY audiophile seeks better than many of the alternative transistors. As the opening paragraph is meant to suggest, it will be the *shape* of the electrical characteristics that I assign much of the responsibility for achieving distortion reduction due to so-called "load-line canceling." Just as the shape of an airfoil causes the lifting effect when air flows over the airfoil, I will also describe the ubiquitous "shape" of something I call the amplifier transfer curve to explain the "sweet spot" effect previously highlighted in an article posted by Nelson Pass. We will also see how the right combination of measurements and analysis will reveal this characteristic shape empirically in a simple amplifier that any DIY audiophile can build. Yep, you really can do it yourself even if you don't have a fancy curve tracer.

I've been intensely curious about the subject of this second installment in my series of articles. Based on numerous posts and emails I've received, the underlying value of the SemiSouth JFET is too much to dismiss. It must be a real and positive effect of audio reproduction that is delivered by a SemiSouth JFET, better than most alternatives currently available, to create this buzz. Furthermore, judging by the many inquiries I have received about different part numbers, it is not just an SJEP120R100 that possesses these properties, although that part seems to remain the most popular. The research I did for this article has led me to conclude that our community's experience-based belief in the "R100" is well founded and rooted in a curious "defect" in the SemiSouth JFET from the point of view of the power switching transistor marketplace. This defect was first pointed out to me by the company's director of marketing long before we had heard from Nelson Pass; namely, the lazy transition from ohmic conduction to saturated conduction (see the "*The fat lady sings*" box later in the article to see what I mean). And while doing my research I have not found a quantitative explanation in other writings for why this defect in a power switch is a valuable property in an audio amplifier. Perhaps one exists better than I have provided here. If so, it will be quickly pointed out to me, I'm sure, by one of you. But, frankly, I feel a lot better about what I think I know about this question after I researched it for more than a year. I hope you will find it as interesting to read and analyze as I did to research and write.

The Shape of Things - form

As a practical matter with airfoils it was necessary long ago to standardize the connection between *shape* and technical *performance*. To achieve this it is first necessary to agree on what we mean by the shape. How will it be described in a non-ambiguous way that everyone can understand? This was done by an ingenuous system called the "NACA airfoil" that penetrates even today down to the semi-popular aviation literature. But what equivalent do we have for transistors used in linear audio amplifiers?

Transistor part numbers are something, if not everything, we could want for this. In the box "*Picking through the jewelry box*" I have highlighted output curves taken experimentally for a number of SemiSouth part numbers. My thinking about this has evolved over time. When I started, I was inclined to discount that the shapes of the output curves were all that identifiable by part number. Instead, I assumed that die-to-die variation would dominate the diversity in the shape of the curves. Therefore, I assumed that the results of load-line canceling would be pretty much specific to an individual part. In other words, I expected that folks would discover that they preferred some parts in their inventory over others and that it would be hard to replace those features without screening parts individually. To a certain extent, this is surely true based on comments I have I have received from DIY audiophiles. The preference for part matching in the community certainly reinforces that belief. But on what objective basis can this point of view be justified?

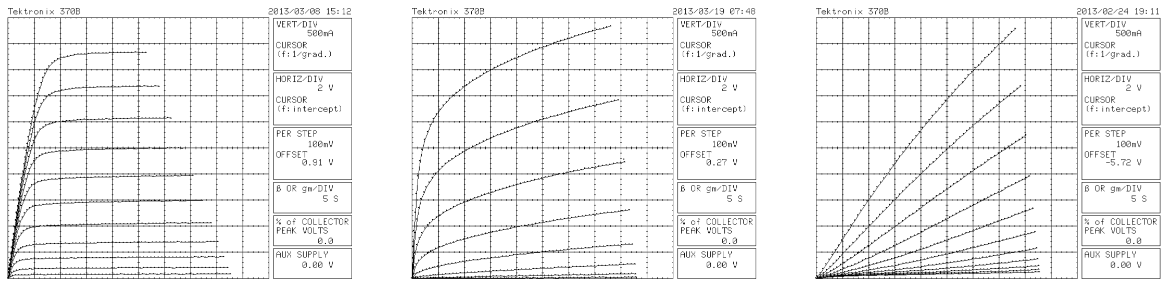

(a) (b) (c)

Picking through the jewelry box. Above are three output curves measured on a Tektronix 370B curve tracer with the same vertical and horizontal resolution (V_{DS} [2 V/division]; I_D [500 mA/division]). In each image the "triode region" (the correct name for the ohmic region in FETs), the saturation region, and the transition between the two appear much different. In (a) we see all three regions; in (b) mostly just the transition region; in (c) it's mostly just the triode region. If you were to look at the same curves with properly scaled resolution, they would look pretty much the same, namely like a pentode as in (a). This was by design, for the power switching applications that these three parts were intended scale by on-resistance and saturation current. But in the audio world, we pick parts not for on-resistance and saturation current, but for gain and the ability of the non-linear output curves to compensate the non-linear transfer curves. To make that possible, the transistor must be able to dissipate enough power to be biased to the correct operating point. In that context, looking at the curves in the way displayed is useful because it reveals the function of the transistor and the linear amplifier containing it: for example, preamp or power stage with high drain resistance or low. The part numbers behind the figures are: (a) SJEP170R550; (b) SJEP065R055; (c) SJDP120R045.

Through the process of researching this article I have concluded that a part number can be a reproducible harbinger of certain expectations. If so, can a particular transistor brand reliably produce an objectively better "sweet spot" than another brand? Answer: If it is a SemiSouth SiC JFET, then probably yes! Delving deeper, will the sweet spot be found at bias points specific to a part number? Answer: again, yes. Well then, if we can find the approximate sweet spot by part number, does the harmonic content of the residual distortion depend upon the part number? Answer: with pure load line canceling without negative feedback, probably not. Instead, fine tuning the bias to get right on the sweet spot has the effect of turning a single-ended amplifier into a balanced push-pull type because of the shape of the amplifier transfer curve. The balance of this article will seek to justify these conclusions. In the process, I will leave you with an analytical framework for understanding my conclusions that you should be able to put to work to optimize DIY audio projects using just about *any* transistor, not just a SemiSouth JFET.

In coming to these conclusions I was helped by reading the opinions and experiences of others on DIYAudio.com. I have also received hints by following the interests of a few especially informed DIY audiophiles who have contacted me about specific part numbers. I was amazed when one clever gent from Norway led me to a whole sleeve of parts I didn't realize existed (SJDP170R1400), which shows how much I paid attention to some of the smaller JFET's that SemiSouth brought to market near the end. Soon thereafter, another wise part scavenger from Seattle contacted me about the R1400 part, and I had to tell him that I was sold out (since then I discovered I have a batch of the same parts in a compact surface mount package, which is the

SJDT170R1400 highlighted at the beginning of this article). But he motivated me to find a small number of another similar part, the SJDP120R340. He got all of those! And later a DIYAudio regular selected a quad of SJEP065R055. While that part has an attractive set of output curves, strictly speaking as a part number it doesn't exist! Not, at least, as a formally released commercial product. It was an engineering build based on high threshold die from the SJDA065R055 production that was commercially released by SemiSouth. In short, I have been humbled by the knowledge displayed by many a DIY Audio man.

The Shape of Things - function

In the box "*Form meet Function*" are images of three aircraft that one Mazzola brother or another have piloted. If you were to ask most rational people what the three have in common, I doubt many would find much to agree upon. The mission of the Boeing 747 is transcontinental transport at very high subsonic speeds at a cost that will make the owners money while satisfying the customers' insistence on reasonable speed, plenty of comfort, and proven safety. The 747 has delivered all of those and more. I remember flying as a passenger in 747's a few times during the 1970's. In each case the pilots squeaked the landings. Radar altimeters, auto-throttles, other semiautomatic landing equipment, and the exceptional skill that comes from being at the top of the airline pilot profession combined to make it look easy. The passengers (including me) spontaneously applauded as though we had just enjoyed a performance of fine live theater. I have not witnessed passengers do that in any other airplane.

The Cessna 172 obviously does not have the mission of a B747. Instead, it performs whatever mission its owners need, often on the "fly" so to speak. Surely it has succeeded at this most general of general aviation roles. It is claimed on Wikipedia that the C172 holds the record as the most produced general aviation aircraft. The vast numbers of types and modifications to this aircraft are proof enough.

The PW-5 sail plane, designed and manufactured in Poland, cannot possibly serve the workhorse missions of the first two aircraft: It hasn't an engine! But what it delivers in aerodynamic performance at "slow" speed the other two aircraft cannot match. Slow is relative; its best glide speed is still highway speed in automobiles. You get a real feel for the speed finishing the landing roll with your body inches above the asphalt runway squeezing the hand lever with all you've got to actuate the barely effective single-wheel brake. A healthy headwind and full spoilers are more effective at slowing the glider down.

(a) (b) (c)

Form meet function. Three aircraft (a) Boeing 747-400; (b) Cessna 172; (c) PZL PW-5 that could hardly be more different in appearance, purpose, and performance. But they all fly. At least one Mazzola brother, motivated by an aviation minded father, has flown each aircraft above. But I'm not the one that flew 747's for United as first officer. That would be my younger brother, the smart one.

There are many differences between all three aircraft, of course, but in terms of function dictating form, the shape of the wings of each aircraft is surely defining of their purpose. The B747 wing is optimized for high speeds at high altitudes, so it must be massively reconfigured with spoilers and flaps to permit safe landing speeds. The C172 wing is not much for promoting speed, but in the hands of a professional, the lift and versatility in the bush-like conditions of short fields, mud, or even water, it performs great. But it also performs well for the beginner, as many first-solo pilots will attest, including me. In contrast, the PW-5 wing has only one purpose, great glide ratio while being strong and light. The PW-5's specified 34:1 glide ratio is not the best among sail planes, but believe me, landing one on a spot requires the spoilers. Otherwise, the bird just doesn't want to give up the air. When low and slow and spoilers stowed it *feels* like it has an engine with a few revs still on the throttle as you glide down the runway in ground effect seemingly forever.

But for all the differences in their wings derived from utterly different missions, one thing is certain: Each aircraft, and any other that actually flies, has a wing that generates more lift than drag. That is why they fly. And by the way, even without the engines running a Cessna 172 and a Boeing 747 will glide for quite a distance if loss of engine power becomes a short term concern for everyone aboard. Check the web for the real life stories of more than one jet airliner with no engines turning that glided to a landing without loss of passengers. For the pilot of the PW-5, this is not news. As the joke goes among those fortunate enough to be glider pilots, *every* landing in a glider is an engine-out emergency. An object that does not generate enough lift to balance gravity is a free falling object. Bombs, so to speak. Consider it a definitive way to sort things out: either your kind of shape falls out of the sky or it doesn't.

That brings me to my point. Is it possible to do analysis to show that a transistor does or does not have a sweet spot? If you look at the output curves in "*Picking through the jewelry box*" we see a selection of three parts. One that has a more definitive transition from the ohmic region to a flatter saturation region, one that has a more gradual transition to a sloped saturation region, and one that might be called the "octopus" of output curves for the spiny way in which the curves evolve, with no strong saturation accept at the lower (and useless) gate voltages. When it comes to the sweet spot I'll go out on a limb and make the following two statements:

- *Every* transistor employed in a linear amplifier with fixed speaker impedance will show a "sweet spot," meaning an operating point where the distortion is measurably less than other operating points nearby.
- *Some* transistors have sweet spots that deliver more output power at a given distortion level; and it's the *shape* of the output curves that determine which ones are better.

In other words, if I drop just about any transistor into an audio amplifier, an operating point can be found that improves the linearity of the amplifier. This operating point is specified in the direct current sense, that is, it can be located by adjusting the bias current of the transistor and the effective power supply voltage. It is also a function of the speaker impedance, but since speaker impedance is generally not adjustable in the sense of the other two variables we will treat it as a fixed value. However, not every transistor type has the right properties to make practical use of the sweet spot. The analogy to the wing is that just about any shaped wing will generate lift, but not every wing can be used for every aircraft. (Duh!) How do we separate those transistors we want to use for operating an amp in a sweet spot from those not worth the effort? One way I suppose is trial error, and I certainly want to take advantage of everything that has been figured out by a community as knowledgeable as the DIY audiophiles. But I also want to know why, and for that reason I have invested more than one year of research (part time) to quantify the features of the transistor that make the difference.

If you scratch the surface of popular explanations for the underlying physics of the airfoil, as is easy to do with the internet, you will surely stumble on a claim that, at least in the past, calculating the aerodynamic forces for a real airfoil is too complicated for practical engineering. But this is changing. Mississippi State University, where I work, is a world leader in unstructured grid generation software. So much so, MSU software is used in the aerospace industry to reduce the need for wind tunnel testing. But long before software and super computers, there were wind tunnels and mathematical expressions. The NACA airfoil allowed real wings to be designed with existing engineering resources, reused appropriately. This may have limited the design space, but it led to practical wings with pre-engineered properties. And all of this was mostly worked out in the first half of the Twentieth Century, barely 30 years from first flight and well before the first electronic computer. NACA numbers give precision to the definition of the shape and a direct link to the performance expected in the application. Lift is assured when an airfoil moves through air. Beyond that, NACA numbers separate airfoils into categories that lead to their selection in applications through engineering.

As it turns out, any transistor is as certain to deliver linearization through load line canceling as an airfoil will deliver lift in an air flow. All that is required is a relationship between the current flowing through the device and the potential across the controlled terminals (drain-to-source in a JFET) that opposes the change in current caused by the signal on the control terminals (gate-to-source in a JFET). Every transistor shows the property of a direct relationship between drain or collector current and drain-source or collector-emitter voltage. While it is most vigorous in the ohmic region of the output curves of the device, the so-called saturated region shows a more modest continuing relationship that leads to the idea of "drain resistance" in the small-signal model. The action of the load and the transistor in the amplifier means that the transistor will compensate for drain current acceleration associated with transistor transfer curve non-linearity by countering with a drop in drain current when the drain-source voltage decreases. The effect increases as the device current is modulated into the ohmic region which leads to the large-signal harmonic distortion process known as clipping. It is a curious fact that a small dose of the effect that eventually becomes linearity killing clipping produces a beneficial compensatory effect called load-line canceling that improves linearity. (When this happens in biology it is called *hormesis*.) But, short of clipping, how well does this compensation work? Just as with the airfoil, the details of the transistor's properties must be consulted to answer the question, and the transistor number is an overall indicator of the properties of the transistor design even though every transistor marked with the same number may not be precisely the same. A transistor number can allow us to select transistors for properties relevant to linearization through load-line canceling just as the NACA numbers tell us details of airfoil performance. But how can we divine these properties? And how do we connect properties to performance in a quantitative way that allows us to achieve the rather challenging job of 0.1% total harmonic distortion with agreeable residual harmonics?

Quantifying load-line canceling

The answer to that question depends on your line of work. Suppose you are a circuit designer for an analog integrated circuit manufacturer. That's a point of view that will likely involve comprehensive circuit design with global feedback and reuse of engineering supplied by the foundry that will make the integrated circuit. That means you use their transistors, not yours. If you are a do-it-yourself audiophile with an interest in solid state, and a follower of Nelson Pass and his point of view about amplifier design, then you value simple circuits with little global feedback. That does not mean your experience is not influenced by negative feedback, because it is (how much is the subject of a complete article that will be the third in this series), but that is not the only thing. You, as a community, have learned to use something about the transistors themselves. In this section I would like to offer my understanding of what it is about the transistor beyond simply common-source gain sacrificed on the altar of negative feedback. It has taken me awhile to arrive at an understanding I'm comfortable sharing with a public audience. After all, while I appreciate the many expressions of respect and interest in my work, I am only human. I have to admit that I started writing this article with some preconceived notions that have definitely evolved. And I have conceived and tested several analytical methods for quantifying the single most important contribution that SemiSouth made to your hobby, namely, the *shape* of the output curves of the SemiSouth vertical channel JFET. Combine this with fixed-impedance speaker properties (admittedly an approximation), and we can explain the significant reduction in harmonic distortion through load-line canceling that I see with SemiSouth JFETs.

Nelson wrote an article about the "sweet spot" [http://passlabs.com/articles/the-sweet-spot] in which the topic of load-line canceling figures importantly. I can say I learned two things from this article that represent important starting points. First, Nelson showed that load-line canceling works for any type of transistor (or vacuum tube for that matter) that might be used in an audio amplifier. Whether it is a MOSFET, or a BJT, or a JFET, they all can deliver reduced harmonic distortion through this process. My conclusion is that the process is inherent to the general characteristics of an electronic valve (transistor or tube), which is that the drain or collector current is dependent upon both the gating potential (V_{GS} or V_{BE}) and the output potential (V_{DS} or V_{CE}). And because of the physics of electronic charge transport (with some interesting exceptions) the drain current increases when either potential increase, which is all the amplifier load needs to produce harmonic cancelation. The second thing I learned is that this effect can add to the improvement in harmonic distortion provided by negative feedback. That came as a revelation to me. Nelson shows this in his article by cleverly using all of the available open-loop gain from each transistor in linearizing feedback that comes from using source or emitter followers. And then he shows that a sweeter spot still can be found by adjusting the bias supply (he says adjusting bias current or speaker impedance will also work and I agree). The additional linearity is assumed to come purely from load-line canceling. My goal is to finish the idea that Nelson started by creating an experimental and analytical procedure that any DIY audiophile with a modest amount of technical preparation and experience can pursue that will lead to solid quantitative results. In other words, I mean to explain a simple way to find the sweet spot for your transistor operating in your amplifier without the guess work.

The technical approach I use in this article is the opposite of Nelson's. Using no intentional negative feedback I present a quantitative way to estimate the change in the *shape* of the amplifier's transfer curve that results from load-line canceling alone. We will isolate the sweet spot in the relevant performance curve the way a pathologist might systematically sort out the various microbes in a sample until the offending microbe is isolated. The relevant performance is revealed not in the transistor's transfer curve that many are familiar with, but in a related curve I like to call the *amplifier transfer curve*. Normally, we talk about the transistor's transfer curve, which is what a curve tracer strives to measure accurately. Plenty of examples of those can be found in my last article "What's the Buzz? – Part I" [http://www.diyaudio.com/forums/pass-labs/229241-semisouth-boiler-room-21.html, post #208]. But the *effective* transfer curve from a distortion perspective is the amplifier's transfer curve that accounts for the fact that the transistor modulates the current in an amplifier driving a finite load, not on a curve tracer. In the box "*What goes up must come down*" I graphically illustrate the difference. In the upper graph the amplifier load-line conspires with the V_{DS} dependence of the output curves to change the shape of the transfer curve. Nelson in his "Sweet-spot" article described this as the give-and-take between the increasing gain caused by the square-law dependence of I_D on V_{GS} and the deceasing gain from the dependence of I_D on V_{DS} revealed by the load line. This is especially obvious at the interception of the amplifier load line and the highest output curve which clearly occurs in the ohmic region of the transistor. That was a choice I made to emphasize the effect, but driving your amplifier that hard would surely be recognized as clipping. We are actually looking for something more subtle (think hormesis), which in this case could come from higher speaker impedance (i.e., shallower slope in the load line).

Nelson used words and distortion plots to prove the point. But you can see graphically the transition of the transfer curve from second order ("square law") shown in the lower left hand graph in the "*What goes up must come down*" box to a higher order trend with an inflection point shown in the lower right hand graph. It is the *inflection point* that represents the possibility for a sweet spot because the change in shape of the curve from upward increasing to leveling off causes the shape of the amplifier's transfer curve to become momentarily straighter than it would otherwise be with zero load. The seemingly ever increasing drain current of the square law reported by the curve tracer turns out not to occur in a real amplifier. Instead, the drain current begins to show signs of "obeying gravity" like an object thrown up into the air that eventually reaches the top of its travel, known as the apogee. In fact, the math for finding the apogee is similar to that for finding the inflection point in the amplifier transfer curve. In the latter, I_D levels off through the influence of transistor saturation caused by the amplifier loading. The effect is much like trying to straighten a coat hanger by hand. You can bend it in all sorts of ways that make it look straighter than it was, but it is still slightly bent. In fact it is impossible for a pure feed-forward process like load-line canceling to perfectly straighten the amplifier transfer curve. That's the realm of negative feedback for reasons fundamental to negative feedback itself. But despite this, two observations can be made pretty much by inspection of the amplifier transfer curve at or near the inflection point:

1. If the amplifier bias point is set at the inflection point, then the acceleration in the change of the drain current due to modulation caused by V_{GS} (i.e., the change in the slope of the transfer curve) will be zero, which is the very definition of a perfectly linear Class A amplifier. Of course, the change in the slope of the transfer curve becomes non-zero as soon as the audio signal being amplified pushes the instantaneous amplifier bias off the inflection point, but near the inflection point this acceleration can be so small that the amplifier will be much more linear than it would otherwise have been. In fact, the

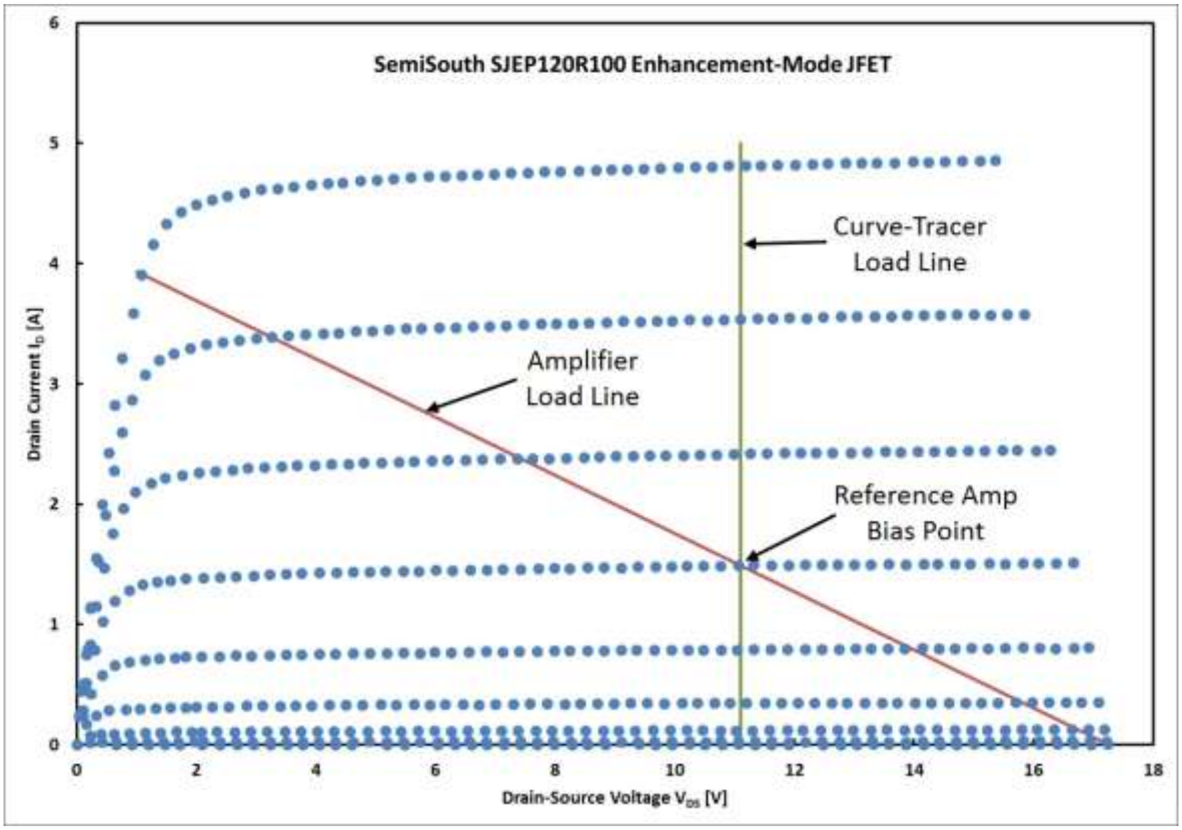

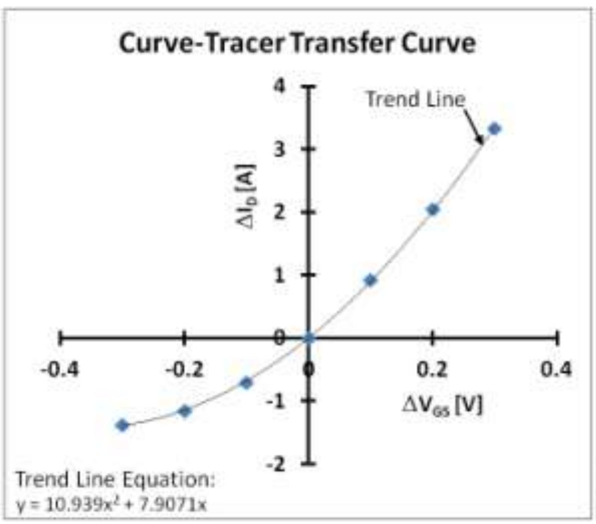

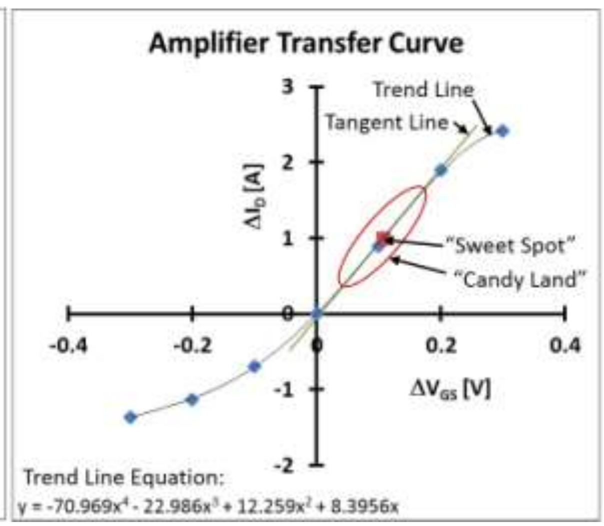

What goes up must come down. The top graph shows the output curves of a SJEP120R100 SemiSouth JFET measured on a Tektronix 370B curve tracer. Superimposed on the output curves are two possible load lines. The vertical load line is approximately what a curve tracer imposes, while the diagonal line is a 4 Ω load line typical of an audio amplifier with an effective power supply voltage of about 17.5 V. Transfer curves can be computed from these two load lines. Because the curve tracer seeks to keep V_{DS} constant the lower graph on the left shows the square-law dependence of the R100 drain current as a function of V_{GS} acting alone. The lower graph on the right shows the dependence of I_D when the change in V_{DS} caused by the modulation of I_D through a non-zero load is included. And, as a result, the shape of the curve is no longer a simple square law. Each blue triangle is the change in I_D and V_{GS} (ΔI_D and ΔV_{GS}) at the intercept of the load line and one output curve measured with respect to the reference amplifier bias point indicated on the graph of the output curves. A trend line for the ΔI_D vs. ΔV_{GS} data was added to each graph, the formula of which is shown at the bottom of each graph.

transfer curve can be so straight near the inflection point that it takes a sensitive instrument to measure the non-linearity. This is why distortion meters report THD on logarithmic scales.

2. At the inflection point the amplifier transfer curve is "push-pull" meaning it has odd symmetry. Therefore, an audio signal amplified by an amp biased at the inflection point will have vanishingly small second-order harmonic distortion even in a single-ended amplifier with a transistor that has a square-law transfer curve. The residual non-linearity of the amplifier transfer curve modulated about the inflection point is mostly third harmonic.

As shown in the "*Running down the second harmonic*" box, both of these predicted behaviors are consistent with my experimental observations on a F2J amplifier (which is single-ended). As the bias point approaches the sweet spot, the second harmonic amplitude will fall, and can even vanish into the noise floor, while the residual third harmonic comes to dominate the much reduced THD. As the bias point passes through the sweet spot, the second harmonic will come roaring back if the transistor can withstand the extra power dissipation.

The conclusion? *The inflection point is the sweet spot*. I have labeled the inflection point, computed using some simple calculus on the polynomial trend line equation representing the amplifier transfer curve, as the "sweet spot." It's kind of like seeing with your very own eyes the Holy Grail you previously only believed to exist. Wow! And, if I can calculate the location of the inflection point by making what are really simple experimental measurements with a transistor of my choice biased in an amplifier of my choice, then *you* should be able to find the sweet spot with much less trial and error.

Why is this important? Here's a reason: There are many sweet spots that can be found in the interactions between an amplifier load (i.e., speaker impedance), the dc power supply voltage, and the transistor output curves. Since some are better than others how do we recognize the best ones? It's the degree of straightness in the amplifier transfer curve near the sweet spot that should catch our eye. If the "sweet spot" is the inflection point, then the straightened portion of the transfer curve near the sweet spot could be called the "candy land." I have marked with an ellipse an example of the "candy land" on the amplifier transfer curve. It should be obvious that the larger the candy land the greater the output power available at the reduced total harmonic distortion available while biased at the sweet spot. This can be seen by plotting the line tangent to the inflection point as shown in the graph of the amplifier transfer curve. This tangent line represents perfect Class A linearity. The trend line which represents the true amplifier transfer curve stays so straight inside the ellipse that the two lines are *almost* the same. Let me reinforce the "almost" qualifier. A sensitive distortion meter will measure the small differences between the two lines which it will report as mostly third harmonic. So the goal of the measurement and analysis procedure discussed in the next section is to help you find the largest candy land practicable with your particular gain device.

Finding the Sweet Spot

Where do we begin searching systematically for the sweet spot? First, it would be helpful to identify the transistor properties that control the relative quality of the linearity expected about the sweet spot. There are a couple of things. First, we want our transistor's output characteristics to show decreasing drain resistance as the bias current increases. This means the slope of each output curve in the saturated region beyond the ohmic portion of the curve is tilting more vertical at each increasing step of the gate source voltage. In fact, it can be shown that if a transistor has constant drain resistance there will be no load-line canceling. Another way to put it is that sometimes it takes fire to fight fire, i.e., we use one non-linearity (the dependence of the drain resistance on gate source voltage) to correct for a second non-linearity (the dependence of the drain current on gate source voltage). This is analogous to the "preemphasis" technique pioneered by Dolby Labs to reduce the noise in audio reproduction from magnetic tape.

Second, we want nice slow transitions from the ohmic region of the output curve to the saturated region of the output curve. The clear existence of two regions of conduction and the necessity of a transition between the two is the feature that makes solid-state transistors "pentode like" in tube speak. It is ironic that by the standards of the power switching community, SemiSouth JFETs have very lazy transitions. So much so, it became a marketing problem when customers pointed out that they like very sharp transitions. It turned out that one feature of a power MOSFET designed for the power switching market is that the slope of the curve in the ohmic region (related to the on-resistance of the power transistor) should be almost constant regardless of the gate-source voltage, and that the knee at the transition between ohmic and saturated conduction should be abrupt. If you are a power switching person, this is what you expect from your field-effect transistor. But if you are a Class A Pass amplifier person, you should want just the opposite because a nice lazy transition extends the size of the "Candy Land" to higher output power. Since SemiSouth was founded on the notion that we would sell to the power switching market, the lazy transition was a minor marketing problem. But I am convinced that the major reason the SJEP120R100 became known as an enabler of better Pass amplifier performance, at least with the single-ended types, is because it does an amazing job of meeting both conditions cited above. The box "*The fat lady sings*" compares two transistors for their output curves. The SemiSouth JFET clearly has the "fatter" curves, i.e., it has the properties better suited to make an extended linear trend about the amplifier's sweet spot.

But whatever! You can make your own judgments about which transistors to go in search of sweet spots. What we need is a search tool. The search tool readily available is the amplifier load line. A load line is superimposed on the output curves of the SJEP120R100 in the "*What goes up must come down*" box. There are two features of every load line that uniquely define it. The first is the slope of the load line which is given by:

$$\frac{\Delta I_D}{\Delta V_{DS}} = -\frac{1}{R_{Load}} \quad , \tag{1}$$

where R_{Load} is the nominal speaker impedance in a working amplifier, which cannot be adjusted arbitrarily. The other parameter that fixes the load line on the output curves is the open-circuit voltage defined as the V_{DS} that would (theoretically) be dropped across the transistor if $I_D = 0$. Mathematically I'll call this voltage the "*x*-axis intercept." The example load line in the box has a slope of about -0.24 S and an *x*-axis intercept of about 17.5 V. In Nelson's sweet-spot article he encourages the reader to go sweet-spot hunting by changing the power supply voltage. That is the basic approach I will suggest here for a couple of reasons:

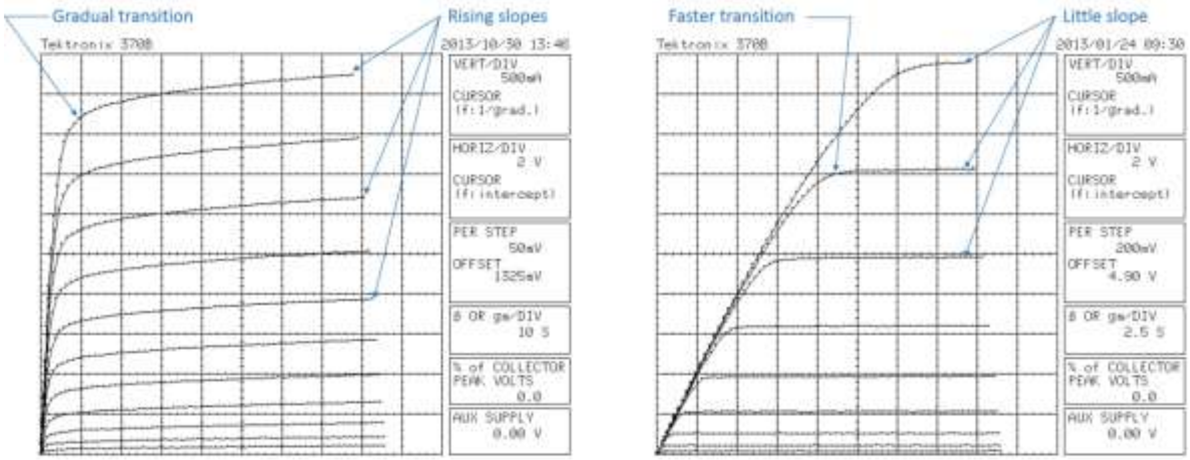

(a) (b)

The fat lady sings. The part numbers behind the figures are: (a) SJEP120R100 (JFET); and (b) 6N90C (MOSFET). The output curves of the JFET show more curvature in the transition between the ohmic mode and the saturated mode of operation. In contrast, the power MOSFET exhibits what the power switching industry has come to expect, namely, nearly constant slope in the ohmic region followed by a swift transition to a nearly constant saturation current (i.e., nearly horizontal line). The "fatter" shape of the output curves of the JFET produce more robust load-line canceling of non-linear distortion.

1. The load resistance in the working amplifier that you want to try out your new found sweet spot is the speaker impedance, which is only changeable in discrete values and in most cases should be selected based on other constraints.
2. While changing the actual power supply voltage in a working amplifier is not much more practical than changing the speaker impedance, changing the *x*-axis intercept voltage in many of the Pass single-ended amps is conveniently done with a bias potentiometer.

The second point is explained in the box "*Amplifier abstraction made easy*" for those that wish the details. But suffice to say we can make a bench top analogue of a working amplifier by testing the transistor with circuit (a) shown in the box. If we do so, we can experimentally determine the output curves of the *amplifier*, rather than the *transistor*. The amplifier output curves "reveal" intuitively the location of the sweet spot better than squinting at the transistor output curves with a load line slashed across it, and it permits visualizing the analytical process for creating amplifier transfer curves where the sweet spot stands out as the inflection point. Getting the amplifier transfer curve is the secret. Once we have that, a simple formula in an Excel spreadsheet will pinpoint the sweet spot with mathematical precision.

At the end of this article I will show you an experimental set up I created for directly measuring the amplifier transfer curve one curve at a time with nothing more than a bench power supply, a digital multimeter, the transistor under test mounted on a heat sink, and a resistor mounted on a heat sink. I will also provide an Excel spreadsheet that will allow you to enter your data and automatically plot it, compute the trend line, and compute the sweet spot from the trend line. What more could you ask for? Well, how about a big picture explanation of what this gets you?

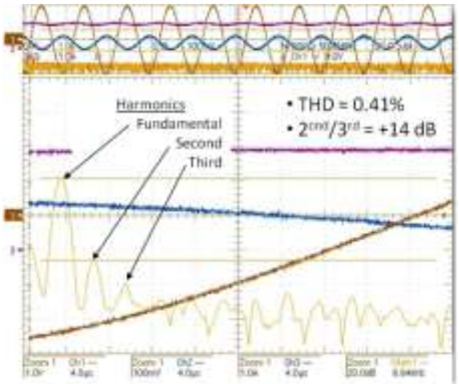

x-axis intercept = 35.4 V

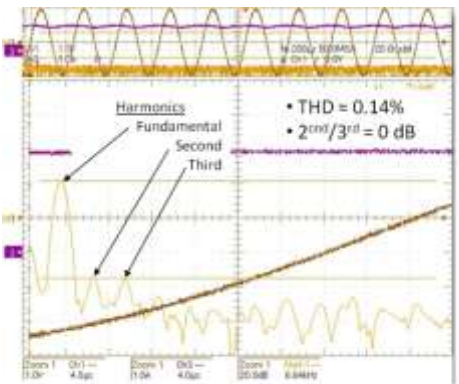

x-axis intercept = 37.5 V

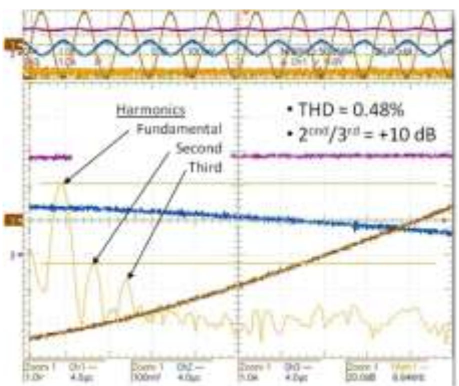

x-axis intercept = 39.5 V

Running down the second harmonic. Measurements are made on a F2J amp with a SemiSouth SJEP120R063 JFET as the gain device, I_{Bias} = 2.9 A, and R_{Load} = 8 Ω. The *x*-axis intercept is computed using equation (3). A Tektronix TDS 540 digital oscilloscope in FFT mode measures the harmonics while the bias potentiometer is adjusted until the second harmonic is minimized. That marks the sweet spot. Above and below 37.5 V the bias point is less sweet.

To do that I refer you to the "*How sweet it is!*" box. There I have fooled the trusty curve tracer into measuring the amplifier output curves rather than the usual transistor output curves. This was done by cleverly adding a load resistor to the drain leg of the transistor under test and shoving the combination into the curve tracer socket as shown in the image. You won't get away with this with the bench level tester I describe later because it uses DC voltages and currents and the resulting power dissipation requires heat sinking both the transistor and the resistive load. In contrast, the curve tracer uses pulsed voltages and currents which allow for at least some of the range of interest to be measured with little bitty resistors and none of the thermal management. The drain current is still measured as usual, and the curve tracer already keeps track of the gate source voltage, but now what it thinks is the drain source voltage is in reality the *x*-axis intercept voltage.

The resulting plot of the output curves shown in the box tells how the shape of things has been changed by load-line canceling. Two things jump out. First, the "ohmic" region is a lot more resistive, meaning the slope is not nearly as vertical. This slope, not surprisingly, is dominated by the resistance of the load, not the resistance of the transistor. This line is the limit where hard clipping will occur because this line tracks the short-circuit current of the amplifier, also known as the *y*-axis intercept current. But below the line are all the output curves. On the right side, the curves are spaced much like the output curves of the transistor, with a gradual increase in the separation between the curves reflecting the square-law of the transistor transfer curve. But if you let your eye wander toward the left, and you look carefully, you should notice that the gaps between some of the curves have become more uniform. Ah ha! That is the approximate location of the candy land! Somewhere in these patches of more uniformly spaced curves is where the sweet spot will be found. Each vertical line represents a different amplifier transfer curve resulting from a different load line. We can plot the combinations of I_D and V_{GS} for each vertical line and look at the resulting transfer curve to find the inflection point. That is the sweet spot. Three

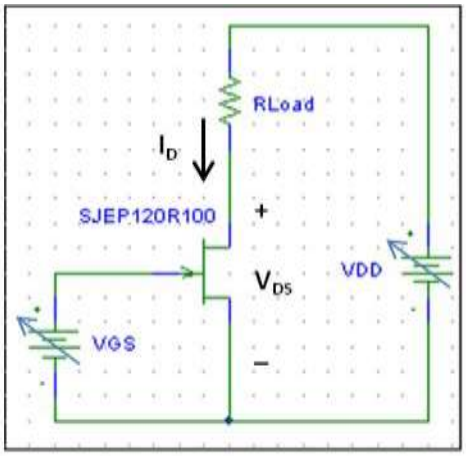

(a)

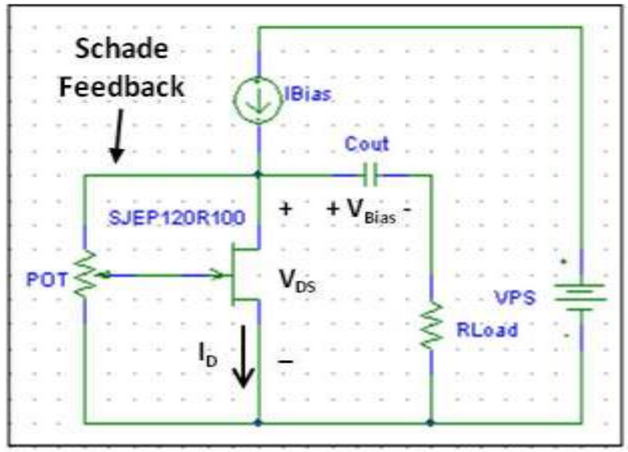

(b)

Amplifier abstraction made easy. Biasing a single-ended Class-A audio amplifier at its most abstract can be reduced to the circuit shown in (a). The circuit is supplied power from a variable voltage source V_{DD} and the current I_D is controlled by a variable voltage source V_{GS}. The fraction of V_{DD} that is dropped across the transistor is V_{DS} which is determined by the load line and the bias current I_D. The load line is easily found by algebraically summing the voltages around the loop that includes V_{DD} and then solving for I_D:

$$I_D = \frac{V_{DD} - V_{DS}}{R_{Load}} \quad . \tag{2}$$

In this case, V_{DD} is the *x*-axis intercept which can be varied by changing the value of V_{DD}. If you take my advice and build circuit (a) as your sweet-spot detector, then changing the *x*-axis intercept is equivalent to adjusting the output of a bench power supply. But what if you want to set up a working amplifier to amplify your favorite audio signal with the benefit of the sweet spot? Looking at, for example, a Pass F1 or a F2 amplifier or the Aleph series, the striking feature is the reliance on current sources for dc biasing. So, when thinking about a working single-ended Class-A audio amp, the circuit shown in (b) is a more relevant abstraction. The advantage of (b) is that the bias current of the gain device is set by the current source through the necessary feature of a Schade feedback circuit. Thus, when no audio signal is present $I_D \approx I_{Bias}$. For example, analysis of the schematic diagram of the F2/F2J reveals that changing I_{Bias} requires changing one or both of two rather imposing high-power resistors (R16 and R17). But it is easy to change the DC bias value of V_{DS}. The potentiometer in the Schade feedback circuit sets this value which is adjustable over a wide range. But how does this install our newly found sweet spot? To do that we need to change the *x*-axis intercept value of the load line in the working amplifier, but how the potentiometer does that is not obvious. Because the current source interposes itself between the power supply and the transistor, we can no longer consider the power supply voltage the same thing as V_{DD} in circuit (a). I have indicated as much by replacing "V_{DD}" with "V_{PS}" in circuit (b). In fact, V_{PS} really should be considered a part of the current source. So, how do I derive the equivalent of equation (2) from circuit (b)? Leaving this to be an exercise "for the student" is an infamous cliché in electrical engineering education, but if one does the analysis, the result is given by equation (3) below:

$$I_D = \frac{V_{Bias} + R_{Load} I_{Bias} - V_{DS}}{R_{Load}} \quad , \tag{3}$$

where $V_{Bias} = V_{DS}$ when no audio signal is applied to the input of the amplifier (known as the "quiescent point") which is the voltage across the output coupling capacitor as shown in (b). The *x*-axis intercept now equals $V_{Bias} + R_{Load} I_{Bias}$. In plain language this means that twiddling the bias pot of the Schade feedback circuit will change V_{Bias}, which will dial the amplifier into the sweet spot for a given I_{Bias} and speaker impedance.

SJEP120R100 inserted in curve tracer socket with 8.3 Ω load.

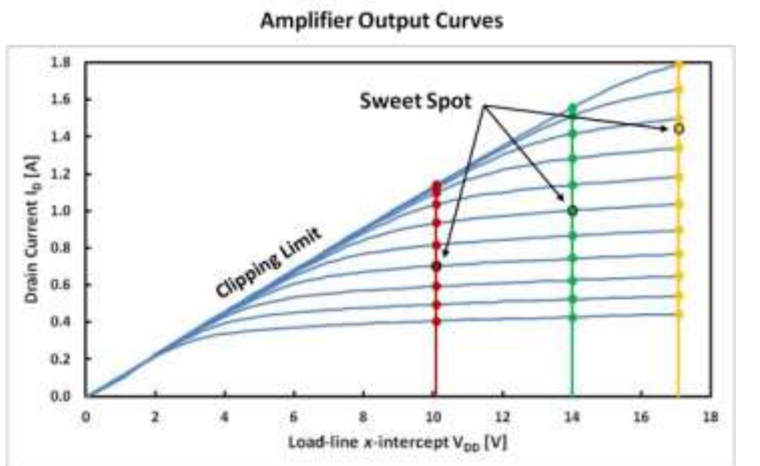

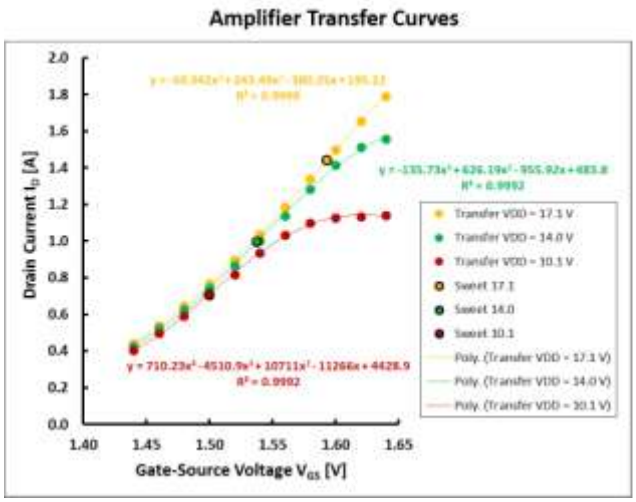

How sweet it is! Many an audiophile may agree with Jackie Gleason, who coined the title of this box to celebrate happy moments in the vintage television show *The Honeymooners*. Finding the sweet spot is equivalent to finding the inflection point in the amplifier transfer curve, which was constructed from the amplifier output curves as follows: (1) Resistors are soldered to the drain leg of the transistor and the combination is inserted in the socket of the curve tracer as shown in the image. This effectively creates circuit (a) from the "*Amplifier abstraction made easy*" box. The curve tracer sweeps what it thinks is V_{DS}, but which in fact is now the *x*-axis intercept voltage. The curve tracer's measurements show the influence of the 8.3 Ω load resistance, including a pronounced sloped current limit which is where hard clipping in an over-driven amplifier will occur. (2) The drain current I_D at the intersection between a vertical line drawn at a particular *x*-axis intercept voltage and one output curve is measured and recorded along with the V_{GS} corresponding to this curve. The curve tracer steps V_{GS} in equal increments and measures one whole output curve at each value of V_{GS}. (3) The measured I_D is plotted as a function of V_{GS} thus forming one amplifier transfer curve. Three different vertical load lines are shown in the graph of the amplifier output curves. The circular dot at each intersection represents the measurement. This same circular dot is plotted in the graph of amplifier transfer curves. (4) From the graph of I_D vs. V_{GS} data it should be clear that the data points provide incomplete information. An equation for a smooth curve is needed to "connect the dots" so that we can do the math to find the inflection point. I use a common form of "professional guessing" called curve fitting to estimate an equation. The spreadsheet program *Excel* has a "trend line" function to automate curve fitting. I fit a polynomial function to each of the three amplifier transfer curves. Each resulting equation is shown color coded to the trend line it corresponds to. A measure of "goodness" of curve fitting is the R^2 value. All three trend lines fit the data well since in each case R^2 almost equals 1. (5) By mathematically analyzing the polynomial trend line, I calculate the inflection point on the transfer curve, which is marked with a circular data point that has a black border. These sweet spots are also marked on the amplifier output curves to draw the eye to the place where the spacing of the output curves is nearly uniform.

amplifier transfer curves taken from the amplifier output curves are drawn in the "*How sweet it is!*" box. It is worth noting that each of them has an inflection point, but as the power supply voltage increases, the bias current where the sweet spot occurs also increases.

But more importantly, as the bias current increases the shape of the amplifier transfer curve near the sweet spot becomes more gradual (i.e., more linear).This straightening of the shape of the amplifier transfer curve supports the widely held belief that larger Class A bias produces better linearity. While dissipating more standby power, the amplifier makes up for it by delivering reduced THD over a greater audio output power range, i.e., the "Candy Land" is wider than at lower bias currents. It all comes down to the point made at the beginning of the article, namely, the shape of the curves matter!

Nagging little problems like power dissipation and excessive temperature rise will eventually get in the way of the theory. The curve tracer approach taken in the "*How sweet it is!*" box, while convenient for visualizing what is happening in the amplifier during load-line canceling, does ignore the crucial variable of transistor temperature, a variable Nelson warned is important in his sweet spot article. This problem will be solved with the practical tester circuit described in the next section. However, it is shown that searching for the sweet spot by adjusting the *x*-axis intercept voltage is feasible in the "*Running down the second harmonic*" box where I illustrate the classic symptoms of the sweet spot in my own experimental investigation of a F2J.

Sweet Spot Test Procedure for the DIY Audiophile

It is ironic that an expensive curve tracer is *not* the best way to locate the sweet spot. Instead, a simple circuit and a digital multimeter is a more reliable way of finding out how to set up the sweet spot in a working audio amplifier. The reason is simple: transistor temperature matters and this is hard to replicate on a curve tracer. Biasing a transistor with an equivalent circuit that mostly reflects what happens in a working amplifier works better because it will require you to dissipate about the same power in the transistor as it does in the working amplifier, and that means the transistor will have to thermally stabilize like it does in the working amplifier.

I have developed a seven step procedure for building and using sweet spot test gear that I believe is within the ability of virtually any DIY audio enthusiast. Of course, I realize that some things may come easier than others, and it is often hard for someone with an electrical engineering background to be understood by those that don't have the same background. But on the other hand, my experience with the DIY audio community is that you all have a wealth of practical experience and a lot to contribute. So, if anything, I expect some of you will improve on what I have to suggest; and surely that's what the on-line forums are for. So, let me list the steps next and refer you to a couple of boxes to deliver additional details.

Seven steps for finding the sweet spot in your amplifier:

1. Acquire the components for circuit (a) in the "*Amplifier abstraction made easy"* box.
   a. An adjustable DC power supply capable of supplying the required bias current (1 to 3 A for an SJEP120R100 or similar transistor). At least 24 V output is needed for 4 Ω speaker impedances or greater, and up to 50 V would be better for the larger speaker impedances.
   b. A power resistor with resistance equal to the desired speaker impedance that is capable of dissipating the power at the bias currents expected, where $P_{\text{dissipated}} = R_{\text{load}} \times I_{\text{Bias}}^2$. I used two 4-Ω power resistors from Ohmite (part number TEH70P4R00JE) in series to simulate an 8-Ω speaker impedance. These resistors come in TO247 packages just like the transistors and are easy to heat sink.
   c. An adjustable DC power supply capable of biasing the gate of the transistor. Instead of a separate supply, the power supply in step 1(a) can be used with a potentiometer to adjust the gate voltage. Just take care to monitor the gate voltage to keep it from changing during the measurements.
   d. A multimeter to measure the power supply voltage (required), the gate-source voltage (required), and the drain source voltage (optional). If you only have one multimeter, no problem. Move it around until you have measured all quantities before changing the bias point.
   e. A multimeter or ammeter to measure the bias current. If you have one with the range to measure the bias current fine. If you don't, the multimeter used in step 1(d) can be used instead to measure the voltage across the power resistor and divide by the resistance to estimate the current.
2. Assemble the amplifier according to the circuit (a) schematic. An example of mine is seen in the box "*Step-by-step we'll get this done*." If you have an ammeter, insert it in series between the power supply and the resistive load, or in series between the resistive load and the drain of the transistor, but not in series with the source of the transistor.
3. Think about a range of *x*-axis intercept voltages that you wish to measure amplifier transfer curves. If you are using an SJEP120R100 or one of its near equivalents, you can look at the amplifier output curves in the "*How sweet it is!*" box to get an idea of your search range. But any transistor should have a sweet spot (just maybe not an attractive "candy land") so maybe you're not testing with a SemiSouth transistor. No problem! Experiment, you have nothing much to lose except the time to make the measurements and the cost of the electricity needed to dissipate the heat at the bias currents.
4. Heat the experimental set up to a stable heat-sink temperature.
   a. Fire up the power supply to your first *x*-axis intercept voltage.
   b. Adjust the gate-source voltage to a drain bias current included in the amplifier transfer curve you are planning to measure (or collector bias current if doing this with a BJT). I would consider the middle of the range but really there is no special value as all you're trying to do is get the transistor to operating temperature.
   c. Let the heat sink sit for a while to thermally stabilize. Since the heat sink temperature will inevitably change during the testing, there is no need to be too exacting about defining thermal equilibrium.

5. Record data along a particular amplifier transfer curve.
    a. Keeping the power supply voltage constant is the main thing to remember to stay on a particular amplifier transfer curve. If the bias current goes up, the power supply voltage may droop a bit. Go ahead and adjust to keep it constant. One thing to keep in mind is that the connecting leads are resistive. Monitoring the power supply voltage from the top of the load resistor to the source or emitter of the transistor will remove the unknown resistance of the lead wires connecting to and from the power supply, thus avoiding one source of error in recording the amplifier transfer curves.
    b. Measure the bias current and the gate-source voltage. Record both numbers as a pair tagged with the power supply voltage. (I have included an Excel spreadsheet in the post with this article that will allow you to enter these measurements into cells in the spreadsheet to automatically analyze the data as described next.)
    c. Adjust the gate source bias voltage to a new drain (or collector) current and re-do step 5(b). Take care to monitor the power dissipation in the transistor (the product of the drain-source voltage and the drain current) so as to not exceed the maximum power dissipation rating of the transistor. In fact, it would be better to stay well below this limit.
    d. Measuring the drain-source voltage is optional, as it is not directly used in constructing the amplifier transfer curve, but it can help you locate the expected sweet spot in your working amplifier later.
    e. After recording a few V_{GS} vs. I_D pairs, you can measure a new amplifier transfer curve by changing the power supply voltage and repeating steps 5(b) and 5(c).
6. Analyze your transfer curves to find the inflection points.
    a. Plot the ordered pairs V_{GS} vs. I_D at each *x*-axis intercept (power supply) voltage on a Cartesian graph with V_{GS} on the horizontal axis and I_D on the vertical axis.
    b. Compute a trend line connecting each data point of each amplifier transfer curve. The trend line should be a polynomial whose preferred order is three. See "*The Devil is in the mathematics*" box if you are interested in finding out why.
    c. Assuming a third order polynomial trend line, solve the equation for the value of V_{GS} at the inflection point using the coefficients derived from the trend line. This analysis is also explained in "*The Devil is in the mathematics*" box and is done for you automatically if you use my example spreadsheet and enter your data correctly in the cells provided.
    d. Compute or estimate I_D from the V_{GS} of step 6(c). This is the inflection point in the amplifier transfer curve which marks the sweet spot at that particular *x*-axis intercept voltage.
    e. Repeat the analysis of steps 6(c) and 6(d) until you have the sweet spots for every amplifier transfer curve you measured. If you don't think any of them are where you want to bias your working amplifier, measure some more amplifier transfer curves. HINT: If your amplifier uses current source biasing, look for inflection points that occur at or near the preset bias current of your working amplifier. If your amplifier uses voltage source biasing, look for amplifier transfer curves that have *x*-axis intercept voltages close to your working amplifier's *effective* power supply voltage (e.g., if you have a split power supply for each channel, then the

sum of the magnitudes of the positive and negative supply potentials may be the effective power supply voltage).

f. Considering the hint in step 6(e), if you can't seem to find an amplifier transfer curve that is right for your working amplifier's biasing situation your transistor may not be well suited for your speaker load. In general, transistor output curves that are flatter are better for higher impedance speakers (a pentode type like transistor (a) in "*Picking through the jewelry box*"). Transistor output curves that are more vertical are better for lower impedance speakers (a triode type like transistor (c) in "*Picking through the jewelry box*").

7. Listen to the sweet spot in your working amplifier.
   a. If your amplifier uses current source biasing then:
      i. Install the transistor under test in the working amplifier.
      ii. Select an amplifier transfer curve that has an inflection point close to the working amplifier's bias current.
      iii. Set V_{DS} of your transistor to the value corresponding to the I_D at the inflection point. This is where the option of recording V_{DS} in step 5(d) comes in handy.
      iv. Let the amplifier warm up and readjust V_{DS} as necessary.
      v. Listen to the amplifier and fine tune to your preferences in residual harmonics as necessary.
   b. If your amplifier uses voltage source biasing then:
      i. Install the transistor under test in the working amplifier.
      ii. Select an amplifier transfer curve with an *x*-axis intercept voltage that corresponds to the working amplifier's effective power supply voltage.
      iii. Set I_D of your transistor to the value corresponding to the inflection point using the amplifier biasing network.
      iv. Let the amplifier warm up and readjust I_D as necessary. Listen to the amplifier and fine tune for your preferences in residual harmonics as necessary.

1 Locate your components

2 Assemble the amplifier

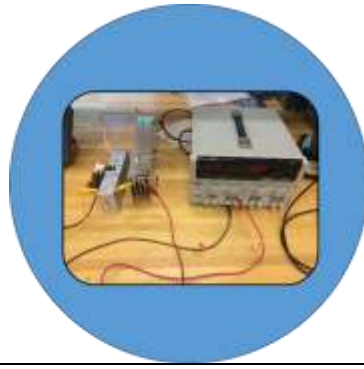

3 Select *x*-axis intercepts

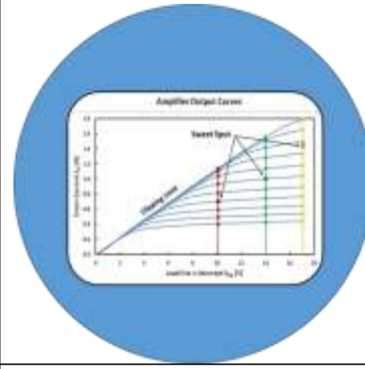

4 Heat up the amplifier

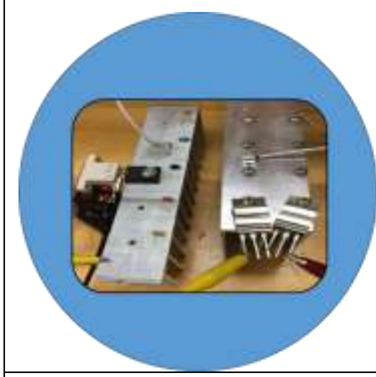

5 Measure transfer curves

6 Find the sweetest spot

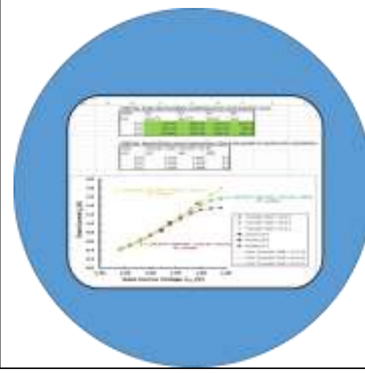

7 Listen to your audio amplifier set up on the sweetest spot

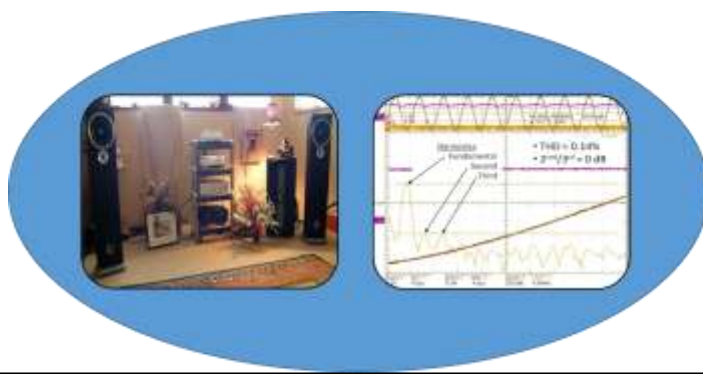

Step by step we'll get this done. The steps needed to track down and listen to your own best "sweet spot." I hope the directions, the suggestions, and the spreadsheet will help you get to your best listening spot without the guess work.

The devil is in the mathematics. There is no reason that a DIY audiophile can't analyze their preferred transistor and speaker load combinations to figure out where the sweet spot is, even those that don't have engineering or scientific backgrounds. Of course, you do have to be able to take ordered pairs of measured data from your test amplifier. I posted with this article an Excel spreadsheet that is designed to make it easy to complete steps 5 and 6 of the instructions. Table #1 in the spreadsheet has columns for you to enter the ordered pairs of V_{GS} and I_D for each amplifier transfer curve denoted by V_{DD} (an optional column is included for V_{DS} if you wish to keep up with it). When you use the columns provided in the spreadsheet the data are automatically transferred to the graph so that you can see the amplifier transfer curve take shape. This is important because I have associated a trend line with each curve that is also automatically calculated and displayed on the graph using color coding to keep track of which trend line goes with which amplifier transfer curve. This is important because in Table #2 you enter the coefficients of the trend lines into the proper rows provided in the spreadsheet. The basic trend line equation is a third order polynomial relating V_{GS} to I_D as shown in equation (4):

$$I_D = a_3 V_{GS}^{\ 3} + a_2 V_{GS}^{\ 2} + a_1 V_{GS} + a_0 \qquad (4)$$

Read off the graph the coefficients a_3, a_2, a_1, and a_0 and enter them in the green cells in Table #2. Table #3 is completed for you automatically using equation (6) as you enter the coefficients by locating the inflection point as the place where the second derivative of I_D with respect to V_{GS} is equal to zero as shown in equation (5):

$$\frac{d^2 I_D}{dV_{GS}^{\ 2}} = 6a_3 V_{GS} + 2a_2 = 0 \qquad (5)$$

Which is then solved for V_{GS} as shown in equation (6):

$$V_{GS} = -a_2/3a_3 \qquad (6)$$

This is the gate source voltage at the inflection point. Plugging V_{GS} from equation (6) back into equation (4) produces the drain current at the inflection point. The coordinates of the sweet spot are at this V_{GS} and I_D. I have set up the spreadsheet assuming your trend lines can be represented by a third order polynomial. If you pick an *x*-axis intercept voltage V_{DD} too small, then it is possible that the inflection point will be too close to hard amplifier clipping which will require a fourth-order polynomial (or greater) to model. I discourage selecting *x*-axis intercepts this low for two reasons. First, the analysis is more difficult and eventually requires some expertise in mathematics. But secondly, and more importantly, because at a sweet spot like this the THD increases too quickly as the audio power is increased. This is why it is well known that sweeter sweet spots come with higher bias currents and power supply voltages. Most people go as high as practical considerations for transistor power dissipation and heat sink temperature will permit.

In Summary…

The bottom line is that the output curves of the SJEP120R100 have enough shape in the right direction to make the relatively small impedance of the speaker adequate to "bend" much of the non-linearity out of the transistor if biased near the special sweet spot. But load-line canceling is a process that depends upon many variables, some of which can be adjusted by you by playing with the bias settings, and others are hard coded into the parts themselves. I hope you will find the step-by-step process laid out in this article useful for discovering the richness of shapes and perhaps you will put them to use in your search for excellence in audio reproduction.

This article is the second part in a three-part series to examine the *buzz* about SemiSouth JFETs in audio. The principal topic in this article was about curves, namely, the importance of the transistor output curves in *shaping* the amplifier output curves. It was shown that the ultimate shape is in the amplifier transfer curve where load-line canceling guarantees an inflection point which is the so-called *sweet spot* we seek. In the next and final article in this series, I'm going to look more closely at the role that negative feedback plays in making beautiful music out of chaos. It influences every amplifier, commercial or DIY, and while it is a topic that can be explored at extraordinary levels of mathematical complexity, the basic impact on audio amplifier design is almost always the same: profound and needed to a certain extent, but also polarizing and prone to overuse. One aspect of Nelson Pass' philosophy as recorded in his writings and his lectures is his ambivalence to the use of negative feedback. On the one hand his amps always have some of it (maybe called one thing or maybe another), but on the other hand his innovation is often revealed by how he sidesteps the use of it in favor of more worldly approaches to solving the time-honored problem of stabilizing and linearizing, while leaving a wee bit of warmth in the final product. Almost like a good mother, Nelson is.

Happy listening!

Mike "Semisouthfan" Mazzola

If you have detailed comments about this article or you would like more information about obtaining SemiSouth JFETs, feel free to contact me at michael.mazzola@impowersystems.com. iMPOWER maintains inventory of SJEP120R100 and R100A, SJEP120R125, SJEP170R550, SJEP120R063 and R063A, SJDP120R085, SJDP120R045, and other specialty SemiSouth JFETs. Transfer curve matching available upon request.

More reading…

"The Sweet Spot" http://passlabs.com/articles/the-sweet-spot

"What's the Buzz? – Part I" http://www.diyaudio.com/forums/pass-labs/229241-semisouth-boiler-room-21.html, post #208