From: http://www.moorepage.net/THD.html

Updated 4-28-2015

About Total Harmonic Distortion Analyzers

THD analyzers are important tools in the design and construction of high performance audio equipment, as well as in the design and construction of oscillators and other test equipment. As Paul Klipsch used to say, "You can't make what you can't measure because you don't know when you've got it made."

How THD analyzers work

Total Harmonic Distortion is measured by using a rejection filter system to remove the fundamental frequency (the first harmonic) of a waveform, usually a sine-wave, and then measuring whatever is left over -- second and higher harmonics, and noise.

Naturally, the measuring system has to have wider bandwidth than the fundamental frequency being measured. For very low distortion signals, say under 0.05% THD, the significant harmonic content will be in the 2nd thru 5th harmonics -- above the 5th, noise tends to dominate. So a measuring bandwidth of 5 to 10 times the fundamental will be needed, with 10 times usually preferred.

Most analog analyzers, like the HP 334A and the Heath IM-5258 for example, basically take the input signal, buffer it from the outside world, adjust its voltage to a convenient level, split it into two paths, sometimes inverting one path to get both non-inverted and inverted signals, and then pass one of the two resulting signals through a narrow-bandwidth filter of some kind -- sometimes a peak (band-pass), sometimes a notch (band-reject). The filter is essential in order to eliminate, as much as possible, the fundamental, while the remaining harmonics and noise are processed for measurement (there are other circuit topologies for measuring THD which I'll discuss below). Then the two signals are summed (if one was inverted) or they are differenced (if they both have the same polarity).

In many distortion analyzers, the fundamental and an inverted exact copy are precisely matched for phase (either precisely in-phase or precisely out) and for level, at which point the summing or differencing eliminates the fundamental and leaves everything else. Feedback is used to sharpen the notch, raising the filter's Q. Depending on the circuit used, the feedback can be negative or positive. Given sufficient feedback, the fundamental will be eliminated but the harmonics will all be passed through with little or no attenuation -- a requirement for accuracy. The notch filter is not infinitely sharp, so some noise signals are attenuated along with the fundamental, but this usually is not meaningful. And if the cancellation is not complete, some residual fundamental can be passed through the filter, setting a floor on the measured minimum signal and maximum resolution. Ultimately, the filter used serves as a notch (band-reject) filter once all the processing is done.

Filter selectivity or "Q" is very important, because most filters have fairly broad peaks or notches, and will ultimately attenuate the 2nd and higher harmonics along with the fundamental, giving inaccurate measurements. A passive Twin-T filter, for example, will attenuate the 2nd harmonic by about 9dB and the 3rd harmonic by about 5dB, with higher harmonics being attenuated proportionately lessâ€”this attenuation represents a severe loss of accuracy unless mitigated in some way.. In active systems, feedback around the filter system sharpens the notch, preventing the harmonic attenuation and yielding trustworthy measurements. General practice is to keep 2nd harmonic attenuation to less than 1dB, and preferably less than 0.5dB, thus having little effect on the measurements of all the harmonics.

The output of this rejection system is then passed to a high-gain, wide-band voltmeter for measurement. Various filters can be used following the rejection system in order to eliminate hum and/or other spurious signals, such as wide-band noise that obscures the actual harmonics of interest. The voltmeter needs sensitivity in the microvolt range when the input signals are below 3VRMS -- a 3V signal with 0.001% THD will have distortion and noise components with an RMS sum of 30 microvolts; obviously, low noise and high gain are essential to the meter.

As mentioned, in a cancellation system, tuning the filter to cancel the fundamental consists of adjusting the amplitude and phase of the filtered signal to exactly match those of the fundamental component of the wideband signal. This means controlling the overall gain of the filtered fundamental and adjusting the phase of the filtered signal. Gain tuning is often done by adjusting the value of an attenuator resistor or a feedback resistor. Phase tuning is usually done by adjusting a filter component, usually a resistor. Automatic tuning is almost essential for measuring very low levels of distortion -- the null is very sharp and hard to achieve and stay in manually.

The natural phase transfer function of filtering the band-limited fundamental makes it possible to use a phase detector to measure and amplify the phase difference between the unfiltered fundamental of the wide-band signal and the filtered fundamental, allowing automatic control of the tuning element. Gain differences between the two fundamental signals are somewhat easier to measure, amplify, and use to adjust the level of the filtered signal. Properly done, an automatic system easily tracks even small differences in amplitude and phase, allowing the system to acquire and hold a deep null of the fundamental. In my experience, auto-tuning can hold nulls of from -100 to -120dB. It is important to realize that the level of any residual fundamental can set the bottom limit of measurable distortion. This means that a fundamental null of -100dB will limit the measured distortion to 0.001%, regardless of how low the actual harmonics are or how low the noise is.

The system described above is how the HP 330 thru 334, all the Heathkits, the Cordell State-Variable, the Tektronix AA 501A, the Sound Technology 1700, and many other analyzers work. There is another way, which is simply to use a traditional notch filter like a Twin-T or a Bridged-T and use positive feedback to sharpen the notch and deepen the null. Of course automatic tuning still offers the needed stability for deep nulls. This is the system used in the ShibaSoku 725 series and the HP 339A, among others. These analyzers often use a Bridged-T filter (with four components) for ease of tuning, while the HP 4333A used a Twin-T, which has six components, which complicates the tuning and switching.

The design issues in THD analyzers are:

1) the fact that optimally tuning them means changes of microvolts in amplitude and changes of milli- or micro-hertz in frequency.
2 measuring signals in the microvolt range while maintaining wide bandwidth.
3) noise.
4) intrinsic analyzer self distortion.

Many forms of filters and amplifiers have been used, each with its own strengths and weaknesses -- here's a summary:

Wien Bridge Filter
The Wien Bridge peak filter (band-pass) is one of the simplest and easiest to tune, which is why it has been used in so many THD analyzers, such as the HP 330 through 334 series, in all of the Heath THD analyzers, several Leader brand analyzers, and others too numerous to mention.

The Wien Bridge uses two resistors and two capacitors -- one resistor and one capacitor in a series leg, and one resistor and one capacitor in a parallel leg, to make a peak filter with a very broad peak. Usually, the two resistors have equal values, as do the two capacitors, which is intrinsically convenient.

The heart of the system (and of most analog analyzers) is that the phase of the filtered fundamental is zero, compared to the input signal, only at the precise center frequency of the filter. The Wien Bridge has an insertion loss of exactly 2/3, or roughly -10dB (-9.54dB), so the wide-band signal also has to be attenuated by the same amount before being combined with the filtered signal.

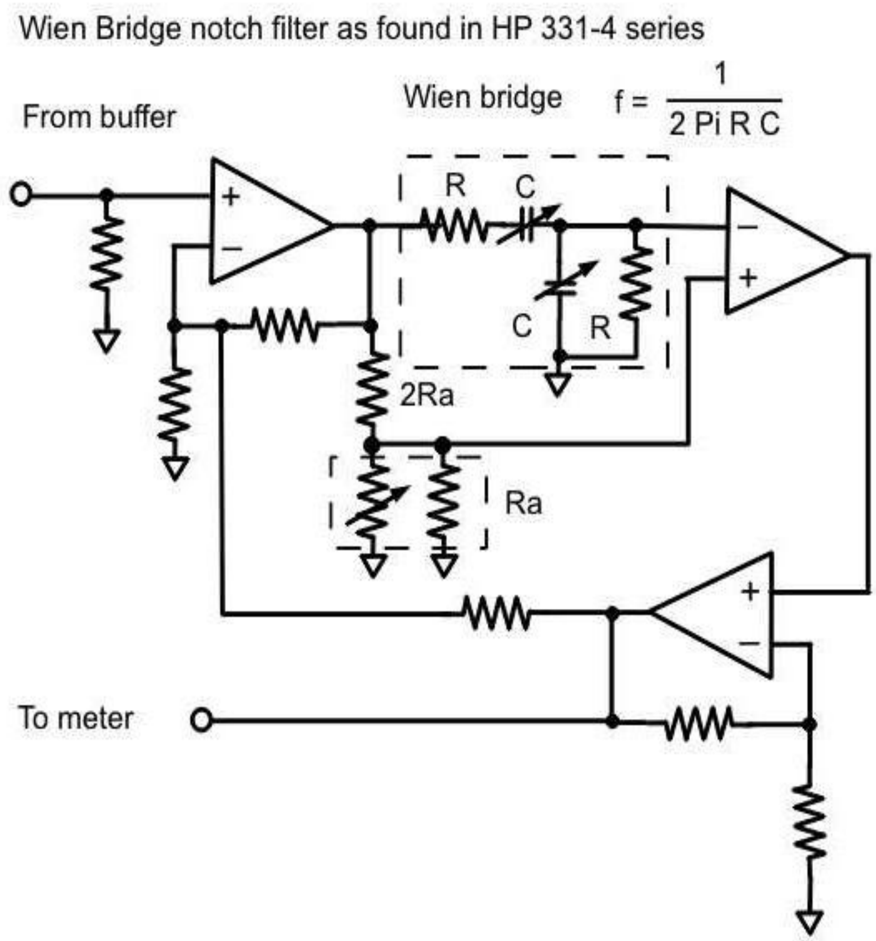

Advantages With phase fine-tuning via small variable resistors fitted as series segments of the total bridge resistances (R, above), and just one variable attenuator resistor needed to fine-tune amplitude (variable part of Ra, above), tuning the Wien Bridge is dead easy from a circuit point of view. The Wien Bridge is easy to auto-tune as well, because both phase and amplitude can be adjusted by voltage-dependent variable resistors, usually photoresistors (CdS cells, also called LDRs) driven by lamps or LEDS, or by varying the channel resistances of JFETs.

The broad peak of the Wien Bridge is significantly sharpened by overall negative feedback around the filter. This sharpening gives a deep null of the fundamentalâ€"the null depth essentially sets the measurement floor of the device if noise is not a limiting factorâ€"while preventing unwanted attenuation of low-order harmonics, especially the second and third, so that the measured values are correct without needing meter correction factors.

The notch depth of the Wien Bridge filter is limited only by the ability of the tuning circuitry to achieve and hold a null; this is limited by the resolution of the variable resistances used in tuning and by the sensitivity and response speed of the amplitude and phase detector systems used for auto tuningâ€"remember, we're talking microvolts and milli- or micro-hertz here.

Disadvantages A limit on performance of the Wien Bridge is, in part, the 10dB insertion loss, which means the rejection amp as a whole has to make up the gain and then some in order to have enough feedback to achieve a narrow and deep notch. This requires an overall main-loop feedback of about 16-20dB for acceptable performance. This means that the overall rejection amp gain with the main loop open needs to be at least 20dB for the feedback; then the 10dB filter loss reduces the signal-to-noise ratio, regardless of whether its loss is made up or not. This situation obviously creates issues with bandwidth and rejection amplifier system THD, not to mention noise.

The challenge in the rejection amplifier system is balancing the moderate signal levels needed for very low THD of the amp stages, especially the input stage, against the S/N ratio losses from levels that are too low -- that's why the insertion loss of the Wien Bridge matters. Note that Wien Bridges can be made from unequal value components -- if one leg has 2C and 1R, then the other has 2R and 1C, for example, then the insertion loss drops from 10dB to 6dB, a nice improvement that is offset by the shallower peak and the slower rate of phase change through the fundamental frequency. I've never seen a THD analyzer built in this way.

Capacitance tuning Most Wien Bridge analyzers are capacitance-tuned. Since large value air-dielectric variable capacitors are also physically large, this means using smaller caps, which in turn means using larger resistors, and at low frequencies, the resistors get very large -- tens of megohms -- which has serious impacts on noise performance. Of course, applying noise filtering to low frequency measurements helps to offset the higher noise from high-Z circuitry, but the problem remains and is substantial.

Resistance tuning The solution to the high-resistance, high-Z problem of capacitance tuning is to do range switching with fixed capacitors, which can easily be as large as needed, and do in-range tuning with either potentiometers or with switched resistors. This advantage is offset for potentiometer tuning systems by the difficulties and expense in making high-precision ganged potentiometers with good tracking. So, designs using switched 1% or better resistors are usually preferred.

The larger issue for me with Wien Bridge systems is that, in my experience, they have quite high residual distortion compared to other bridges, which significantly limits their performance. In practice, they have 10 to 30dB higher residual THD+noise floors.

Bridged-T Filter

One form of the Bridged-T notch (band-reject) filter works in an active circuit that is similar to the way that the Wien Bridge is used, with feedback sharpening the notch -- but in this case, the feedback is positive.

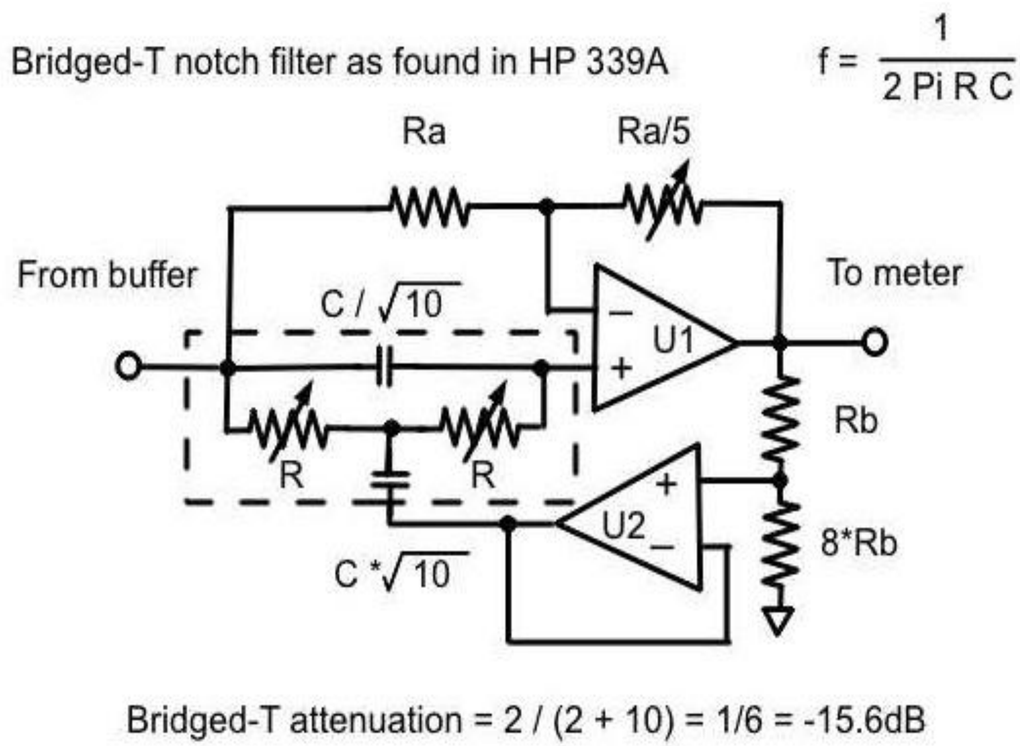

Like the Wien Bridge, the Bridged-T uses two resistors and two capacitors, but they are used to make a notch filter rather than a peak filter,with the advantage that the slope of the phase change through the fundamental frequency is quite steep, compared to the Wien Bridge, making auto-tuning sharper and more sensitive. As in the Wien bridge, range switching can be either by the resistors or the capacitors, with the other then being used for in-range tuning. Unlike the Wien Bridge, two of the components, either the resistors or the capacitors, have to be unequal in value, often in a 10:1 ratio for practical reasons of switching and impedance.

1-9-2011 -- In looking at this page again, I noticed that I had forgotten to mention the role of U2 in the circuit above. U2 provides positive feedback, whose level is set by the ratio of Rb and 8*Rb, the divider resistors in the input leg of U2. The ratio shown sharpens the notch enough to only attenuate the 2nd harmonic by about 0.3dB, an essentially negligible amount. If the value of Rb is too small, the circuit becomes unstable; too large, and the attenuation of the low-order harmonics becomes too great and THD readings are inaccurate.

This arrangement is useful for capacitance-tuned analyzers but, like the Wien Bridge, does result in large resistor values at low frequencies for the typical relatively small capacitances of ganged air-variable tuning caps.

Resistance tuning In resistor-tuned form, the Bridged-T is easy to tune, with auto-tuning via LDRs, which I'm sure was one of the reasons that HP used it in the 339A analyzer (it also is the filter type for the 339's oscillator section) and ShibaSoku chose it for the 725. Note that the bridge capacitor is the smaller element and the pillar capacitor is the larger one.

Capacitance tuning Like the Wien Bridge, if the in-range tuning is done with equal air-variable capacitors, the range resistor values are very large at low frequencies, and noise is again a problem. This is why commercial analyzers like the HP 339A use unequal capacitors for range switching and equal switched resistors (with a vernier) for in-range tuning. Note that, opposite to the resistance-tuned case, in capacitor tuning, the filter's bridge resistance is the larger element and the pillar resistance is the smaller one -- same basic calculations, but the location is reversed.

Disadvantages As in the Wien Bridge, the major disadvantage is insertion loss. I simulated the Bridged-T filter circuit of the HP 339A in LTspice. I found that the circuit can easily achieve null depths of over 100dB, if the tuning resistors have a few milliohms resolution. This level of performance is achieved with capacitance range switching and a capacitor ratio of 10:1 for the two bridge caps. The wide-band signal is attenuated by essentially the same amount as the bridge's attenuation of the fundamental, which amounts to about 6 times, or -15.6dB, with the exact value determined by component tolerances. This means a fairly high insertion loss -- more than the Wien Bridge -- which means more noise -- but the sharper phase tuning is a nice trade-off.

A deeper notch (before overall feedback is applied) can be attained with a larger cap ratio, but that also means, unfortunately, more attenuation of the wide-band signal, resulting in a poorer signal-to-noise ratio. HP definitely optimized the overall design of the 339A, using a 10:1 cap ratio in the rejection filter section and a 100:1 ratio in the oscillator section, and it's performance is very good, if not great, by today's standards -- I now have one.

But the real disadvantage to both the Bridged-T and Wien Bridge filter systems, in my experience, is that they both end up with a relatively high 2nd harmonic signal component in the output of the notch filter system. In the case of the HP 339A, this 2nd H. spike is what sets the measurement floor of 0.001%, and it is what limits the resolution of the HP 331-4 series to about 0.01%. I believe that these distortions come from the non-inverting topology of the amplifier, and are the artifacts of common-mode failure in the amp. An all-inverting amplifier system would be highly preferred.

Twin-T Filter

The Twin-T is a notch filter that is ideal for a THD analyzer, since it essentially has zero insertion loss, plus it has the big advantage over the Bridged-T of an intrinsically very deep null, needing active amplification only to sharpen the notch. Here, a FET-input opamp U1 buffers the filter and U2 provides the active positive feedback. There is no need for gain adjustment as seen in the Bridged-T thanks to zero insertion loss. Notch bandwidth, or "Q," is controlled by the ratio of the Rb resistors -- too sharp and it's hard or impossible to tune; too wide and harmonic attenuation is excessive.

1-9-2011 -- 10*Rb is probably too much feedback; a fixed ratio of about 8:1 for the two resistors works fine -- no need to fuss with a pot.

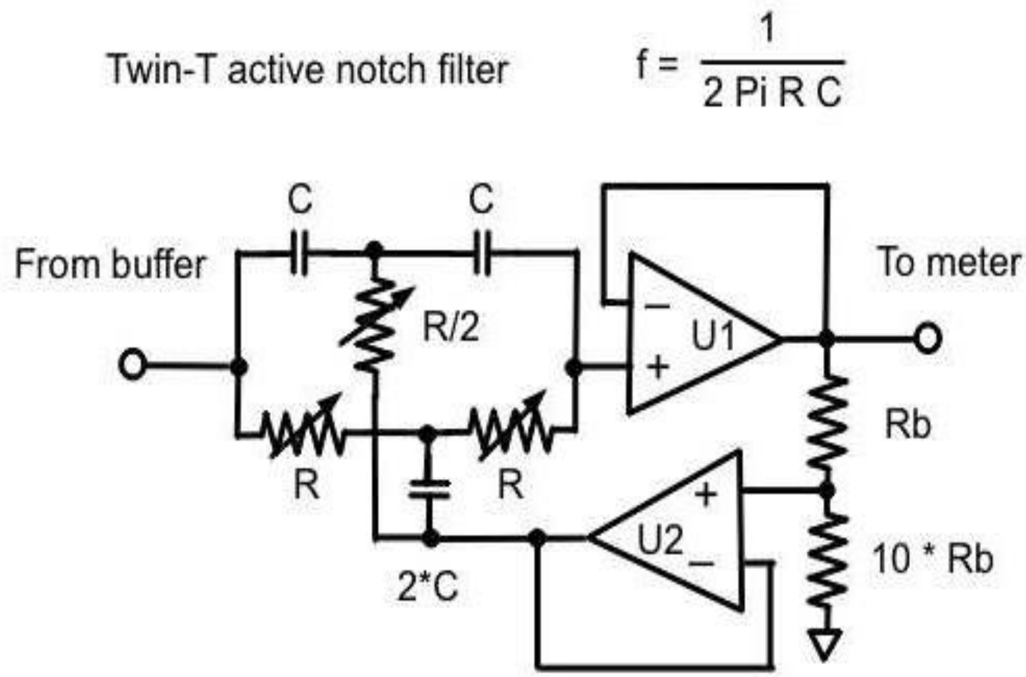

Advantages As shown in the example above, the Twin-T is easily adapted to an active filter design, with feedback to sharpen the notch and thus make correction factors unnecessary. Two sets of switched resistors (large pots don't work well for high resolution due to "graininess" and high temperature sensitivity) and a small precision pot in each leg (with scaled values) can take the place of the Rs shown above in order to yield high resolution; and with great care and even greater patience, such an analyzer is capable of extremely high levels of performance, well into the range of a few parts per million or even much better. See my Active Twin-T notch filter page for details.

Disadvantages But the Twin-T has the huge disadvantage of being made up of three capacitors and three resistors, and they also have unequal values, which makes range switching complicated and makes tuning extremely slow and tedious. Auto-tuning is doubly difficult, and as far as I know, only HP has made a commercial analyzer, the 4333A, that uses a Twin-T filter, and which has a somewhat complicated auto-tuning system. I've never seen or used one of these and don't know how good it is, but I suspect it has high performance.

Despite the disadvantages, detecting the relative phases of the two branches of the fundamental in the bridge is quite straightforward, since the signals at the centerpoints of the two Ts have signals that are 90° apart at precise tuning. Phase detectors can easily sense this exact difference and adjust an LDR in the resistive leg with a lamp or LED, and can adjust an opamp configured as a variable negative impedanceâ€”that is, a capacitorâ€”in the capacitance leg, again with LDR, to achieve precise tuning.

State-variable Filter
The filter type of choice? For an analog analyzer, the state-variable (phase shift) filter may have no equal, since it has no insertion loss, is easy to tune, can be relatively low-Z for best noise performance and lends itself to precise auto-tuning. HP used it in the 8903 series, as did Bob Cordell in his terrific THD Analyzer, and it has been used in many others as well, including Tektronix and Sound Technology units.

It can be capacitance or resistance tuned, with the usual caveats about large resistors for capacitor tuning, so switched-resistor tuning is generally preferred. 1kHz THD floors of 0.0003% (3ppm) or better are achievable, making these hard to beat.

State variable notch filter

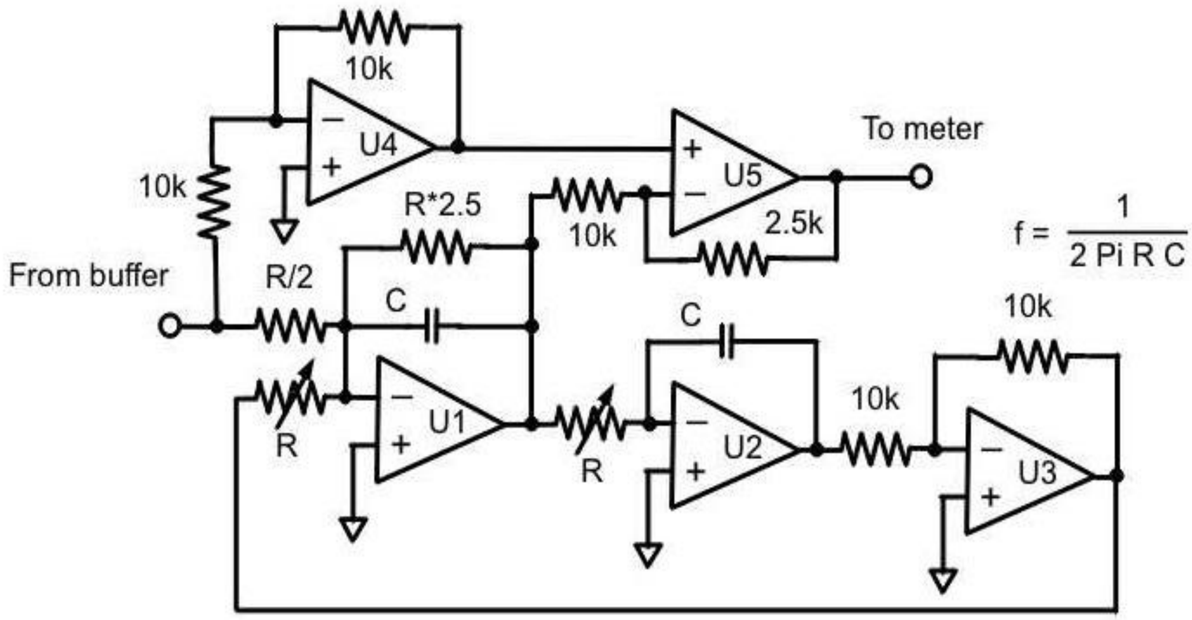

Some comments are in order about this circuit. The two variable Rs and the two Cs together tune the filter at integrators U1 and U2. The output from U1 is the band-pass signal, and the output from U2 is a low-pass signal, if one is needed for anything. U3 is a unity-gain inverter to get the overall phase correct. Unexpectedly, changing either of U3's 10k gain-setting resistors will alter the tuning *frequency,* which provides a clue to auto-tuning the phase differential between the broadband signal's fundamental and the band-pass filtered fundamental. U4 is a unity-gain inverter to drive U5 with the correct polarity wide-band signal.

Resistors R/2 and R*2.5 set the notch sharpening gain, in this case about 14dB, yielding a 2nd harmonic attenuation of a bit more than 0.3dB. U5 is a differential amplifier that algebraically sums the wide-band input signal from U4 and the band-pass signal from U1 and the filter. The two signals to U5 are in phase, meaning that U5 differences them, and so the fundamental is subtracted, leaving only the distortion and noise to be sent to the meter. U5's gain-setting resistors control the amplitude of the differential between the two inputs to U5, again indicating one place where auto-tuning for amplitude can be applied. In essence, U5 reduces the gain in nearly precisely the proportion that U1's gain setting resistors raise it, providing a very deep null. Obviously, changing the ratio between R/2 and R*2.5 or between U5's two resistors, 10k and 2.5k, or between U4's two gain-setting 10k resistors, will change the signal balance at U5, thus changing the null depth.

It only occurred to me after drawing the circuit above that it would be better to put U4 on the output of U1 and invert the band-pass signal, leaving the wide-band signal with the cleanest signal path possible. The band-pass filtering helps to keep the band-pass signal clean, so nonlinearity in U4 will tend to have less effect on the overall result, because we only really care about the fundamental coming from the band-pass filter. Note that the only non-inverting stage is U5, the diff amp, which helps to keep common-mode errors low, but not insignificant -- if distortion residuals will be a problem, here is where they will happen.

Computer software-based analyzers
Digital signal processing is very attractive for distortion analysis, thanks to very good software performance. Most computers have some sound input capability, and very good external or internal "sound cards" are available, so often only a software application (several good ones are free) is needed to do analysis. Software analyzers work by converting the analog input signal to a digital signal using an analog-to-digital converter (ADC). Once in digital form, all of the processing is done in floating-point math at relatively high resolution, so the limits on performance are only the linearity, bandwidth, and signal-to-noise ratio of the input amplifier, and the linearity, bandwidth, and resolution of the ADC. Generally, software systems for distortion analysis work as spectrum analyzers with an amplitude vs. frequency display.

Advantages at least 16-bit resolution is common in nearly all PC sound cards and on-board systems, In general, these chipsets yield 1kHz THD values in the 0.005-0.01% rangeâ€”very good considering the low cost, once the computer is paid for. Some sound cards and on-board chipsets offer higher performance, with 24-bit resolution. These systems often yield 1kHz THD floors in the 0.001-0.002% range, which is quite good, and some offer results in the 2-3 parts-per-million area. Higher performance yet is found in some dedicated digital analyzer systems, such as those from Audio Precision, with measurements better than those from most (or perhaps all) analog analyzers.

Disadvantages The glaring deficiency of nearly all PC-based solutions is bandwidth. The Nyquist Limit constrains digital systems from measuring signals higher than 1/2 the sampling frequency of the ADC. Many PCs have 16-bit, 44.1kHz (or 48kHz) ADCs, meaning that the highest signal component measurable is around 20-22kHz, although 192kHz, 24-bit capabilities are becoming common. This in turn means a maximum fundamental frequency of from 2kHz to 5kHz, depending on the harmonic content of the signal. The best onboard systems typically have 96kHz or 192kHz sampling, with a bandwidth of around 45kHz to 90kHz, allowing examination of fundamentals up to 8 or even 20kHz for lower order harmonics.

Most users will want more measurement bandwidth. Plug-in sound cards and outboard sound systems can offer 24-bit resolution with 96kHz sampling being common, and with 192kHz sampling increasingly available, meaning that fundamentals of up to 10 or 20kHz can be measured, again depending on harmonic content. However, systems with very good linearity are still quite expensive, often in the same price area as used high-performance analog analyzers that have similar THD floors and better bandwidth.

Another disadvantage is that in general, the display of harmonics in such analyzers is a spectrum display, meaning the harmonics are shown as spikes, and the user has to add their individual levels up as the root of their mean squares to arrive at a THD value, which is tedious. Some software systems do the math of computing the RMS sum of the components, which is very handy. Also, software solutions generally do not provide selectable low-pass noise filtering, a feature which is useful in arriving at trustworthy THD values and for seeing low-level distortion components.

Finally, most PC sound inputs have a limited level range, meaning some external level adjustment is needed to handle a wide range of signal amplitudes. This is a relatively minor issue, but needs to be considered, since maximizing signal-to-noise ratio is essential for best performance.

1-9-2011 -- In fairness, a good way to extend the amplitude sensitivity of a computer based analyzer is to pre-filter the input signal with an active Twin-T filter like the one described above. Reducing the level of the fundamental by 40 or 60dB will relieve the spectrum analyzer's input amp and ADC from needing huge dynamic range and extremely low intrinsic THD, leaving only its signal-to-noise ratio as a measurement concern. Providing that you can accurately measure the amplitude of the fundamental, and that you know the precise 0dB reference level of the spectrum analyzer, you can easily determine the dB level of harmonics referred to the level of the fundamental. If the analyzer is reading the 3rd harmonic at -105dB, the 0dB reference level of the analyzer is 1VRMS, and the oscillator is putting out a 5VRMS signal, then the 3rd harmonic is actually at -119dB relative to the fundamental.

Using my PC's "Intel High Definition Audio" chipset sound input and the ARTA spectrum analyzer software, I was able to see a THD residual of my state-variable oscillator, IG-18 #2, as being around 0.0007 to 0.0012%, depending on a variety of factors relating to level and measurement bandwidth. By using the active Twin-T filter between the oscillator and the sound input, and by nulling the fundamental to -80dB, I was able to see that the only significant harmonic was the 3rd, at 0.0005%, with everything else below -120dB or in the "grass" of the analyzer display. With the active filter, the only limitations to the PC system are the maximum frequency of fundamental and the frequencies of the harmonics that are accurately measurable with the overall measurement bandwidth.

3-10-2011 -- But a concern remains, and that is overall linearity. I recently measured the linearity of the system using a precision 40dB attenuator (1ppm accuracy) driven from the very accurate attenuator in my HP 652A Test Oscillator. My PC system has very good linearity down to -100dBu, then begins to compress rapidly below that, so that a signal at -110dB displays as -105dB. Compression in that direction means that the results are worst-case, which is comforting. Now, -100dB is nothing to sneeze at, but every such system needs to be checked for its linearity deviation, especially the 16-bit ADCs. This loss of linearity has effects on the evaluation of analyzer null depth among other things. See below.

5-5-2011 -- Using the very high spectral resolution of the ARTA software, I was able to separate out the oscillator signal from adjacent noise spikes and determine that the linearity is actually much better than first noted above. I found that significant compression didn't begin until below -110dBV, and amounted to about +2dB at -120dB. This is much better than I expected.

6-19-2012 -- Real resolution has come to my bench via the Active Twin-T notch filter described on my webpage about it, and coupling that with an E-MU 0204 USB 24-bit/192kHz sound module. This combination has given me a measurement bandwidth of about 90kHz and superb linearity approaching -140dBV relative to a 0dBV (1VRMS) input level. Meaningful measurements of 10kHz signals are practical, and useful measurements of even 20kHz signals can be made if they are very low distortion (i.e. have only low-order harmonics) to begin with. Highly recommended.

1-18-2011 --
Notes on some specific THD analyzers

HP 330 series (and various old Heath analyzers with two-digit model numbers)
These still show up on auction sites. They are vacuum tube units of the Wien Bridge type that are physically large and which have limited resolution, with a maximum full scale sensitivity of 0.3% (or 1% for some Heath and other tube unitsâ€”and even on some older solid-state units from Leader, Meguro, Marconi, and others). Tuning range is typically 20Hz to 20kHz. Regardless of age, they can be useful and can be improvedâ€”whether that's really practical depends on your frame of mind. Once, in leaner times, I replaced the tubes in an HP 330D with cascode amplifiers made up of JFETs and high-voltage NPN transistors, and was able to get residual THD measurements in the 0.01% area. That unit served well for years, but was a bear to tune since these analyzers didn't have auto tuning. Size and weight are the main barriers to using them as a springboard for greening, but they are treasure-troves of parts.

HP 331A thru 334A series (and Heath IM-5258)
These are solid-state, discrete-transistor designsâ€”no opamps. They are Wien Bridge units with very wide rangeâ€”in the HP units, up to 600kHz fundamentals!!â€”and wide measuring bandwidth, above 3MHz except on the most-sensitive range. They offer low noise, and, in the 333 and 334 (and Heath unit), auto tuning. They employ average-responding metering. They have either a 400Hz passive LC high pass filter that is very good, or a passive 30kHz low pass LC filter, also very good, but not both filters. They desperately need expanded low pass filtering for better performance in the audio mid-band. Max full scale sensitivity is 0.1%, fully usable to 0.01% and below, except that their residual THD floor is around 0.01% (I've been told that the Heath has a higher floorâ€”I've not used one). In the HP units I've used, this floor is dominated by 2nd harmonic distortion, which comes from the bridge amplifier/notch feedback system. I've never checked the depth of the fundamental nullâ€—I really should; but I have good reason to believe that it is better than -90dB and may approach -100dB. Except for their high residual THD floor, and somewhat insensitive metering (300uV/0.1% full scale), the low noise and nulling of these analyzers may make them useful for measurements to 0.003% or even below. I think they may be candidates for greening.

HP 4333A
This is an analyzer onlyâ€”no oscillator is built-in. It was the last one from HP to use discrete transistor circuitry, although it does have a sprinkling of opamps in locations non-critical for bandwidth, noise, or residual THD. As I've mentioned elsewhere, I've never seen or used one of these analyzers, but I do have a PDF copy of the manual. Effectively, the 4333 bridges between the 331-4 and the 339. It has some auto-tuning features and uses a Twin-T filter with fairly elaborate and clever tuning circuits. It specs out closely to the 339, but does not have built-in low-pass filters. It's voltmeter has a sensitivity of 100uVRMS full-scale, with less than 10uV residual noise over a bandwidth of more than 200kHz, so it is pretty quiet. The resulting maximum full-scale THD is 0.01%, so based on noise alone, the THD floor would be under 0.001% -- the spec is less than 0.0018%. The metering is average-responding, calibrated in the RMS of a sine wave. Probably a reliable performer, but certainly less desirable than a 339.

Sound Technology 1700 series
These solid-state units are often for sale on the auction sites and I judge them a good buy at the prices they currently sell forâ€”better in every way than the HP 331-334 and not much more expensive. They have Wien Bridge oscillators and phase-shift/state-variable notch filters, with balanced input and output, and offer max full scale sensitivity of 0.01%, easily readable to below 0.001%. They were the standard for high-performance measurements for quite a while. Some models included intermodulation

measurement capability. The metering is average-responding. I used one for several years, and the residual measurement floor, as I remember it, was in the 0.0012% area. The built-in oscillators have very low THD, using a combination of JFET and LDR AGC, and cover the frequency range from roughly 10Hz to 100kHz (or was it 20Hz to 200kHz?). The analyzers offer both 400Hz high-pass and 80kHz low-pass filtering to improve sensitivity and resolution. Despite the inconvenience of their relatively large size, they are still fine units as-is, and may offer potential for even higher performance, depending on actual nulling capability and noise levels. Sound Tech made extensive use of Harris 2600 series wide-band, low-noise opamps in the design of the 1700.

HP 339A
This instrument, like the ST 1700, has both oscillator and analyzer sections, and has been a favorite of broadcast engineers and techs, audiophiles, and audio service techs for years. It offers true RMS metering with a VU mode if wanted. It has a very good Bridged-T oscillator with JFET AGCâ€”mine has 1kHz THD well under 0.0002%â€”and it has a Bridged-T analyzer section that is comparable in performance to that in the Sound Technology units. Like ST, HP made extensive use of the Harris 2625 opamp, which has absolutely terrific performanceâ€”it's a shame that they are no longer in production. Few current opamps have the 2625's combination of virtues and HP put them to good use. The metering system is true RMS, a great advantage for accuracy.

The 339A I have at the moment has a null depth of around -105dB of greater than -110dB, based on the PC system's linearity test, setting an absolute THD floor of less than 0.0003%. Mine has a residual THD of a bit less than 0.001% with 30kHz low pass filter and 400Hz high pass filters engaged. If a 10kHz low pass filter were available, the 1kHz residual might drop another dB or two, but the real measurement floor is set to around 0.001%, as noted above, by the strong second harmonic spike in the notch amplifier.

It isn't clear to me that the null floor or residuals can be any lower in these units, but if they can, then they may be capable of 0.0003-5% floors. While they are not fully automatic, like the 8903 series, they are easy to use, thanks in part to auto set level, and to simultaneous tuning of the oscillator and analyzer. On the negative side, you rarely see HP instruments from this era with knobs that aren't broken or damagedâ€”not HP's finest hour, mechanically. Because of their construction, I find them hard to work on, but it would be fun to get better performance out of one.

My experience is that the low-pass filter amps are quite noisy, and replacing the quad opamp -- I think it's an LM324 -- that is central to them would be a good idea. I have discovered that the 2nd H. of the oscillator can be minimized by "tuning" one of the two 2.00k resistors at the base of the oscillator's JFET AGC, thereby improving residual THD; and a friend has suggested that replacing the quad op-amp in the AGC circuit with a quieter and faster unit can significantly reduce the THD above 10kHz. This may also require replacing or adding some capacitors.

HP 8903A/B/E
The 8903A has oscillator and analyzer sections with unbalanced output and input. The 8903B has state-variable oscillator and analyzer systems, and has balanced out and in, very useful for measuring bridge-type amplifiers or balanced gear. The 8903E is an analyzer only, with balanced in. These units offer true RMS digital display metering and a very usable frequency counter display, and employ computerized operation with fully automatic level setting and tuningâ€”just connect the signal, push the DISTN button, and wait a few seconds. Their most popular application seems to be in radio production testing and servicing. They offer SINAD measurements, and the fully automatic operation makes them great units for final test or the service bench. I have an 8903E on hand at the moment, and the null floor,

taking measuring system non-linearity into account, is greater than 120dB at 1kHzâ€"much better than 0.0003%.

However, from a high performance viewpoint, these units, which should be naturals, are hampered by a couple of things. One is that the 400Hz high pass filter, if fittedâ€"it and several other special filter types are optionalâ€"is in the signal chain *ahead* of the notch filter, adding a pair of opamps to the chain and contributing significant additional THD and noiseâ€"this filter really should have been after the notch filter, as in the 339A, which is where the low pass filters of 80kHz and 30kHz are.

A second issue is that there are five opamps in the signal chain ahead of the notch filter, even when no special filters are usedâ€"seven or eight if a filter is usedâ€"and this raises the residual THD floor significantly. These opamps, all NE5534s, do balanced to single-ended conversion and auto set level functions, and programmable-gain functions for auto level setâ€"and the programmable-gain amps are not by-passable without interfering with the computer operation of the analyzer.

A third issue is that using the monitor output for post-processing is difficult, because the output level does not have a fixed reference levelâ€"it varies with the gain-setting done by the microprocessor and the various programmable gain amps.

As with the 339, I'm not sure if a deeper null is achievable with the 8903 or how best to reduce the residual THD of the amplifier chain. Better amps than the 5534 do exist, but it isn't clear that using, say, AD797s or LME49710s or OPA134s or LT1122s would actually help. I may have to find out. The 8903E on my bench measures the oscillator in my HP 339A at 0.0009%, with the 30kHz low pass filter in the circuit, and I think that's the best this analyzer can do at present. Having a 10kHz low pass filter available would certainly help. Substantially modding the 8903 would not be easy because of the microprocessor based operation, but it isn't altogether impossible to make constructive changes....

Tektronix AA 501A
I recently (April 2015) had an opportunity to use a Tektronix AA 501A Distortion Analyzer. This fully automatic unit is very easy to use and, except for lacking a frequency display, is fully comparable to the HP 8903A/B/E units in performance and operation.

It looks to me like the notch filter is a Bainter filter followed by a non-inverting summing amp for the notching. And the result is a strong 2nd H., just as with the HP units. The particular one I looked at had a THD residual at 1kHz of 0.0011%, with a 2nd H. level of -105dB, as seen from the Function Output. Its null was deep at -120dB. This analyzer depends utterly on the NE5534 opamp (in singles and duals) and so perhaps a better summing amp, one with less common mode error (an OPA1641?) would improve this 2nd H. issue.

ShibaSoku 725
I recently (October 2014) had one of these on the bench, and it is as good an analyzer as I think exists at this timeâ€"like the HP 8903E, it is an anlyzer only, no oscillator is included. Like the HP 8903E, it offers automatic operation and a digital frequency readout. It uses an active Bridged-T notch filter like the one in the HP 339A, but after that, everything is different.

This instrument uses a 12-bit ADC to digitize and process the post-notch signals, providing a variety of options for analysis, and is capable of resolving THDs of less than 1ppm, 0.0001%! In addition to supplying TRMS metering, it offers several options for average-responding metering, and several LP and

HP filter options. It also provides analysis of the levels of individual harmonics from 2nd through 5th, with resolutions into the ppb (yes, parts-per-billion) range.

I could not determine the self-distortion of the 725C that I got to useâ€"I have no oscillator with low enough distortion to see where the floor actually is. The noise levels of this unit are certainly low enoughâ€"well below -140dBuâ€" to imagine an actual THD+N resolution of perhaps -125 or -130dBu.

What about Boonton, Krone-Hite, and other analyzers?
Great question. The Boonton and Krone-Hite units that I am aware of are basically much like the HP 8903 series, offering microprocessor-based automatic operation, at least 20Hz to 100kHz operating range, RMS metering, digital frequency display, and 1kHz THD resolutions in the 0.001% area. The similarities of specs and performance make me think that all of these units were built to meet some telecom industry or governmental requirements, and this probably includes the ShibaSoku as well, even though it far outperforms all of the others.

From: http://www.moorepage.net/Twin-T.html

Most recent update 12-16-2014
Active Twin-T notch filter
A path to high-resolution distortion analysis

Update — I got some emails....
I designed and built this Active Twin-T filter because I was looking into the distortion performance of ultra-low-distortion oscillators and I wanted to be able to evaluate them throughout the audio range, at least to 20kHz, and also to higher frequencies if possible. So the design described here is fully tunable from 20Hz to 20kHz.

Great if you want to evaluate oscillators, but what about other stuff? I have received emails from folks who just want to test other devices, like audio preamps and power amps, and they wanted a somewhat simpler design, with perhaps a few fixed frequencies. Fine, I answered, that can be done, and this article really has all the information buried in it needed to make such a change. But that got me thinking, what's really needed for this?

Here's the situation — This Active Twin-T notch filter works in conjunction with some kind of Analog-to-Digital-Converter (ADC) and spectrum analyzer software on a PC. I happen to use two different ADCs: an EMU-0204, and a QuantAsylum QA-400. Both of these units have very good ADC performance but they also include very good Digital-to-Analog-Converters (DACs) too, and these DACs let them output very low distortion sine-wave signals.

That capability is of no use in evaluating oscillators, so I didn't dwell on it previously. But for testing audio components, this is extremely useful. Note well that these capabilities are also already built into the sudio functions of most modern computers and in general they work very well, if not quite as well as the outboard components I just mentioned.

How low is low? Modern computers now have audio systems that offer 24-bit resolution in both ADCs and DACs, at sample rates as high as 96kHz, and most go to 192kHz. The result is ADC spectrum analysis capabilities that can resolve individual distortion component levels in the mid-band audio range below -100dB (0.001%) and most can even go to -120dB (0.0001%) or more.

On the DAC side, the midband (1kHz) products can be well under -90dB, and many can go under -100dB. For most audio testing, the performance of these combined units is more than good enough. Using a spectrum analysis package like ARTA (see link below) lets you connect the DAC output to the ADC input in "loopback" mode to evaluate the best possible performance your computer can deliver. You can see the individual harmonic components, and the package can calculate the Total Harmonic Distortion plus Noise (THD+N) for you automatically.

Do you need the Active Twin-T? Only if your gear is truly state-of-the-art — you know, SOTA. I'm not trying to discourage you from building this notch filter, but it's a lot of work to go to if you really don't need it. But, say you do need it — do you need all of its capabilities?

Practical realities of testing Any piece of audio gear worth its salt will have distortion at 100Hz or even at 20Hz about as low as at 1kHz, except for possible power supply limitations in power amps or transformer issues in tube gear; and measurements at 20Hz are often hard to make due to power-line related harmonics and noise. So do you need to test down there with the Twin-T? I believe the DAC/ADC testing will give you the info you need at low frequencies, if you need to test there at all, even with SOTA gear — if your results with the Twin-T are great at 1kHz, no reason to do more than check with the basic DAC/ADC audio system at lower frequencies.

At higher frequencies, the issues are a bit different. Above 20kHz, the performance of the test oscillator and/or ADC input buffer amp will begin to be an issue, even if they are very good, so at 50kHz or 100kHz, just use the audio system DAC/ADC combo. What's left is testing between 1kHz and 20kHz.

So, clearly, testing at (or near) 1kHz is an important qualifier and quantifier of quality in audio gear. From it, results can be inferred about performance in general. Now a couple of options crop up: 1) You can make a Twin-T that works at 1kHz and at 10kHz, which involves (as described in the article below) switching capacitors and using one set (fixed and variable) of resistor values. Or 2) you can test at 1kHz and 20kHz, which again involves switching caps. Or 3) you can test at 2kHz and 20kHz, which is similar to the first option. Or 4) you can test at 1kHz, 10kHz, and 20kHz, the more complicate build, switching both cap sets and resistor sets. I personally like option 1. But many audio designers have followed Bob Cordell's lead and look at the 1kHz to 20kHz distortion ratio as a measure of goodness; for that you need option 4.

Now, on to the build...

Background

Distortion analyzers fall into two main groups -- Total Harmonic Distortion (THD) Analyzers and Spectrum Analyzers (also including Wave Analyzers). The THD Analyzers are represented by units like the famous HP 339A and units from Sound Technology, like the 1701A, and from Tektronix, like the TM AA501A. Spectrum analyzers (and wave analyzers) for audio use are represented by units like the HP 3580A, HP 3581A and the HP 3586A, the Tektronix 7L5, and more recently by units like AP One from Audio Precision; other instrument makers also have units for audio use, like Rhode & Schwartz, and others. All such analyzers use high-Q filters, either band-pass or band-reject, in order to separate the fundamental of a complex signal from its harmonic constituents, which are seen to be non-linearities or impurities in the signal; or to tune to a very narrow band of frequencies in order to isolate and measure the level of one particular frequency.

Broadly speaking, the THD analyzers tune out the fundamental and measure some or all of everything in the signal that's left, whereas wave analyzers and spectrum analyzers tune sharply to individual frequency components to measure their frequency and level. Wave analyzers can only look at one spectral component at a time, but many can plot the levels of spectral components by sweeping the hi-Q tuned filter over the range of frequencies of interest. Spectrum analyzers inherently sweep the signal range of interest to make a plot of amplitude(s) vs. frequency.

The thing is, such sophisticated instruments can be expensive and may or may not serve to reveal extremely low levels of non-linearity in high-quality audio gear. I encountered this problem directly when I wanted to measure the distortion of an oscillator I was modifying, IG-18 #1. I had on hand an old but very serviceable HP 334A THD Analyzer. Unfortunately, it couldn't measure distortion below about 0.01% at 1kHz, and I was pretty sure the oscillator was better than that. As Paul Klipsch used to say, quoting Barnard's Law, "You can't make what you can't measure, 'cause you don't know when you've got it made."

Modern technology has helped folks like me a lot, with computers commonly having pretty decent digital audio capabilities built-in. Naturally, software engineers have made a number of very useful tools to take advantage of these capabilities, including spectrum analyzer software. But I knew from experience that even very good PC sound capabilities, either on-board, as plug-in sound cards, or as USB modules, didn't have the performance to make the kind of measurements I was interested in -- measurements of THD below 0.002%. I tried my project oscillator with a USB sound adaptor from Griffin Technologies called the "iMic" and got very good, but not great results. Was the problem the oscillator, the PC, the analyzer software, or all three?

I built another oscillator that I was very sure had great performance, a state-variable design by Bob Cordell. I built that into a second Heath IG-18 chassis, IG-18 #2, the BIG-18. The PC and the iMic measured it at below 0.005%, but the results were equivocal -- noisy and lots of spuriae. I was sure the oscillator was better, and that the problem was the dynamic range of the hardware. Then I discovered that my PC has 24-bit, 96kHz digital sound capability on-board (sound of head hitting wall...), so I settled on using the wonderful ARTA Audio Analysis software. which includes a high-resolution spectrum analyzer function. ARTA and the PC took IG-18 #2 down to THD under 0.002% at 1kHz; that's very good, but was the oscillator actually better? I had to know.

I reasoned that if I combined a notch filter with the on-board sound capability and the ARTA software, I could extend the measurement range of the computer to see the distortion products at low levels. Would it work?

The Twin-T filter

I needed a way to measure very low levels of distortion that was easy and relatively cheap. I wasn't ready then to buy a high-precision THD analyzer or a wave analyzer -- I thought I could cheap out. Years ago I had built a passive Twin-T notch filter tuned to 1kHz to make quick and dirty measurements of amps and oscillators. The Twin-T is the most attractive notch filter from a performance point of view -- it can (with care) be tuned for a notch depth of over 100dB, though the tuning process is a little (!!!) finicky. The trouble is, a) they are hard to tune over a range of frequencies, b) they require numerous adjustable parts for best performance, and c) they attenuate the harmonics of interest to varying degrees, so that you can't easily make a level measurement with an AC voltmeter or scope and know how much correction factor to apply.

The typical passive Twin-T filter attenuates the 2nd Harmonic of the tuned frequency by about 9.4dB. That's a lot, and the 3rd H. is attenuated by over 5dB. Accounting for variations in just those two harmonics alone makes precision measurements impossible. The reason for this is that although the notch of the twin-t is very deep, the filter's bandwidth at the -3dB points is too wide; in other words, the Q is too low. Clearly, some form of sharpening of the filter's frequency band is needed.

Googling "twin-t filter," for example, will lead you to a number of articles about implementing an active twin-t. The general form of an Active Twin-T filter is shown below. Here, a FET-input opamp U1 buffers the filter and U2 provides the active positive feedback. There is no need for gain adjustment thanks to zero insertion loss. Notch bandwidth, or "Q," is controlled by the ratio of the two Rb resistors -- too sharp and it's hard to tune; too wide and harmonic attenuation is excessive. The 10*Rb value shown is approximate -- an exact ratio will be given below in the schematic.

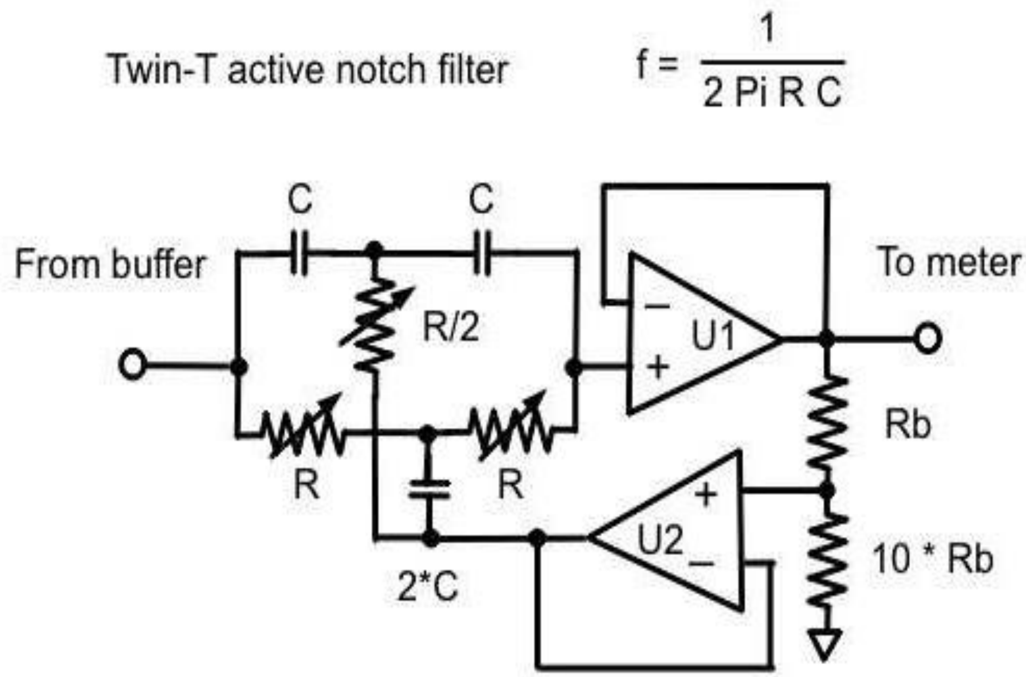

Advantages As shown in the example above, the Twin-T is easily adapted to an active filter design, with feedback to sharpen the notch and thus make correction factors unnecessary. Two sets of switched resistors (large pots don't work well for high resolution due to graininess and large temperature sensitivity) and a small precision pot in each leg (with scaled values) can take the place of the Rs shown below in order to yield high resolution, and with great care and even greater patience, such an analyzer is capable of extremely high levels of performance, well into the range of a few parts per million or even much better.

I consider the Twin-T to be the best overall filter type for an active design, despite its relative complexity, because feedback is used only to raise the filter Q -- in other feedback filter forms, which use algebraic summation, feedback is used to tune for both phase and amplitude matching, and my experience has been that this raises low-order distortion products. In fact, all of the summing analyzers I've looked at have fairly high level 2nd Harmonic distortion in the null amplifier. The HP 331-334 series for example have a distortion floor of 0.01% (-80dB) which is due entirely to the level of the 2nd H in the Wien bridge notch filter. Similarly, the HP 339A and the HP 8903 series of analyzers, outstanding instruments, have a measurement floor of about 0.001% (-100dB), again due entirely to the level of the 2nd H. of the summing notch filters they use.

I first encountered this form of the Active Twin-T notch filter in an article by John L. Linsley-Hood in Wireless World magazine in the late 1970's. That design stuck with me, so I decided to try and make it work with some of the new quiet, low-distortion opamps that are available, like the Burr-Brown OPA134. I breadboarded a circuit, as described in IG-18 #2. It worked very well, but it was single-frequency only. I needed to be able to tune it. I built the circuit using switched capacitors for the range tuning, and potentiometers for the amplitude tuning elements. It worked, but the carbon composition pots I used were prone to temperature problems, noise, and drift.

It works ... mostly

I made some measurements and discovered that not only was it hard to tune and impossible to hold on a null, the noise floor of the whole system was hiding distortion components. Checking amplitude linearity with a precision attenuator revealed that levels below -100dB were being compressed by the PC system and giving poor results. On the good side, the filter's attenuation of the 2nd H. was only 0.5dB and the 3rd H. was less than 0.2dB, so I no longer had to worry about accurate measurements of harmonic levels.

I decided on a rebuild. The rebuild would have three ranges to cover the 20Hz-20kHz audio band -- 20Hz-200Hz, 200Hz-2kHz, 2kHz-20kHz, giving the coarse tuning a "20" to "200" range. It would have switched metal-film resistors for null tuning stability and accuracy, and it would have an additional amp for switching in 20dB of gain to the filter output to get wider dynamic range -- a solid 120dB. The unit would have battery power for great noise immunity.

I was already into this active filter for the cost of some pots from Radio Shack, the opamps (which I had anyway), and a 3 pole, 4 position switch. I could spend a little more -- the second oscillator had cost *quite* a bit more than that. I went to a Michael's hobby supply, which is a nationwide chain, and got a great little 8" x 6" x 3" thin steel box with a lid to put everything in for $5 (I've since discovered that Michael's only have these boxes around Christmas, if they have them at all -- I got lucky). I ordered two 4-pole, 23-position switches from an eBay supplier in Hong Kong -- these switches are used in volume control attenuators for high quality analog audio preamps and such -- these weren't cheap, but they're very nice switches. These let me use 1% metal film resistors as the tuning elements for two of the three levels of tuning precision. This arrangement yields great tunability and stability. I had a couple of 500 ohm ten-turn pots on hand, bought years ago from a friend for just such an eventuality.These provide the tertiary level of tuning.

Resistors

I had to work out what sizes of resistors to use for the switches. I wanted the range capacitors to be easy values to buy and to sort for tolerance. I figured that caps with "10" values would be cheapest and easiest. Given the tuning frequency calculation, a pair of 10nF bridge caps and a 20nF (two 10s in parallel) pillar cap would tune to 1kHz with a total of 15.9k ohms for each of the bridge Rs and half that or 7.95k for the pillar R. Downscale, 200Hz would need about 79.6k ohms in the bridge Rs and 39.8k in the pillar R. At the far end, 2kHz, the bridge Rs would be 7.96k and the pillar R 3.98k.

With the 23-position coarse and medium switches, and the ten-turn pots for the fine tuning, the two bridge Rs and the pillar R can be divided up for high resolution. The three sections of each branch are in series. If the medium and fine Rs can go to zero, then a minimum R for each branch of the coarse tuning at the high end of the range can be a bit under the needed R values. It was convenient to use a coarse R resistor in each branch as a limiting R in series with the switched steps, which sets the maximum tuned frequency a little over the "200" end of the scale, and raised the total series R to give the other end a frequency of just a little under the"20" end of the scale.

The nicest thing would be to have the resistors scale to a log ratio so that the frequency steps would be spread evenly around the panel and knob. But that meant a strange assortment of resistors, and at low frequencies, fine tuning that wasn't fine enough, with the opposite problem at the other end of the scale, where the finest resistor granularity would be too coarse. So, I took the easy way out and used equal resistors for each step, resulting in a scale that isn't pretty, but works. The coarse tuning uses 3.24k ohms for each step of the bridge Rs and 1.62k ohms for each step of the pillar R. This is handy because these are standard 1% metal film resistor sizes -- and the two extra series resistors in each branch make the highest possible frequency about "245", with the lowest frequency about "19."

For the medium, second level tuning, I needed resistor values that would total to somewhat more than a single step of the coarse tuning switch. I chose 165 ohms for the bridge Rs (23 * 165 = 3.795k) and half that or 82.5 ohms for the pillar R -- again these are standard 1% values. The 500 ohm pots were good as is and could be shunted if necessary for more resolution.

Capacitors

The only remaining problems were sorting the caps and adjusting for variances in cap value. I bought a few 5% caps in the 100nF, 10nF and 1nF sizes so that I could select for tolerance. I have a nice little 3-1/2-digit cap meter, and at these values it could get me close to 1% tolerances, matching the tolerances of the resistors. Even so, some variation might upset the cart. So I put a 1k pot in series with each of the three branches, and fiddled until I got a good null at both ends of the range and between ranges; you might have to add some resistance to one of the pots to get everything to work reliably.

Here's the circuit diagram -- DDB, a contributor to the DIYAudio website has made a readable schematic for the filter. You can expand the image in most browsers for greater detail, or save the image to your desktop and expand it in any JPG viewer for easy viewing:

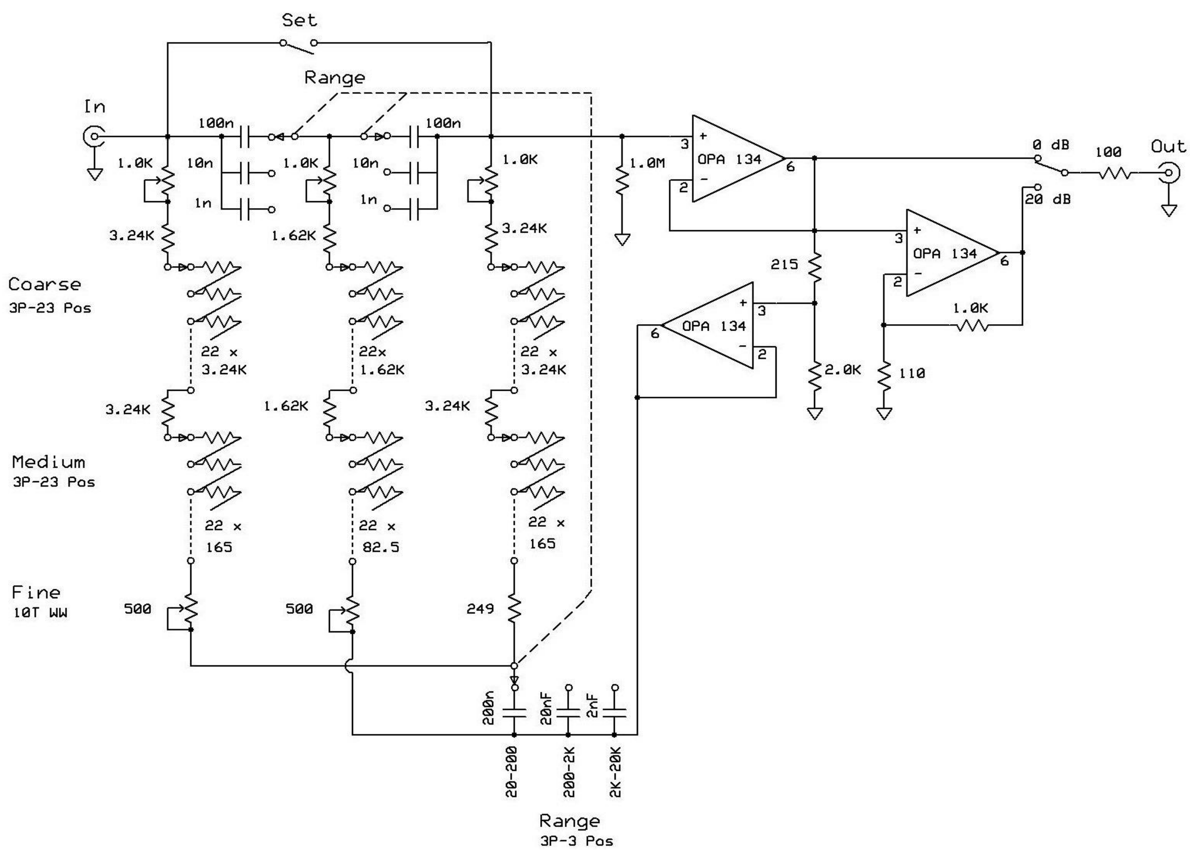

The coarse step values are 3.24k and 1.62k, and the medium step values are 165 and 82.5. These are all standard 1% values.

But Patrick Turner raised the question of using less precise (and less expensive) parts. Resistors can be selected from the 5% series to 1% or better with almost any 3-1/2 digit DMM -- absolute accuracy isn't important, just the relative value of one to the next, and maintaining a very close 2:1 ratio for the bridge vs. pillar values. Convenient 5% resistor values are 3k and 1.5k for the coarse switch, and 150 and 75 for the medium switch. Carbon film resistors work fine in this application, but metal films are just better all around.

My personal opinion is that sorting parts is a pain, and given that 1% metal films are readily available at reasonable cost, buying 10 or 20 times as many 5% carbon film resistors to get the ones you need ends up not being a big saving.

One last issue is the fine tuning pots. The wire-wound 500 ohm pots I used are a little grainy for achieving nulls better than 100dB, but they're what I had -- just a couple of 200 ohm carbon pots may be all that's needed. When you're using a spectrum analyzer, you just need to get the fundamental largely out of the way -- 40dB seems to be good and 60dB of null is plenty. For reading with a meter, however, you want the fundamental gone, and it may take two stages of fine-tuning pots to hold nulls of more than -100dB, if you actually can hold that deep a null long enough. You might need carbon pots of 500 ohms and 100 ohms in series. With the current set-up, I can hold 100dB nulls for a short time, say 10 or 20 seconds -- long enough to read a meter. I can hold -110dB for a second or two -- not long enough for a meter to settle and get read.

How well does it work?
A great question and I have a great answer -- very well. Here's a spectrum of the oscillator in an HP 339:

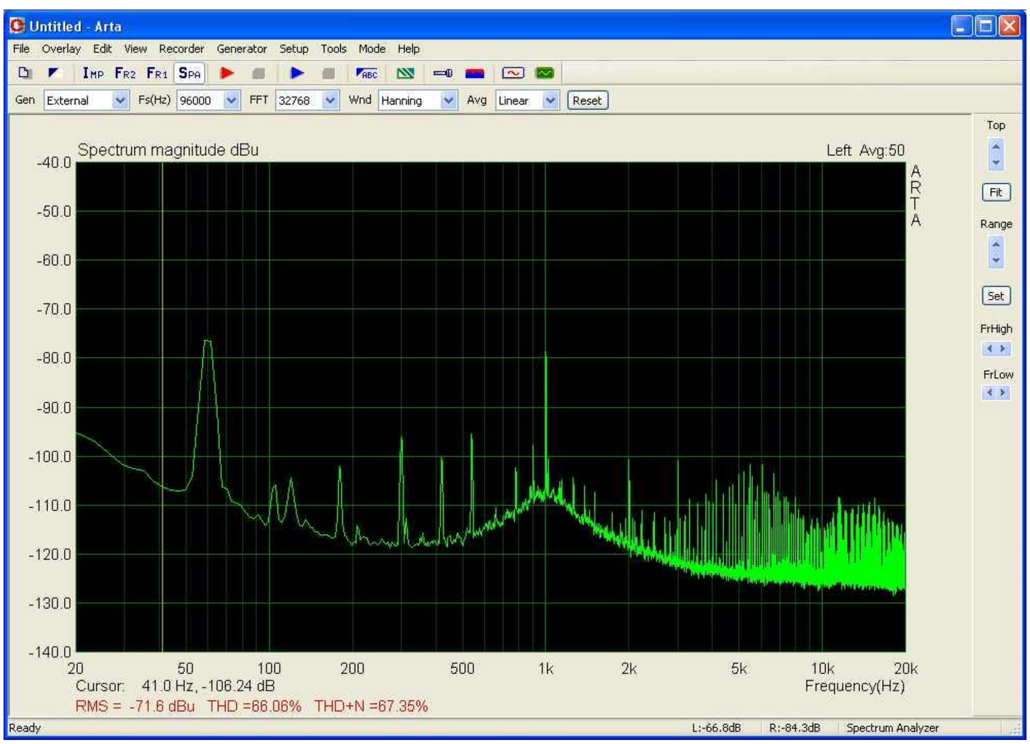

Spectrum of HP 339A oscillator output at 0dBu at 1kHz, nulled with Active Twin-T;
Y-axis values are actually 20dB lower, due to added gain in the filter

Note that the null is at an actual -100dBu, and that the 2nd and 3rd harmonics are actually at -120dB each. The hump in the noise floor at the null frequency is caused by the loss of feedback at the null, and accurately traces the shape of the noise curve for this amplifier and filter configuration.

Here are some pictures of the construction:

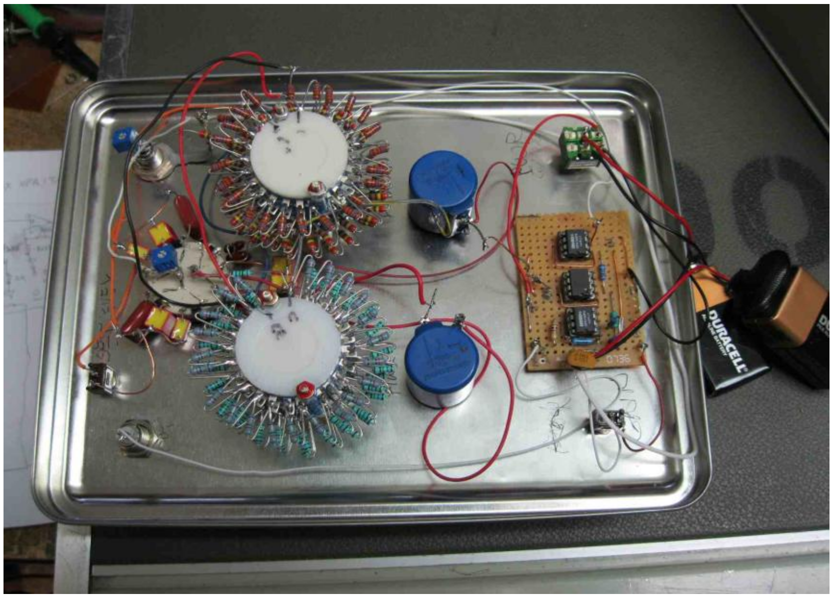

Inside top panel of the filter showing the arrangement of parts

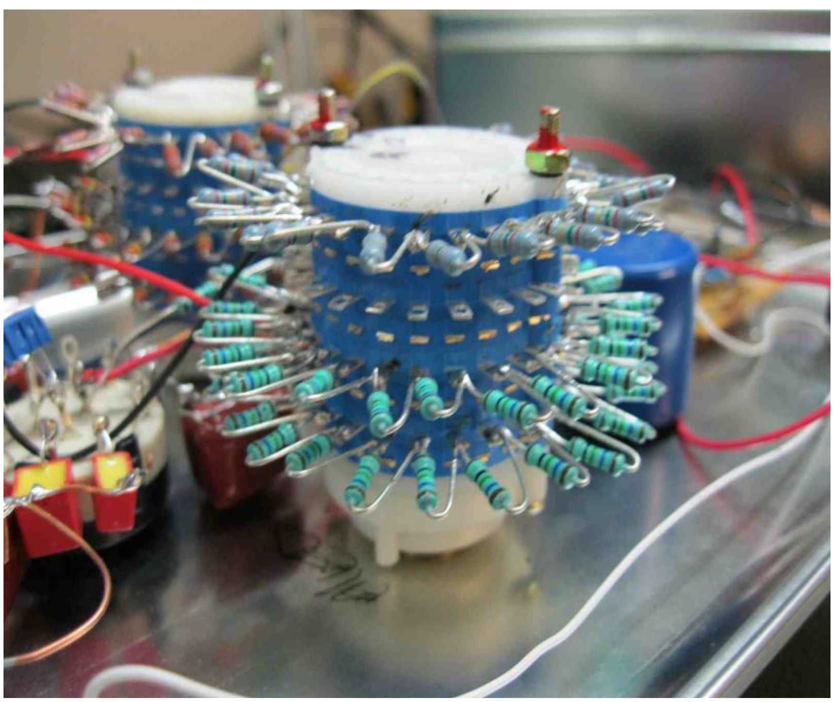

Close up of one of the 23-position switches -- three of its four poles are used

The top of the unit

I hope you're impressed by what a relatively simple active filter and some ordinary computer hardware and very good software can do. This is performance that can't be matched by an analog THD analyzer, and it works every bit as well at 10kHz as it does at 1kHz. Of course the bandwidth of the computer's ADC sampling system needs to be high enough to let you actually see important distortion products, and that is the primary limitation of this combination. My PC does 24-bit, 96kHz sampling, so I can get useful spectral products out to around 46kHz -- enough for working at 10kHz but not for 20kHz. If I had a 24- or 30-bit, 192kHz sound capability then the audio band would be fully covered.

Actual measurements It needs to be said that it is a pain to have to mentally or physically calculate the RMS sum of the products; this filter could be improved by adding low frequency and high frequency stop band filtering for reading results with an RMS meter -- and that's doable, but reading those results also requires a lot more gain to make those 1ppm products big enough for most meters to measure -- say another 60dB of gain. I'm OK with just looking at the spectra and estimating the sum of the products, or just considering them individually. It's very nice to know which products are big and which aren't, so I find the spectrum display very useful.

However, ARTA offers a solution -- you can tell it to calculate the RMS sum of the just distortion components, as well as the RMS sum of THD+noise, too. If you null the fundamental by precisely 60dB, then read the ARTA calculated values and divide by 1000, you have values equivalent to meter readings from a THD meter -- except that there's no noise filtering, of course. ARTA does let you filter out low-frequency hum and noise from the calculated values even though they are not removed from the spectrum display. 60dB is a good value for the null because it's high enough above the noise to keep the fundamental amplitude steady, and it's low enough to ensure that the PC sound system's residuals won't be a problem. And if you add gain from the filter's +20dB amp, be sure to include that too.

As importantly, ARTA does give you a crude but effective way to control the measurement bandwidth -- change the sampling frequency -- I often want a 10kHz bandwidth for measuring the distortion+noise of 1kHz fundamentals, so as to exclude high-frequency noise that isn't relevant but still skews the numbers. I just change the sampling frequency to 22kHz, and the ARTA calculated values are cut off at about 10kHz; for a 20kHz bandwidth I elect to sample at 44.1 or 48kHz. The results have proven to be very close to RMS measured values with similar bandwidth.

This amplifier filter system is quiet enough that the post-filter amp gain could be raised even more if the gear you're working with can benefit from the added resolution. Unlike with a THD analyzer, the spectrum analyzer's measurement floor isn't determined by the depth of the null of the fundamental, so you really can see extremely low levels of distortion.

Update 4-13-2011 -- Larry Burk asked me a good question about an attenuator for the input, for the case of testing power amps at high levels. This circuit likes to be driven from a relatively low-Z source -- the output of a power amp is good. So is the output of an opamp. The circuit is somewhat sensitive to source Z, both in level and Q. For example, at 1kHz, a 600 ohm source causes a drop in the upper passband level of 0.28dB and softens the Q somewhat. As a general rule, the lowest resistance at the input will be about 3/4 of one bridge arm's total resistance. At 1kHz, for example, one bridge arm resistance is about 15.9k, making the input resistance roughly 12k. As frequency decreases, the input resistance increases significantly, so at the bottom end of a range, the resistance will be as much as 5X higher.

For high signal levels like from big power amps, I recommend using a two-resistor voltage divider with a bottom leg of 100 ohms, and an upper leg suitable to get the division ratio you need -- but it may be important to add some capacitance across the bottom leg resistor to compensate for the residual shunt C around the upper leg resistor, making a frequency-compensated attenuator that keeps the frequency response flat out to 100kHz. Try to keep filter set-level outputs in the 0.5 to 1VRMS range when using a PC spectrum analyzer. As long as the voltage divider is reasonably precise, you can monitor the input to the attenuator with a voltmeter and you'll know the input level to the filter by definition.
###

Update 12-21-2011 -- Up until recently, I had been using the on-board sound input capability of my PC to do the Analog-to-Digital Conversion (ADC) work. But I was limited in upper bandwidth by its maximum sampling frequency of 96kHz, and the linearity was a little suspect. I wanted to look at higher frequency products. Having an ADC with 192kHz sampling would allow reasonable estimates of at least low-order distortion products out to perhaps 80kHz, letting me look at fundamentals as high as 20kHz.

After looking at relatively inexpensive "sound cards," I chose an E-MU 0204 USB adapter. It turned out to have lower noise and distortion than my on-board system.

Update 6-7-2012 --
I've had lingering questions about linearity of the EMU 0204's ADC. The measurements made with the Active Twin-T notch filter can't be any better than the linearity of the ADC used with it. I have unwittingly been plagued by ARTA choosing the wrong driver for the EMU 0204 on start-up; the idiot using the software didn't notice that this was happening. This was causing all kinds of strangeness. So I made sure that the ASIO driver was loaded in ARTA and connected my precision HP 353A attenuator to the HP 339A oscillator output, with the output of the 353A connected to the EMU's line input, giving me

a combined attenuation range of over -150dB -- that's plenty. The 353A is rated for a maximum error of +/- 0.5dB over its whole range, and I know that the attenuator in the 339A is better than 0.2dB over its range.

I set the ARTA software to a reference level of 0dBV using a source of precisely 1VRMS (it's good to have an accurate AC voltmeter -- min is an HP 3458A 8-1/2 digit DMM) at 1kHz from the 339A oscillator. Then I stepped the attenuation in -10dB steps. Here are my results, in short form. I'll annotate each spectrum plot below it. The first is at a signal level of -60dB:

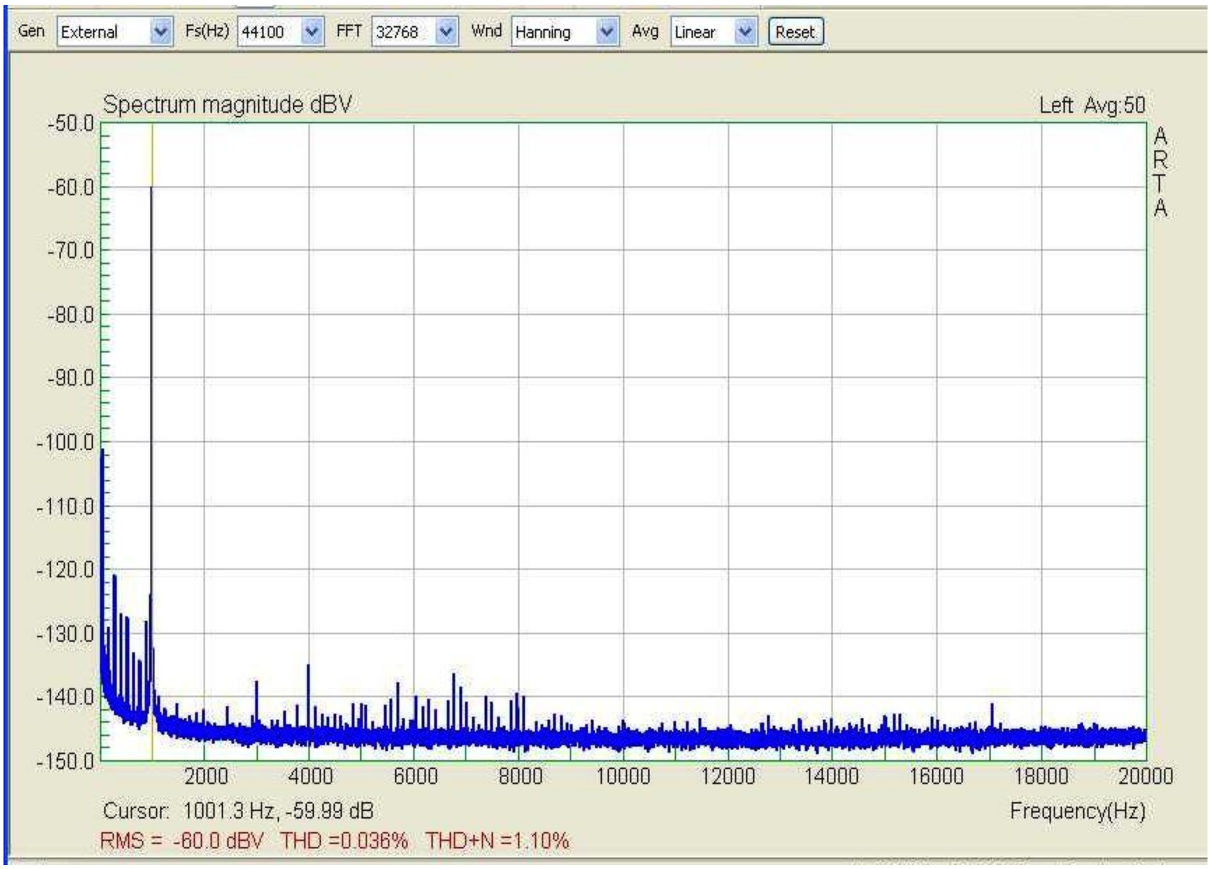

This is a wide frequency range plot that clearly shows the noise floor at the given sampling frequency and FFT sample size. This noise floor is the combined floor of the oscillator and the ADC. Note the cursor value displayed -- I've set the cursor in each plot to show the exact center of the oscillator frequency -- I had to tune the oscillator to hit the center of the band displayed by aRTA for best results, but this isn't crucial. I averaged each plot for 50 samples to reduce jitter and noise effects.

Here is the plot for -120dBV:

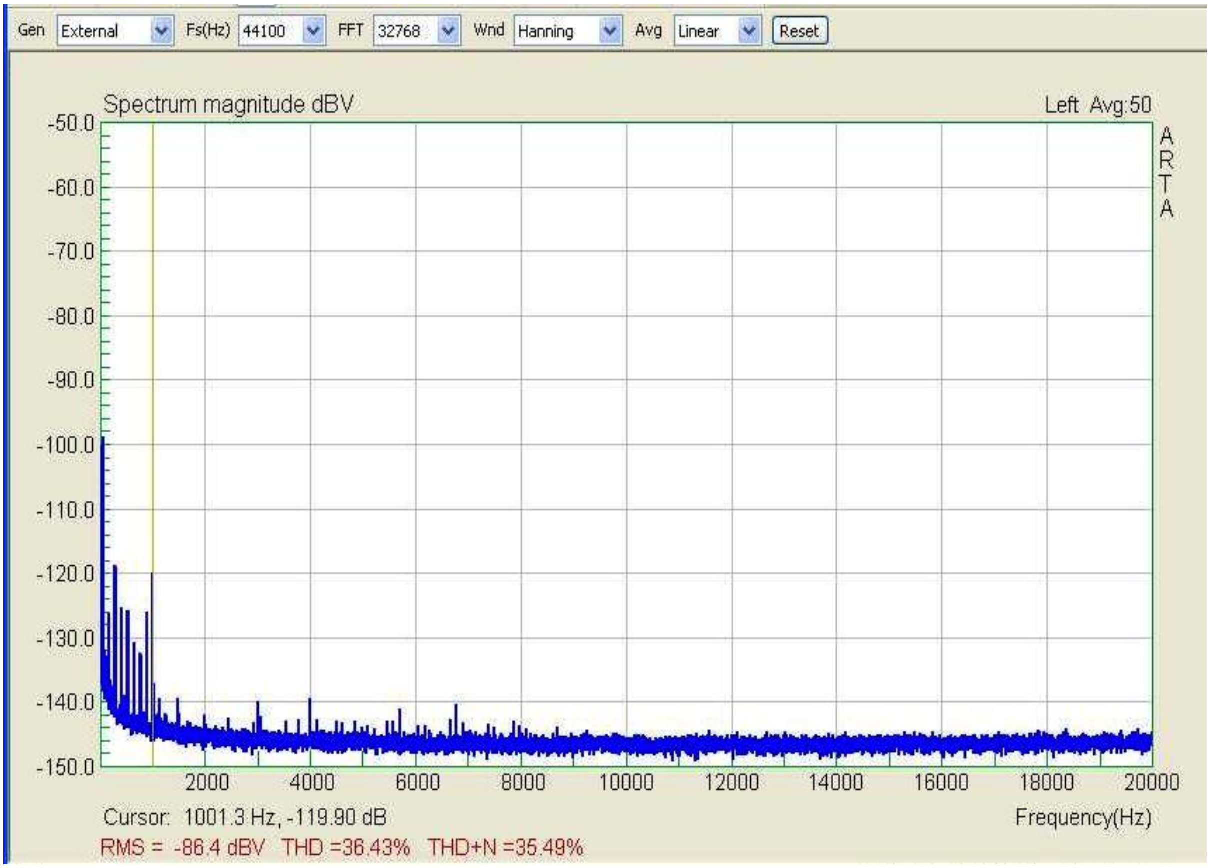

Note the +0.1dB error -- I don't know if that's attenuator error or ADC error.

Here is the same plot up close and personal:

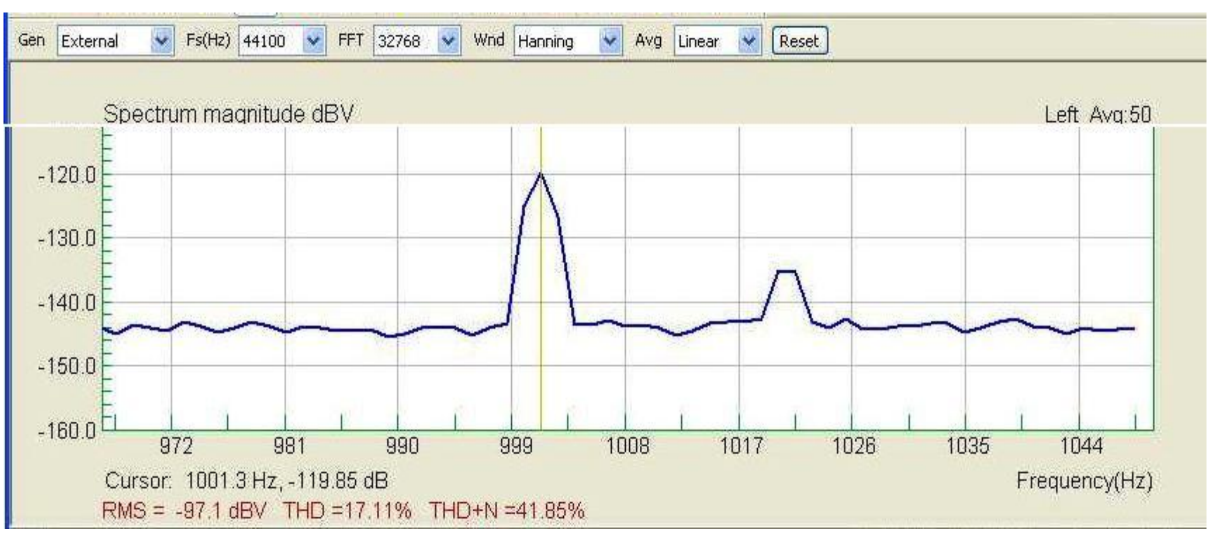

Now the error is seen to be +0.15dB.

Here's the plot for -130dBV:

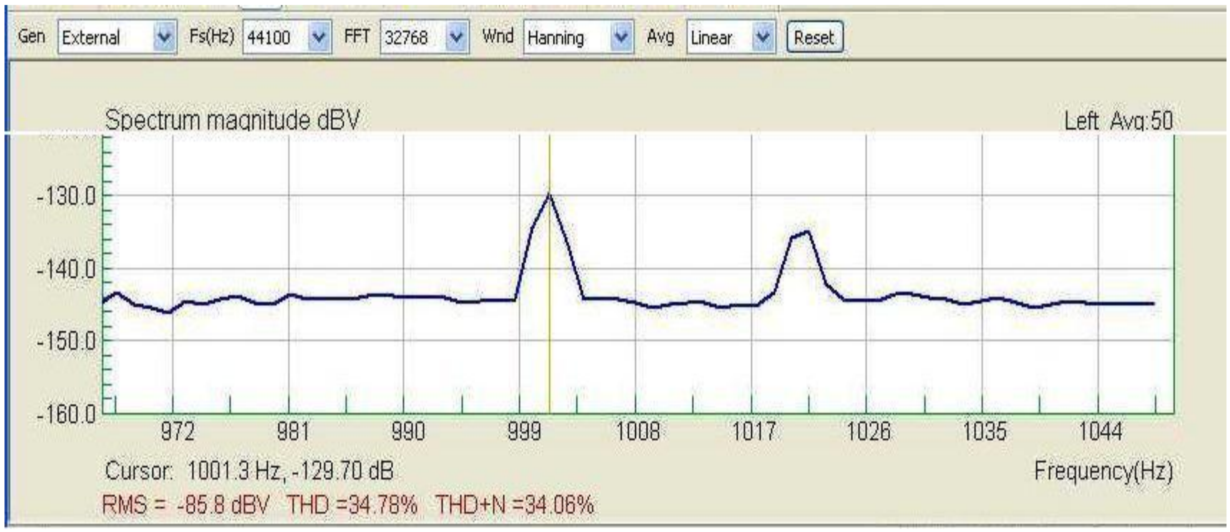

The error is seen to be +0.3dB.

Finally, just to impress, here's a broad spectrum plot at -140dBV:

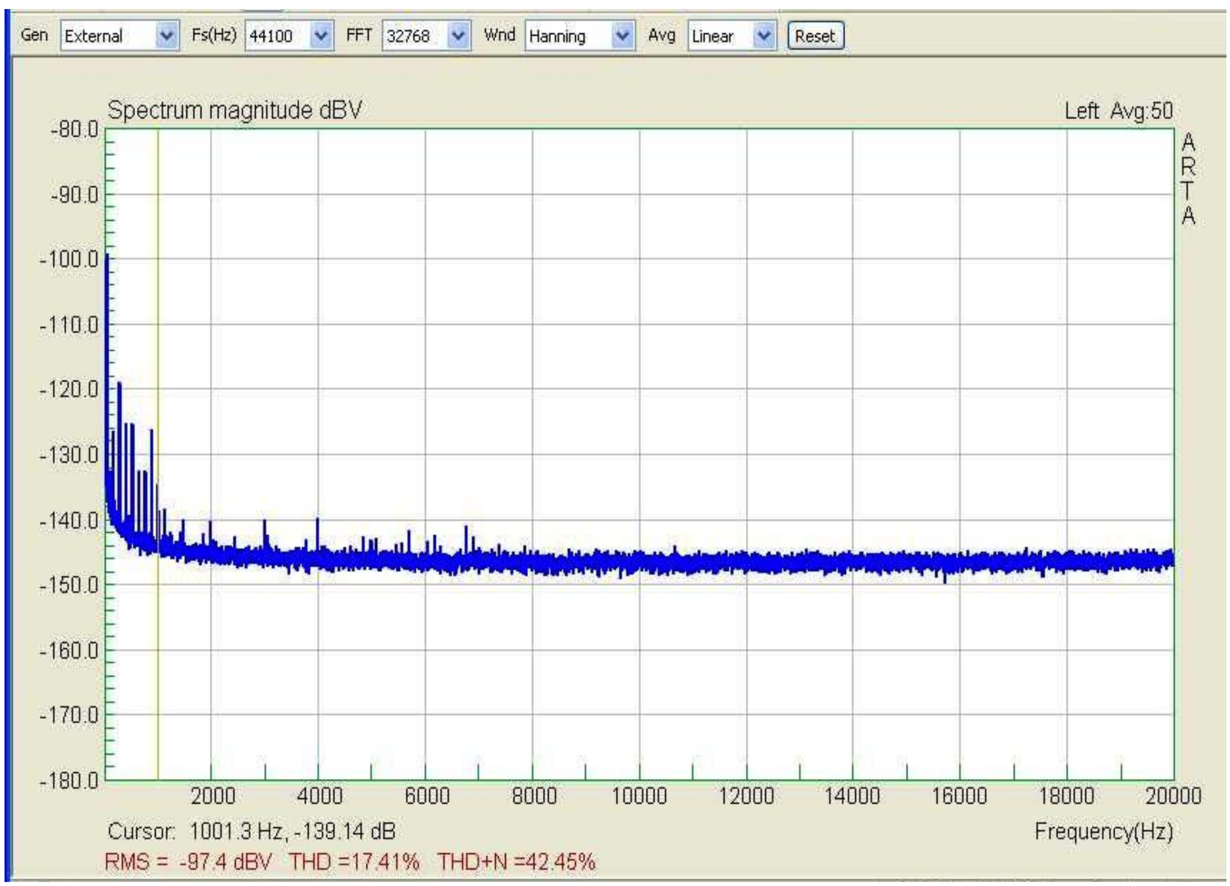

Note that the small spike of the oscillator is obscured by the cursor and is between two larger noise spikes. ARTA says the error is less than +1dBV.

Now it is possible that the combined errors of the two attenuators and the non-linearity of the ADC combine to null each other out. If so, that's a happy result and I'll take it.
###

I've done some evaluation of the EMU, and have begun to ponder the question: Where is the residual distortion coming from? In my test setup, there are three sources of distortion -- the oscillator, the Twin-T filter, and the EMU.

I don't have a distortionless oscillator. Using the digital oscillator in the ARTA analyzer program, fed thru the EMU's DAC, isn't the solution -- distortion is low, but not low enough. My HP 239 oscillator has low distortion, as does the oscillator in the HP 339 Analyzer -- they are the lowest distortion oscillators I have. But how low is their distortion? To properly evaluate the residual distortion of the Twin-T filter and of the EMU ADC, I should have an oscillator with distortion products lower than -140dB, and lower if possible. Well, that's not going to happen right now, although my friend David Barber is working on it. All I can say is, "stay tuned."

What I have done is make some spectra of the 10kHz output from the HP 239, using various combinations of level settings and the other gear. The descriptions follow each spectrum. I chose 10kHz for examination because, although the oscillator has lower distortion at 1kHz, so do the Twin-T and the EMU. 10kHz stresses everything just a bit more. I chose to use 1VRMS as the highest output from the 239, because that is the level that the EMU is good with at moderate gain in its input amplifier. With these settings, the EMU clips at an input of a bit over 3dBV, so 0dBV is comfortably under the full-scale input, but high enough to assure lowest ADC distortion. I tried to find the best distortion+noise vs. level trade-off, but that's really hard to evaluate when everything in the signal chain is near its lowest levels anyway. This, as Paul Klipsch used to say, is "milking mice."

What can be expected of the EMU, given its 24-bit resolution? that's a dynamic range of 144dB! But noise and distortion seriously limit that maximum performance, including contributions from the input amplifier. It's an open question as to what can be expected here from any of the gear.

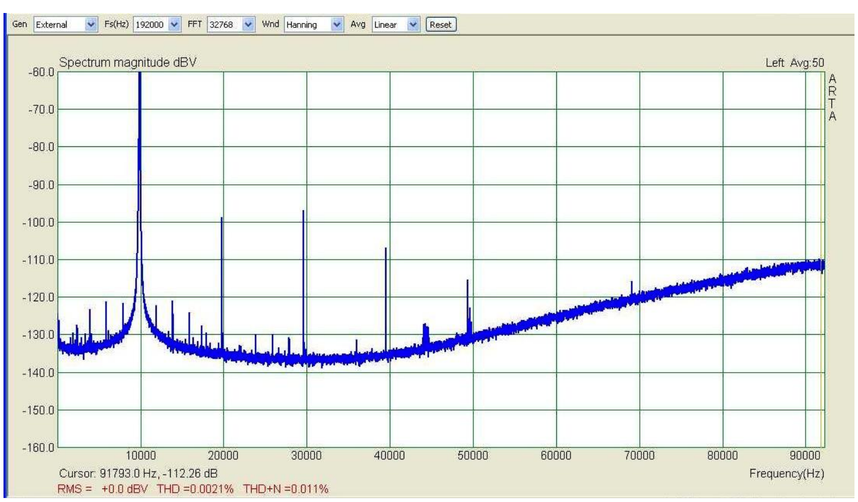

This is the 239 fed directly into the EMU. The indicated spectrum levels are referred to 0dBV = 1VRMS. The 239's attenuator is set to 1V, but that's its level into a 600 ohm load, so the variable knob is adjusted

-6dB for 1V out. The 239's output feeds a passive 30k ohm attenuator pot (set very near full output), which feeds the EMU. The 2nd and 3rd harmonics are near -100dB, which is the distortion of the EMU's total system -- input amp plus ADC -- at this input level; the HP 239 is lower. So, these distortion levels are the floor for the E-MU at near-full-scale input.

Note the rise in the noise floor as frequency increases. I speculate that this is an artifact of the noise-shaping filtering of the E-MU's sigma-delta converter. ARTA's calculated THD of 0.0021% (-94dB) looks to be right on compared to the product levels.

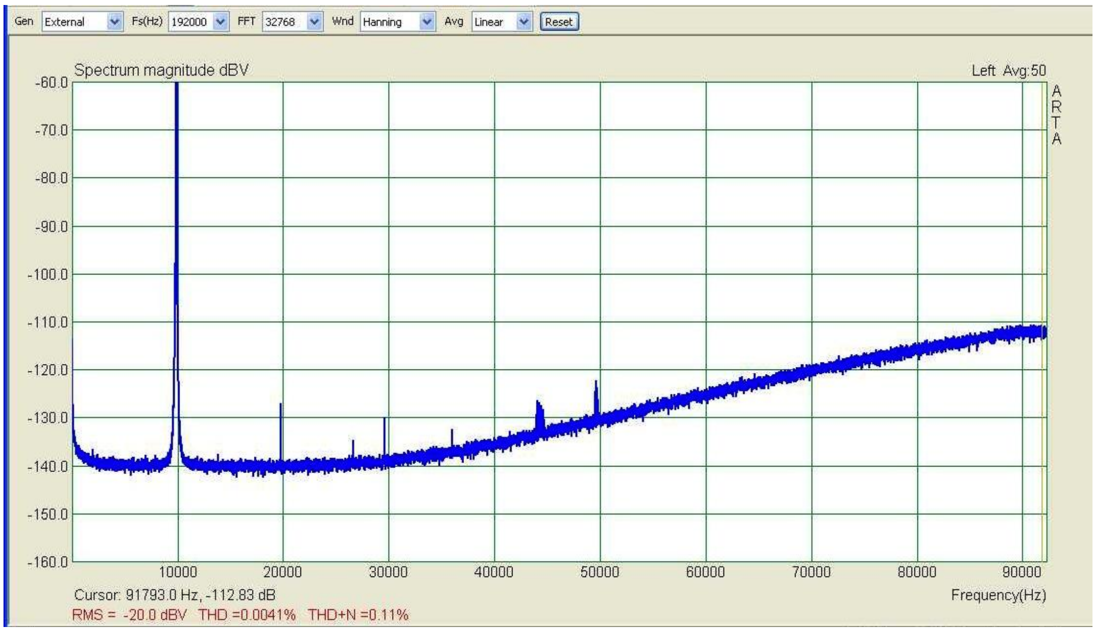

This is the same as 1 except that the output attenuator of the 239 is set to 0.1V, a drop of 20dB in the input to the EMU. The distortion levels have decreased by almost 30dB, not the expected 20dB, which is counter-intuitive, because the distortion of ADCs goes up as level goes down. So, this could be the lowered distortion of the EMU's input amp at the 20dB lower level, bringing us near the floor for the EMU as a whole -- but we also could be closing in on the oscillator's distortion.

ARTA's calculated THD of 0.0041% looks all wrong compared to the levels of the products on-screen -- I have no explanation. Everything above 40kHz seems to be spurious junk -- there is some 50kHz 5th harmonic signal there but it's being modulated by line and/or other noise, and its true level is not visible.

The next question is: Do the Twin-T's amplifiers add distortion? So first, a spectrum of the Twin-T in set-level mode, then a spectrum in filter mode:

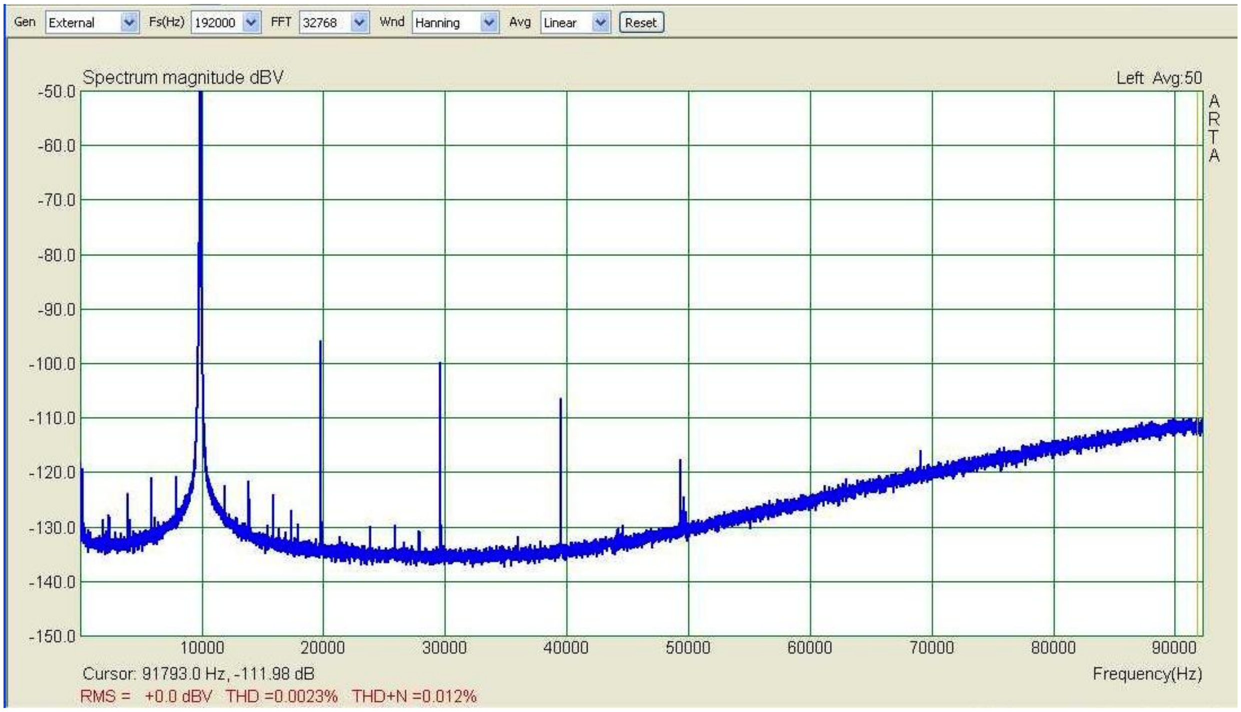

This is the same as Spectrum 1 with the Twin-T in the chain. The Twin-T is in set-level mode, so there is no filtering. There are slight differences in the relative levels of the 2nd and 3rd harmonic products compared to 1, but the overall distortion is about the same. So the Twin-T isn't adding much, if anything.

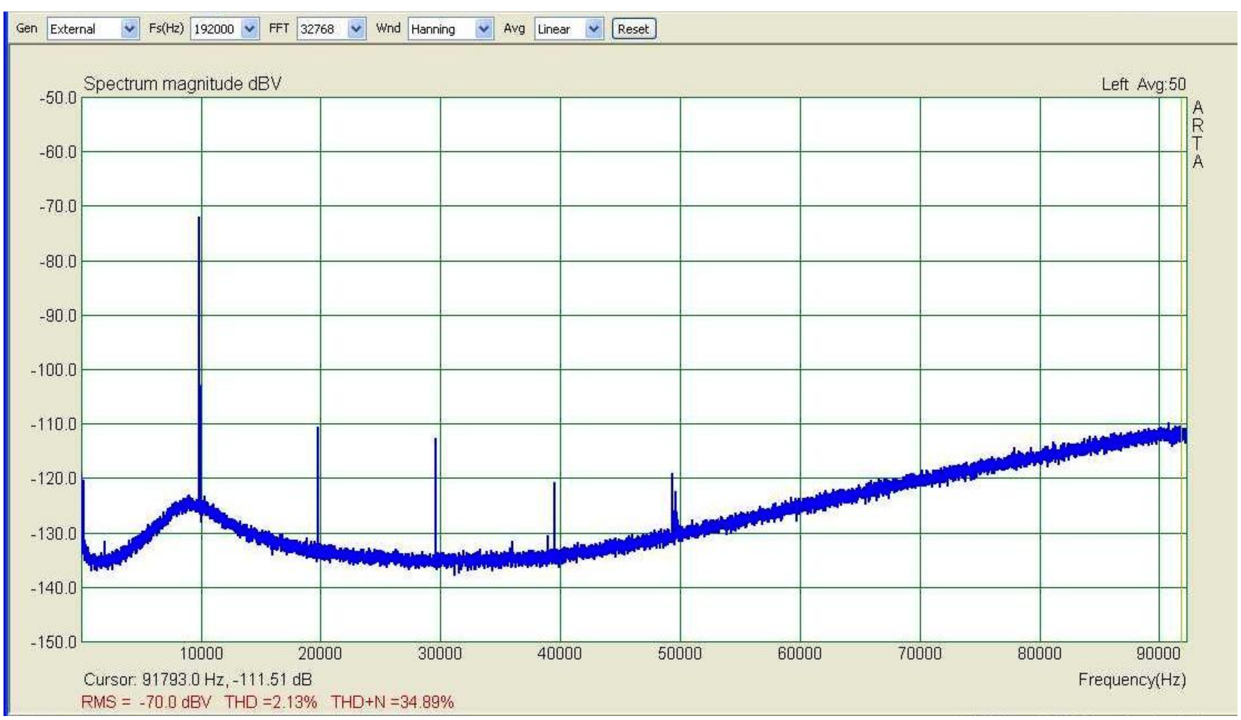

Here, the Twin-T is in filter mode, as can be seen by the low level of the fundamental. The distortion product levels are referred to the 0dBV level. The low fundamental makes ARTA's calculation of the THD hard to refer to the 0dBV level -- but it calculates out about right for a THD of about 0.0005% -- pretty good for a 10kHz oscillator. What matters are the levels of the products compared to the levels in Spectrum 3. The levels have dropped by 10dB or more, which means that the EMU's distortion has gone down, but not nearly as much as in Spectrum 2, indicating that we are seeing either the residual of the HP 239 oscillator or the residual of the Twin-T filter, or both. Which is it?

Spectrum 5 --

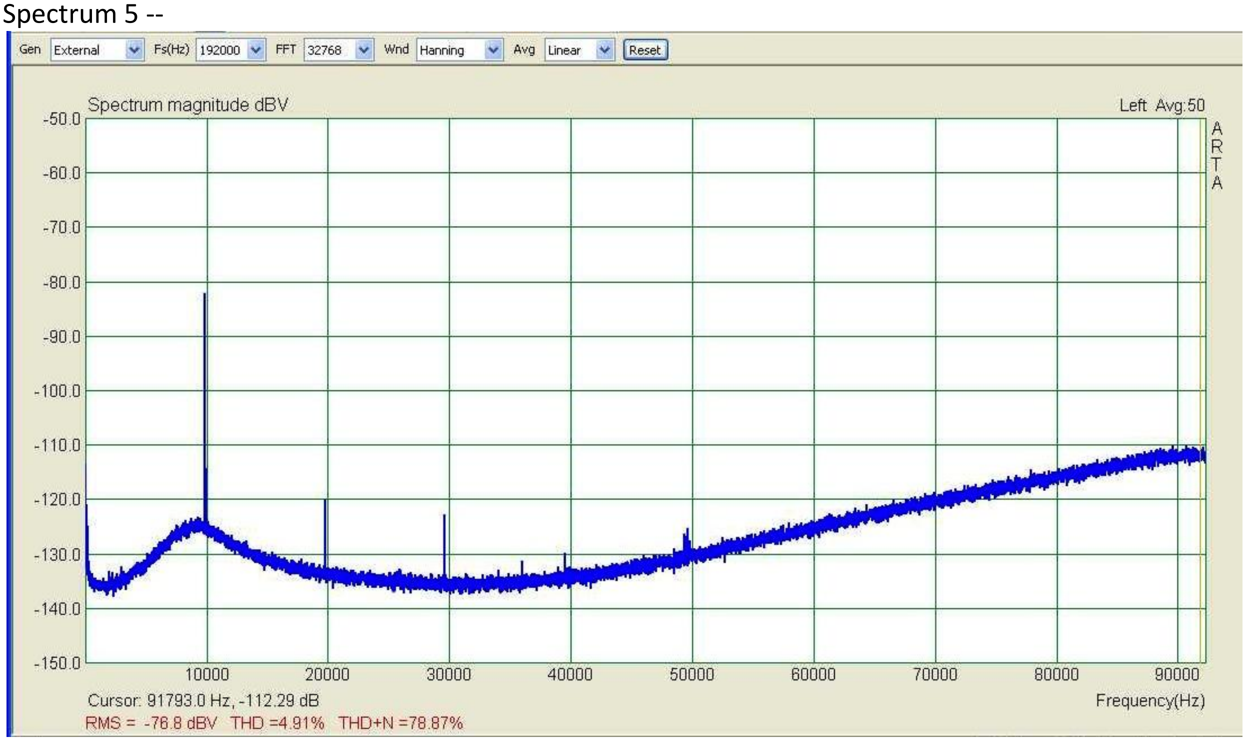

This is the same as 4, but the oscillator's output attenuator is dropped 10dB -- this does not change the distortion relative to the level, it just changes the level. If the distortion is primarily from the Twin-T, then the lower level should result in more than a 10dB drop in the levels of the products. If the distortion is primarily from the oscillator, then the products should drop by the same amount as the attenuation of the output -- 10dB. As can be seen, the products do drop almost exactly 10dB, indicating that we are seeing the 10kHz distortion of the HP 239.

Given the outstanding linearity of the EMU's ADC, this result means that the Active Twin-T and the EMU 0204 together offer an accurate view of very low distortion levels -- in the area of -130dB or lower -- in oscillators and other electronic gear.

Update 10-28-2012 --
Simpler configuration for a few frequencies
I've had some requests for more information about parts values and possible modifications, and I've been thinking about the usability of the filter. I find I mainly use it at 1kHz and at 10kHz, with other frequencies only very rarely. The design is easily simplified to do just a few frequencies with much less elaborate switching of the tuning resistors. A possible penalty is not having much tuning range, but most oscillators that will be used as sources can be tuned as well. Having four frequencies, for example, would require twelve 1% resistors, at most twelve capacitors (but probably fewer, depending on the

frequencies chosen) and three 10-turn wirewound pots -- these pots should have values 5% or less of the R value in the frequency formula.

The detailed description above covers the formula for tuning frequency, so choosing resistors and capacitors is straightforward. In general, the resistor value R shouldn't be under about 2kohms, but higher is better to avoid source loading; unless the cap value C is then too small -- this is a bigger issue at high frequencies, as a little work with a spreadsheet or calculator will show. A C value that is too small will be significantly changed by the inevitable stray C of the parts and wiring, making use at 50 or 100kHz, for example, a challenge.

###