(Yet Another) Tracking Generator for the HP 8566B Spectrum Analyzer

By Joseph Haas, KEØFF
8/28/2016
joeh-*at*-rollanet-*dot*-org

Introduction

I recently acquired an HP 8566B Spectrum Analyzer (S/A) for a bargain price. I was familiar with the analyzer, having worked with that specific unit before, so I knew it was in reasonable shape. For those not familiar with this model, it covers 100Hz to 22 GHz and is a workhorse in the EMI/EMC test lab venue. The wide frequency range makes it an excellent general purpose analyzer, covering literally audio to microwaves. If you are only going to own ONE spectrum analyzer, the HP8566B should be on your short list.

As capable as this analyzer is, it is not without its drawbacks. For one, it is massive. It is comprised of an RF unit and a video unit. Together, these weigh in at a whopping 112 lbs (51 kilos). This doesn't even come close to "transportable" much less portable. It is also large: with the rack-ears, it measures in at 19" (48cm) wide and over 26" (66 cm) deep including the rear-cables. I decided to purchase a lift table to support the new addition to my home lab since bench-mounting was not at all attractive.

Figure 1. The HP-8566B

Most notably, the HP 8566B lacks something that many modern analyzers seem to generally have included (or at least provisioned for): a tracking generator (T/G). A tracking generator provides a test signal that tunes in sync with the analyzer input frequency. This allows circuits and devices to have their frequency response "swept" across a frequency range. The result of the combination of S/A and T/G is the beginnings of a scalar network analyzer (an SNA measures magnitude only, no phase), which can be an indispensable tool for testing/tuning filters, amplifiers, and other RF components. While such a setup neglects phase, when one considers the cost of a true network analyzer, the S/A-T/G combination is a very cost effective substitute which can offer valuable data about the frequency dependent behavior of a device under test (DUT).

The T/G is not the only way to produce a SNA using an S/A. Another popular method is to use a sweep oscillator. These are generally dedicated pieces of test equipment, but they can be constructed using any frequency source that can be electrically tuned. Two methods are generally employed: open loop, and slaved. For the open loop method, the sweep oscillator is constantly cycled while the spectrum analyzer is in max-hold mode (generally, the S/A sweeps relatively quickly, and the sweep oscillator sweeps relatively slowly). In this way, the two pieces of equipment asynchronously cycle through the sweep range (ideally, they would both have the same start-stop frequencies, but this is not required) eventually filling in the S/A span with the scalar response of the DUT. The slaved mode couples a span signal from the S/A (a roughly saw-tooth analog output that varies from some low voltage to some higher voltage as the S/A sweeps across its span, provided by some analyzers) to drive the sweep oscillator frequency.

These methods are used often (in fact, I employed the open-loop method to investigate the performance of some of my T/G components [(4)]) and can produce results similar to those of a tracking generator, and generally require test setups that are less involved than a construction of a T/G. The drawbacks are that the open-loop method can take an inordinate amount of time to fill in a full trace (as much as several minutes for a single trace, which makes it very time-consuming to perform adjust-and-measure operations), and the slaved method is usually only good to spans of a couple hundred MHz or less and not all S/A devices provide the span signal.

This document describes the journey that was inspired by my desire to add tracking generator functionality to my HP8566B. It turns out that others have followed this path before me. Their efforts were very helpful and are noted in the bibliography. It also turns out that the fundamentals of tracking generator design are not all that difficult to understand. This all hinges on the provision of a critical signal from the S/A: the 1st LO output.

The Beginning

Most spectrum analyzers are basically just a super-heterodyne receiver that is tunable across a (typically wide) range of frequencies. Rather than outputting a demodulated audio or data signal, the base-band detected signal is fed to an oscilloscope-style display which is swept in concert with the tuned-frequency sweep. The result is an amplitude vs. frequency display of the signal spectrum – essentially the Fourier Transform of the time-varying input. Modern analyzers (and even ones like the HP 8566B) are somewhat more involved in that they generally digitize the spectral result for post-processing and display, but all S/A designs begin with this basic structure.

Thus, the signal chain of an S/A (see Figure 2) consists of a swept local oscillator (LO) that is mixed with the input signal to reach the 1st IF. The 1st LO, by definition, sweeps in step with the tuned frequency of the S/A input, offset by the 1st IF frequency. In order to produce a tracking output, one simply needs to mix the 1st LO (hereafter referred to as the SALO) with another LO (hereafter referred to as the TGLO) to “add-back” the 1st IF frequency, resulting in a signal frequency that is located exactly at the S/A tune point at any given moment.

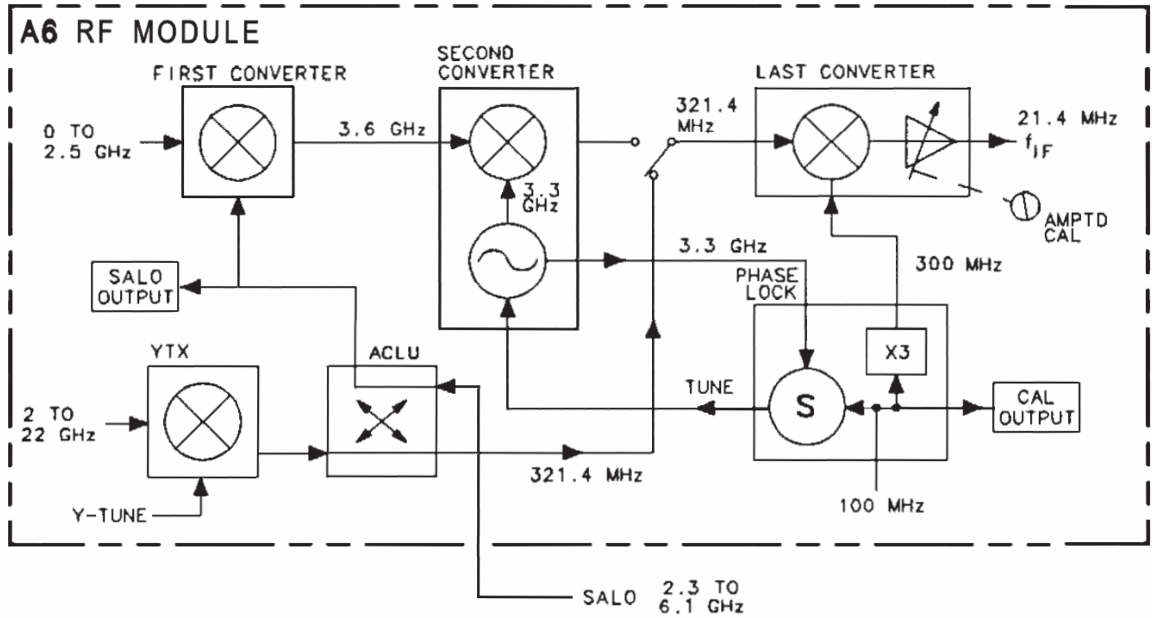

Figure 2. Simplified Block Diagram of the HP-8566B

Of course, nothing is ever as simple as it first appears. In the case of the 8566B, and many other wide-band spectrum analyzers, the entire bandwidth of the analyzer is not swept in a single throw. In addition, the variances introduced by the various band-pass filters employed between the S/A input and detector can introduce issues. For the 8566B, the input scan ranges are broken up in to several bands. The focus of this project is on the first band: 100 Hz to 2.5 GHz. For the first band, the S/A uses a 3.6214 GHz 1st IF (this information is presented in the detailed block diagram for the 8566B).

All that is needed to produce a T/G output is to mix the SALO with a TGLO of 3.6214 GHz. The output of this mixer (which is down-converting for this configuration) will produce a signal that is at the S/A input frequency as long as the 8566B is set such that its start and stop frequencies are both below 2.5 GHz. The mixer will also produce an image as well and other possible product results. However, as long as these other frequencies are not excessive, the filtering of the S/A will generally be sufficient to ignore them. Of course, when considering that an active DUT might be involved, these off-frequency products may not be as easy to sweep under the rug.

The Details

There are several subtleties that influence the detailed design of the T/G hardware. A partial list follows: the span of the analyzer, the LO frequencies that are used, the level of the LO output, the isolation of the SALO output, and fine-tuning of the TGLO to account for variances in the narrow-band filter settings of the S/A. Of these, the isolation of the SALO from the T/G mixer is generally the most important design parameter if one wishes to achieve the highest possible dynamic range of the T/G response.

The SALO output is generally derived from a directional coupler inside the S/A (this coupler is not depicted in the simplified diagram of Figure 2). The coupler samples a small amount of the LO energy just before injection into the 1st mixer. This diverted LO energy is routed to the SALO output connector. While the coupler is directional, the amount of directional isolation is not generally enough to completely reject any non-desired signals that might try to travel INTO the SALO output.

The effects of this issue can be easily observed by simply sweeping the S/A over the full span with the T/G mixer connected to the SALO, and with no S/A input (see below). The resulting plot will depict the degree of coupling of stray mixer products and other interfering signals into the SALO connection. Any result that is above the normal S/A noise floor (without the T/G mixer connected) is indicative of back-coupling which has the effect of reducing the dynamic range of the S/A-T/G combination.

Dynamic range is important only when there isn't enough of it to spare. This is generally determined by the type of measurement and the resulting measurement setup. Because of this, the best result is to have as much dynamic range as possible to keep from being "boxed-in". Generally, the minimum S/A-T/G dynamic range is dictated by the maximum attenuation that one might expect to see in a DUT. If the DUT attenuation at a given frequency exceeds the dynamic range available at that frequency, the true behavior of the DUT in that region will be obscured by the artificially raised noise floor of the S/A. Generally speaking, for a +10 dBm T/G, the best dynamic range one can hope for is on the order of 70 dB or so (this will vary based on analyzer, span, and filter settings).

To address this, it is good design practice to include a significant degree of isolation and shielding into the SALO portion of the T/G signal chain to reduce the effect of mixer products or other signals that might find their way to the SALO connector. A terminated circulator, A.K.A. an isolator, is a good candidate but most offer only 20 or 30 dB of isolation across a particular bandwidth. This isolation also generally drops off significantly outside the specified bandwidth.

Others who have attempted this tended to focus on an isolator that covers the SALO range. However, the isolator actually needs to cover from the S/A input-low-frequency up to beyond the upper limit of the SALO. In this case, the ideal isolator coverage would need to be from 100 Hz to above 6.3 GHz. Isolators that cover such a wide bandwidth and to such a low bottom frequency may well exist, but only as very expensive components in very expensive test equipment. Thus, while an isolator that covers the SALO range is good practice, other means of isolation must be employed to improve the T/G performance.

A high-isolation amplifier stage can offer an additional 30-50 dB of isolation. However, not all amplifiers feature high isolation. For this design, a MiniCircuits GVA-62+ was employed with attenuators to produce a nominal 35 dB isolator with essentially no gain. An attenuator at the output of the amplifier also has the benefit of presenting a better broadband 50Ω match to the T/G mixer, helping to reduce spurious mixer products.

To confirm that this approach has merit, the T/G mixer was tested with an isolator and a removable attenuator installed between SALO and the T/G mixer. For this test, the mixer RF port was connected to the SALO via a 6SD61 isolator. The LO port was connected to the PLL and supplied with a 3.6124 GHz TGLO signal. The T/G mixer IF port was terminated with a 50 ohm load and the S/A was set to plot from 1 KHz to 2.5 GHz (the S/A input was left open).

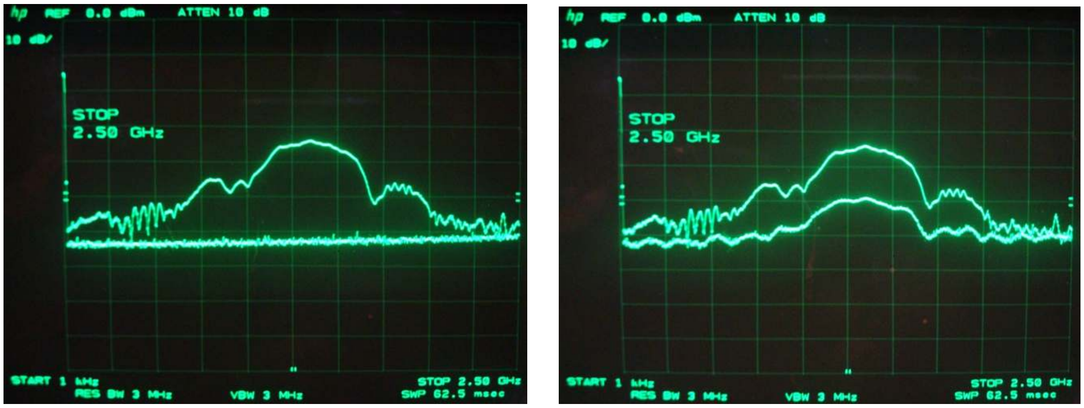

(a) (b)

Figure 3. At (a), the S/A response with the T/G mixer active (top) and inactive (bottom). At (b), the S/A response with no attenuation between the isolator and T/G mixer (top) and with 10 dB of attenuation (bottom). All traces are max-hold.

As the plots of Figure 3 demonstrate, the maximum degradation (the difference between the undisturbed noise floor and that with the SALO connected to the T/G mixer) experienced by the S/A is easily observed and is almost -30 dBc between 1.3 and 1.6 GHz. Next, a 10 dB attenuator was added between the isolator and mixer. The addition of just 10 dB of attenuation reduces the interaction to about -45 dBc. This suggests that the addition of over 30 dB of isolation should decrease the interaction such that nearly all of the latent S/A dynamic range is available. Taming this effect may also make it possible to do most non-T/G measurements with the T/G connected and turned on.

To address the fine-tuning requirements that are expected at narrow RBW/VBW S/A combinations, a PLL synthesizer makes a good choice for the TGLO source. After some research, the ADF4351 was chosen. This PLL has an integrated VCO and the device offers a wide range of output frequencies: 35 to 4400 MHz. The ADF-4351 is an I/M/N device with a programmable reference divider, so it offers a great degree of freedom for the TGLO frequency selection.

To help control the PLL frequency stability, the 8566B 10MHz output is used as the REF input to the PLL board. This allows the PLL to lock to the same reference that the S/A uses, which greatly reduces many of the drift issues that can be present with a separate, un-correlated, REF source.

The Feature Creep

Originally, only a single band was desired for the TG. However, based on the suggestions of Kerry Wong's attempt [2], it was decided that the 2^{nd} 8566B band might also be supported by this effort. This requires a slightly different TGLO setup since the 8566B's 2^{nd} band uses an IF of 321.4 MHz. The 2^{nd} band of the 8566B covers 2.0 to 5.8 GHz using the fundamental of the SALO. This means that the SALO covers 2.3214 GHz – 6.1214 GHz and the TGLO needs to be 321.4 MHz for this band.

As for the subsequent bands, the 3^{rd} band of the 8566B covers 5.6 – 12.5 GHz using the 2^{nd} harmonic of the SALO and would require a TGLO of 2.6393 GHz. The 4^{th} band covers 12.3 – 18.8 GHz using the 3^{rd} harmonic and would require a TGLO of 8.0928667 GHz. Finally, the 5^{th} band covers 18.4 – 22.0 GHz using the 4^{th} harmonic and would require a TGLO of 13.71965 GHz. Of the harmonic bands, only the 2nd and 3rd could be directly supported by the ADF4351.

The decision as to whether to support the higher bands, and how many of them might be practically supported would impact the T/G design in a couple of ways. First, the choice of mixer and output amplifier is driven by the frequency range needed. Second, filtering of the T/G mixer output becomes more complicated. While filtering is likely not needed to support passive DUTs, it is desired that this T/G be able to support active DUTs to the greatest extent practical. To do this, filtering of the T/G mixer is needed to reduce spurious signals that might interact with an active device, and more band/SALO combinations = more filters needed, plus the overhead to switch between them.

Using fixed frequency filters represents a compromise between cost and performance. The use of relatively low-cost, fixed high-pass and low-pass filters limits the effectiveness of the filtering. In truth, a tunable bandpass filter is what should be employed here. However, cost is the main limiting factor. Surplus filters for this exact function are available, but the cost puts this option out of reach for the moment. However, the modular construction means that one could be added later with a minimal re-design effort.

The mixer also represents a cost (and availability) problem. Broad-band mixers are not generally available as new COTS components. By broad-band, I mean REALLY broad-band, with RF/LO ports that have useable response from 300 MHz up to 12,000 MHz (or more) and an IF port that is good from DC to 12,000 MHz. The reality of the mixer-conundrum is that availability is the bigger problem. Matching over such large bandwidths is difficult to accomplish. Even devices that have similar ranges in other respects, rarely go down below 1 GHz on the RF/LO ports. The mixer chosen for this project could be used in the 2nd band, if it were wired differently (using the IF port for the SALO port, more on this later). Still, this limits the TG to just 8 GHz (close, but not quite close enough, to the desired 12 GHz target).

The Derived Requirements

While the elements of the T/G could be reasonably cobbled together when needed, it was desired that this design be a permanent fixture attached to the 8566B. This requirement suggests a chassis to contain the various sub-systems. Since the 8566B can be rack-mounted, and this particular unit had the mounting ears to that effect, it was decided to use a 19" rack chassis profile for the T/G.

To keep cable lengths appropriately short, the proper location for the T/G chassis is under the 8566B. To facilitate this, a 19" rail structure (Figure 4) was added to the lift cart. The width of the cart was almost exactly the proper opening width, so the task was relatively simple and used standard angle and square-U channel aluminum stock. This rail structure allows the 8566B to be suspended just above the T/G chassis so that the appropriate connections can be conveniently located on the T/G front panel. Since the S/A is independently suspended, the T/G can also be easily removed if needed.

Figure 4. HP 8566B lit-cart rack-mount structure

When I first worked with this particular 8566B, I produced a GPIB to Serial interface computer using an SiLabs evaluation board set for their C8051F120 processor (were I to do this today, I'd likely use an ARM-M4 microcontroller). The GPIB interface computer was included with the 8566B, so it was available as a control point for the ADF-4351 PLL and other T/G features. The T/G chassis would thus be the perfect location for this peripheral.

The chassis itself was surplus from my repeater controller days and consisted of a 17"x13"x3" sheet metal aluminum chassis with a top-lid and a blank 19" rack panel that was given up from another surplus chassis I had lying around. While the area seemed ample at the beginning, it soon became clear that space would have to be managed carefully to make sure everything would fit with room for future enhancements.

A 12V/5V switching supply in need of a home provided the power source. The GPIB interface processor and power supply take up almost 1/3 of the floor space in the chassis. While the T/G RF subassemblies are relatively small, once attached with minimal length connections, they end up occupying more space than anticipated.

Another oft used tool in RF investigations is that of a signal generator. These devices produce a fixed frequency output at a particular signal level to stimulate a DUT at a single frequency. In order to support this future extension, and also to provide a user interface (UI) for other tool features, an LCD and keypad were added to the design. The software support of the LCD was provided in another tool that uses similar interface SW as the GPIB computer, so it wasn't a difficult SW task to add this feature. However, the GPIB computer's I/O was not efficiently mapped when it was created, so this complicated the integration of the new I/O devices.

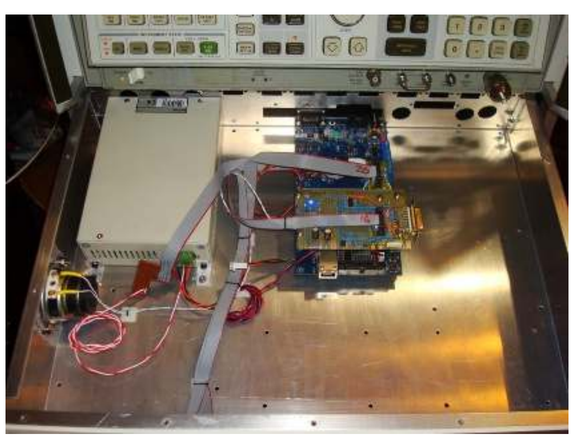

Figure 5. T/G Chassis with power supply and GPIB controller installed

A sloped sheet-metal enclosure was fashioned to support the LCD and keypad to provide an ergonomic presentation. The flat model was designed in AutoCAD and transferred to a piece of 40mil scrap aluminum. Once the flat shape was cut, a makeshift brake was used to bend the flat sheet into the final shape. Press-fit nuts were applied to allow the enclosure to be easily attached to the T/G front panel. The result is somewhat rough around the edges, but not half bad given the tools available.

Figure 6. T/G front panel (minus RF connectors). The SALO output can be seen just to the left of the main S/A input

With these elements in place, the platform for the T/G was well established. All that remained was to add the RF hardware.

The RF Hardware

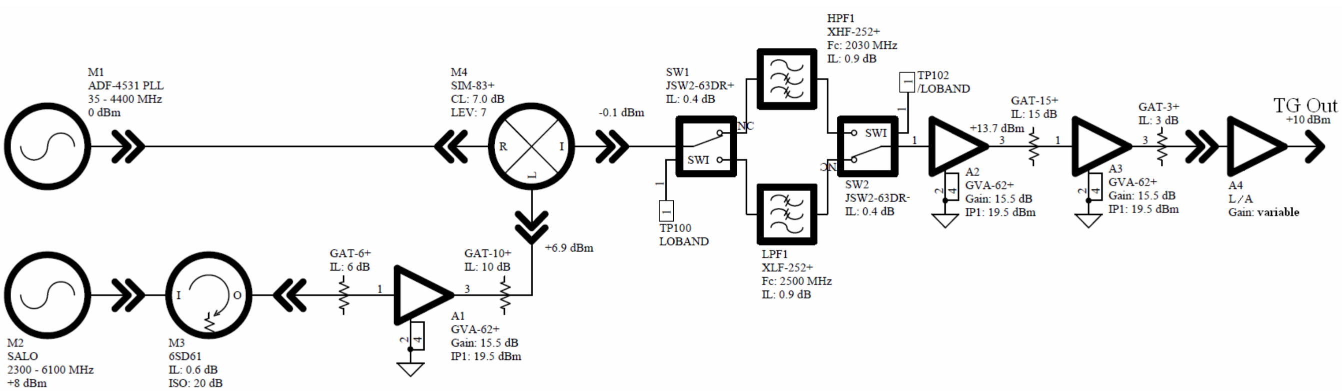

Figure 7. Block Diagram for the T/G signal chain

Figure 7 illustrates a block diagram of the major components of the T/G RF circuit. Not shown are the bias circuits and inter-stage AC coupling capacitors. The mixer was constructed on a small piece of Rogers-Duroid material with edge-mount SMA connectors for each of the ports. The isolator (M3) is a self-contained unit with SMA connectors and the ISO-AMP (A1) was a custom PCB constructed on 0.062” FR-4 that was included with another PCB design. Similarly, the filters and final amp stages were constructed using another custom PCB that was part of the same panel as the ISO-AMP PCB.

The isolator, M3, is intended to connect directly to the SALO output connector using a M-M, right-angle adapter. A short piece of RG-402 with right-angle connectors routes the SALO cable into the T/G chassis via a feed-through hole in the front panel. A bulkhead connector would be a better aesthetic approach, but this method is simpler, cheaper, and has a bit less loss than if there was an additional set of connectors in the coax connection. In addition, this method provides mechanical stress relief that would otherwise be difficult to achieve for a bulkhead connector setup.

Once inside the chassis, the SALO connection attaches to the ISO-AMP board, which is connected to the mixer board. The PLL connects to the opposite side of the mixer board and the IF output of the mixer connects to the T/G Final amplifier board. Power and the filter band select signal complete the RF hardware connections.

A control signal from the GPIB computer drives SW1/SW2 to effect the desired filter setting. The low-pass filter is used for the 1^{st} band (100 Hz to 2.5 GHz), while the high-pass filter is used for all other bands.

The ISO-AMP features input and output pads and a GVA-62+ amplifier with a nominal 20 dB of isolation from output to input. The pads help limit the signal level (no amplification is needed at this point in the signal chain, isolation is the primary requirement for this stage). While it would be best to place all of the attenuation at the output of the amplifier (this improves the noise figure of the stage) the input attenuator is needed to keep the amplifier well under its IP1 level. Doing so reduces the mixer spur levels.

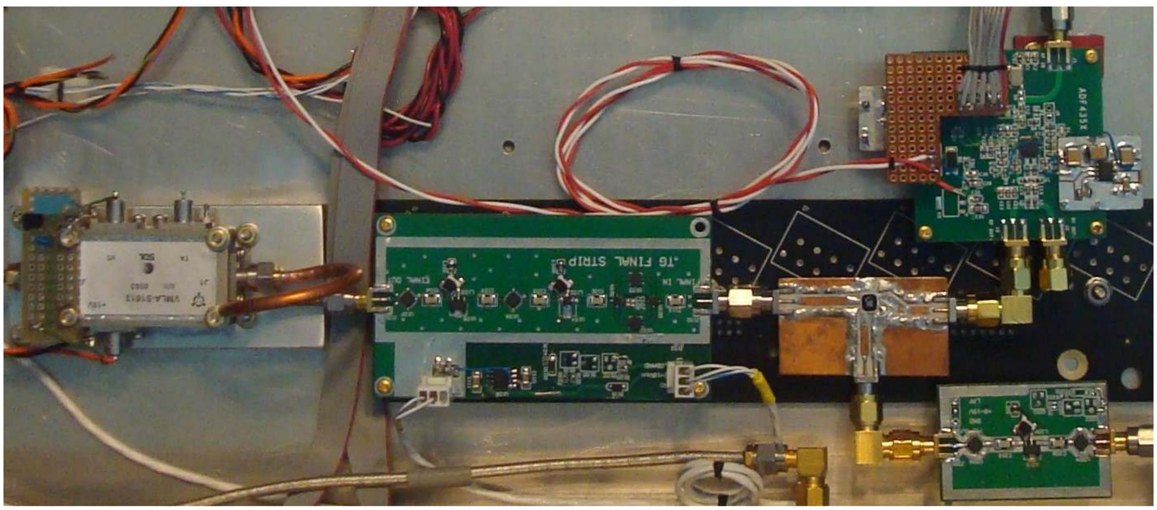

Figure 8. T/G RF subassemblies. From the upper right clockwise: TGLO source, ISO-AMP, MIXER, Final strip, and Levelling Amp.

Figure 8 illustrates the final placement of the RF sub-assemblies. This arrangement places the ISO-AMP input right next to the access hole for the coax connection to the SALO isolator, which will be located outside the T/G chassis. The T/G output connector is located at the bottom-right-of-center of this image. This layout provides sufficient room for the RF subassemblies and places them at or reasonably close to the requisite front panel connector locations.

The levelling amp (L/A) is a modestly priced item purchased off of e-bay that is intended to produce a constant level output given a widely varying input. This greatly improves the behavior of the T/G by producing a nearly flat output across the T/G sweep range (give or take a dB). A heat-spreader and heatsink (the chassis) are needed for this part. With the thermal configuration as shown, a case temperature rise of only about 6° C is observed for the L/A.

Future multi-band support will require a switched mixer arrangement (assuming that a broad-band mixer is not found instead). This will likely be accomplished by fabricating a multi-level mixer assembly with RF switchgear to select the mixer appropriate for the desired band. The multi-level assembly will increase the assembly complexity, but will allow the upgrade mixer to occupy the same footprint as the existing mixer.

Performance

The “finished” T/G shows relatively good performance. Figure 9 illustrates a normalized SNA plot of a MiniCircuits PLP-1000+ LPF for two different DUT configurations. This figure shows that the SNA dynamic range increases with greater attenuation applied to

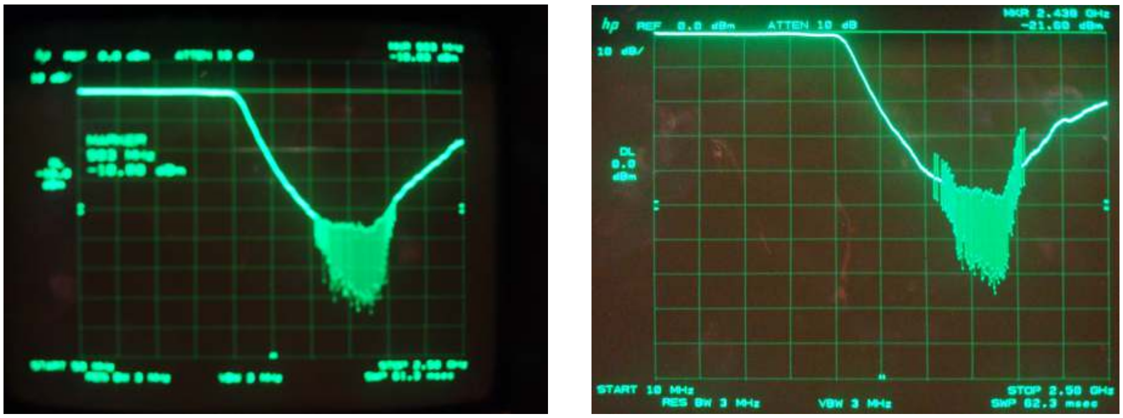

(a) (b)

Figure 9. Normalized SNA plot of a MiniCircuits PLP-1000+ LPF with (a) 16dB of attenuation in the DUT chain, and (b) 10 dB of attenuation.

The T/G output. Reducing the attenuation provides more signal and this would imply greater dynamic range. However, what the data of Figure 9 suggests is that this is contrary to the observed result: greater attenuation appears to provide greater dynamic range. The attenuation also helps present a better broad-band 50Ω match to the DUT-S/A input chain which makes it easier to de-embed the DUT setup with normalization.

However, one can only increase the attenuation to a certain point before the undisturbed noise floor of the S/A begins to dominate. Thus, further improvement would have to be achieved by increased isolation to the T/G mixer. The open circuits and close proximity present in this setup make it possible for signals to couple into the sensitive SALO branch of the TG mixer chain. Isolation by shielding offers a decent chance of improving D/R and is best applied to the ISO-AMP and mixer modules. Each module was built with provisions for shielding and application of shielding materials is easily accomplished.

Conclusion

With a DR of 50 dB or so, the S/A-T/G combination accomplished here is not perfect, but this represents a useable tool with some room for improvement. While still a work in progress, the T/G produces good performance and has proven very useful for examining filter and insertion loss features of several sub-projects.

Bibliography:

1) "A 2.5-GHz Tracking Generator for the HP 8566A/B Spectrum Analyzer", 2004, John Miles, KE5FX, http://www.ke5fx.com/tg8566.html

2) "Building a Tracking Generator", January 10, 2016, Kerry D. Wong, http://www.kerrywong.com/2016/01/10/building-a-tracking-generator/

3) "Troubleshooting and Repair Manual, Volume 1, HP 85660B Spectrum Analyzer RF Section – HP Part No. 85660-90210", December 1990, Hewlett-Packard, Co.

4) "Connecting the HP 5086-7023 YTO", 3/14/2016, Joseph Haas, http://www.rollanet.org/~joeh/projects/Connecting_the_HP_5086-7023_YTO.pdf

