An IF Module for Wide-Bandwidth Signals

The HP 70911A IF module provides the HP 71910A receiver with wideband demodulation and variable bandwidths up to 100 MHz, while maintaining the gain accuracy of a spectrum analyzer.

by Robert J. Armantrout, Terrence R. Noe, Christopher E. Stewart, and Leonard M. Weber

The HP 70911A IF module provides much of the new functionality present in the HP 71910A microwave receiver. From the start, the primary design goal of the HP 70911A was to overcome the 3-MHz IF bandwidth limitation of existing Modular Measurement System (MMS) spectrum analyzers. At a minimum, we wanted a tenfold increase in bandwidth, but really hoped to achieve 100 MHz. Although bandwidth was the major design focus, several other goals were also important, including:

- Accurate gain
- Variable bandwidths
- Pulse detection
- Direct connection to demodulators
- FM demodulation
- I-Q down-conversion.

Of these goals only the first three are usually considered in spectrum analyzer IF design. The others were based on the need to better address the more complex signals employed in modern communication and radar systems.

Given the range of bandwidths required, previous spectrum analyzer IF design work has concentrated on center frequencies of 3 or 21.4 MHz. Obtaining the accuracy and stability of both gain and bandwidth required even at these IF frequencies has always been challenging. While there are a number of well-understood design alternatives and approaches available for 21.4-MHz and 3-MHz IFs, they did not exist for the 321.4-MHz center frequency chosen for the HP 70911A. Because of this some degradation of accuracy and stability was anticipated, and the design team was anxious to minimize any such degradation.

Fortunately, advances in both components and fabrication techniques were underway that were applicable to the needs of the project. The increasing availability of wide-bandwidth RF components in surface mount packages and the growing internal repertoire of surface mount manufacturing expertise suggested that the performance goals could be achieved without the need for internal microcircuit developments.

The resulting design makes extensive use of surface mount technology to achieve the goal of 100-MHz bandwidth at the 321.4-MHz center frequency while maintaining the excellent gain accuracy and stability expected of spectrum analyzers. In addition, optional down-conversion and demodulation features extend the utility for wide-bandwidth signals with complex modulations.

Fig. 1 shows the major internal functional blocks that make up the HP 70911A. A detailed discussion of the design considerations for these blocks is given below. Note that the module is partitioned into standard and option sections. An option cardcage, similar to that offered in the HP 859xE Series spectrum analyzers, provides a standard interface for all options.

Variable-Bandwidth Design

The following discussion is divided into three parts. The first part gives some background about the design of variable-bandwidth filters. The second part describes an alternative design that was considered and proven for 1-MHz-to-10-MHz bandwidths, but not included in the final product release. The final part discusses the design of the 10-MHz-to-100-MHz bandwidths of the HP 70911A.

Background. To serve as background material for describing variable-bandwidth filter design, the design approach used in the HP 70903A IF module is described here. The HP 70903A was the predecessor of the HP 70911A and used the synchronously tuned class of filters.

Synchronously tuned filters consist of several poles with the same center frequency and Q* with buffering between the stages. There are several advantages to using this particular topology, foremost being the excellent pulse response of these filters. This response allows for fast sweep speeds on a spectrum analyzer. Since we are trying to create a continuously variable bandwidth over a large adjustment range, it is also important to have a filter that can be easily adjusted. Synchronously tuned filters are easy to tune and are tolerant of a slight misalignment in different stages. Also, unlike other bandpass topologies, the Q of each stage is less than the final required filter Q.

To make these stages variable-bandwidth, a series resistance is added to reduce the Q of each of the individual stages. The individual stages look like the circuit in Fig. 2. The bandwidth of this circuit is given by the following equation:

$$BW_{stage} = (R_s + R_p)/(2\pi C R_s R_p)$$

* Here Q is filter quality factor, not quadrature as in I-Q modulation.

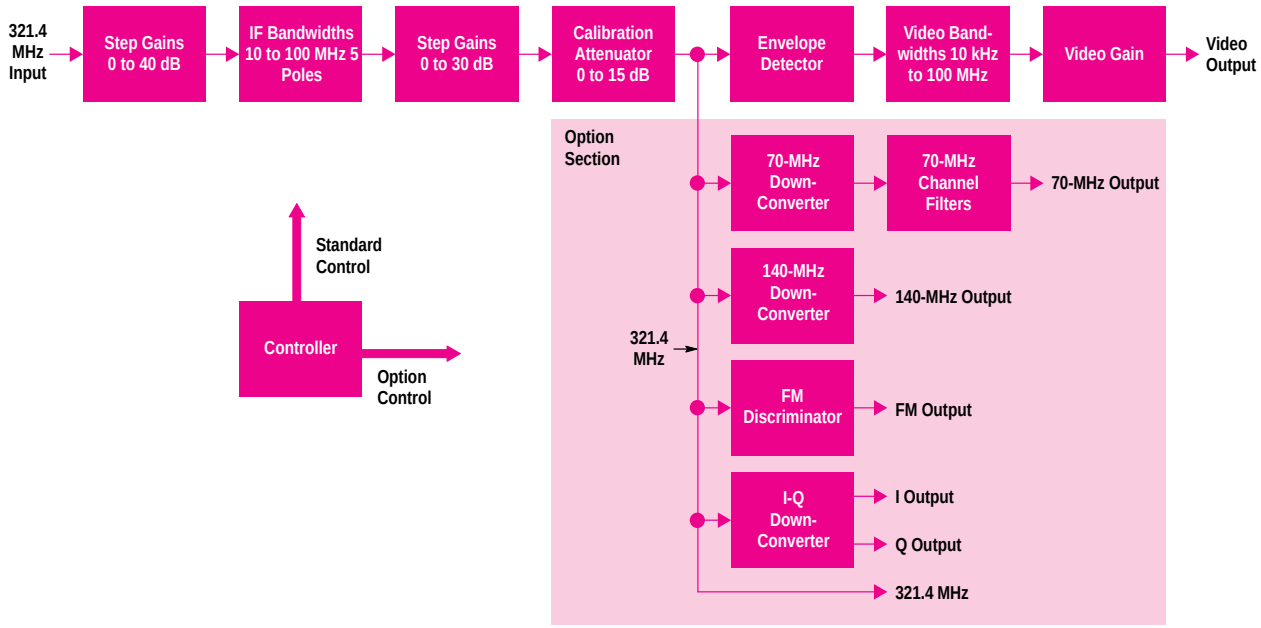

Fig. 1. Block diagram of the HP 70911A IF module.

where R_p is the equivalent parallel resistance across the tank circuit and R_s is the series Q-reducing resistance. By adjusting R_s, the bandwidth can be adjusted continuously. R_p is the combination of the input impedance of the buffer stage and the equivalent parallel resistance of the tank circuit.

By cascading several of these individual stages, a synchronously tuned filter with the desired bandwidth can be created. The equation for the bandwidth of an n-stage synchronously tuned filter is:

$$BW_{total} = BW_{stage}\sqrt{\left(2^{\frac{1}{n}} - 1\right)}$$

The typical HP spectrum analyzer has four or five stages in a synchronously tuned filter, which results in individual stage bandwidths of 2.3 to 2.6 times the overall filter bandwidth.

To implement a continuously variable synchronously tuned filter, the series resistance is created by using p-i-n diodes as variable resistors. The p-i-n diodes used are optimized as current-controlled RF resistors. The RF resistance varies with forward bias current according to the following relationship:

$$R = aI^{(-b)}$$

where a and b are constants and I is the forward bias current in the diode.

This resistance characteristic holds for frequencies above the low-frequency limit, which is set by the minority carrier lifetime of the p-i-n device. Below that frequency the devices behave like ordinary p-n junction devices and rectify the signal. This results in distortion effects that can limit the dynamic range of the filter. The recommended operating frequency is ten times the low-frequency limit, which is given by the following equation:

$$f_{min} = \frac{1}{2}\pi\tau$$

where τ is the carrier lifetime. To minimize the distortion effects from rectification, often several p-i-n diodes are used in series to minimize the signal voltage across each individual diode (see Fig. 3).

This topology depends on a low impedance driving the p-i-n diodes and a high impedance buffering the tank circuit. Typically an FET buffer amplifier is used as the amplifier at the output of each stage because of its high input impedance. Care must be taken in the design of this amplifier to avoid distortion problems caused by the large signal voltage across the tank circuit. Keeping the nonlinear junction capacitance of the FET buffer amplifier small compared to

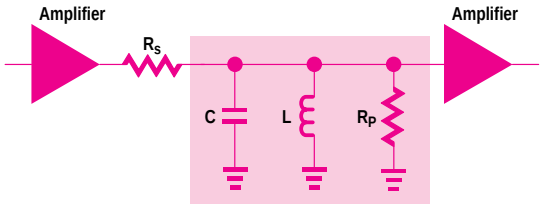

Fig. 2. RLC tank circuit with a series resistance (R_s) for adjustment. This circuit represents one stage of a synchronously tuned filter.

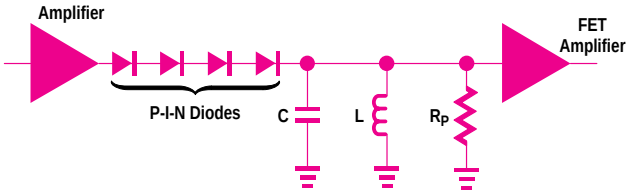

Fig. 3. RLC tank circuit with p-i-n diodes in place of a series resistance.

the capacitance of the overall tank circuit minimizes these distortion effects.

Adjusting the current in the p-i-n diodes can provide continuously variable bandwidths over a large range. Usually a digital-to-analog converter (DAC) is used to control the current in the p-i-n diodes and allow for setting different bandwidths.

This method of varying the bandwidth of the filters works very well with one slight problem. The series resistance in combination with the parallel resistance across the tank circuit creates a voltage divider. Varying the series impedance into the tank circuit not only changes the filter bandwidth, but also changes the loss through the filter as well. This amplitude change is an undesirable side effect. Several methods have been used to compensate for this change in amplitude.

One of the methods that has been patented by Hewlett-Packard uses feed-forward compensation (see Fig. 4). This method has several advantages over previous schemes that rely on feedback for amplitude compensation. The idea is to sum the proper signal at the output node to offset the drop across the series resistance element.

By summing a properly scaled version of the input signal back into the output node with a compensation resistor R_c, the voltage drop across Rs can be canceled. By setting $K = 1+R_c/R_p$, the voltage at the output node is always equal to V_{in}, independent of R_s.

Since R_p is determined by the Q of the tank circuit and the input impedance of the FET buffer amplifier, it does not vary with bandwidth. Thus R_c can be adjusted for each pole of the filter to compensate for amplitude variations. Variations in R_p over temperature can be compensated by using a thermistor in the R_c circuit to cancel their effect.

Summing a scaled version of V_{in} into the output node without introducing significant amounts of noise and distortion is accomplished in some HP IF circuits with a transformer circuit. By adding a primary winding to the tank inductor a transformer is created with a one-to-four turns ratio (Fig. 5). This sets the value of K to be four and determines the value of R_c for a given R_p as: $R_c = 3R_p$. Using a transformer with a one-to-four turns ratio yields an impedance transformation of 1 to 16. Thus, a resistor on the primary side of the transformer looks like 16 times the resistance from the secondary side. Feeding the primary side of the circuit from V_{in} through a compensation resistor requires a resistance of:

$$R_c = 3R_p/16.$$

Cascading several of these stages together implements a synchronously tuned filter that has a continuously variable bandwidth and no change in amplitude.

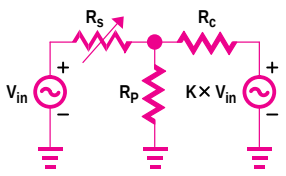

Fig. 4. Topology for a feed-forward amplitude compensation circuit.

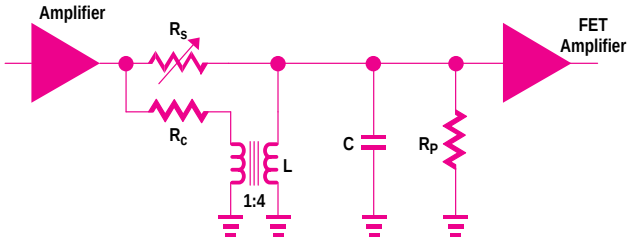

Fig. 5. Feed-forward amplitude compensated RLC tank circuit.

The HP 70903A uses four of the stages shown in Fig. 5 to implement bandpass filters with bandwidths adjustable from 100 kHz to 3 MHz at a center frequency of 21.4 MHz.

Design for 1-MHz-to-10-MHz Bandwidths. In the HP 70911A IF module we originally wanted to have continuously variable bandwidths down to 1 MHz at a center frequency of 321.4 MHz. This required an overall Q of 321.4. Even with several cascaded stages in a synchronously tuned configuration, the individual poles still required a loaded Q greater than 120. To achieve a loaded Q this high requires a device that behaves as a resonant circuit with a much higher unloaded Q. At a center frequency of 321.4 MHz there are very few choices of resonators that can achieve a Q this high. Given the size constraints of fitting on a PC board inside an MMS module, the possible solutions to this design problem were limited.

Some of the traditional choices for high-Q resonators in this frequency range include helical resonators and transmission line resonators. The size of either of these choices was the biggest obstacle to their use in the HP 70911A module. A new resonator technology was found that met all of the constraints. This resonator is a quarter-wavelength shorted coaxial transmission line formed from a high-dielectric-constant ceramic material. The dielectric constant of the ceramic is approximately 90.5, which yields a length of less than 1 inch at 321.4 MHz for a quarter-wavelength resonator. The coaxial resonators are formed with a square outer conductor 0.238 inch on a side and a circular inner conductor of 0.095-inch diameter. These dimensions are small enough to mount four of these resonators on a single printed circuit board with the appropriate circuitry to create a four-pole synchronously tuned filter. The unloaded Q of these ceramic coaxial resonators at 321.4 MHz is around 220.

A shorted transmission line (T_L) behaves like a parallel RLC resonant circuit at a center frequency corresponding to a quarter wavelength of line. An equivalent RLC lumped-element model for this circuit can be calculated by matching the slope of the reactance change with the frequency of the transmission line circuit at resonance to an equivalent RLC circuit (Fig. 6). The equivalent parallel resistance can be calculated from the Q of the resonator.

To implement a synchronously tuned filter all of the stages need to be aligned to exactly the same center frequency. By adding an adjustable capacitance in parallel with the shorted transmission line the stages can be pulled into alignment with the center frequency. This requires that the resonant frequency of the resonator be higher than the final required center frequency because the added parallel capacitance will lower the resonant frequency.

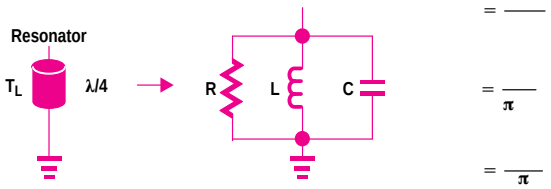

Fig. 6. Equivalent circuit for a ceramic resonator.

The resonator chosen for the HP 70911A investigation was cut to a length that corresponded to approximately 360 MHz so that it could be pulled into alignment at 321.4 MHz. Using varactor diodes for the parallel capacitance allows the alignment of all of the center frequencies using a DAC under automated computer control.

For a square transmission line with a round center conductor the characteristic impedance of the line can be approximated by the following formula:[1]

$$Z_0 \cong \frac{60}{\sqrt{\epsilon_r}} \ln\left[1.079\frac{w}{d}\right] \text{ohms}$$

where w is the width of the square transmission line, d is the diameter of the coaxial element center conductor, and ϵ_r is the relative permittivity of the dielectric.

From the dimensions given above for coaxial resonators, Z_0 is calculated to be approximately 6.3 ohms. Using the formulas given for R, L, and C in Fig. 6, the equivalent circuit of the resonator looks like Fig. 7.

To implement a four-pole synchronously tuned filter, the final Q of each stage needs to be 140 to meet the final desired bandwidth of 1 MHz. This implies a total parallel equivalent resistance of 1004 ohms. Since the resonator parallel resistance is only 1765 ohms, the total impedance of the circuit that buffers each stage must be greater than 2327 ohms. It is a challenging design task to generate a buffer stage with that high an impedance at a frequency of 321.4 MHz. To attain a maximum bandwidth of 10 MHz the equivalent parallel resistance needs to be 100.4 ohms.

The circuit topology used for the 10-MHz to 100-MHz bandwidths, which is discussed in the next section, worked well at the lower Q levels, but was unable to provide the high impedance necessary for the minimum bandwidth of 1 MHz. To attain the high impedance needed, a GaAs FET buffer stage is used across the resonator (see Fig. 8). The driver stage is a common-base configuration so the output impedance level can be set high enough to be stepped up by a tapped-capacitor transformer circuit, which is similar to the 10-MHz-to-100-MHz bandwidth circuit. The varactor diodes used to vary the capacitive taps have a tuning range of approximately 10 to 1.

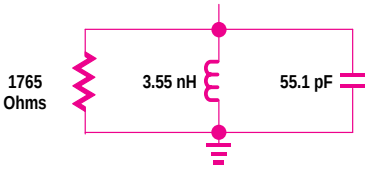

Fig. 7. Circuit values for the ceramic resonator's equivalent circuit.

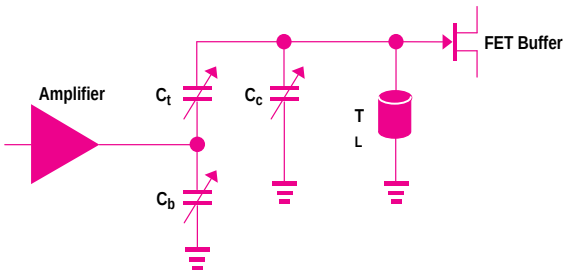

Fig. 8. Resonator with circuitry for bandwidth and center frequency tuning.

For a tapped-capacitor transformer the effective turns ratio is given by: $N = C_b/C_t + 1$. The impedance ratio varies with N^2. This impedance ratio provides the required bandwidth range but there is a drawback. The tapped-capacitor transformer also steps up the signal voltage at the input of the FET amplifier. This leads to distortion problems. The solution was to step the voltage back down with a fixed-ratio tapped-capacitor transformer (see Fig. 9). This keeps the voltage at the FET down to a level that keeps the distortion within allowable limits.

Varying C_t and C_b can set the desired bandwidth from 2.3 MHz to 23 MHz for each pole. C_c is used to adjust the center frequency to 321.4 MHz for each pole. Since the effective capacitance across the resonator changes as the tap capacitors are varied, the center frequency needs to be readjusted as the bandwidth is varied. This is accomplished with varactor diodes driven by DACs and a lookup table containing the appropriate voltage settings for each bandwidth in 10% increments over the entire range of bandwidths. Cascading four of these stages as a synchronously tuned bandpass filter yields an overall bandwidth of 1 MHz to 10 MHz.

Design for 10-MHz-to-100-MHz Bandwidths. The dynamic range limitations of the resolution bandwidth filter design approaches described above meant that they would not work for the HP 70911A. A different approach was needed. A synchronous type of filter was still desired because synchronous filters have low group delay variation. This is a requirement for good pulse fidelity, which was one of the goals for the HP 70911A. A five-resonator synchronous filter was chosen for the shape factor requirements and the range of bandwidth desired. These are two conflicting requirements because, unlike other filter types, increasing the number of resonators in a synchronous filter decreases the required Q of the individual resonators. For the required maximum

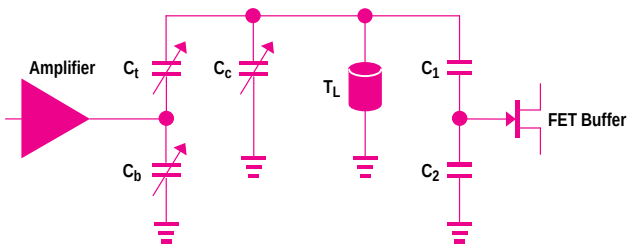

Fig. 9. 1-MHz-to-10-MHz bandwidth filter stage.

bandwidth of 100 MHz at a resonant frequency of 321.4 MHz, the fractional bandwidth of the composite filter is over 30%, and with five resonators, each tank would have a fractional bandwidth of over 80% of its center frequency.

A variable resonator with low insertion loss and low distortion was needed. Existing variable-bandwidth filters changed the Q of the resonator by varying its load resistance. For a five-resonator synchronous filter:

$$Q_{section} = Q_{overall} \times \sqrt{2^{1/5} - 1}$$

For the 10-MHz bandwidth,

$$Q_{section} = (321.4/10) \times \sqrt{2^{1/5} - 1} = 12.39$$

and for the 100-MHz bandwidth,

$$Q_{section} = (321.4/100) \times \sqrt{2^{1/5} - 1} = 1.24$$

For a single resonator, the bandwidth would be 26 MHz for the composite filter to have a 10-MHz bandwidth and 260 MHz for a 100-MHz setting. That means that a parallel resonator with an impedance of about 35 ohms at resonance would need to see a parallel resistance of between 45 and 450 ohms.

One of the ways the Q was changed in previous variable-bandwidth filters was to change the loading on the resonator with p-i-n diodes. A current source drove a series of p-i-n diodes connected to the top node of the resonator, which was connected to a high-impedance amplifier.

This is a good solution since p-i-n diodes act like inexpensive electronically controllable RF resistors. Distortion in p-i-n diodes can be reduced by putting a lot of them in series and using the same bias current. This method was tried but there was a problem. For the narrow bandwidths, a large RF voltage is present at the top node of the resonator. When this voltage is applied to the gate of a FET or the base of a bipolar junction transistor, the junction capacitance is varied by the RF voltage, causing distortion. At 21.4-MHz or 3-MHz center frequencies where this scheme has been used, the change in impedance because of this parasitic varactor is not significant. At 321.4 MHz the degradation in the third-order intercept is too great given the aggressive goals of the HP 70911A.

It seemed wise at 321.4 MHz to avoid high impedances, high RF voltages, and noise-figure-degrading p-i-n diodes. Transforming our characteristic impedance of 50 ohms up and then down using reactive transformations would allow us to avoid high-impedance amplifiers and p-i-n diodes. A capacitive transformer could be implemented with varactors to give us the desired continuous bandwidth variation. However, reference texts suggest that capacitive transformers should be used in cases where the resonators are only operated up to 20% bandwidth. In the HP 70911A, the resonators need to operate up to 81% bandwidth. It seemed like there was little hope of getting this scheme to work, but it was tried anyway.

With this topology the only place that there would be high RF voltages is at the top node of the resonator. Since there were going to be varactors at that node, there was concern about distortion. This was solved by putting the varactor diodes in a back-to-back configuration so that there would

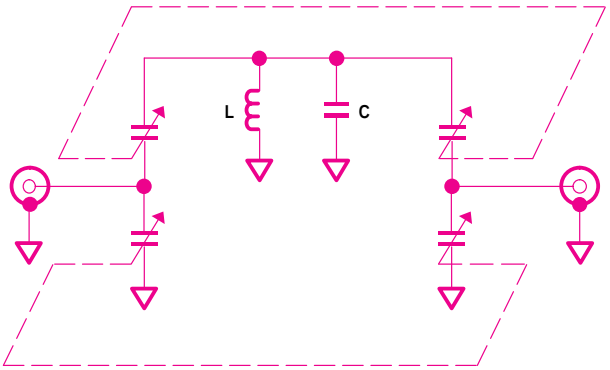

Fig. 10. Resonator for 10-MHz-to-100-MHz bandwidth. The variable capacitors are varactor diodes.

be some cancellation of the effect of the RF voltage (see Fig. 10). In this circuit, when the upper varactor increases in capacitance because of a positive swing of the RF voltage, the lower varactor decreases its capacitance, canceling out the change. Thus, the distortion problem was minimized.

The main effect of trying for over 80% bandwidth with capacitive taps is a nonideal filter shape (Fig. 11). At the wider bandwidth settings the upper tap capacitors are much larger than the lower tap capacitors. The circuit resembles a high-pass filter and doesn't have the ideal resonator rejection above resonance. This can be compensated by adding series inductors that will resonate with the upper tap capacitors (Fig. 12). The bandwidth of these outer resonators is high enough for the maximum bandwidth desired. As the main resonator bandwidth is decreased the outer resonator shifts up in frequency because the upper tap capacitance decreases. This shift does not cause trouble since the outer resonator has a bandwidth that is high regardless of the tap setting because of its 50-ohm loading on one port and variable loading on the other port.

The main resonator impedance was chosen to be 35 ohms at resonance so that for the widest bandwidths the Q-reducing resistance required was greater than 25 ohms (50 ohms at the input in parallel with 50 ohms at the output). Once that

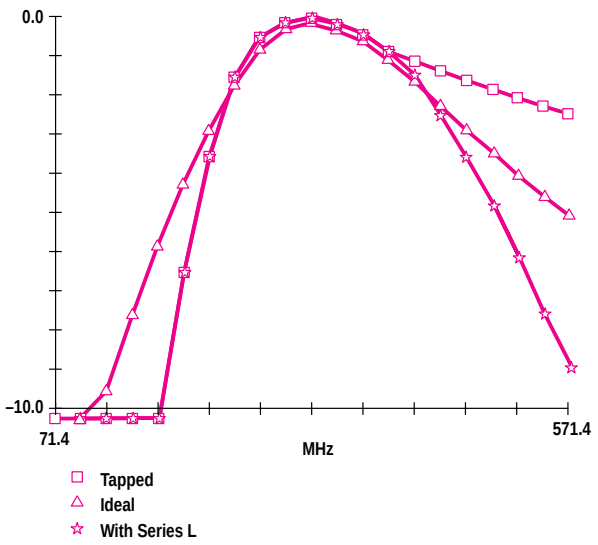

Fig. 11. The nonideal filter shape that results from using capacitance to achieve over 80% bandwidth.

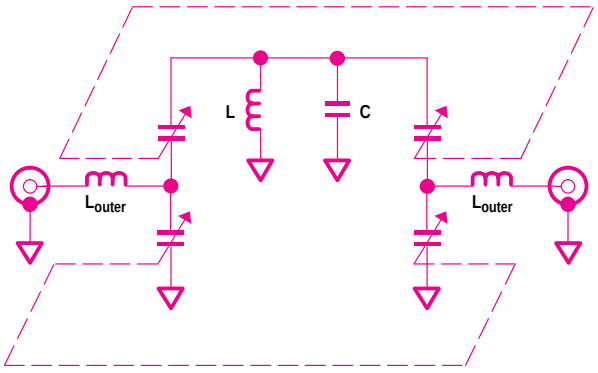

Fig. 12. Resonator compensated by adding series inductors that will resonate with the upper tap capacitors.

was decided, the values for L and C were easy to calculate. One of the complications of using the tapped capacitors is that the equivalent capacitance in shunt with the tank inductor changes with the bandwidth. This problem is solved by using DACs to control the voltages of all the varactors. A lot of calibration ROM space is required to support this circuit topology. All five resonator circuits have the lower and upper tap varactors ganged together (see Fig. 13). The shunt tank capacitors are connected to separate DAC outputs allowing independent control of the center frequency of each resonator.

The resonator circuit shown in Fig. 13 is used in the HP 70911A. The insertion loss for this circuit is less than 6 dB for the 26-MHz setting and about 1 dB at 260-MHz bandwidth. The third-order intercept point is about +29 dBm referred to the output for all settings. Group delay variation is less than half a nanosecond in wide mode and about 3 ns for the narrow-bandwidth setting.

Accurate Gain

The gain accuracy of the HP 70911A IF module depends on the gain of the seven step gains and the five filter poles and the accuracy of the calibration attenuator. How gain accuracy is achieved in each of these elements is discussed below.

Calibration Attenuator. The calibration attenuator is used during self-calibration of the HP 71910A receiver. The customer performs receiver self-calibration periodically to ensure that the receiver meets all of its specifications. This procedure measures and corrects several aspects of receiver performance. Among other things, it measures the gain of the step gain and attenuator stages and measures and corrects displayed linearity errors in the linear detector.

Since the calibration attenuator is used as a reference standard against which other parts of the receiver are measured, it is essential that the attenuator yield accurate and stable gain over the receiver's specified 0-to-55°C operating temperature range. Over this range, and over the attenuator's 0-dB-to-13-dB attenuation range, accuracy is guaranteed within 0.3 dB at 321.4 MHz.

At these frequencies, variable attenuators are traditionally designed using semiconductors with bias dependent resistivity. Examples would be p-i-n diodes with a current dependent resistance or GaAs FETs with a resistance that depends on gate voltage. Unfortunately, these types of attenuators do not demonstrate the required temperature stability. For this reason, the calibration attenuator was designed as a series of

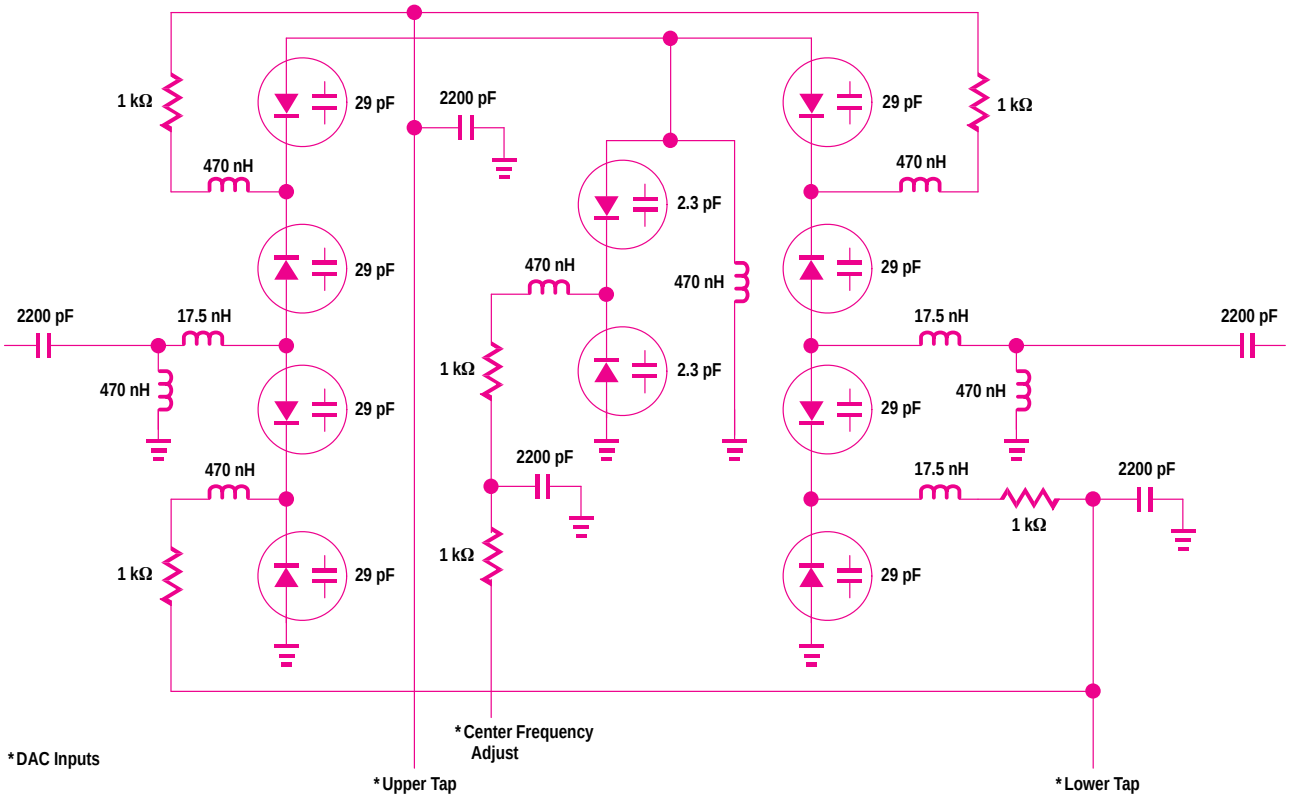

Fig. 13. Circuitry for one pole of the IF bandwidth filter showing the upper and lower tap varactors ganged together.

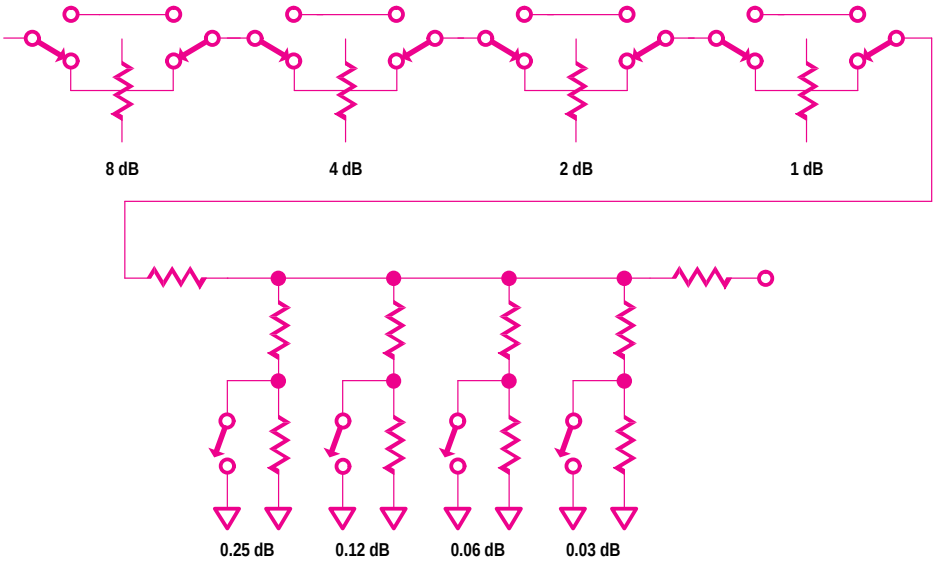

Fig. 14. The calibration attenuator is designed as a series of switchable attenuator sections.

fixed switchable attenuator sections (Fig. 14). The 1-dB through 8-dB attenuator stages are pi attenuators made with surface mount thick-film resistors. The 0.25-dB through 0.03-dB attenuator stages could not be designed as pi attenuators because the resistance values required for these very low attenuation values would not be achievable at 321.4 MHz.

Instead of trying to figure out a way to build a 0.03-dB attenuator, we built a 6-dB tee attenuator with an attenuation we could vary slightly. This was done by changing the resistance of the shunt element of the 6-dB attenuator. By switching around small resistors in series with much larger ones, very small attenuation steps can be realized. Changing only the shunt element in this attenuator does cause the attenuator's return loss to vary across its 0.5-dB attenuation range, but this effect is small enough to be acceptable.

With standard 1% tolerance resistors, the attenuation accuracy of this circuit will not be exact enough without alignment. During alignment of the HP 70911A, each 1-dB calibration attenuator step is measured and corrected to the desired value by turning on the appropriate combination of small attenuator steps. This alignment data is then stored in EEPROM.

Step Gains. The purpose of step gains is to substitute a known fixed gain ahead of the detector to enable accurate measurement of low-level signals. The ideal step gain has a 0-dB gain state and a 10-dB gain state. The implementation in the HP 70911A is shown in Fig. 15. The 0-dB (bypass) path actually has approximately 2 dB of loss, while the 10-dB (gain) path has approximately an 8-dB gain. The goal of the circuit is to make the gain difference between the 0-dB and 10-dB states exactly 10 dB. The variable attenuator in the gain path allows the gain to be trimmed to achieve this accurate gain difference. During alignment the DAC values required to trim the gain are determined for each of the step gains from measurements made at 0, 25, and 55°C. These DAC values are stored in EEPROM tables which are consulted by the module firmware during operation. As mentioned above, the calibration attenuator is used during calibration to measure the actual gain step value. In addition, because the calibration attenuator is accurate to within 0.3 dB, it can be used in conjunction with the step gain to provide accurate 1-dB gain steps over most of the 70-dB gain range.

Filter Pole Gain. As discussed above, bandwidth variation is obtained with a controlled variation of the Q of the filter pole. Because of this, the gain of the filter pole also varies with bandwidth. It is necessary to compensate for this gain variation if the module gain is to be accurate for all bandwidths. Since bandwidths are in 10% steps (10, 11, 12.1, ...), there are a finite number of bandwidths for which gain compensation is required. Associated with each filter pole is a programmable gain block (see Fig. 16). This gain block is used to provide the necessary gain compensation. The DAC values for this compensation are determined during alignment and stored in EEPROM tables, which are consulted by the firmware each time the bandwidth is changed.

In addition to controlling the nominal gain of the filter pole these programmable gain blocks also play a role in temperature compensation of the overall gain of the module. Gain drift with temperature is most troublesome during warmup. For this reason, the temperature of the module is monitored during warmup and the temperature value is used to adjust the gain to keep the output levels relatively constant. The warmup period is defined as the first hour after the module

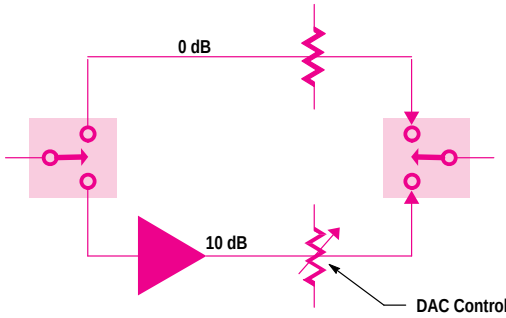

Fig. 15. A block diagram of the step gains in the HP 70911A.

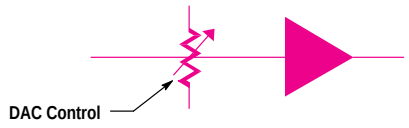

Fig. 16. A representation of a programmable gain block.

is powered on. During this period the temperature is measured once per minute and the rate of change is used to determine the size of the gain adjustment required. After the warmup period, the gain is stable for small changes in temperature so this compensation mechanism is disabled.

The module firmware orchestrates gain changes based on sampling a temperature sensor voltage with an ADC. The ADC values are used to calculate the gain change based on the following equation:

$$\frac{V_t - V_n \; G_p}{V_{55} - V_n}$$

where V_t is the voltage representation for the current temperature, V_n and V_{55} represent the voltage values for 25 and 55°C respectively, and G_p is the peak gain change between 25 and 55°C for each bandwidth. G_p is determined during alignment.

The gain change calculated is used to index into an EEPROM table to determine the DAC value necessary to achieve the desired gain. The DAC-value-versus-gain relationship is determined and stored during factory alignment.

Pulse Detection

The linear detector allows the receiver's user to recover AM and pulse modulation from the input signal. It strips the carrier from the input signal and leaves only the envelope (Fig. 17). The resulting envelope information can then be displayed on an oscilloscope, allowing the user to analyze the modulation or transient characteristics of the input signal.

The key performance specifications for the detector are bandwidth, dynamic range, and pulse fidelity. We would like the detector bandwidth to be much wider than the IF module's bandpass filters so that it does not limit the IF module's bandwidth. The bandpass filters have a maximum bandwidth of 100 MHz, which is equivalent to 50 MHz after detection. The detector is guaranteed to have at least twice this bandwidth, or 100 MHz. Dynamic range is a measure of the linearity of the detector. This is measured by changing the input RF voltage in 1-dB steps and measuring the resulting change in the dc output voltage. Ideally, it should also change by 1 dB. Our specification guarantees that over a 26-dB range, this change will be accurate within 3%.

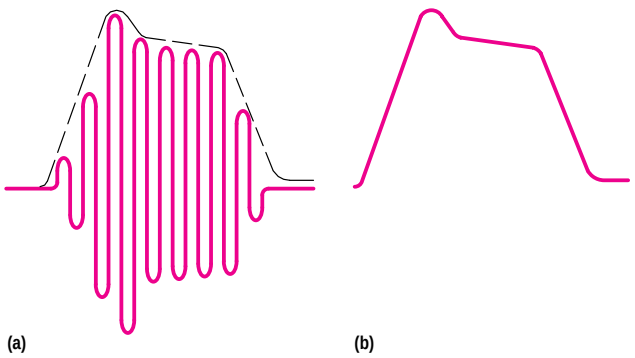

Fig. 17. (a) Input to the linear detector. (b) Output from the linear detector after the carrier is stripped off.

Previous linear detectors in HP spectrum analyzers have achieved this performance, but at much lower IF frequencies of 10.7 or 21.4 MHz. Achieving this performance at 321.4 MHz was the most challenging aspect of this design. A schematic of this circuit is shown in Fig. 18. Q1 is a common-base buffer stage that drives Q2, which is the detector transistor. Q2 and CR1 each act as half-wave rectifiers. Positive half cycles of Q1's output current flow through CR1 to ground. Negative half cycles flow through Q2's emitter and collector and develop a voltage across R1, the load resistor.

The fundamental linearity problem is that the input impedance of Q2 varies dramatically with signal level. With no input signal, Q2 is biased at 120 µA. This yields a dc resistance looking into the emitter of 217 ohms. At full-scale output, the dc emitter current is 10 mA, reducing the resistance to 2.6 ohms. This load resistance is in parallel with several parasitic loads (Fig. 19). Among these parasitic loads are Q1's output capacitance, Q1's collector bias network, the parasitic capacitance of the printed circuit board, and the capacitance of Q2's base-emitter junction. At high signal levels, Q2's input resistance is low, and essentially all of Q1's output current is delivered to the desired load. At low signal levels, Q2's input resistance is high, and the parasitic elements tend to shunt current away from the desired load. This variable current shunting degrades the linearity, so good linearity requires that these parasitic elements load the circuit as little as possible.

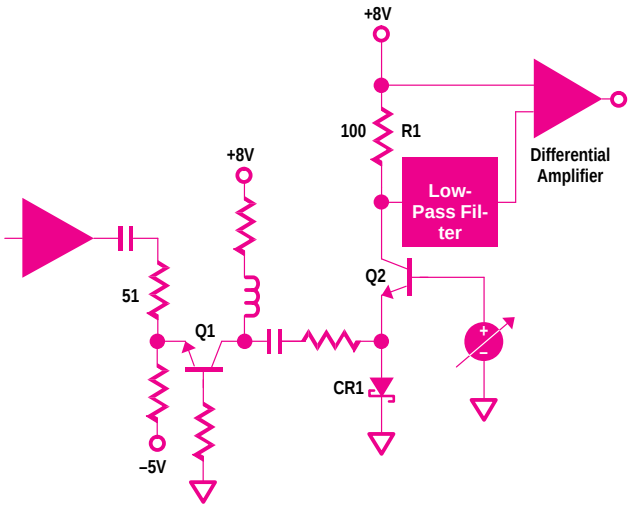

Fig. 18. Linear detector circuit.

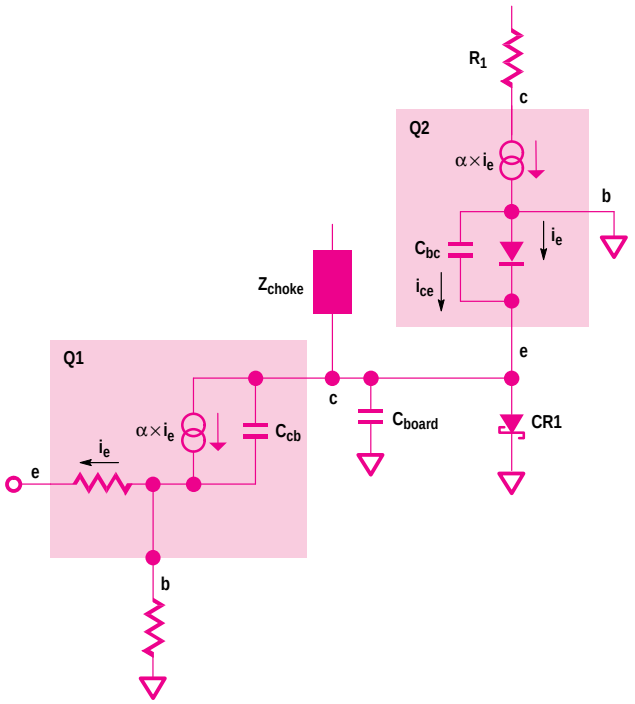

Fig. 19. Linear detector equivalent circuit.

Q1's output capacitance is minimized by using a common base configuration. Also, a microwave transistor is used because of its low capacitance. The load impedance of the collector bias network is maximized by careful design of the bias network. The printed circuit board layout is also carefully designed to make nodal capacitance as small as possible without sacrificing manufacturability.

The selection of the right transistor for Q2 was perhaps the most critical component of the design. We needed to use a microwave transistor to get the low junction capacitance we wanted. We needed two things from this transistor: a low base-emitter capacitance and, if possible, a capacitance that decreases linearly with decreasing collector current. We wanted this relationship between capacitance and current because if capacitance decreases linearly with current, then that capacitance will not degrade the detector's linearity. This is because the junction's capacitive reactance will increase as its resistance increases, and the fraction of current "stolen" by the capacitor will not vary with signal level. Since this shunting effect is independent of signal level, it will not degrade linearity.

It rarely happens, but sometimes semiconductor physics decides to give you just what you'd like. This is one of those cases. To a first-order approximation the base-emitter capacitance of a bipolar transistor is linearly proportional to bias current, at least at moderate current levels. Even better, we could easily extract this information from a transistor's data sheet curves. Low base-emitter capacitance is roughly equivalent to high f_T (transition frequency). A capacitance proportional to bias current will reveal itself as a curve of f_T that is flat versus bias current. Theory suggests, and experiment demonstrated, that the best detector transistors are those that have a high and relatively constant f_T over their entire operating current range. Unfortunately, most microwave transistor data sheets do not give f_T curves over the 100:1 range of bias currents that we wanted. Fortunately for us, we have a lot of data books and found some microwave transistors that met our needs. As expected, the transistors with the best f_T curves yielded the most linear detectors. Typical linearity error for the detector we selected is shown in Fig. 20.

The detector's output current flows across R1, generating a 1-volt drop at the maximum input level. Since the other end of R1 is tied to the 8-volt supply, it is necessary to use a differential amplifier to reference the signal to ground. A simpler approach would have been to tie R1 to ground instead of +8 volts and to tie CR1's cathode to −8 volts. This would have eliminated the need for a differential amplifier. But this would have made it difficult to achieve good pulse fidelity.

Achieving good pulse performance can be hard even with nominally linear circuits, but it is particularly difficult to do with inherently nonlinear ones like detectors. These circuits can exhibit overshoot, droop, or both on any time scale (microseconds to seconds) if their bias networks are not designed correctly. If the bias networks exhibit significant impedance at virtually any frequency below hundreds of MHz, the bias voltages in the detector can vary with the input signal, causing imperfections in the detector transient response. For this reason it seemed risky to try to build a good enough bypass network that could have presented CR1's cathode with uniformly low impedance across a broad frequency range. Rather than accept this risk, we chose to ground CR1's cathode and accept the complexity of a differential amplifier to recover the detected voltage.

The differential amplifier is integrated with a low-pass filter that removes the 321.4-MHz component from the half-wave rectified voltage across R1. This is an elliptic low-pass filter with a 200-MHz corner frequency. Even though elliptic filters have notoriously poor pulse response, we can use one here. We can do this because of the bandwidth limitation imposed on the input signal by the IF module's bandpass filters. The elliptic filter's bandwidth is four times higher than the effective postdetection bandwidth of the IF module's resolution bandwidth filters. Since these filters prevent the higher-frequency components from reaching the elliptic filter, only very low levels of ringing are observed in the detected output. We were able to demonstrate this by simulating the pulse response of the resolution bandwidth filters cascaded with an elliptic detector filter. As a result, we avoided the

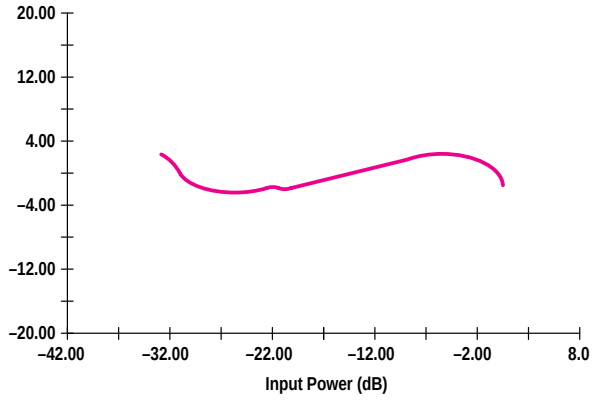

Fig. 20. Detector linearity error.

need for a more complex full-wave rectifier with its inherent carrier suppression.

Standard IF Outputs

For direct connection to commercial demodulators a 70-MHz or 140-MHz IF output is required. The HP 70911A offers either or both of these outputs as options.

A simplified block diagram for these options is shown in Fig. 21. Both down-converters use the 321.4-MHz local oscillator circuits (described later). This design has a VCO tuning range sufficient for both down-converters. The LO frequency for the 70-MHz down-converter is 391.4 MHz and the LO frequency for the 140-MHz down-converter is 461.4 MHz. These LO signals are applied to a mixer which has some buffering in front of it and is followed by an optional filter and an amplifier. Image rejection filtering is not part of the design since it is assumed that the variable-bandwidth filters are in the upstream path. The output filter is used to confine the output bandwidth to a specified amount.

The same basic design is used for both down-converters. The key difference is that in the 140-MHz design a pad follows the mixer, whereas in the 70-MHz design there is a diplexer at the mixer output, which provides a good out-of-band impedance match. The 70-MHz design has also been made available as a special option for the HP 859XE Series spectrum analyzers.

Channel Filters

The channel filters option provides an electronically switchable bank of five bandpass filters and variable gain that can be used at 70-MHz, 140-MHz, or 160-MHz center frequencies.

The input of the board goes to each filter cell through a series of GaAs switches and well-isolated stripline 50-ohm printed circuit board traces. The cells are large enough for a standard-size printed circuit board-mounted filter. The machined aluminum shield has pockets on the bottom to keep the signal pins of the filter isolated from each other. There is also a through path available for bypassing the filters. After the switching network, there is a p-i-n diode attenuator that allows continuous electronic amplitude control. Next, there is a high-dynamic-range, wide-bandwidth amplifier. The amplifier also provides temperature compensation for the gain of the board. The compensation is done by using a thermistor to vary the current in a p-i-n diode which varies the emitter degeneration impedance with temperature.

The excellent isolation, wide bandwidth, and variable gain make the channel filters a flexible option for any of the standard IF outputs.

FM Outputs

The FM discriminator generates an output voltage that is linearly proportional to the frequency of the input signal. It is used to demodulate wideband frequency modulated signals such as those found in satellite television links or chirp radars.

The key performance specification for the FM discriminator is linearity. Ideally, the frequency-input-to-voltage-output transfer function should be a straight line. Our goal was to make the maximum error from a straight line less than 1% of the full-scale output across the 40-MHz deviation range of the demodulator. The techniques used in this design were driven primarily by that goal.

Many different types of circuits have been designed to do FM demodulation. There are Foster-Seely discriminators, ratio detectors, phase-locked and frequency-locked demodulators, and slope detectors. Digital techniques, which count the zero crossings of the input signal and extract the frequency information mathematically, offer the promise of the highest linearity. These techniques are used in the HP 5371A,[2] the HP 53301A, and other modulation-domain analyzers from Hewlett-Packard. Although they achieve excellent linearity, these products are large and expensive and certainly would not fit on a single 4-inch-by-7-inch card in the HP 70911A. For these reasons, it was necessary to pursue a different approach.

Two analog demodulators seemed to offer the best potential for high linearity across a broad band: a pulse count demodulator and a time-delay discriminator. A pulse count demodulator (Fig. 22) generates a fixed-length output pulse at every zero crossing of the input signal. Since higher-frequency signals have more zero crossings, the output pulses occur more frequently. As a result, the dc average value of the output pulse train is higher for higher-frequency inputs. The low-pass filter placed after the pulse generator filters out

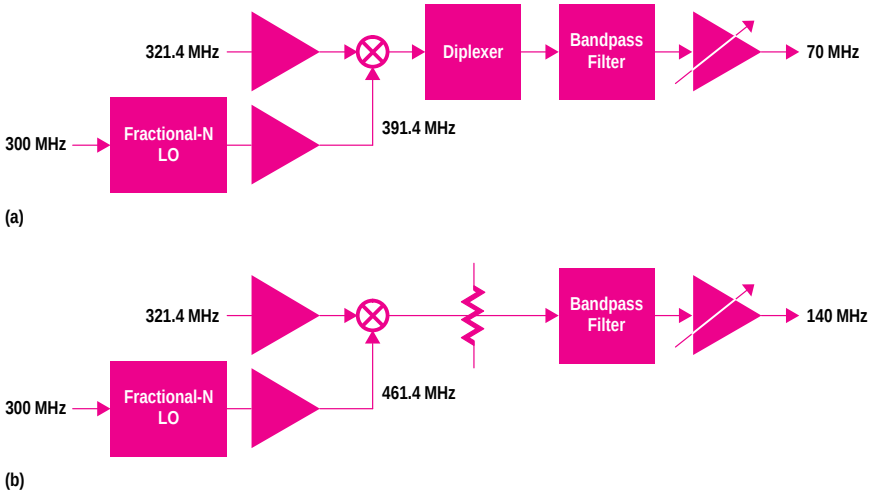

Fig. 21. Simplified block diagram of the output options for direct connection to commercial demodulators. (a) 70-MHz IF output. (b) 140-MHz IF output.

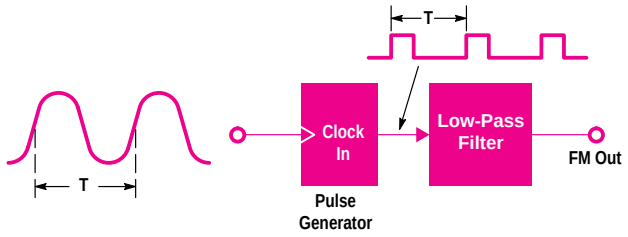

Fig. 22. Pulse count discriminator.

the carrier frequency component of the pulse train, leaving a dc value linearly proportional to the input frequency. This linear conversion of input frequency to output voltage is just what we needed to build a linear discriminator.

This type of demodulator can be implemented very simply and inexpensively by using a retriggerable one-shot timer to generate the output pulses. It does, however, have disadvantages in our application. For one, very narrow pulses would be required to make this work at 321.4 MHz. Also, these simple one-shot timers tend to have noisy outputs because of variations in the width of the output pulse.

A time-delay discriminator works by converting the input signal's frequency modulation into phase modulation (Fig. 23). A delay line delays the input signal by a fixed amount of time. A phase detector on the delay line output compares the phase of the input signal against the phase of the time-delayed version of the input. Since the phase of a high-frequency signal changes more rapidly than the phase of a low-frequency signal, the phase difference between the two inputs to the phase detector will increase linearly with frequency. The output voltage of the phase detector is proportional to this phase difference and thus, proportional to the frequency of the input signal. Typically, the length of the delay line is chosen so that the signal will be delayed 90 degrees at the center frequency of the discriminator. This gives zero volts dc output at the center frequency and centers the output in the middle of the phase detector's transfer function. This inherently linear conversion of frequency to phase seemed to make this type of circuit a logical candidate for our application.

However, this type of discriminator posed two potential disadvantages for our application. First, this discriminator is inherently limited in the maximum frequency deviation and the maximum modulation rate it can handle. Typical phase detectors only behave well when the phase difference between the inputs varies by less than 180 degrees. Because phase difference is proportional to input frequency, the maximum frequency deviation the discriminator can handle is limited. Also, the sensitivity inherently rolls off at high modulation rates. In other words, as the input frequency starts to vary more quickly, the level of the demodulated output will start to drop. The longer the delay line, the lower the modulation rate at which this will occur. In our case, we need to demodulate broad frequency deviations and as a result the maximum delay line length is limited by deviation requirements and not modulation rate needs.

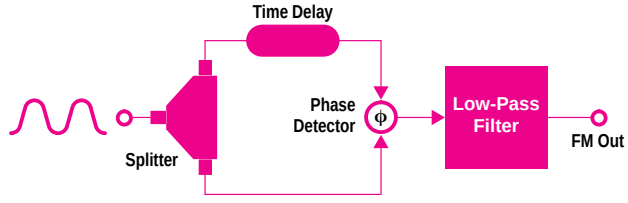

Fig. 23. Time-delay discriminator.

The second disadvantage of the time-delay discriminator is based on phase detector characteristics. Our high IF of 321.4 MHz would suggest using a double-balanced mixer as a phase detector. Conventional double-balanced mixers are designed to work with a sinusoidal RF port drive. The result is that the mixer output voltage varies sinusoidally with the phase difference between the LO and the RF waveforms (Fig. 24). Therefore, it is only linear if the phase difference between the input signals does not vary much from 90 degrees. Since we wanted good linearity, that meant a short delay line. Unfortunately, the shorter the delay line, the lower the sensitivity of the discriminator. Short delay lines mean low phase shifts and therefore low output voltages. For good signal-to-noise ratio, we wanted to maximize the time delay.

A double-balanced mixer has a sinusoidal transfer function because its RF input voltage is sinusoidal. Ideally, if its inputs are square waves, the transfer function would be linear over a 180-degree range. However, generating very fast square waves is hard, and the mixer would need a very broadband dc-coupled IF to work well. Fortunately, there is a type of double-balanced mixer that meets these requirements: the exclusive-OR gate. An ideal double-balanced mixer generates its IF by inverting the RF waveform whenever the amplitude of the LO crosses zero (Fig. 25). This is exactly what a digital exclusive-OR gate does with logic-level inputs. Thus, with this characteristic an exclusive-OR gate can be used as a double-balanced mixer.

Because of our high IF and broad frequency range, we needed to use very fast logic circuitry if we wanted this to work. Motorola's ECLinPS Lite family of emitter-coupled logic turned out to be perfect for our application. These logic gates come individually packaged in eight-pin small-outline ICs and feature rise times under 300 picoseconds. The fast, square pulses generated by this logic are perfect for making a very linear phase detector.

When logic gates are as small and fast as these, it's only natural to use them wherever you can. In the end almost all the functions on the board including limiting amplifiers, mixers,

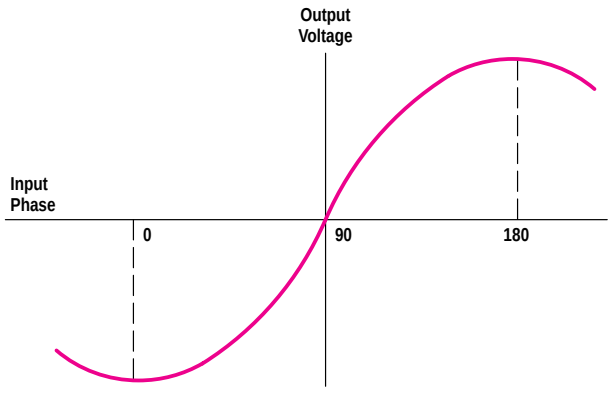

Fig. 24. Transfer function for a double-balanced mixer.

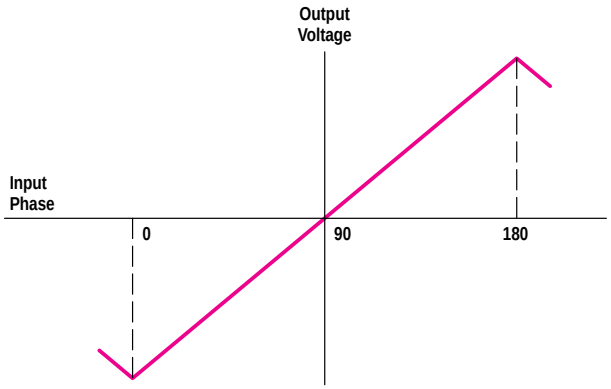

Fig. 25. Exclusive-OR transfer function.

and phase detectors were implemented using these RF logic gates. Because of the nature of FM modulation and demodulation, logic parts work well in frequency modulation applications.

Fig. 26 shows a block diagram of the FM discriminator. The 321.4-MHz input is applied to a limiting amplifier. The limiting amplifier is a high-gain stage that turns the incoming FM sine wave into a constant-level frequency modulated square wave. Given that our signal processing is done with logic parts, we obviously needed something like this to convert the input into ECL levels.

An ideal FM demodulator is insensitive to amplitude modulation of the input signal. The output voltage should not change at all when the input amplitude varies. Limiting amplifiers are used to achieve this. They have high gain and clip the level of the output signal at a predefined level. Our limiting amplifiers are implemented with ECL line receivers, which are differential-input high-gain amplifiers with ECL-level outputs. Their high gain and hard limiting allow the FM demodulator to work properly with inputs as low as −30 dBm.

The output of the limiting amplifier is a square wave with a nominal center frequency of 321.4 MHz. This is mixed against a 250-MHz LO to a lower frequency of 71.4 MHz, where the actual demodulation takes place. Originally, the intent was to do the demodulation at 321.4 MHz. As we better understood the problems we faced in trying to achieve good FM linearity, it became clear that using a lower frequency would produce better results. At a lower frequency, the period of the IF is longer. The rise time of the parts used does not change, so overall the square waves are "squarer." Our analysis of the time-delay discriminator showed that it was perfectly linear, but this is true only if the square waves are perfect.

The use of small surface mount logic parts enabled us to design compact LO generation and frequency conversion circuitry. The 250-MHz LO is derived from the 300-MHz reference frequency available in the HP 70911A. The 300-MHz signal is converted to ECL levels by a limiting amplifier. The 300-MHz reference clocks a prescaler, which divides the input frequency by six to produce a 50-MHz output. The 300-MHz and 50-MHz ECL square waves are then applied to the inputs of an ECL exclusive-OR gate. This gate performs as a double-balanced mixer, producing 250-MHz and 350-MHz outputs. The 250-MHz output is selected by a bandpass filter. This filter is ac coupled, so a limiting amplifier is placed on the output to convert the 250-MHz LO back to ECL levels.

The 250-MHz LO and the 321.4-MHz hard-limited input signal are then applied to another exclusive-OR gate. This gate is also used as a double-balanced mixer, producing outputs at 71.4 MHz and 571.4 MHz. The 71.4-Hz output is selected with a low-pass filter. The entire LO synthesis and frequency conversion circuitry occupies only 3 in^2.

The use of exclusive-OR gates and square waves, as opposed to traditional diode mixers and sine waves, has a surprising consequence. As noted earlier, a traditional diode mixer has a sinusoidal phase-to-voltage transfer characteristic. As a result, the IF out of an ideal diode ring mixer with sinusoidal inputs is another sine wave. In contrast, the logic level mixers we use here have a triangular transfer function. As a result, the IF output of these mixers is a triangular, rather than a sinusoidal waveform. In our case, we don't care whether it's sinusoidal or triangular, because we immediately convert the IF to a square wave with another limiting amplifier.

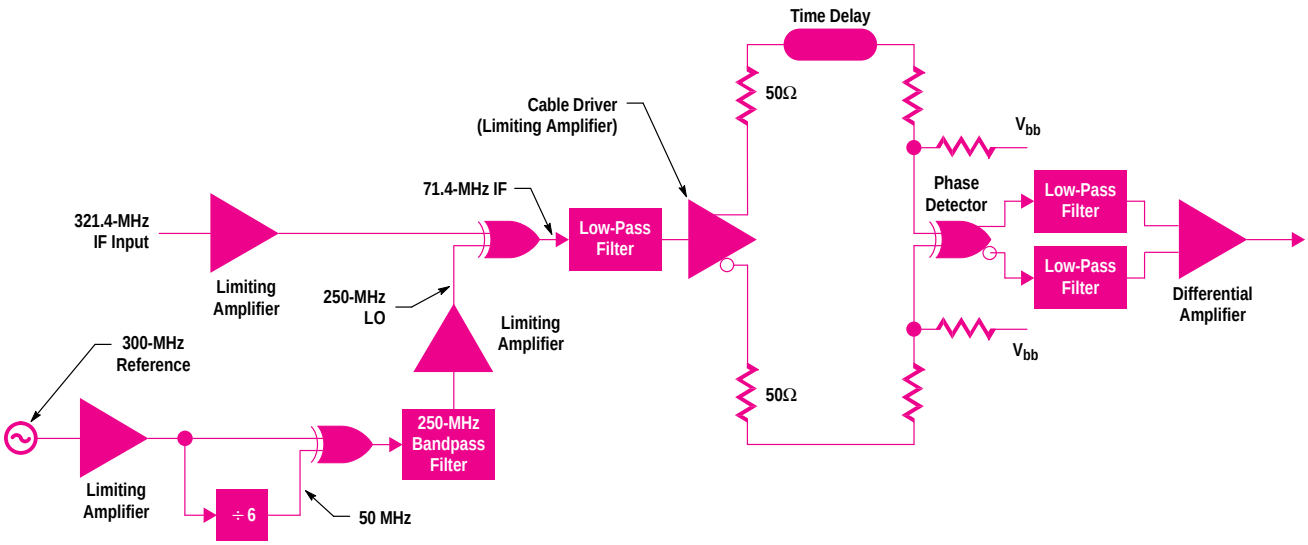

Fig. 26. FM discriminator block diagram.

The 71.4-MHz IF is next applied to the discriminator circuitry. The IF is applied to a special ECL cable driver IC which produces ECL-type outputs that have a larger than usual voltage swing. This large voltage swing allows us to place a series 50-ohm resistor on the output, cutting the signal amplitude in half. The resistor increases the output impedance of the gate to around 50 ohms, which turns out to be crucial to achieving good discriminator linearity.

The outputs of the cable driver follow two different paths. One output goes directly to the input of the phase detector. The other output goes to a delay line. This delay line is a 19-in-long 50-ohm stripline trace embedded in the middle of the printed circuit board. Numerous bends and turns squeeze it into a 1-in-by-3-in area. The board, made out of HP FR4, has a dielectric constant of about 4.5, yielding a 3.5-ns delay. This delay produces a 90-degree phase shift at the center frequency of 71.4 MHz.

The delayed and undelayed signals now meet at the phase detector, which is another exclusive-OR gate. The square waves are applied to the high-impedance input of the phase detector through a 50-ohm matching pad. The attenuation value of this pad is critical to good discriminator linearity. As mentioned earlier, good square waves are critical for good FM linearity. The attenuation value chosen strikes a balance between two "desquaring" mechanisms.

If the attenuation value is small, the delay line output will not be isolated from the 1-pF input capacitance of the phase detector. This capacitor degrades the return loss of the delay line's load. When a pulse emerges from the delay line output and hits a poor impedance match, some of the pulse's energy is reflected back into the delay line. It then travels backwards through the line to emerge at the delay line input 3.5-ns later. When the pulse reemerges here, it hits the output of the cable driver. This incident voltage disturbs the bias of the cable driver output transistors, and as a result causes disturbances in the shape of the new square wave that the cable driver is trying to generate. The degraded shape of the square wave degrades the FM discriminator's linearity. As it turns out, this effect is worst when the reflected pulse arrives at the cable driver at the edge of a new pulse. Unfortunately, this inherently occurs at the frequency where the delay line has 90 degrees of phase shift—right in the center of the passband. This effect is seen as the linearity ripple in the center of the passband (Fig. 27).

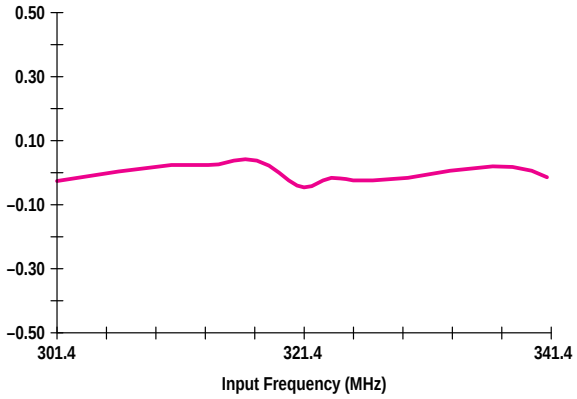

Fig. 27. Discriminator linearity error (percent of full scale).

The way to minimize this problem is to use a matching pad to isolate the delay line from the capacitance of the phase detector's input. The attenuation can only be so large because excessive attenuation introduces other problems. The attenuator reduces the voltage swing to less than ECL levels. As a result, the phase detector must provide gain to produce ECL levels at its output. Unlike the ECL line receivers, these exclusive-OR gates have a relatively low gain of 12 dB. So, with low-level inputs, the output pulses of the phase detector start to look less square. This manifests itself as the broad, slow droop in the linearity curve. In the end, an attenuation value of 4 dB was chosen as a reasonable compromise between these two linearity degrading mechanisms.

In the frequency domain, the phase detector can be thought of as producing a dc voltage proportional to the phase difference between its outputs. Looking at it in the time domain is also interesting. The two inputs to the phase detector are square waves with a fixed time delay of 3.5 ns between them. As a result, the phase detector produces output pulses of fixed 3.5-ns width. As the input frequency changes, these pulses occur more frequently, but the pulse width remains the same. This is also exactly how the pulse count demodulator works! So as it turns out, by using a linear phase detector our time delay discriminator turns out to be equivalent to a pulse count demodulator. It works as well as it does because using a delay line to fix the output pulse width is more stable than the RC time constant of a simpler implementation.

The phase detector outputs are applied to low-pass filters to remove the ac component of the pulse train. These filters have a 12-MHz bandwidth that sets the maximum frequency modulation rate the discriminator can respond to. Since the phase detector has differential outputs, a differential amplifier is used after the filters. The differential amplifier removes the dc offset inherent in the ECL level output of the phase detector. Further gain after the differential amplifier is used to give a 1-volt swing for a 40-MHz change in input frequency. The maximum frequency deviation the FM discriminator can handle is limited by the drive capability of this amplifier, rather than the discriminator circuitry itself. It has been verified experimentally that the discriminator will respond to as much as 100 MHz of deviation with essentially nondegraded linearity. A switchable amplifier provides a higher-sensitivity setting, giving a 1-volt swing for a 10-MHz frequency change.

I-Q Outputs

The I-Q down-converter (Fig. 28) recovers the in-phase and quadrature components of the input signal. The IF input, with a nominal center frequency of 321.4 MHz, is mixed against a 321.4-MHz local oscillator. This creates an IF with a nominal center frequency of zero hertz, or dc. The output bandwidth extends from −50 MHz to +50 MHz.

The input signal is split into two paths. Each of these paths goes to the RF port of a mixer. The 321.4-MHz LO is applied to the LO ports of both mixers. The LO input to one of these mixers is shifted by 90 degrees. The IF outputs are low-pass filtered to remove the image frequency, then amplified and sent to the front panel of the HP 70911A.

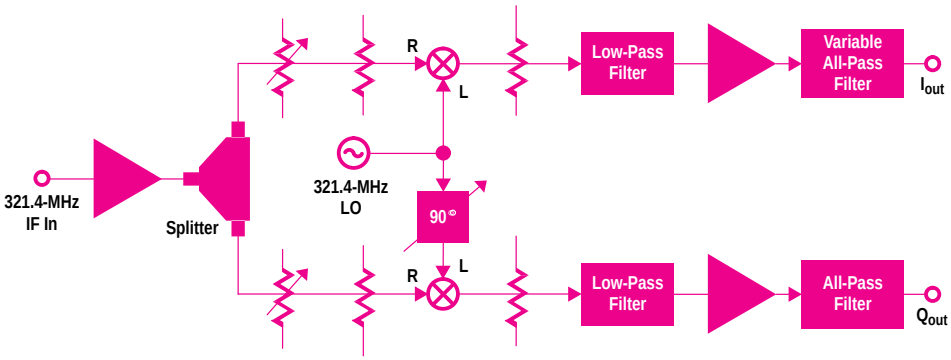

Fig. 28. I-Q demodulator block diagram.

321.4-MHz LO. The 321.4-MHz LO produces a synthesized signal that is locked to the 300-MHz reference signal and level stabilized (Fig. 29). The LO has a VCO that runs at twice the output frequency (642.8 MHz). The reason for running at this frequency is based on the availability of a 600-MHz-to-1000-MHz VCO design that has proven to have good phase noise and has been in use for some time. The VCO output is buffered and split into two paths: the main signal path and the phase-locked loop path. The phase-locked loop path goes from the splitter to a pad-amplifier-pad combination to maintain reverse isolation from the prescaler. The prescaler divides the 642.8-MHz signal by 32, 33, 36 or 37. The divide number is controlled by an HP synthesizer IC that implements the fractional-N division. The output of the synthesizer IC is equal to 300 MHz/160 = 1.875 MHz when the VCO is phase-locked. This signal goes to one input of a phase detector. The phase detector output is low-pass filtered, summed, and fed to an integrator and loop filter. This is where the synthesizer IC's noise is filtered. The noise comes from the method of fractional-N synthesis used in the IC. This noise is designed to be well outside the few kilohertz of bandwidth of the phase-locked loop where it is easy to filter.

The main signal path goes from the splitter to a divide-by-two IC. This is an ECL part that is biased in the middle of its threshold to allow for ac coupling of the 642.8-MHz VCO signal. The output of the divider is 321.4 MHz which is then input to an amplifier and resistive splitter. The splitter outputs are fed to the last gain stages of the board. These amplifiers are run well into compression to get a constant output power. The amplifier outputs are combined with a 3-dB splitter/combiner and then aggressively low-pass filtered to reject the harmonics produced by the limiting action.

I-Q Down-Converter. Two key performance specifications for an I-Q demodulator are amplitude balance and phase balance. Amplitude imbalance is the gain difference between the I and Q output ports. Ideally this gain difference should be zero across the 100-MHz input bandwidth of the demodulator. Phase imbalance is a measure of the error in the phase shift between the I and Q outputs. Ideally this phase shift should be 90 degrees across the input bandwidth.

The amplitude and phase balance of the demodulator are both factory adjusted for best performance. Variable p-i-n diode attenuators in the I and Q RF paths allow the gain of the two channels to be adjusted independently. The 90-degree phase shifter on the LO is also adjustable and is used to align quadrature. These adjustments allow us to align the channels very closely. The mixers used are purchased as a matched set, with specified gain and phase matching across our passband.

The difficult part was maintaining this balance across the entire 100-MHz input bandwidth of the demodulator. If the frequency responses of the two channels differ even slightly, amplitude and phase balance will be degraded. For this reason, we tried to make the two channels as symmetrical as possible and as flat as possible. The printed circuit board layout of the RF paths for both channels is identical so that any board parasitics will be the same for both channels. The IF circuitry was designed to be as broadband as possible. For example, the IF low-pass filters have a corner frequency of 175 MHz, even though they only need to pass frequencies as high as 50 MHz. The corner frequency was placed this high to minimize the filter's phase shift in the 50-MHz passband. Our IF amplifiers are fast op amps that provide over

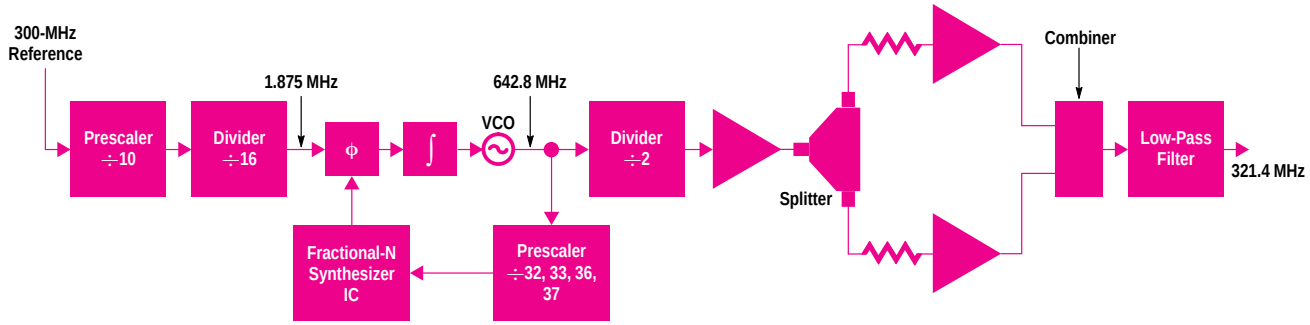

Fig. 29. Block diagram of the 321.4-MHz LO.

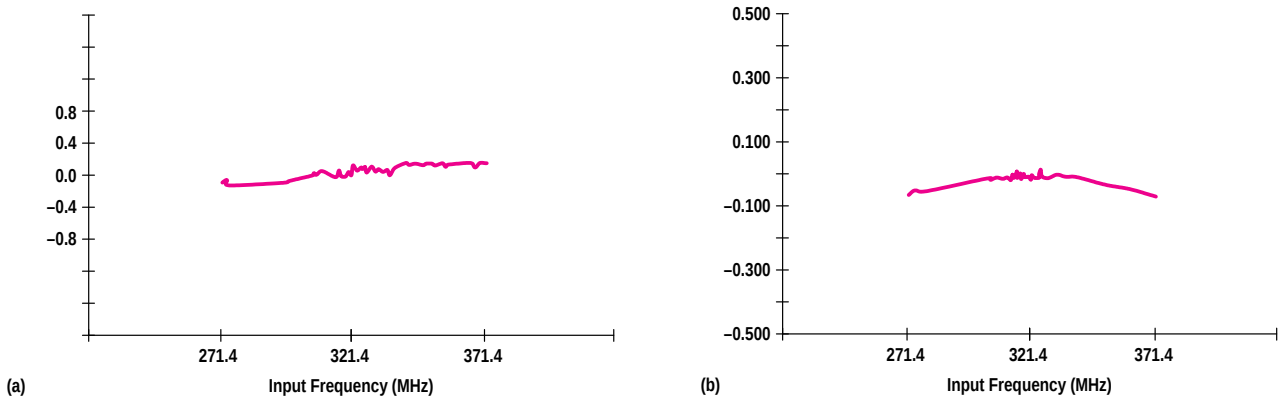

Fig. 30. (a) I-Q quadrature phase error. (b) I-Q amplitude imbalance.

200 MHz of bandwidth. These op amps are also used to minimize the phase shift in the 50-MHz passband. If these parts have significant phase shift, then there are likely to be significant phase shift differences between the two channels, and phase balance will be degraded. Representative performance for the I-Q demodulator is shown in Fig. 30.

To achieve the best phase balance across our bandwidth, an adjustable all-pass filter is used on the I-channel output. The phase shift versus frequency of this circuit is adjustable. It is used to compensate for mismatches between the channels in phase shift as a function of output frequency.

Acknowledgments

The authors would like to thank John Fisher for his support and guidance on this project. Others who made significant contributions to this project include Greg Quintana for controller and firmware design, Bill Walkowski for market research and product definition, and Gil Strand for alignment methodology and production test development. The authors would also like to thank the entire management team for their support during development.

References

1. H. Riblet, "An Accurate Approximation of the Impedance of a Circular Cylinder, Concentric with an External Square Tube," *IEEE Transactions on Microwave Theory and Techniques*, Vol. MTT-31, Oct 1993, pp. 841-844.

2. *Hewlett-Packard Journal*, February 1989, Vol. 40, no. 1, pp. 6-41.