HP CaLan: A Cable System Tester that Is Accurate Even in the Presence of Ingress

Daniel D. Van Winkle

Today, cable system operators have to deal with bidirectional traffic from sources such as pay-per-view television, high-speed Internet access, and two-way telephony. A cable testing system is described that can handle bidirectional traffic even with RF noise (ingress) on the return path.

The deregulation of the telecommunications industry has resulted in an unprecedented surge of effort to implement new telecommunications services. The cable television industry, in particular, is pushing hard to provide customers with two-way telephony, fast Internet service, and interactive video programming. As they prepare to implement and maintain two-way communications over their cable networks, cable industry engineers and providers of cable TV test equipment are facing a whole new set of challenges.

This article describes the background and design of the new HP CaLan 3010H and 3010R sweep/ingress analyzer (see **Figure 1**), which enables cable television providers to do return path alignment in two-way cable systems and to troubleshoot the system quickly regardless of the presence of RF noise (ingress).

Daniel D. Van Winkle

Dan Van Winkle is a hardware design engineer at HP's Microwave Instruments Division and is currently working on the HP Broadband Series Test System. Since joining HP in 1990 he has worked as a production engineer in the modular measurement system and microelectronic departments. He received a BSEL (electronics engineering) degree in 1984 from the California Polytechnic State University in San Luis Obispo, California, and an MSEE degree in 1988 from Santa Clara University. Born in Greenbrae, California, Dan is married and has two children. In his free time he enjoys outdoors sports such as hiking, mountain biking, and skiing.

Figure 1

The sweep/ingress analyzer consists of the HP 3010H headend unit and the HP 3010R field unit.

Background

Cable systems have historically been one-directional. Signals from sources such as satellite feeds, community access programming from local studios, tape machines with prerecorded programming and commercial insertion, and various antennas for both local and distant broadcast television signals are received at a central cable office (headend). Here they are appropriately combined and rebroadcast to various branches over the cable system until they finally reach the subscriber's home.

A typical cable system has trunk lines that feed bridge amplifiers. The amplifiers feed distribution lines, which in turn feed the flexible cable drops to the subscriber's home. Approximately 40% of the system's cable footage is in the distribution system and 45% is in the flexible drops to the home.[1] A cable system can be as simple as a small headend with a single trunk line and only a couple of distribution lines, or it can have multiple trunk lines, each with multiple distribution lines and possibly a remote headend connected to the central headend through a midsplit super trunk. **Figure 2** shows a portion of a typical conventional cable system.

Figure 2

Conventional low-frequency coaxial cable system.

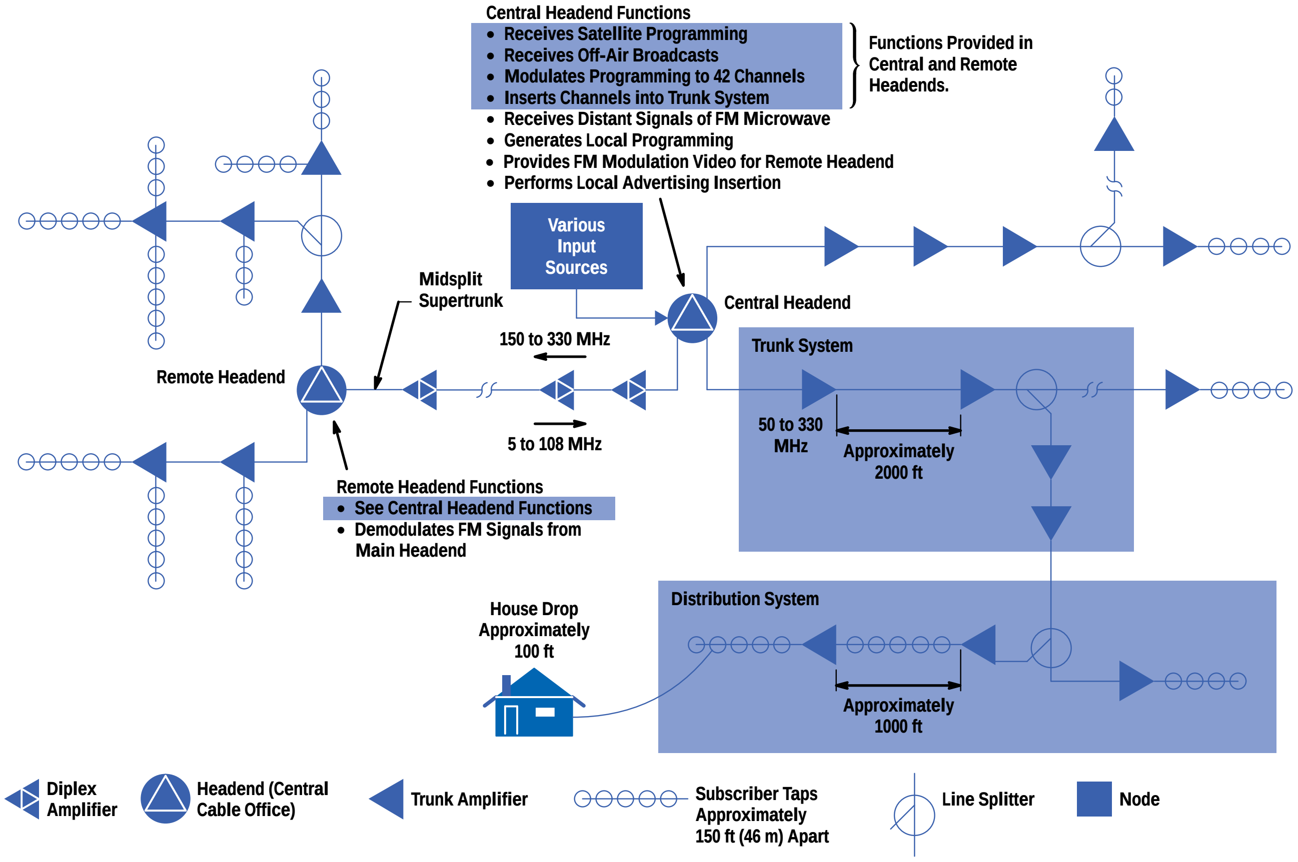

Glossary

Ingress. Any undesired signals present on the cable system. These signals usually come from inside the home, but can also come from ham radio operators, car ignitions, and other sources.

Signal Level Meter (SLM). A signal level meter is used to measure signal levels in a cable system. It is typically a tuned receiver that is calibrated to measure a fairly wide dynamic range.

Forward and Return Pilot. The forward pilot is a modulated fixed-frequency carrier (set by the cable operator) that is used for forward communication in a cable system. Forward communications travel from the headend to the remote site. The return pilot is a modulated fixed-frequency carrier (set by the cable operator) that is used for return communications in a cable system. Return communications are from the remote site.

Headend. The headend in a cable system is the location where signals are gathered from various sources (off-air, community access, tape machines, and so on). The headend is also the central distribution point for the entire cable system. In upgraded two-way systems, the headend will likely contain some type of file server connected to the Internet through a T1 line.

Remote Unit. The remote unit as defined for the HP 3010H and HP 3010R sweep system is the unit that is taken into the field to do forward and return sweeping. It is handheld and battery operated.

Return Sweep. The return sweep is the magnitude response versus frequency of the entire (or partial) cascade of all the amplifiers and cable lengths in a cable system in the return direction. Return sweep (and forward sweep) are used as alignment tools for cable systems.

Many cable systems are upgrading their plants* from low-frequency coaxial cable systems to hybrid fiber/coaxial cable distribution systems. In hybrid fiber/coaxial systems (also called fiber backbone systems) the system is divided into several smaller cable systems with amplifier cascades limited to four to six amplifiers. Each of the smaller cable systems is fed with a fiber link to the headend (see **Figure 3**).

---

* The entire cable system is referred to as the “cable plant.”

Figure 3

Hybrid fiber/coaxial cable distribution system.

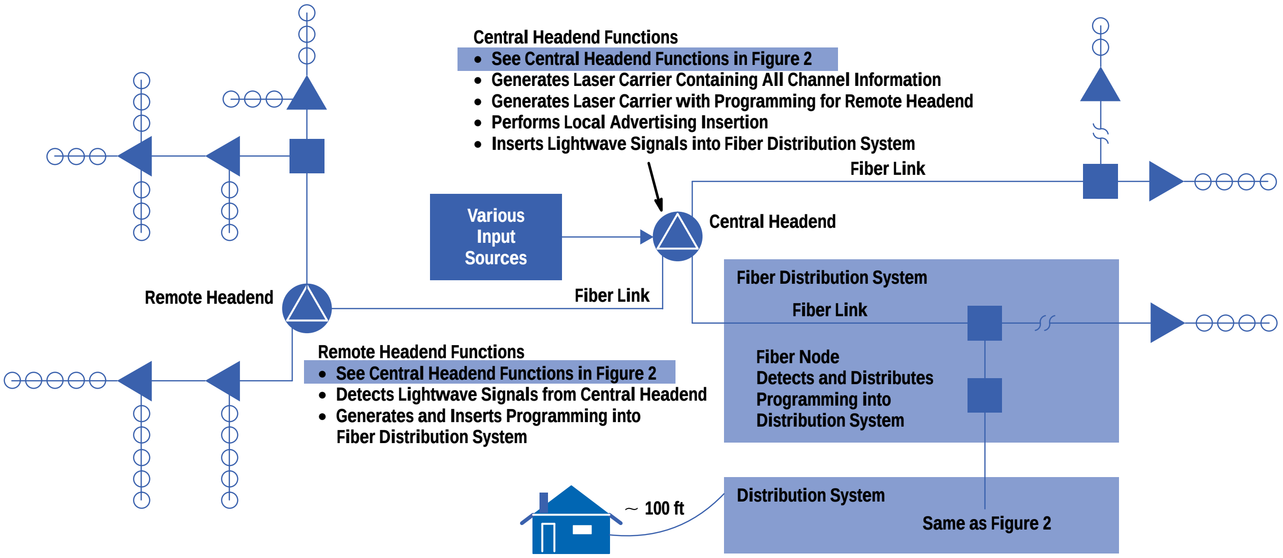

Figure 4

Modern two-way cable system.

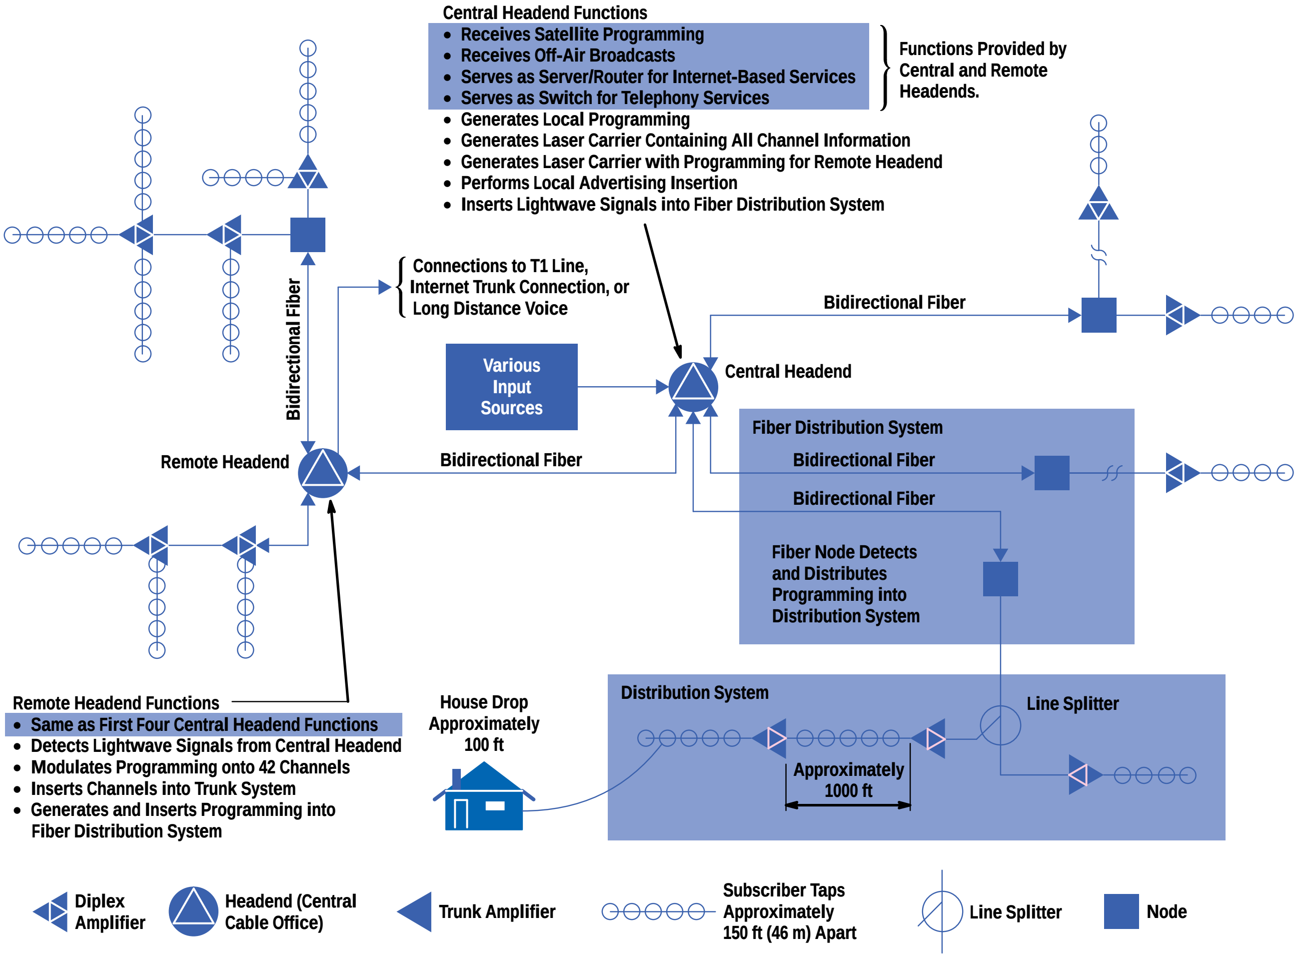

The advantages of using hybrid fiber/coaxial systems include significantly lower vulnerability to amplifier outages, reduced bandwidth restrictions, lower noise buildup (because fewer amplifiers are in series), and greatly reduced ingress. The greatly reduced ingress makes a two-way system practical. Another major benefit of the fiber backbone approach is that implementation cost is relatively low.

Many cable operators are also upgrading their plants to be bidirectional. Not only will the system transmit signals from the headend, but signals from the subscribers' homes will also be transmitted from various in-home devices back to the headend. **Figure 4** shows a generic return path system.

The return path allows new services, such as audio and video telephony, video-on-demand, and Internet access, to be offered. As the cable operators bring up these systems, they are finding that the return path (from the home to the headend) is causing some new and interesting challenges. One-way cable systems have a tree-like structure, in that information (modulated video carriers) is fed from the root up through the trunk, into individual branches, and finally into the leaves (subscribers' homes). On the other hand, two-way systems have a river-like structure, in that every house drop is like a small creek that feeds into a larger river and then pours into a lake (headend). Any debris (noise) that comes from small creeks will also end up being dumped into the lake (headend). In a large system, this noise power can be enough to cause significant interference. Since most cable operators plan on using the low-frequency (5 to 40 MHz) end of the cable system, and since most of the noise generated in homes is in this range, careful attention must be paid to establishing and maintaining the return path. As mentioned above, a hybrid fiber/coaxial distribution system can help significantly in reducing return-path ingress.

The Development Process

The development process for the HP CaLan 3010H and 3010R sweep/ingress analyzer proceeded in three phases. First, we developed a proposal and model of an "ideal" cable testing system. Second, we created the design concept for the new analyzer by making a list of hardware and software tasks and performing a system analysis. Finally, using the output from the previous two phases, we developed the product.

The Proposed System

The HP CaLan 2010/3010 product line has historically been used by cable operators to sweep and align their forward path systems. The HP Calan 2010A was essentially a signal level meter (SLM) that allowed cable operators to measure carrier levels and perform some FCC testing. The HP 3010A (the first instrument in the HP 3010 family) added the capability to perform a forward sweep like a scalar network analyzer. For example, to perform a forward sweep, an HP 1777 integrated sweep transmitter was required in the headend (see **Figure 5a**). This system worked by stepping the HP 1777 source across the 5-MHz-to-1-GHz band at a power level below the visual and audio carriers. (Guardbanding was used to prevent the sweep source from causing unwanted interference.) The HP 1777 was pulsed for approximately 5 to 10 μs at each frequency step. The HP 3010A then widened its receiver time window to catch the signal from the HP 1777. Careful timing of the source (HP 1777) and the receiver (HP 3010R) was required to maintain synchronization. Also, because the signal from the HP 1777 passed through the entire cable system, the resultant signal level received by the HP 3010A showed the swept amplitude response of the entire cable plant from the headend to the remote site where the system was being tested.

Because cable companies are working rapidly to upgrade their return path systems, an urgent measurement opportunity has emerged: *return sweep*. To perform a return sweep, the remote unit has to contain a frequency source with enough frequency range to cover the entire return bandwidth, and the headend unit has to contain a receiver. Because the existing products (HP 3010A and HP 1777) did not contain the necessary hardware, the design team proposed a scheme in which a sweep source would be added to the remote unit and a new headend unit would be created. The headend unit would have the same functionality as the remote unit but would be capable of acting as a central control point for the system by coordinating the sweeps of multiple remote units. A block diagram of this scheme is shown in **Figure 5b**.

The Design Concept

After evaluating the system shown in **Figure 5b**, the design team decided that to build these new instruments they needed to:

- Create a sweep/ingress analyzer that includes a sweep source and maintains the functionality of the original HP 3010A
- Move the design to a process that is compatible with surface mount technology and use preferred parts where possible
- Move the design to a less expensive substrate process (The original HP 3010A RF board was Teflon laminated onto FR4.)
- Develop firmware changes to create a communications protocol that would allow communication between a remote unit and a headend unit over the system being measured.
- Correct any of the current RF board yield issues, including:
  - The tuning range issues in the first LO caused the 4.4-GHz-to-5.6-GHz oscillator to have chronic problems with dropping out at the low end of the frequency range at higher temperatures.

Figure 5

Setup for a forward sweep (a) and a return sweep (b).

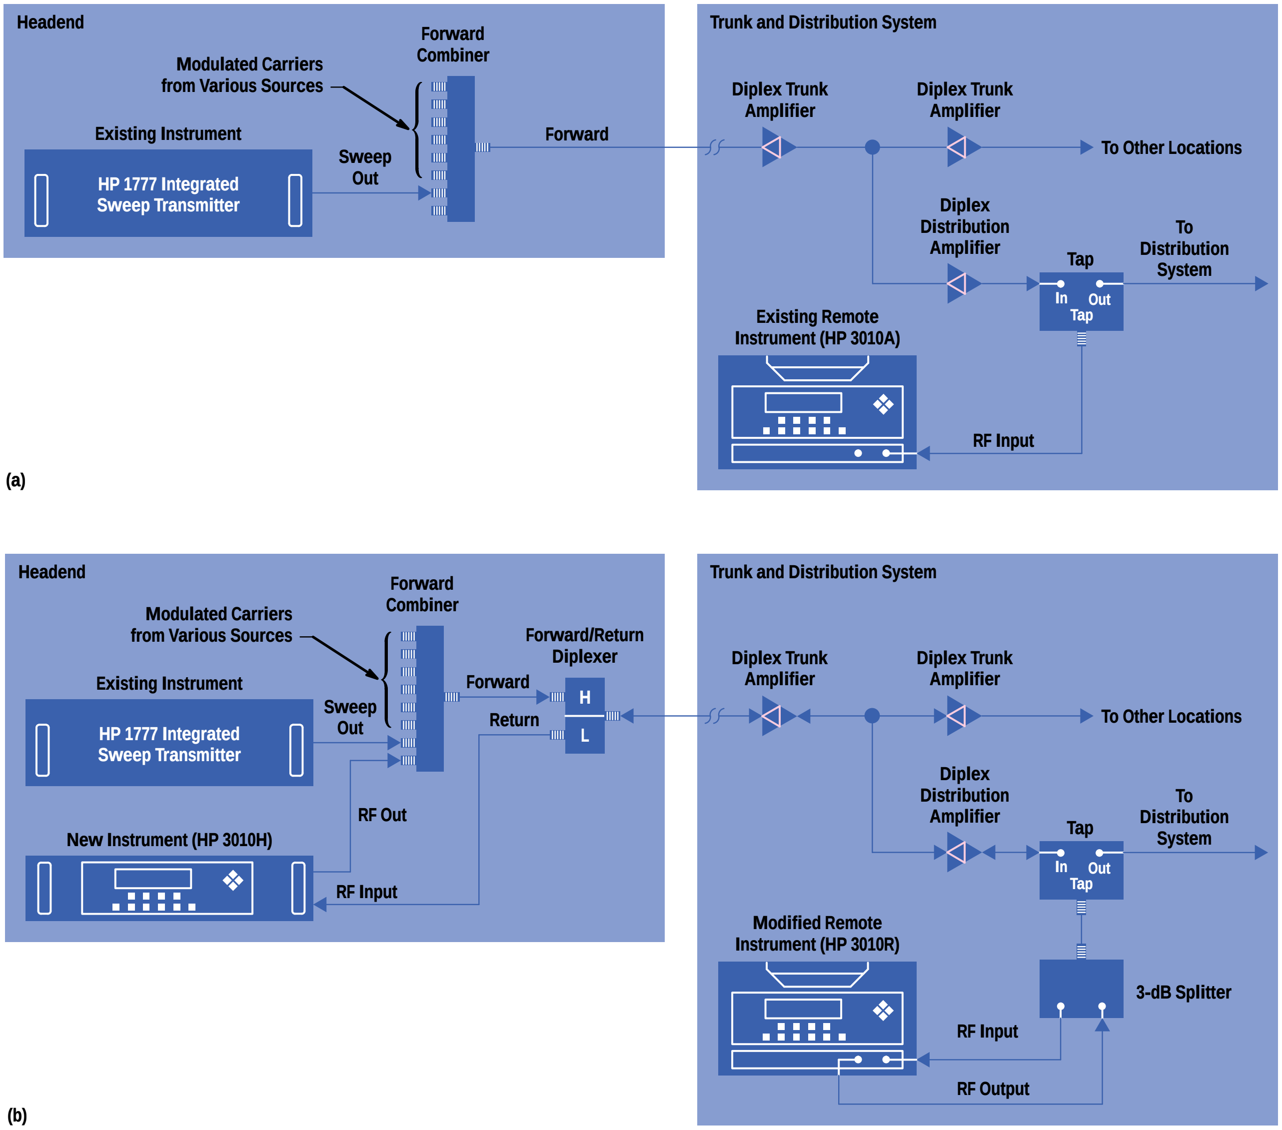

- The first mixer input flatness was very dependent on the positioning of the dual-diode mixer package.
- The 20-dB preamp was very dependent on transistor f_T (cutoff frequency), and since devices tend to vary from lot to lot, the preamp often required hand-selected parts.
- The 20-dB input match required tuning.

At the beginning of the project, the required hardware changes to the original HP 3010A block diagram (see **Figure 6**) were not readily obvious. We considered several schemes for the sweep source portion of the new remote unit. We finally settled on an approach that is based upon a tracking generator concept used in most spectrum analyzers. In a spectrum analyzer, a tracking generator is implemented by operating the RF conversion in reverse. That is, instead of converting a

Figure 6

Original HP 3010A block diagram.

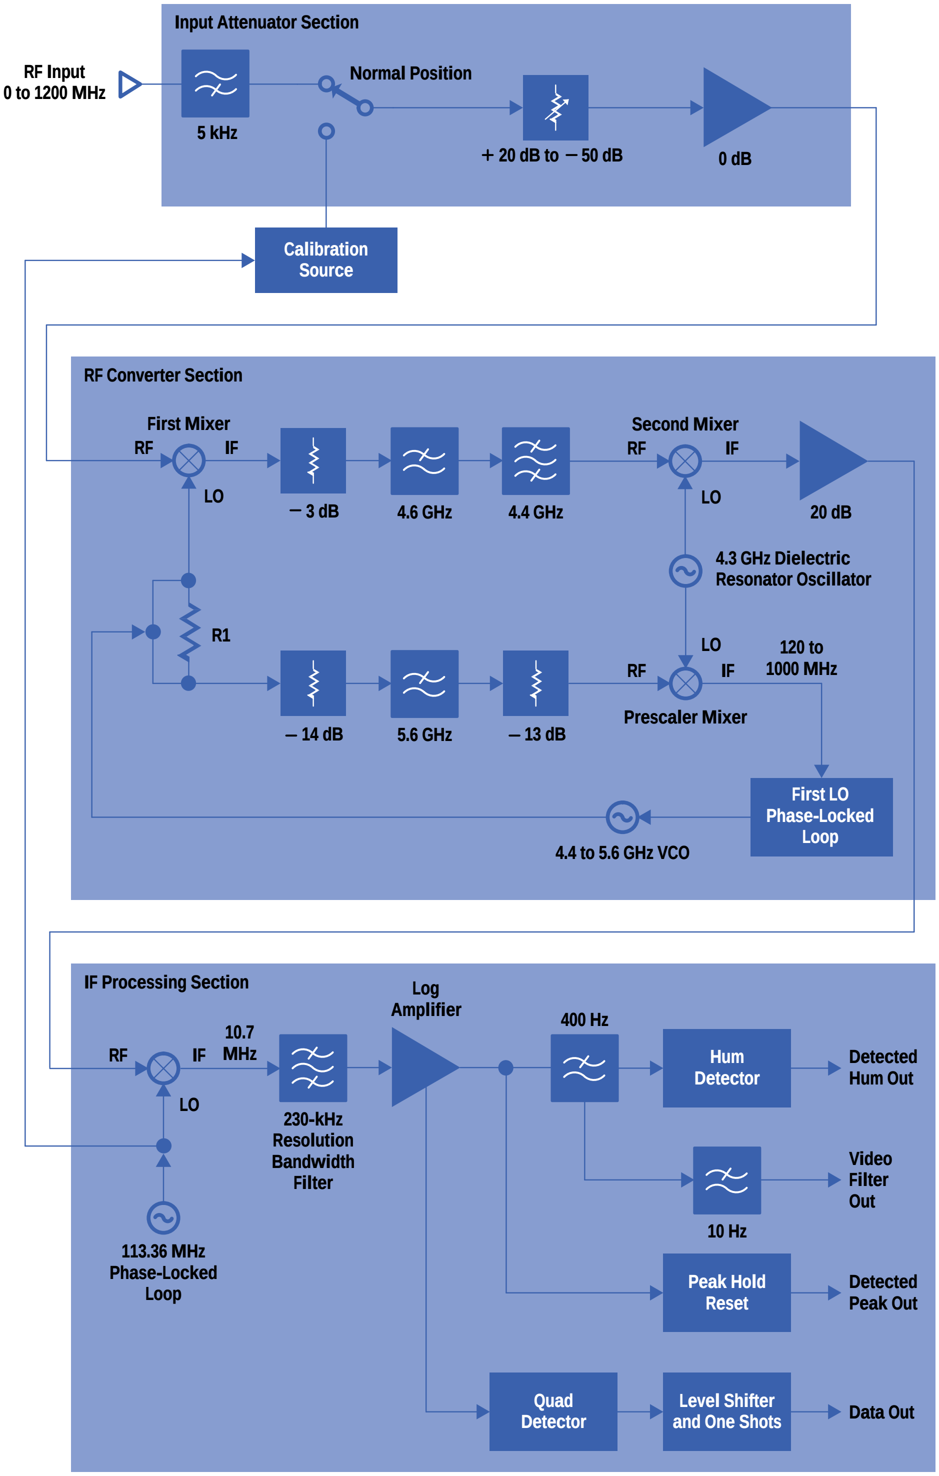

broad input frequency range into a fixed IF using superheterodyne techniques, a broad output range is created by operating a superheterodyne receiver in reverse (see **Figure 7**).

Figure 7

Generic spectrum analyzer tracking generator.

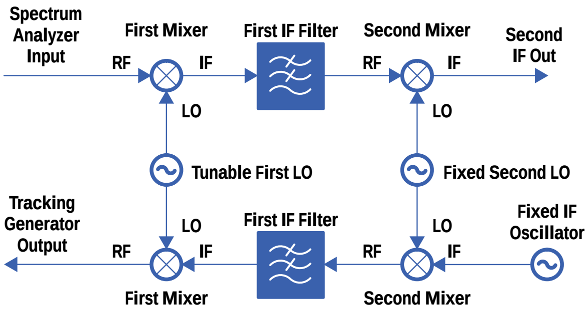

The tracking generator approach appeared to be really promising. However, to continue to use the unit as a receiver, we added switches around the RF converter section to switch between transmit and receive. The final configuration is shown in **Figure 8**.

This approach operates as follows. In the receive mode, the signal enters through the RF input and the attenuator section, is passed through the first transmit/receive switch, and is injected into the first mixer. Based upon the reference level, the attenuators are set to keep the input power at the first mixer at approximately −30 dBm to −40 dBm. The first mixer upconverts the signal to a fixed IF of 4.454 GHz, which is filtered and passed on to the second mixer where it is downconverted to 124.0568 MHz. The 124.0568-MHz signal is passed through the second transmit/receive switch to the IF section where it is amplified and mixed a second time with 113.3568 MHz and downconverted to the final IF of 10.7 MHz. The signal is amplified again, passed through the resolution bandwidth filter, and applied to the log amplifier. The log amplifier output can either be filtered through a 400-Hz low-pass filter, or it can go directly to the peak detector. The output of the peak detector is sent to the main processor board, which contains an analog-to-digital converter (ADC).

In the transmit mode, the two transmit/receive switches are reversed, and the 124.0568 MHz phase-locked loop oscillator is turned on. The 124.0568 MHz oscillator signal is passed into the second transmit/receive switch and injected into the second mixer where it is upconverted to 4.454 GHz, filtered, and injected into the first mixer. Based upon the frequency of the first LO, the first mixer, operating in reverse, downconverts this fixed IF to a signal in the range of 5 MHz to 1 GHz. The resulting signal is then applied through the first transmit/receive switch to the output ALC (automatic level control) amplifier, which provides approximately 20 dB of ALC range to allow for slope correction and vernier 1-dB adjustments of the output power. The signal is then passed through an output attenuator which can be set to 0, 10, 20 or 30 dB, giving the unit an overall output power control range from 10 dBmV to 50 dBmV.

Overall Design

With the concept defined, detailed hardware development could be started. Five main hardware tasks had to be resolved. From the concept definition, two more design tasks were created: a 124.0568-MHz phase-locked loop oscillator and a leveled sweep amplifier (which would be adjustable from 10 dBmV to 50 dBmV).

Since one of the main priorities of this project was time to market, breadboarding was critical to start firmware development and identify unforeseen implementation problems early. The breadboards ended up being very solid and, in fact, are still being used by the firmware engineers to develop further enhancements to the product.

Figure 8

Final block diagram for the new HP 3010.

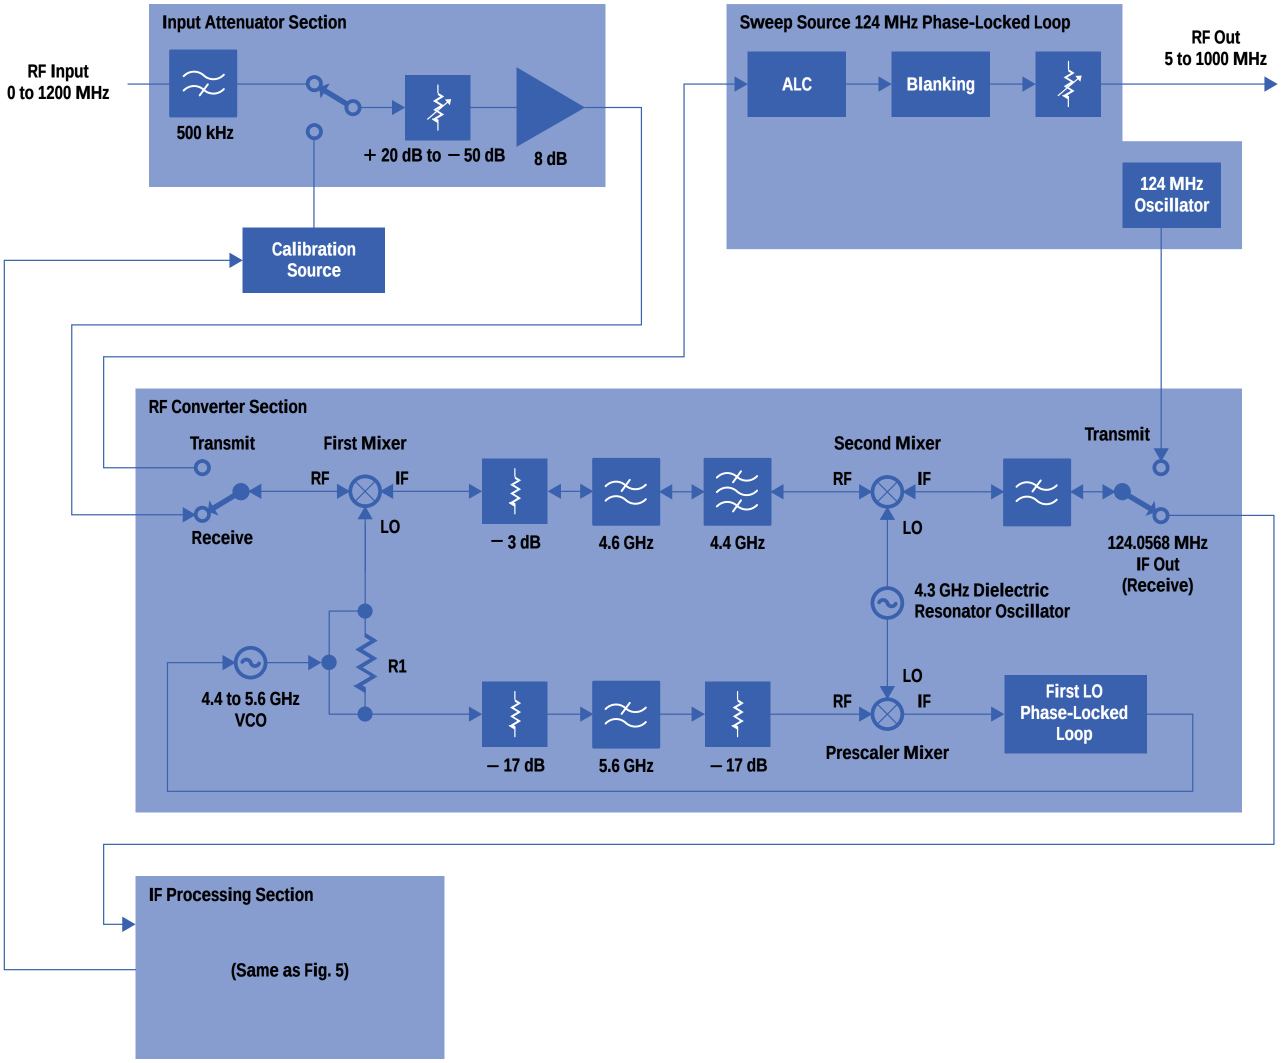

The next step in the hardware design was to convert the board from the fairly expensive low-loss material to a fairly inexpensive GeTek* material. The circuits that were most likely to cause problems were those that caused problems in the original design: the first mixer and the first local oscillator.

First Mixer

A symbolic schematic of the original first mixer is shown in **Figure 9**. The main function of the first mixer is to take the input signal (5 MHz to 1 GHz) and upconvert it to the fixed first IF of 4.454 GHz by multiplying it by the first LO frequency. For an upconversion at 5 MHz, the first LO must be 4.459 GHz, and for 1 GHz, it must be 5.454 GHz.

---

* GeTek is a fairly low-loss, woven-glass type of material that is much easier to process than Teflon-based material laminated to FR4. Because GeTek is easier to fabricate, the raw board is significantly less expensive than one made with Teflon.

Figure 9

The original configuration of the first mixer used in the HP 3010A.

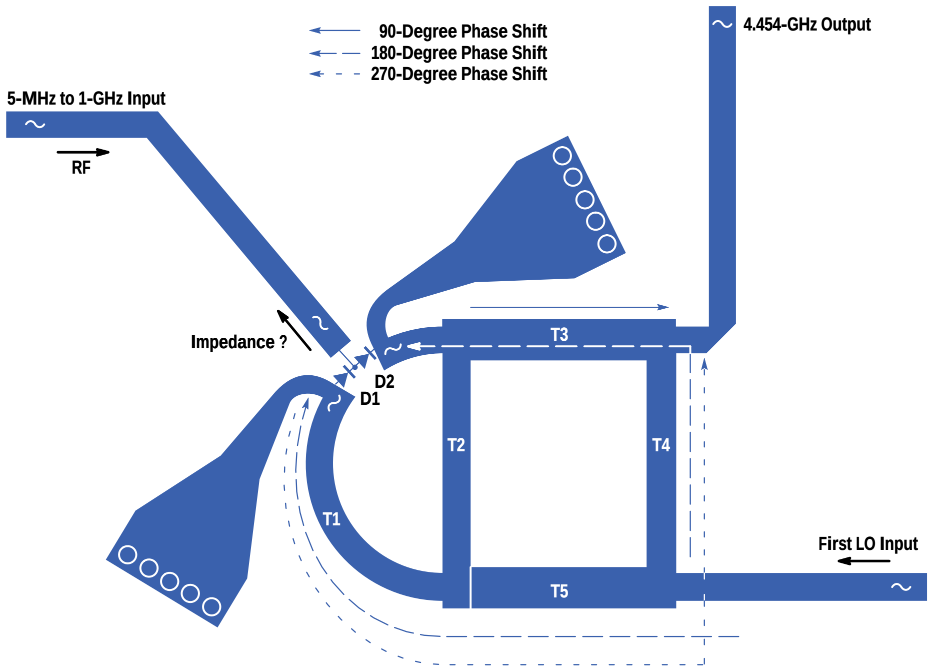

The mixer in **Figure 9** works as follows:

- During the first half of the LO cycle, D1 is forward biased and D2 is reverse biased. A forward-biased Schottky diode looks simply like a resistor, so the RF signal flows through T1 into the microstrip hybrid, emerging at the T3/T4 junction of the hybrid shifted by 270 degrees. Because a reverse-biased Schottky diode looks like an extremely small capacitor, an insignificant amount of the RF power flows into the T2/T3 junction.
- During the second half of the LO cycle, D2 is forward biased and D1 is reverse biased so that now the RF signal flows into the hybrid at the T2/T3 junction, emerging at the T3/T4 junction shifted by 90 degrees.

The net signal coming out of the T3/T4 junction can be seen to be a multiplication of the LO signal and the RF signal. That is, the phase reversal of 90 and 270 degrees of the RF signal is essentially the same as multiplying the RF signal by a unit LO pulse train with an amplitude of ± 1. **Figure 10a** shows an overlay of an RF input signal of 1 GHz and the first LO signal of 4.5 GHz. The solid trace in **Figure 10b** shows the voltage at the anode of D1. The dashed trace in **Figure 10b** shows the voltage at the cathode of D2. As these signals are phase shifted through the hybrid, the resultant trace at the T3/T4 junction is shown in **Figure 10c**. This is equivalent to multiplying the two traces shown in **Figure 10d**.

This multiplication performed in the time domain results in a translation in the frequency domain and the net result is the sum and difference of the RF and first LO frequencies. More specifically, if the input frequency is f_{rf} and the first LO frequency is f_{lo}, mixing produces two new frequencies: $f_{hi} = (f_{lo} + f_{rf})$ and $f_{low} = (f_{lo} - f_{rf})$. In our case we are interested in the f_{hi} so we simply filter the f_{low} with the first IF bandpass filter.

It was not clear, at first, what was causing the conversion loss to be so sensitive to the position of the mixer diode pair. A closer look at the layout shown in **Figure 9** and the functional description reveals that when D1 is forward biased the impedance looking into the common leg is extremely important. In fact, it would be best if it were a short circuit at the LO frequency, and conversely, with an open circuit, the mixer would hardly function. Looking out of the mixer RF input

Figure 10

Waveforms appearing at the junctions of the first mixer shown in Figure 9. (a) An overlay of an RF input signal of 1 GHz and the first LO signal of 4.5 GHz. (b) The voltage at the T1 junction (anode of D1) and the cathode of D2 (T2/T3 junction). (c) The phase-shifted signals at the junction of T3/T4. (d) The two signals that could be multiplied together to produce the trace shown in Figure. 10c.

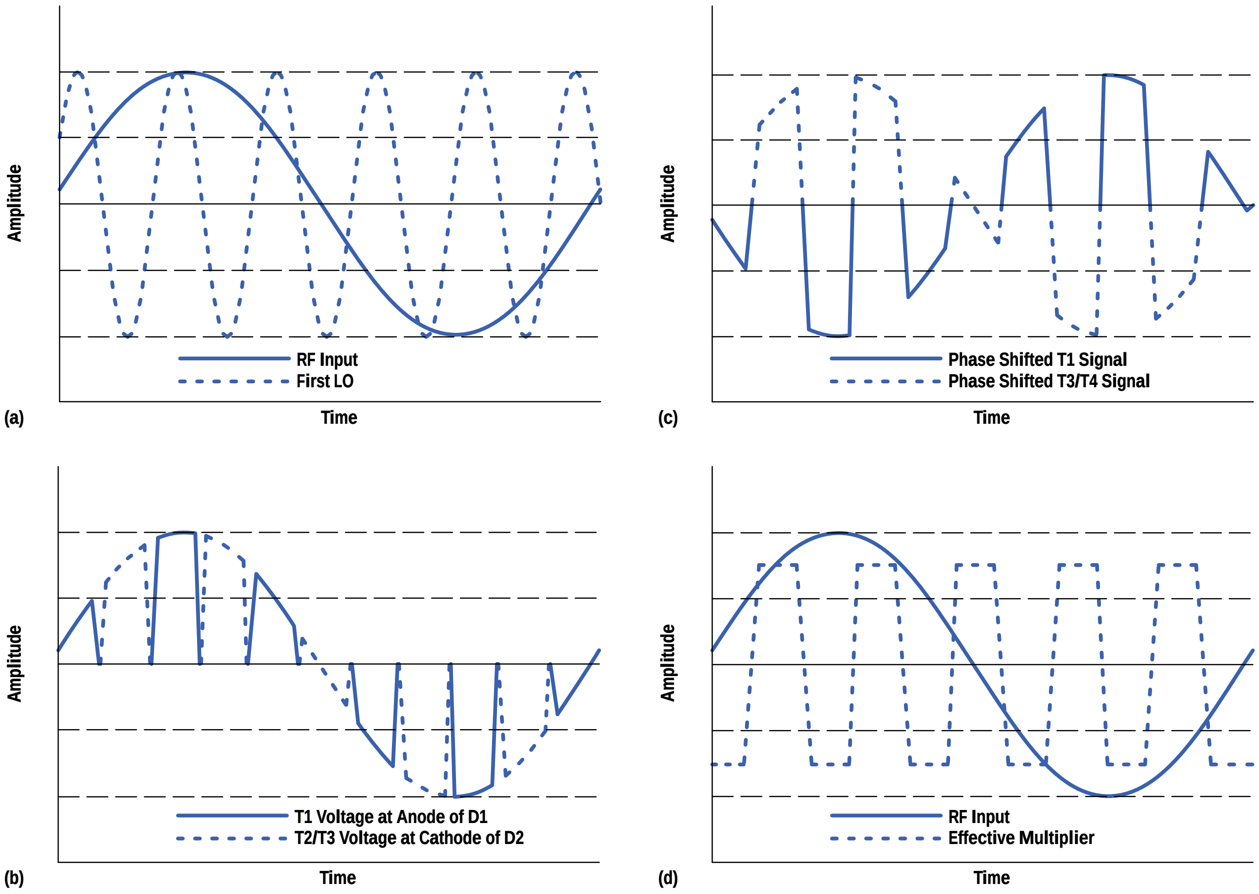

leg, it was extremely hard to tell what that impedance would be at 4.4 GHz or 5.6 GHz, and it was even harder to control. That leg is fed by a low-frequency (5-MHz-to-1-GHz) 8-dB preamp and its output impedance is not easily defined or controlled at 4.4 GHz. What was needed at that point was an ac short over the LO range (4.4 to 5.5 GHz) that would pass a signal in the 5-MHz-to-1-GHz RF input range. This was a prime application for a radial stub. After conducting simulation we decided to add a pair of radial stubs to the mixer, resulting in the layout shown in **Figure 11**.

The new mixer shows a much better conversion loss (approximately 2 dB) and much better flatness across the band. We also found that the new mixer is less susceptible to LO power variations, and when operating in the tracking generator mode, shows approximately a 10-dB improvement in harmonic performance. A breadboard of the mixer was built and tested on the new GeTek material and was found to perform quite well, with the measured results closely matching those of the simulation.

First Local Oscillator

The first local oscillator, which had been a chronic problem circuit on the original board material, needed to cover the 4.4 to 5.6 GHz range. This was perhaps the single most risky circuit to try to fabricate on a lossy dielectric material.

Figure 11

A modification to the configuration shown in Figure 9. This is the configuration used in the new HP 3010H.

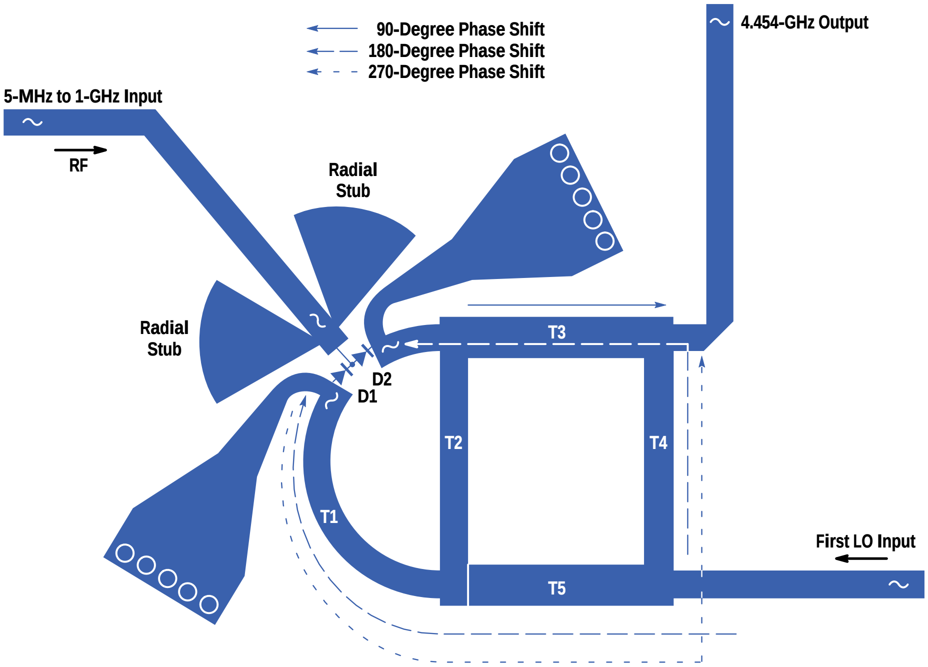

Because oscillators require fairly high Q resonators so that the negative impedance can overcome the lossy portion of the resonator, going to a higher loss material seemed like it could only make the problems worse.

We simulated the oscillator in its original configuration to understand the nature of the original circuit's problem with the low end of the frequency range dropping out of oscillation when subjected to heat. Also, the oscillator had been susceptible to *moding* (multiple oscillation modes) in the past. The oscillator was a negative-resistance oscillator with a microstrip stub and a varactor as its resonator. The simulation revealed several things about the oscillator that gave insight into the cause of the recurring problem at the low end of the frequency range. The absolute value of the negative resistance was barely enough to overcome the losses of the resonator, and as the active device's parameters moved with increased temperature, the oscillator often dropped out of oscillation. To rectify this problem, the circuit was modified by adding a small capacitor on the collector of the oscillator transistor. The added capacitance had the effect of moving the negative resistance to a higher absolute value, thereby enabling the oscillator to overcome the losses of the resonator. Furthermore, the resonator was adjusted to more closely center on the negative resistance of the active device.

Second Mixer and Dielectric Resonator Filter

The design team originally thought that the second mixer and dielectric resonator filter would not need many changes other than scaling the line widths and lengths. However, because the second LO is loaded by the second mixer and dielectric resonator, the line lengths become critical (see **Figure 12**).

To prevent loading the dielectric resonator oscillator (DRO) at its center frequency (4.33 GHz), the combined electrical length from reference plane 2 to reference plane 6 should be a multiple of 180 electrical degrees (or half wavelengths). In addition, it is best that the electrical length from reference plane 2 to the mixer diode reference plane 4 be a multiple of 180 degrees at the IF frequency of 4.454 GHz. This is to prevent loading the mixer diode and to create a maximum voltage

Figure 12

The second mixer.

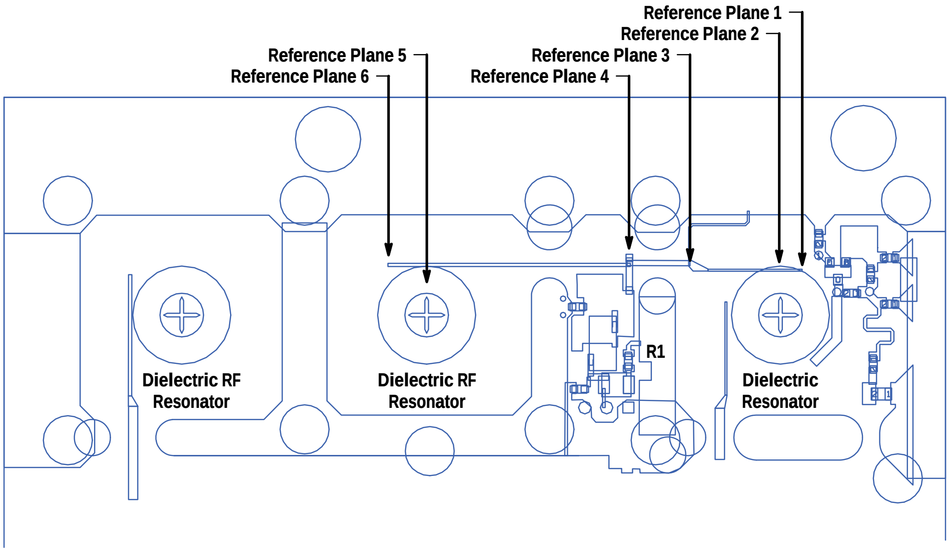

swing across it. Finally, the electrical length from reference plane 1 to reference plane 5 had to be a multiple of 180 degrees at 4.454 GHz (the first IF frequency) to prevent loading the dielectric RF resonator (DRF) by the DRO. The result is that the overall length of the long interconnecting microstrip line between the mixer and oscillator has to be some multiple of half wavelengths (such as 0.5, 1.0, or 1.5).

The original Teflon-based design called for this circuit to have approximately one wavelength from the DRF reference plane (reference 5) to the DRO reference plane (reference 2). This is made up of 0.5 wavelength from reference plane 5 to reference plane 4 and 0.5 wavelength from reference plane 2 to reference plane 4. Since GeTek has a higher dielectric constant than the Teflon substrate used in the old design, and since higher dielectric constants result in shorter wavelengths, it was physically impossible to leave the length of the line at one wavelength. Unfortunately, we didn't realize this until after the first board turn.

Much of the health of this circuit is based on the amount of mixer bias voltage coming from the average dc current induced in the mixer diode by the second LO. In the first turn of the new board, the mixer bias voltages across R1 were approximately 30 to 40 mV. In comparison, the original board had a mixer bias voltage of approximately 100 mV. After this mistake was realized, and an extra 180 degrees of electrical length from reference plane 4 to reference 5 was added to the line, the mixer bias voltage was actually increased to an average of 110 mV.

124.0568-MHz Oscillator

The design of the 124.0568-MHz oscillator was fairly straightforward. It was implemented as a classical phase-locked loop design using a Motorola MC145150 phase-locked loop IC, having a reference frequency of 2.7648 MHz. Since all the data exchanged between the headend and the remotes employs phase shift keying (PSK) modulation of the carrier, this capability was included in the 124.0568-MHz oscillator design. Since frequency equals the derivative of phase with respect to time, a step in phase is equivalent to an impulse in frequency. This phase shift keying modulation was implemented by impulsing the tuning line of the oscillator at a frequency outside the loop bandwidth of the oscillator. Since the VCO's output frequency is a function of voltage, we can accomplish a step in phase by applying an impulse of voltage through a high-pass differentiator to the tuning input of the oscillator (see **Figure 13**).

Figure 13

PSK modulation generation.

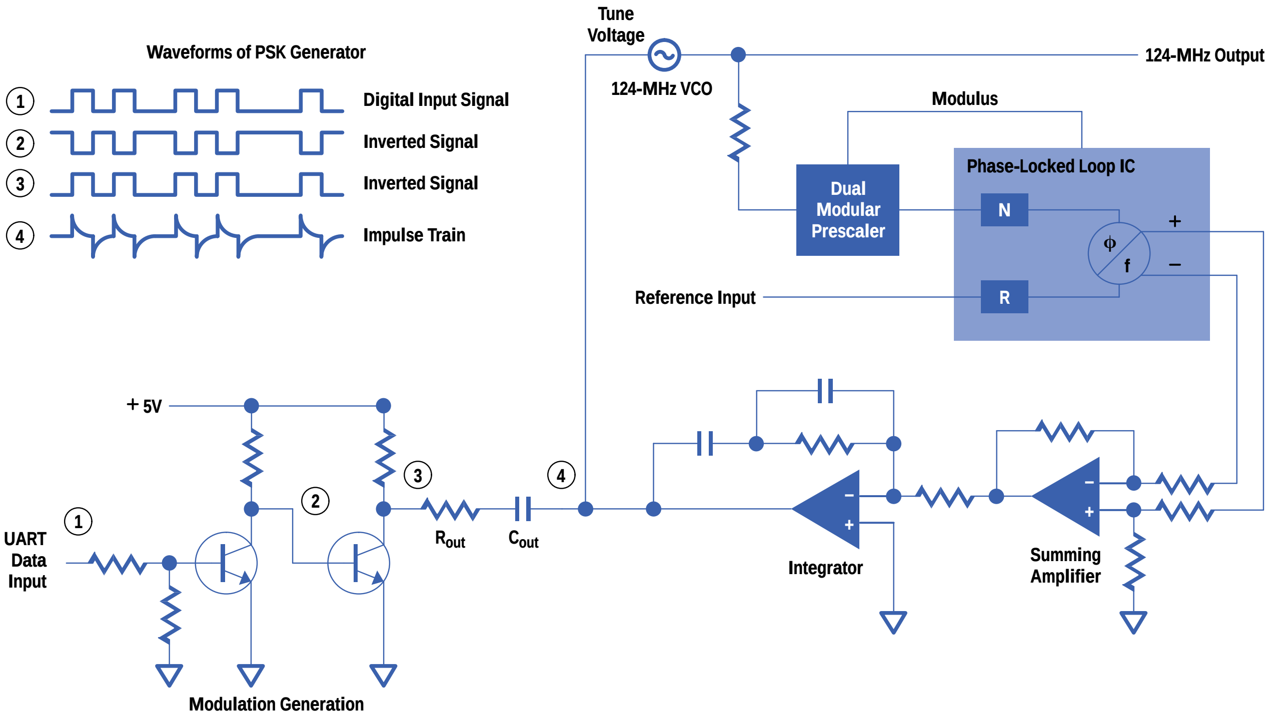

Output ALC Amplifier

The output ALC amplifier was implemented using classical design techniques. This included a variable gain amplifier (HP IVA05208) on the input to provide the variable gain required for leveling. The signal was then passed through two fixed-gain amplifiers and detected using a classical diode detector with a reference diode for temperature stability.

Input Attenuator Match

Since the HP 3010H and HP 3010R sweep/ingress analyzer is used as a signal level meter, a large part of its accuracy depends on the input match. The specification for input return loss of the instrument is 18 dB. The RF module has a specification of 23 dB.* The previous design had some difficulty meeting this specification, and often the modules had to be manually adjusted to reach the desired performance. One goal of this project was to eliminate manual tuning to achieve the input match. This instrument is required for use in cable systems having a 75-ohm characteristic impedance. All of the switches and relays that were available for the instrument design were designed for 50-ohm systems (75-ohm relays and GaAs switches do not exist).

To use these parts in a 75-ohm system required some unconventional matching techniques. In a 75-ohm system, a 50-ohm transmission line (if it is short enough) tends to function like a parallel capacitor. For a superheterodyne receiver, like the HP 3010, it is desirable to have a low-pass filter in the input. By using 50-ohm parts as shunt capacitors in a classical inductor and capacitor low-pass filter structure and high-impedance transmission lines as inductors, we were able to achieve good input return loss and proper filtering. At the initial production release, the input return loss did not quite

---

* The RF module has a tighter specification than the instrument because there is a cable and two connectors between the module output and the instrument output. These cables and connectors degrade the return loss of the module.

meet specifications. However, with some minor repeatable modifications, the new RF assembly aligned much faster than the original RF assembly. The complete RF converter board is shown in **Figure 14**.

Figure 14

RF converter board.

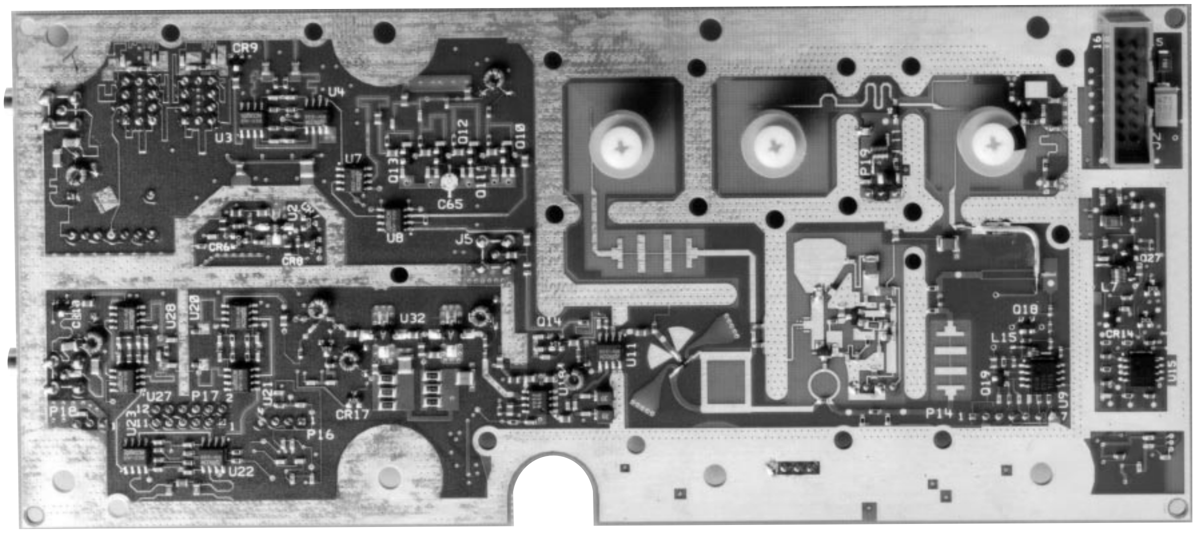

Communication Protocol Development

Another major part of this project was the development of firmware. Three major features were added to the existing firmware:

- Reverse sweep and communications protocol
- Ingress detection
- Digital-channel power measurement.

Reverse Sweep

The primary effort of the firmware development was to create an algorithm to perform a return sweep. Return sweep is used to generate a sweep response plot of the return (subscriber's-home-to-headend) position of the cable system (see **Figure 15**). With this sweep response, technicians can align their amplifiers for the best possible response.

On initial inspection, the problem of communication between a remote unit and a headend unit seemed like a fairly simple proposition, with the communication going something like:

- Headend to Remote: "Hey, are you there?"
- Remote to Headend: "Yes, I'm here."
- Headend to Remote: "OK, what do you want to do?"
- Remote to Headend: "I want to sweep."
- Headend to Remote: "OK, go sweep."
- Remote to Headend: "OK, sync now."

Figure 15

Typical return sweep response.

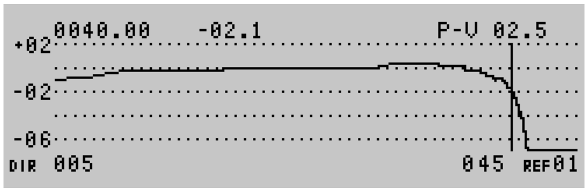

The headend and remote would be synchronized, and the remote would then send out pulses at various frequencies until the end of sweep was reached. At that point, the headend could simply process the data and send it back to the remote. It is easy to see that a simple one-remote system would be easy to program.

Trying to handle multiple users (remotes) with this technique might result in the following communication dialogue:

- Headend to Remotes: “Hey, are you there?”
- Remote 1 to Headend: “Yes, I’m here.”
- Remote 2 to Headend: “Yes, I’m here.”
- Remote 3 to Headend: “Yes, I’m here.”
- Remote 4 to Headend: “Yes, I’m here.”
- Remote 5 to Headend: “Yes, I’m here.”
- Remote 6 to Headend: “Yes, I’m here.”
- Remote 7 to Headend: “Yes, I’m here.”
- Headend to Remotes: “Wait! Wait! Wait! I can’t understand you all at once.”

At this point, communication would break down.

In our development, we chose to take a more civilized approach. Each remote unit is assigned a serial number to enable the headend to identify and communicate with it. In addition, we decided there would be two states for the headend and remote units: *connected* and *unconnected*. In the connected state the headend and remote units are in sweep loops. The headend can be connected to many remotes. It keeps a list that is used to control when each remote is run through the sweep loop. The unconnected state would indicate for the headend that it is not connected to any remotes, and for any remote that it is not connected to the headend.

New users (remotes) would be polled by repeating messages broadcast on the forward pilot from the headend. The forward pilot is a frequency that is set aside in the cable spectrum for communication. The HP 3010R uses a forward and return pilot to complete the communication loop.

When new users are first connected to the cable system, they respond to the new user poll by returning (over the return pilot) their serial number and a cyclic redundancy checksum (CRC) of the sweep table stored in a remote’s memory.* The sweep table checksum response is important because if it does not match the headend’s sweep table checksum when the unit begins to sweep, the headend and remote units will not be synchronized, resulting in bad sweep data. Once a new user has responded to a new user poll and received a sweep message from the controlling headend, it is then considered connected and is added to the headend’s queue.

---

* Sweep tables are used in the HP 3010 SLM system to keep the sweep signals out of the way of visual, aural, and, in some cases, digital carriers.

Since it is possible for more than one unit to respond to a new user poll simultaneously, the headend can use four allotted time slots to listen for responses. To create random time slot selection, each remote in the not-connected state that attempts to respond will choose a time slot based on a pseudorandom number generated from its internal clock. Since the likelihood of any two remotes being set to the same time is very slim, a fairly equal probability exists for a given remote to respond in any one of the four time slots. Because the headend can only respond to one time slot at a time, the first remote unit's time slot that contains a message will be the next unit to be connected. To conclude the connection, any responding remote that receives a sweep message from the headend will be considered connected.

Ingress Detection

Since the influx of noise (ingress) is a problem, we wanted to add a feature that would allow remote field technicians to see the level of ingress being received at the headend (or at a hub site). In addition, it is important for field technicians to be able to see this ingress even if the remote unit can no longer communicate with the headend. To accomplish this, we decided to have two modes of ingress operation: *return spectrum* and *broadcast ingress* (see **Figure 16**).

Figure 16

Ingress detection.

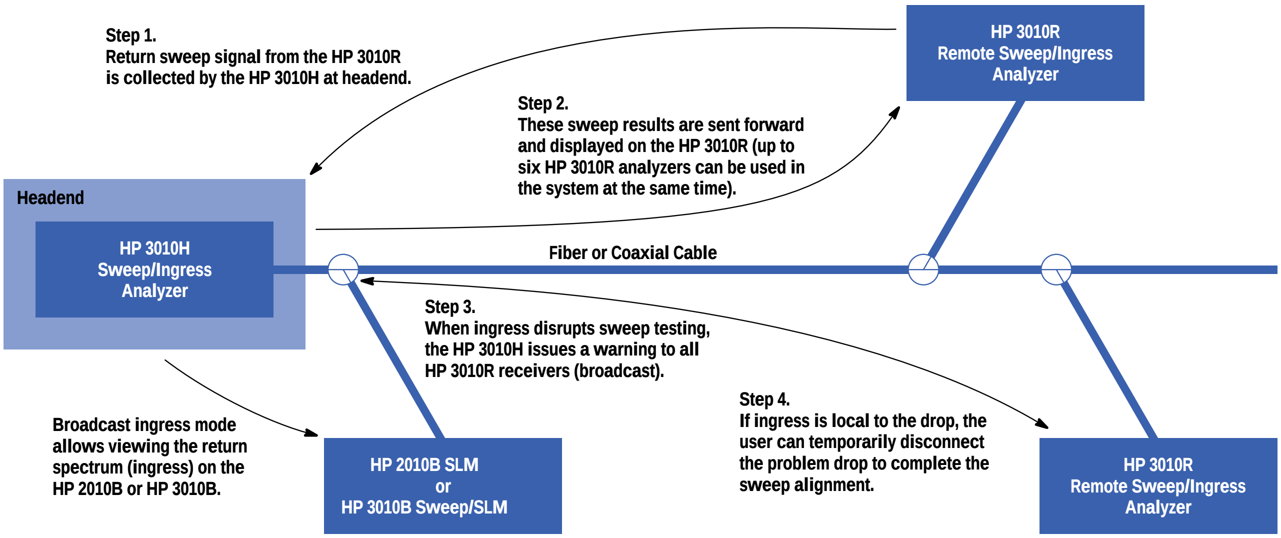

Return spectrum, as the name implies, is nothing more than a snapshot of the return spectrum taken by the headend and sent to the remote unit over the forward path. The return spectrum measurement, like the return sweep measurement, is a demand measurement (the remote unit demands the measurement). It is very similar to the return sweep measurement, except that the remote unit does not send sweep pulses. As a consequence, no synchronization pulse is required. When the remote unit demands a return spectrum measurement (during the synchronization message) it waits in the listen mode. The headend then performs its usual new user poll, but instead of listening for the sweep pulse from the remote unit, it measures the noise ingress. Upon completing the return spectrum measurement, the headend sends a data message to the remote unit that requested the return spectrum measurement.

Broadcast ingress, on the other hand, is not a demand measurement. Broadcast ingress measurements can only occur if the headend is in continuous mode, or if it receives noise above a certain level while trying to receive a message from the remote unit. When this occurs, the headend sends a broadcast message containing the return spectrum data to all remotes. When the remote units receive this message, they add a softkey to their display and store the data in an array.

The user can then view the data at any time by simply pressing the softkey. When the headend is set to continuous ingress mode, it sends a broadcast ingress message every other sweep cycle.

Digital Channel Power Measurement Algorithm

Because many cable operators are now beginning to offer digital services, and since most digital services are broadcast using some sort of digital modulation scheme (such as QAM, QPSK, and so on) that results in a broadband noise-like signal, it is no longer an easy matter to measure channel power. Digital signals tend to look like noise pedestals in the frequency domain (see **Figure 17**). A new algorithm has been added to the HP 2010 and HP 3010 firmware to enable the user to measure digital channel power easily and accurately.

Figure 17

Digital channels.

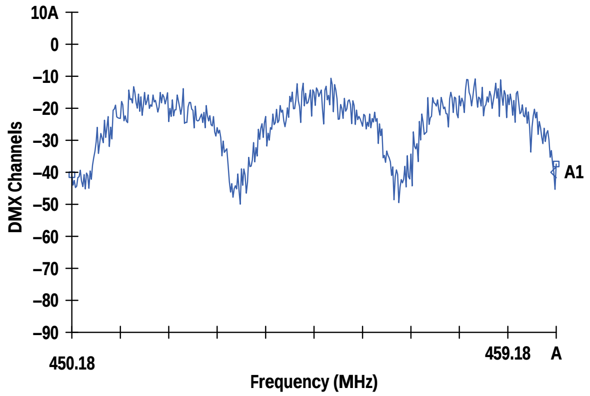

An algorithm developed for the HP 8591C spectrum analyzer has been leveraged into the HP CaLan 2010 and 3010 signal level meters. This algorithm is based on the fact that, at any given center frequency, the power level detected at the output of the log amplifier is a time sample of the total power contained within the resolution bandwidth of the receiver. That is, it equals the bandwidth multiplied by the noise power density in that band (watts/Hz). As an example, with a noise power density of -100 dBm/Hz, the total power contained within a 230-kHz bandwidth would be -46.383 dBm. (This calculation must be performed with linear power. To do the conversion in dB, $-100 + 10\log(230 \times 10^3) = 46.383$.) Since the resolution bandwidth of the HP 2010 and HP 3010 modules is approximately 230 kHz, to measure an 8-MHz channel requires at least 35 steps (8 MHz/0.230 MHz). To calculate the total power in a channel, the receiver is tuned to the low end of the channel and then stepped across the entire band in steps narrower than the bandwidth of the filter. At each frequency step, the power is converted to linear, averaged five times, weighted by the overlap in the step size, and added to the other frequency data. After the final frequency is reached, the total power is converted back to dBmV for display.

The accuracy of this measurement depends on two major factors: resolution bandwidth accuracy and the ability of the detector to operate in sample mode. Since the shape of the resolution bandwidth filter is not easy to model mathematically, it is necessary to define the filter's noise equivalent bandwidth. This can be defined as the effective ideal bandwidth of the filter. It can be calculated by finding the area under the normalized curve of the filter. For example, an ideal brick wall filter would have a normalized amplitude response of one and a bandwidth of $\Delta\omega$. Thus, the brick wall filter has a noise equivalent bandwidth of $\Delta\omega$.

Since the resolution bandwidth filter is not consistent from unit to unit, we devised a way to calibrate the noise equivalent bandwidth of the HP CaLan 2010 and 3010 modules. This is accomplished by turning on the internal calibrator and stepping the first LO in fine increments around the center frequency. The responses are then summed in a trapezoidal integration to give the resultant noise equivalent bandwidth.

For accurate measurement of a noise-like signal, it is necessary for the detector to be in sample mode. This is very important because it is impossible to know what the peak to average power level ratio is for any given type of modulation scheme. Fortunately, it is possible to run the detector in a sample mode. The final algorithm gives results that correlate well with an rms power meter and with the HP 8591C digital channel power downloadable program.

Other firmware changes included upgrades for printer support and improvements to the user interface to accommodate the newly added features.

Conclusion

The new HP CaLan 2010 and 3010R and 3010H sweep/ingress analyzers were completed quickly to fill the urgent needs of cable system operators to sweep and align the forward and return paths of their cable systems. The final product is easy and intuitive to use and includes many system improvements. This was achieved by setting and maintaining clear priorities from the very beginning. In the order of precedence, these priorities were schedule, quality and reliability, performance, and cost. By focusing on the priorities, we were able to introduce a new product within a very short time.

Acknowledgments

Many people were instrumental in getting this product to market in what may be record time for the Microwave Instruments Division. The lab team made up of Bill Morgan (overall project leader), Brenda Coomes (firmware project leader), Bill Scales (digital design), Cinda Craven (mechanical hardware design), Marv Milholland (remote firmware design), Chin-Min Wang (headend and digital power firmware design), and John Mullan (user interface and printer support firmware design) were all instrumental in the creation of this product. The documentation team made up of Jerry Green and Al Barnes did an excellent job of creating a comprehensive set of manuals. The manufacturing team including Dave Wolbeck (hardware support and RF board test development), Jim Stock (hardware support and RF board test development), Dennis Nishi (production engineering RF board), Ken Silk (production engineering instrument test development), Phil Melman (production engineering support for everything else), Adam Pirog (instrument line technician and all-around troubleshooter), Kim Holliday (industrial design), Charlie Wolfe (production engineering mechanical engineer), and Patty Hicks (line supervisor) were instrumental in reaching our shipment and pilot run goals. The new product introduction group consisting of Vic Sallee (new product introduction purchasing), James Ellison (design drafter), Steve Comins (SMTC test support), Linda Morrison (new product introduction coordinator), and Julie Silk (SMTC liaison) were also quite helpful in pulling things together.

Reference

1. W. S. Ciciora, *Cable Television in the United States—An Overview*, Cable Television Labs, Inc. 1995 (***http://www.cablelabs.com/C-IT/ciciora.html***).

Online Information

For more information about the product described in this article, take a look at the information located at the following URLs:

- http://www.tmo.hp.com/tmo/datasheets/English/HPCaLan.html
- http://www.tmo.hp.com/tmo/datasheets/English/HP3010H.html
- http://www.tmo.hp.com/tmo/datasheets/English/HP3010R.html
- http://www.tmo.hp.com/tmo/datasheets/English/HP3010B.html
- http://www.tmo.hp.com/tmo/datasheets/English/HP2010B.html

▶ Go to Journal Home Page