LESSONS FOR THE INTERACTIVE RETURN SYSTEM

Dr. Kerry D. LaViolette
Philips Broadband Networks, Incorporated

Mr. Tom Staniec
The Excalibur Group, a Time Warner Company

Abstract

Activating and using the CATV return system as part of a profitable communications network has inspired tremendous work, discussion, and demonstration. This activity level tends to concentrate on known problems with subsystem components but often loses direction toward a complete system architecture. This paper presents laboratory results for a full system implementation as well as for subsystem components. Then data on impairments from real implementations is presented. Together these allow the larger system view while focusing on key areas for success.

INTRODUCTION

The interactive return system is the emerging technical challenge in CATV architecture. The conceptualization and implementation of a functional system requires a thorough understanding of the desired goals, the subsystem components, and the interaction between these components. In addition, the ability to overcome external impairments (like ingress) and the alignment/maintenance of the interactive network pose significant challenges.

This paper concentrates on the most prevalent modulation implementation in the return system. It is no surprise that Quadrature Phase Shift Keying (QPSK) data transmission is the modulation scheme of choice for most current implementations. This modulation scheme is extremely tolerant to impairments based on its low modulation order and use of FM techniques for data transmission [1]. Indeed, 20 T1-rate (1.544 Mbits/sec) carriers can easily fit within a 10-40 MHz bandsplit providing essentially a is a telephone line for every home on a node without any blocking! Therefore QPSK is not only robust, but also provides enough bandwidth for today's applications.

This paper begins by showing test results of return amplifiers as independent units while also introducing to the testing methods, equipment, and graphs that are used. The next section looks at return lasers, both Fabry-Perot (FP) and Distributed Feedback (DFB) type. Following this is a look at a cascade performance built with currently available equipment. The cascade not only shows the complete system performance, it also shows results over temperature.

The next two sections are devoted to practical impairments and draw on experience from actual field trials. Both sections focus on ingress with the first as a discussion of actual field experience. The next looks at potential mitigation strategies, followed with a conclusion on the lessons learned.

RETURN AMPLIFIER PERFORMANCE

Before showing the performance of a return amplifier, the testing methods need to be explored. Figure 1 shows the standard test set along with the second-order and third-order products that fall within the 19 MHz test signal. Twenty-one T1 QPSK modulators are used as inputs to the device under test. These modulators are spaced from 5 to 40 MHz at 1.75 MHz spacing, while the tested frequency is at 19 MHz. The frequency plan is graphically shown in figure 2. The level into the device under test (DUT) is adjusted by attenuators and (if needed) gain blocks, while the output is filtered and adjusted to keep the input to the QPSK demodulator within its AGC range. The system is fed by a Bit Error Rate Test set (BERT) at 1.544 Mbit/s with a 2^{23}-1 PsuedoNoise (PN) sequence. A PN sequence of bits simulates a random pattern, but eventually repeats. In this case the pattern repeats every 2^{23}-1 bits. The QPSK modulator/demodulator pair, at the 19 MHz test frequency, was specifically selected because it operates within 1 dB of theoretical performance. *Therefore, the device or system is being characterized rather than modulator/demodulator design.* In addition, any gain blocks used were also tested to make sure they were not impacting the performance of the DUT.

When performing this characterization, two factors are considered. The first is noise performance. At low input levels the DUT noise will produce errors, referred to as the noise side or left side of the curve. The second consideration is high input levels. At high input levels the distortions generated will fall under the tested carrier eventually rising enough to cause errors. For QPSK and items tested in this paper this behavior happens only after the DUT goes into compression or clips. This is referred to as compression or the right side of the curve. These two curves create a washtub graph that shows the input dynamic range that can be used for error-free performance. For this paper Bit Error Rates (BER) below 1×10^{-8} are considered error-free. This error-free dynamic range is critical because it sets the levels of operation that overcomes noise and ingress while maintaining margin for environmental and system related impairments.

This technique was applied to a typical return amplifier, and the results are shown in figure 3. Keep in mind that this graph is similar to following graphs so the descriptions, provided here, need only to be stated once. The X-axis of the graph is RF power (dBmV) per channel as presented to the device input. The device input is a valued reference when aligning the return path [2] and will be used for all the following graphs.

The washtub curve is clearly evident. The noise side has two different lines plotted. The first is with just one QPSK (the 19 MHz test channel) present, while the other has all 21 QPSK channels on. For this device there is no discernible difference between the two conditions and since it is a thermal noise limited device, this similarity is expected. At $1x10^{-8}$ the dynamic input range is nearly 68 dB. This is an extremely wide operating window; however, it will be narrowed by noise funneling in a cascade situation. This narrowing will be shown in the system test section.

Also, it should be noted again that errors are not observed on the right side curve until the device is in compression. The compression sets the absolute upper limit that the input of the amplifier can tolerate and along with any input padding determines the highest level available in the coax. This in turn sets the carrier to ingress level. Input padding is a recent approach to improve tolerance to ingress. Aligning the system for unity gain by padding at the input of the following amplifier allows the highest RF levels possible in the coaxial cable, thus maximizing the signal to ingress level.

RETURN LASER PERFORMANCE

This section shows the performance of typical lasers and/or transmitters that can be acquired for systems today. Discussion continues over which type of device, either an FP or DFB, should be used. The comparison that was made in [3] still applies. A DFB can inherently give better performance in terms of link length and overall operating widow. However, this performance is not free. Coolerless DFBs have been sampled to the market place and may narrow the cost differential. The choice, as always, is performance versus cost and can only be decided when one knows the system needs, the upgrade path, and the upgrade timetable. This section will show the relative performance of these devices and point out some practical considerations that have not been previously mentioned.

Figure 4 shows the performance of a typical FP transmitter with the same testing methodology mentioned in the previous section. The device was biased at 0.5 mW and transmitted over 6 dB of singlemode fiber. Although a return receiver was used, it plays no part in the performance and, therefore, will be regarded as transparent for QPSK data. However, the return receiver does have a role: its gain adjustment is used to maintain the proper level for operating in the demodulator's AGC range.

As we look at the figure we see a well defined operating window of 30 dB. The right side of the curve was determined by laser clipping while the left side of the curve is determined by laser noise. Also, note that again there is no discernible difference between single-channel and multiple-channel operation. This noise is sporadic in nature. Note that the slope of the curve is not steep as would be expected for white noise. This sporadic noise is caused by laser chaos when Rayleigh backscatter from the fiber plant is presented at the FP device[4].

Clearly, this transmitter performance is poorer than the single return amplifier. However, 30 dB is an extremely wide operating window and should be sufficient for practical systems, as will be shown in a following section. It should be noted for future reference that this transmitter design has only a minor shift over temperature for QPSK data.

Since this performance is lower than that for a return amplifier, methods to improve it were considered. One method was to ask a supplier to change the design parameters of the FP. It must be recalled that these devices were designed and optimized for ON/OFF digital performance and although this allows "technology piggybacking," it may be a suboptimal design for QPSK carriers. Figure 5 shows a redesigned device with two sets of curves. The first has an output of 0.5 mW. This device' 10 dB improvement over the device in figure 4 is significant and is on the noise side. The slopes of the curves are now extremely steep, so white noise performance has replaced sporadic noise performance.

A practical consideration is shown in the second set of curves. Another "improved" FP laser was biased at 2 mW and then tested over 6 dB of fiber. This device was rated to 4 mW, and the thought was to increase the right side of the curve by increasing the clipping point. Although this step was accomplished, the noise side of the curve was extremely degraded. The dynamic input window on this device is 30 dB, 10 dB less than the 0.5 mW improved FP! The additional current used to bias the device 6 dB higher has allowed additional laser modes. These additional modes increase the laser mode partition noise, as they compete for gain, causing the degradation. This outcome demonstrates the importance of defining and characterizing both sides of the washtub curve.

A return DFB is shown in figure 6. This device was biased at 2 mW and is both cooled and isolated. The transmission distance was 6 dB. The device performs as expected with an operating window of 45 dB, which is 15 dB better than the typical FP in figure 4 but only 5 dB better than the improved FP, biased at 0.5 mW, in figure 5. The noise curve is white noise limited, and there is no difference between a single carrier and a multiple carrier.

A practical matter is illustrated by the DFB in figure 7. This device was operated over 11 dB of fiber and performs as expected with regard to clipping. However, the noise side shows an interesting dichotomy. There is a large difference between single-carrier performance and multiple-carrier performance. In fact multiple-carrier performance is 15 dB better! The supposition is that the device has a quieter noise level when it is operated under a RF load; multiple carrier operation is, therefore, better. However, in practical systems the dynamic input range would be limited to 30 dB which is the difference between the single-carrier noise performance and clipping. This reasoning assumes that the system allows dynamic carrier allocation and, therefore, there will be times when only a single carrier is operating. This can be mitigated by forcing some carriers to always be operational; however, these carriers take away useful information capacity. Again, we see the importance of defining both sides of the washtub curve and not assuming what is a worst-case scenario.

CASCADE RESULTS

Using the return amplifiers and standard FP transmitter from the previous sections, a cascade was built. The cascade, shown in figure 8, has 5 amplifiers in cascade with additional noise added to simulate 50 amplifiers in the node. The return RF cascade was aligned for unity gain and a flat response. Its characteristic curve, at room temperature, is shown in figure 9. This figure shows that the cascade's dynamic input window is 50 dB, which is 18 dB less than a single return amplifier, with the change due almost entirely to noise funneling.

Since the input window is still larger than that of the FP transmitter and when compared to figure 4 has different RF levels for compression, alignment for unity gain from the last return amplifier to the return transmitter did not make sense. Therefore, alignment was made so that the compression levels were the same. This was facilitated by padding the input to the return laser transmitter. With this alignment the entire transmitter dynamic range is available while using the upper portion of the cascade. **Therefore, the signal-to-ingress level is maximized in the RF section where ingress occurs.** Due to the very nature of fiber transmission, no additional ingress occurs in the glass, so the carrier-to-ingress level is set before the optical transmitter.

The performance of the cascade at room temperature is shown in figure 10 and is referenced to the input of the first return amplifier. Due to the alignment described in the previous paragraph, the full system performs with a 30 dB window. The dynamic input range's width, as expected, is the same as the FP transmitter and is still extremely wide, providing room for system operation. The system was cycled over temperatures from -35 °C to +55 °C.

The resultant worst-case performance is shown in figure 11. The dynamic input range is reduced to 20 dB. An investigation into this reduction showed that the laser transmitter performance was nearly stationary over the entire temperature range and contributed very little. However, the changes in gain and cable length in the RF section were significant. The right side of the curve is set at low temperature, where gain is the highest and therefore the laser clipping point is reached more quickly when referenced to the cascade input. The right side of the curve is set at high temperature, where the additional cable lowers the signal level at the transmitter, bringing it closer to the noise.

In summary, the cascade has a dynamic input window of 20 dB over temperature. This window is still very large, and system operation can be robust. Work in thermal gain control, raising the RF amplifiers compression point, and widening the return transmitter's input range will make the system even more robust.

DISCUSSION OF INGRESS

Ingress, the unwanted by-product of poor shielding integrity in coaxial networks, poses a problem that has to be reconciled before a network can function as a reliable communications platform. This is really a very simplistic statement about a fairly complex problem. Fifteen years of history return networks proves that you can have an operational network that will support telecommunications. However, what is required are a sound view of the problem's principally origin and a steadfast attitude on how to eliminate, mitigate or diminish its impact on the network.

First, how does ingress get into a network? By empirical observation, 70% of all ingress comes directly out of subscribers' homes. That number could be higher, or possibly lower (doubtful but possible), as indicated in [5]. Further, the drop cable from the tap port to the ground block at the house contributes another 25% to the problem. An interesting study, done by CableLabs originally for digital network quantification of high-order modulation schemes, dealt with RF drop shielding as measured in the 88 to 108 MHz FM band. A number of conclusions were drawn from study, but the most significant is up to 60% of the drops hanging in the air today have shielding effectiveness of less than 50 dB! This is a statistic with significant implications. It points to signal leakage in the return portion of the plant where previously problems have not been seen. The remaining 5% of ingress problems are in the physical hard coaxial plant and come from critters, craft and catastrophes.

Examining the home and the subscriber link, one issue is immediately apparent. By Federal Communication Commission rule, the MSO has no right to deal with the internal house wiring once it has crossed the wall boundary of home unless it causes a leakage problem. That mandate puts the operator in a difficult position: how to deal with a problem coming out of the subscribers home but which does not affect forward network operation, cause a Cumulative Leakage Index (CLI) problem, or impact anything the customer maybe doing internally to the home. If the customer has a paid up and current bill, the operator can be and *has been* denied access and be left without recourse.

More recent work by CableLabs shows that RF shielding is a major source of high-level impulsive strikes in a network. Figure 12 is an extreme example of this problem. The spectrum analyzer print shows the effect that a thermostat, used in an apartment complex's electric baseboard heating, has on a return network. Obviously the picture is not a pretty one to start with and is compounded by the intermittent nature of the problem.

The above deals with problems internal to the home, but the greatest problem in the network is the subscriber. Contemplate how many times your system personnel have had to troubleshoot direct pick-up problems, poor video performance and numerous other subscriber complaints, which subscribers caused themselves on there own. A sobering thought considering this is an obvious chink in your network armor.

The drop cable from the tap to the home suffers from many of the same problems, but the subscriber is not the culprit. The crafts person, cable age and weather are the greatest threats in this area. These should be manageable because a bad drop can be repaired; unfortunately, it is less often managed than it is found by accident. CLI and leakage ride outs are touted as the equalizer. However, a large percentage of systems never get ridden because they are off easement. Thus, problems show up as a "blip" on leakage test equipment screen. Roughly speaking, the drop wire from the pole throughout the house covers as many miles as your hard coaxial plant. In a very large number of cases this "plant" never gets measured.

Additionally, fly-overs are promoted as a way to feel secure. Contemplate the following: a recent "blip" was closely investigated and a leak of 1000 μV/m was found about 400 feet off the road easement. It was caused by a maintenance man "making" a five-way splitter to serve new television sets. During the same time, a fly-over took place and no leak was registered. Why? The building had a steel roof and structure. The ingress caused by the "splitter" was causing a problem in the return network. This situation took place in a system which operators believe is a tight network because of CLI and a their fly-over results.

A closer look at ingress in the forward system plant shows that signal leakage, in the air navigational portion of the band, is in large part created by short wavelength slot aperture antennas. Return ingress is made up of a long wire antenna that, particularly in the low end of the return band, acts like AM radio. It is well known that attaching a piece of wire from an AM receiver to a cold wire pipe increases your reception capability. The same thing happens with the shield on coaxial cable and house electrical wiring. This implies that you are dealing with over- the-air radiated signals and ground-wave radiated signals. A look at the construction of an AM radio station or a short-wave broadcaster such as Voice of America further illustrates the situation. The result is an opportunistic signal that gets into the system by poor shielding characteristics. By interpolation, 95% of all problems come from the coaxial cable from the tap down through the home where the shielding integrity is the worst and where CLI is not performed with any consistency. Thus, it is no mystery that a CLI that can meet or exceed specification yet not give you a clean return system. This is a simplified view of the nature of ingress but one that is accurate.

PRACTICAL SYSTEM OPERATION

Fixing the ingress problem and setting up the network are directly related to each other. If the network is set up improperly, it may not be operating as far above the noise floor and ingress as it could. What options are available in the network to overcome noise and ingress? There are three possible methods. First, talk louder than the ingress giving yourself headroom above the problem. Second, repair the ingress to knock it down, thereby increasing the headroom. The third choice is to use a combination of the previous two. All of the above assume a modulation scheme that has a good chance of surviving in the return environment.

As discussed in the previous section, the nature of ingress largely centers on poor shielding. Objectively, there are two ways to deal with the problem: filter it or fix it. Filtering is the single most effective method to clean up and control the problem. Fixing the problem does not ensure it is gone but only momentarily mitigates it. As can be seen, using this approach never takes control of the network. No control leaves the network continuously vulnerable to whatever happens behind the closed doors of the subscribers' homes.

Utilizing the filter approach in number two or three above has the effect of either killing ingress in total or controlling (the lossy-filter approach) the total level allowed into the return network. The best possible position for either type of filtering is at the tap port. Positioning the filter there removes the need to repair all ingress problems in a node caused by the drop system through to the house. Interestingly, while it is often stated that the return network there is 35 MHz of bandwidth, without some type of ingress mitigation strategy, a system may have only 20 MHz or less that is functional. That significantly impacts the total number of services that the network can provide.

What can be done to increase the bandwidth and provide a quality of service that will provide long-term success? The best method is a combination of approaches. Fundamentally, it must be recognized that network must be owned by the provider and protected from all interlopers who are either ill-informed, accidental or malicious. Second, control of ingress is paramount and works in concert with the above strategy. Once the above are realized, then operational methods can be devised. Primarily, the return network should be set up to reference flat inputs to the return amplifiers and at the highest level possible without operating in compression. By input-referencing the return amplifiers, as opposed to the forward amplifier output-referencing, maximum carrier-to-noise and distortion can be achieved. By properly selecting a modulation scheme, a network can be established that is capable of providing high-performance telecommunications.

Logic indicates that increasing the output-level capability of the communications device in or on the home will provide a way to build more headroom and margin into the network. While desirable, that solution opens the door to the problem in drops with less than 50 dB RF shielding. In short, a signal leakage problem will arise in a frequency band where it was never experienced in the past. A recent letter[6] points out the shielding problem, only in reverse. With amateur radio transceivers having a 0.5 μV/m receiver sensitivity and high transmitter output power, a significant problem can materialize both from ingress and egress in the return network. The proper operation of the network requires control, maximized input reference levels, minimized ingress levels and resilient RF modulation schemes.

CONCLUSIONS

What lessons can be gleaned from this discussion? First, today's subsystem components, when configured in a typical system, allow a wide operating window for return system performance, this includes FP transmitters as well as current return amplifiers. Current work in thermal control, amplifier-compression point, transmitter design (both house modulator and laser) will only strengthen this system. Second, in order to operate a network, the provider must own the network. This ownership responsibility must then include must be taken by a combination of "protective" strategies, such as filtering, maintenance and the provision of gateway access. Indeed, a fundamental shift in system design is required for network operation. The ground-up strategy, that has been so successful, for broadcast video must be changed to a network-down strategy, thus concentrating on the most critical issues without allowing subsystem "solutions" to adversely affect network operation.

REFERENCES

[1] LaViolette, et. al., "System and Design Considerations for the Cable Return System," Proceedings of the NCTA, May 1994.

[2] Staniec, "Return Path: Perspective from Field Trial Experience," Cablelabs Seminar, HFC Return Path: Paving the Way, Dallas Feb. 1996.

[3] LaViolette, "Fiber by Design," Proceedings of the NCTA, May 1995.

[4] LaViolette, "The Impact of Rayleigh Backscatter Induced Noise on QPSK Transmission with Fabry-Perot Lasers," IEEE Photonic Technology Letters, May 1996.

[5] Cable Television Laboratories (CableLabs), "Return Path Characterization."

[6] QST Magazine, March 1996

Test Diagram

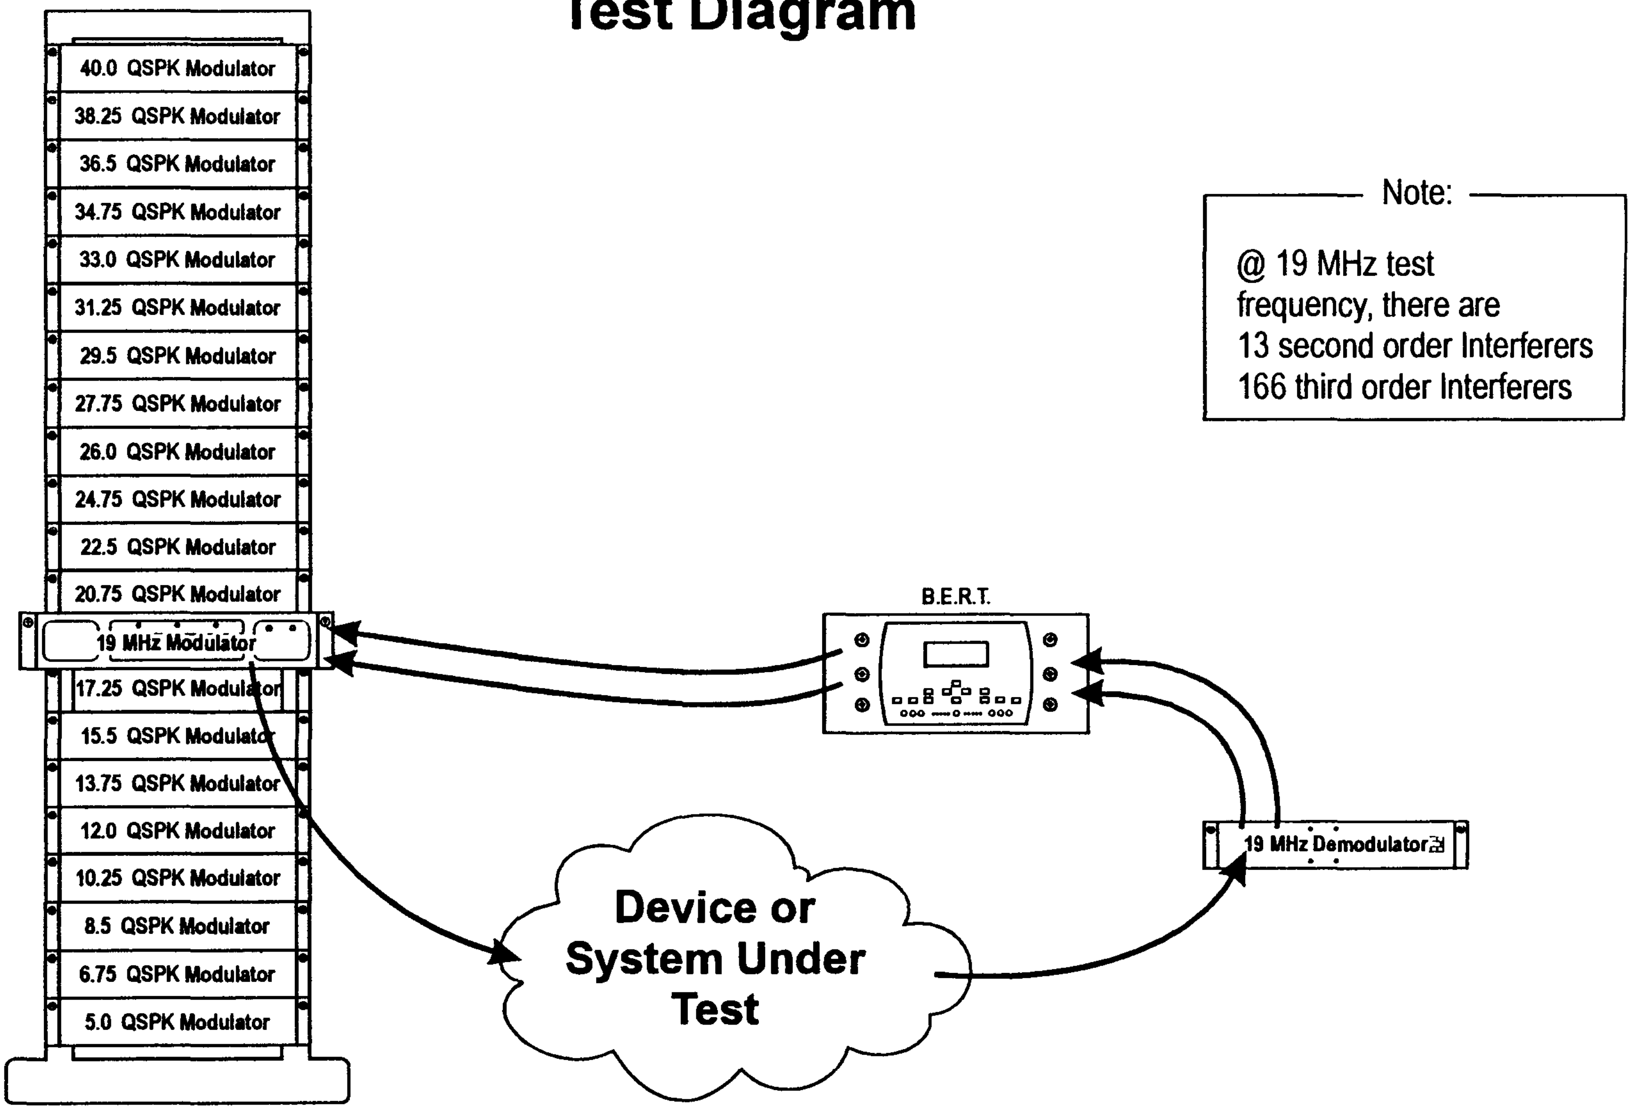

Figure 1

Frequency Plan

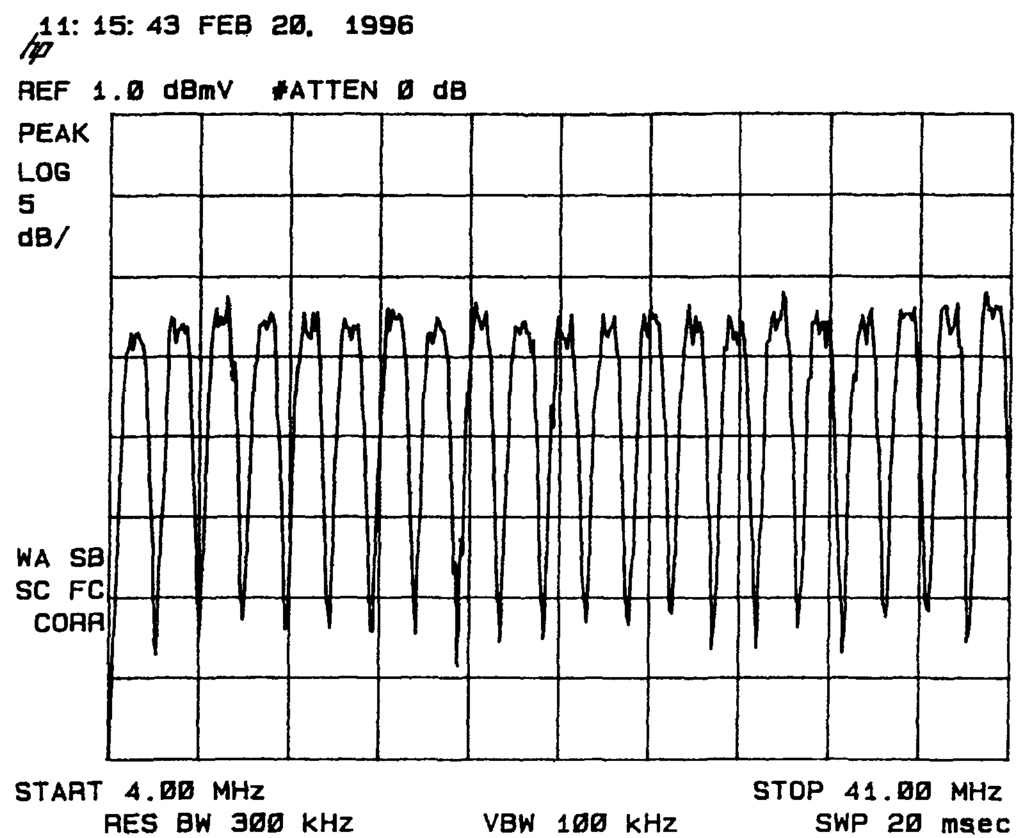

Figure 2

RETURN AMPLIFER TYPICAL PERFORMANCE

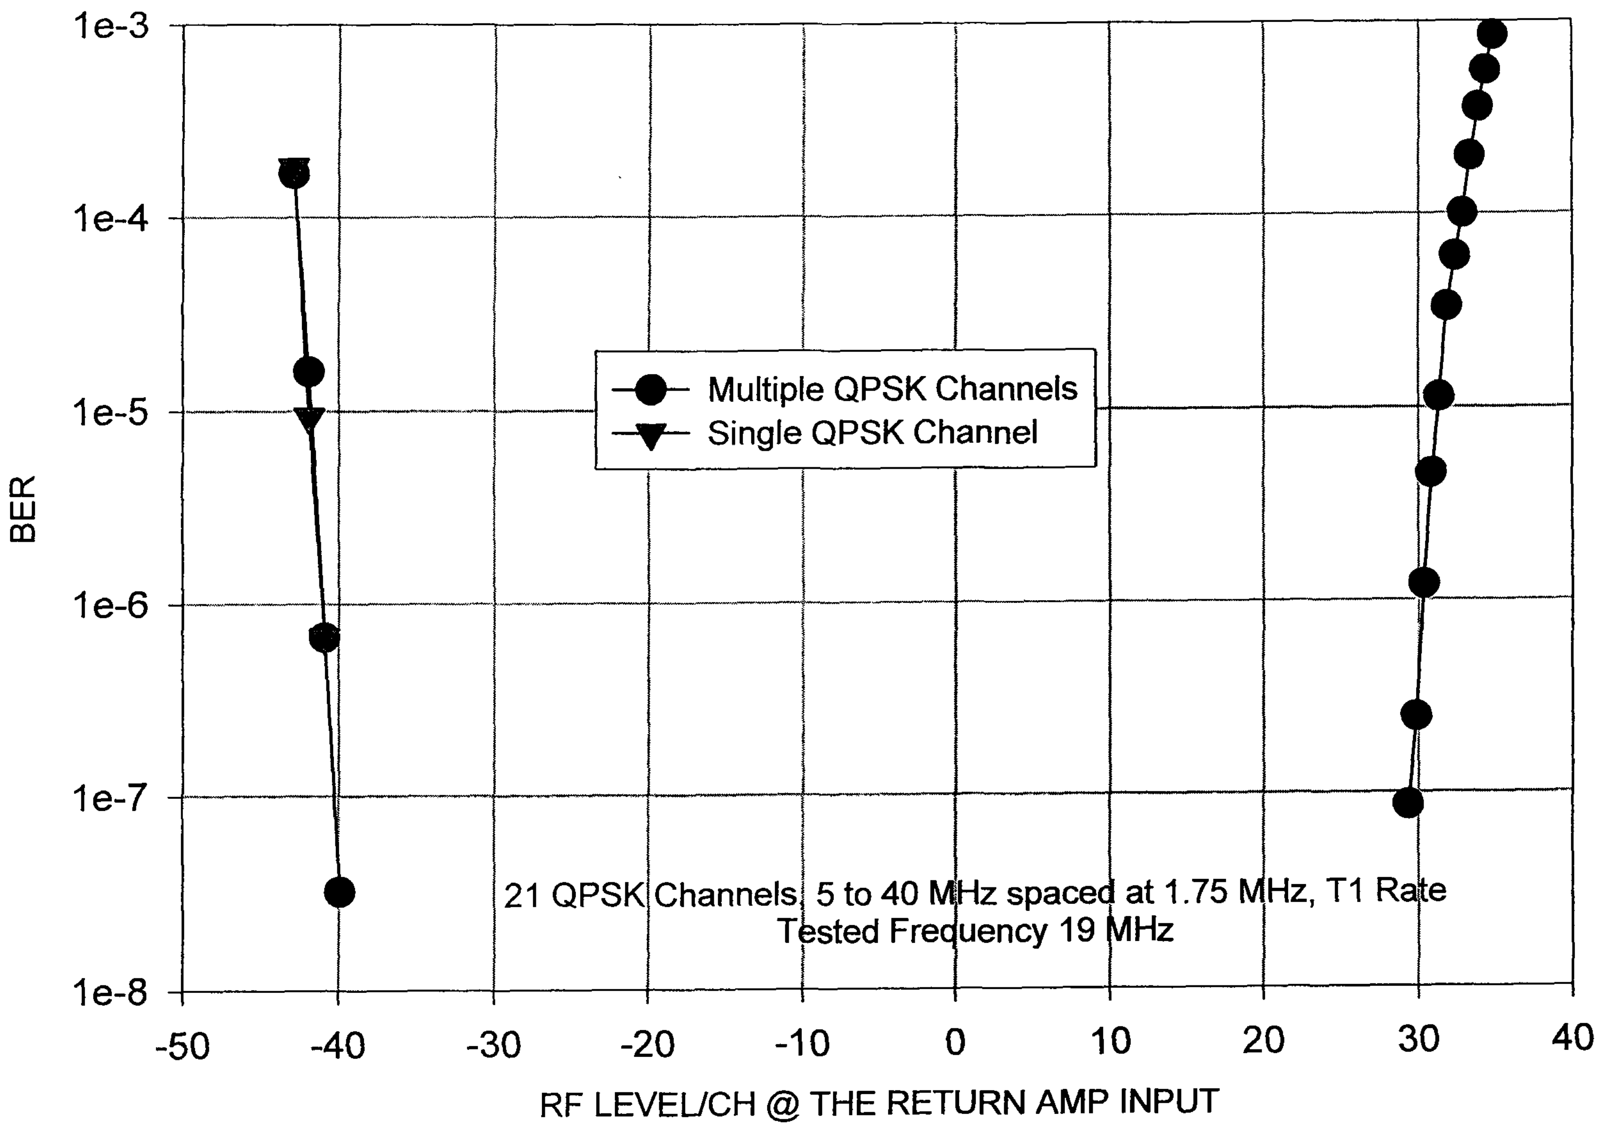

FIGURE 3

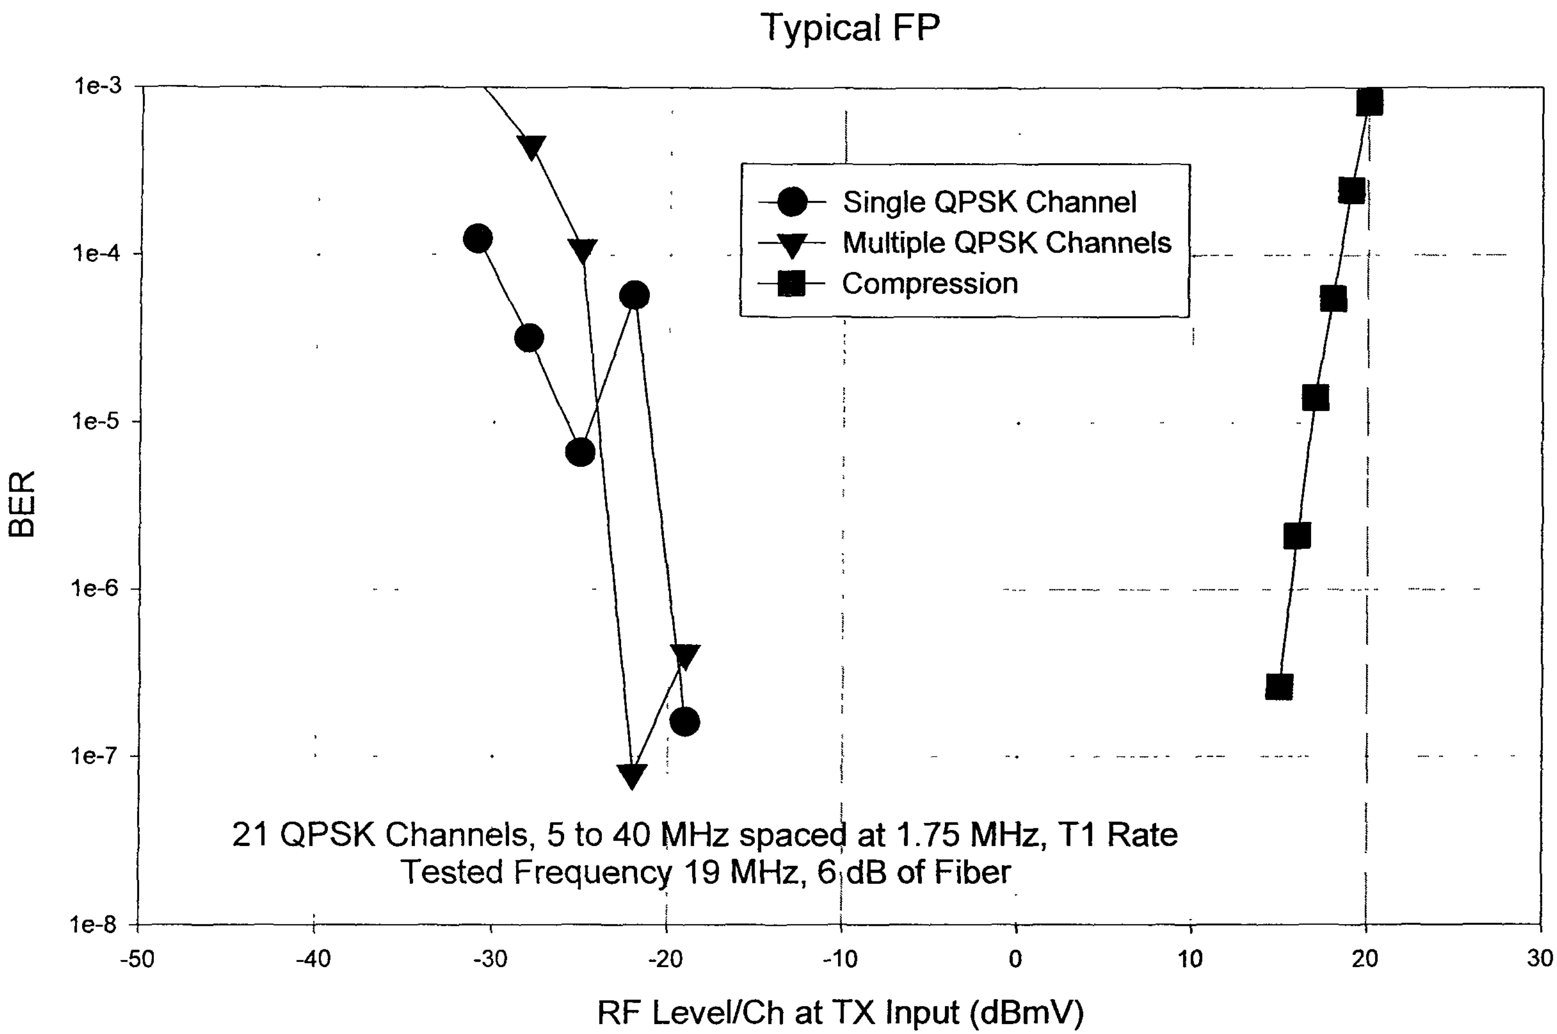
Typical FP
1e-3
1e-4
1e-5
1e-6
1e-7
1e-8
BER
Single QPSK Channel
Multiple QPSK Channels
Compression
21 QPSK Channels, 5 to 40 MHz spaced at 1.75 MHz, T1 Rate
Tested Frequency 19 MHz, 6 dB of Fiber
-50
-40
-30
-20
-10
0
10
20
30
RF Level/Ch at TX Input (dBmV)

FIGURE 4

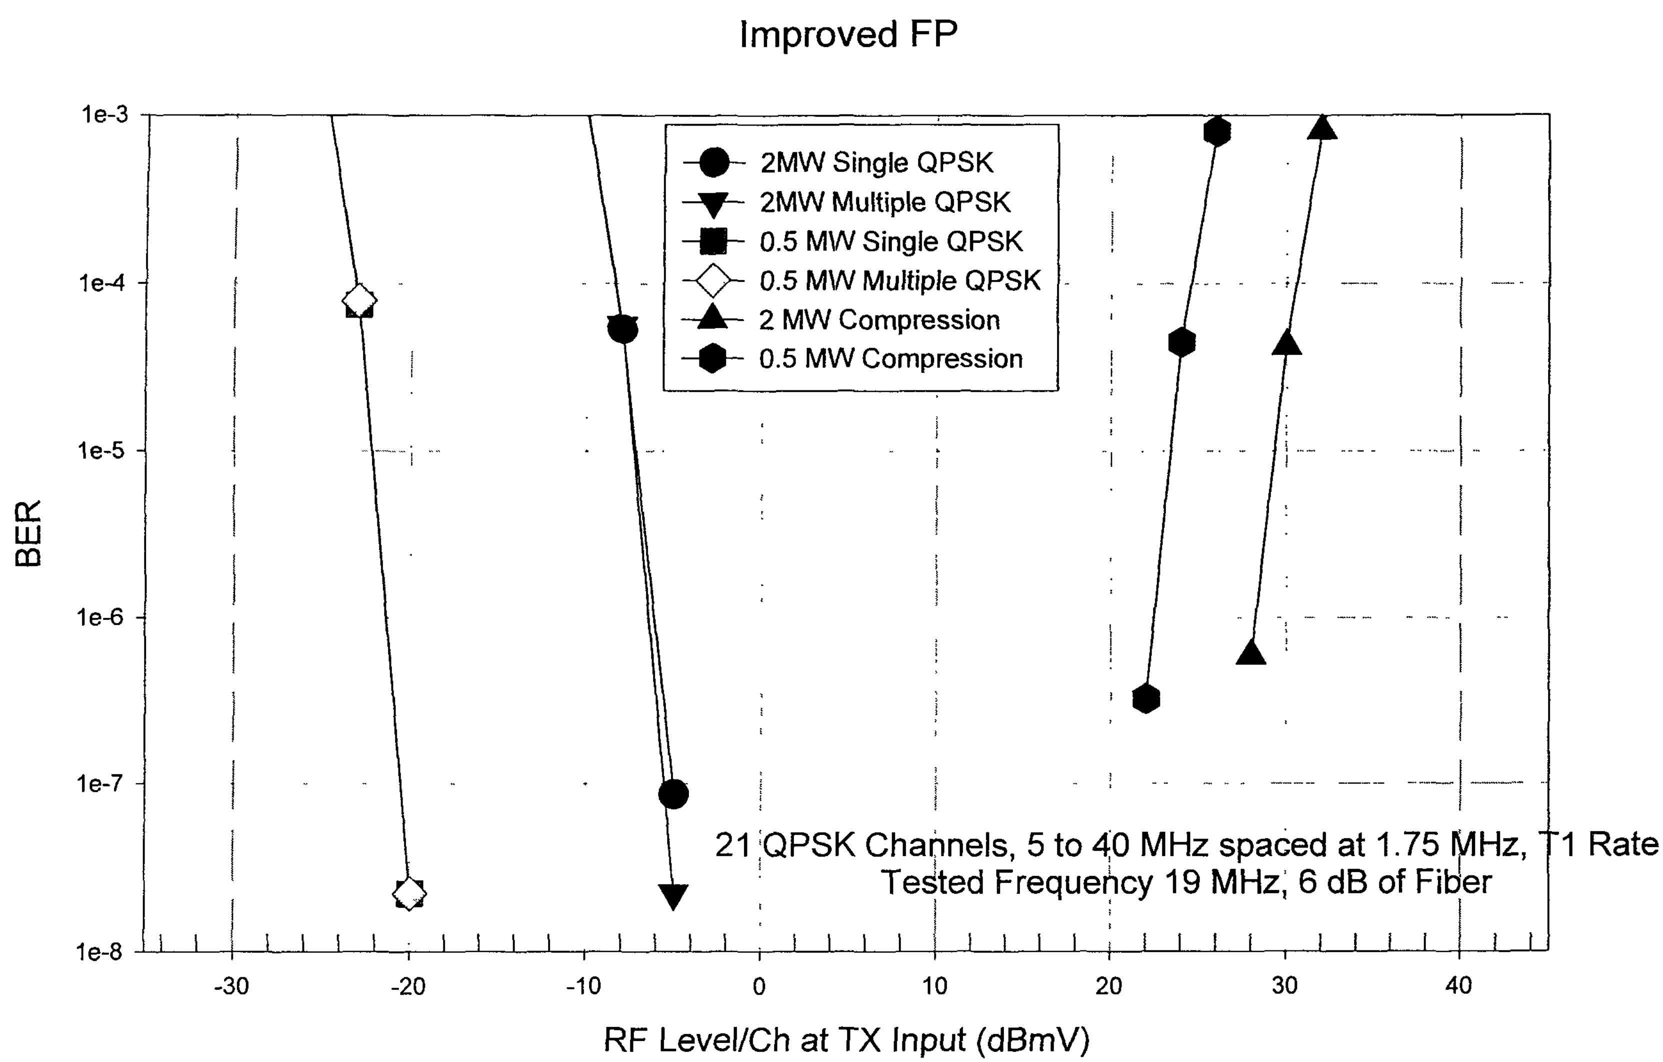
Improved FP
1e-3
1e-4
1e-5
1e-6
1e-7
1e-8
BER
2MW Single QPSK
2MW Multiple QPSK
0.5 MW Single QPSK
0.5 MW Multiple QPSK
2 MW Compression
0.5 MW Compression
21 QPSK Channels, 5 to 40 MHz spaced at 1.75 MHz, T1 Rate
Tested Frequency 19 MHz, 6 dB of Fiber
-30
-20
-10
0
10
20
30
40
RF Level/Ch at TX Input (dBmV)

FIGURE 5

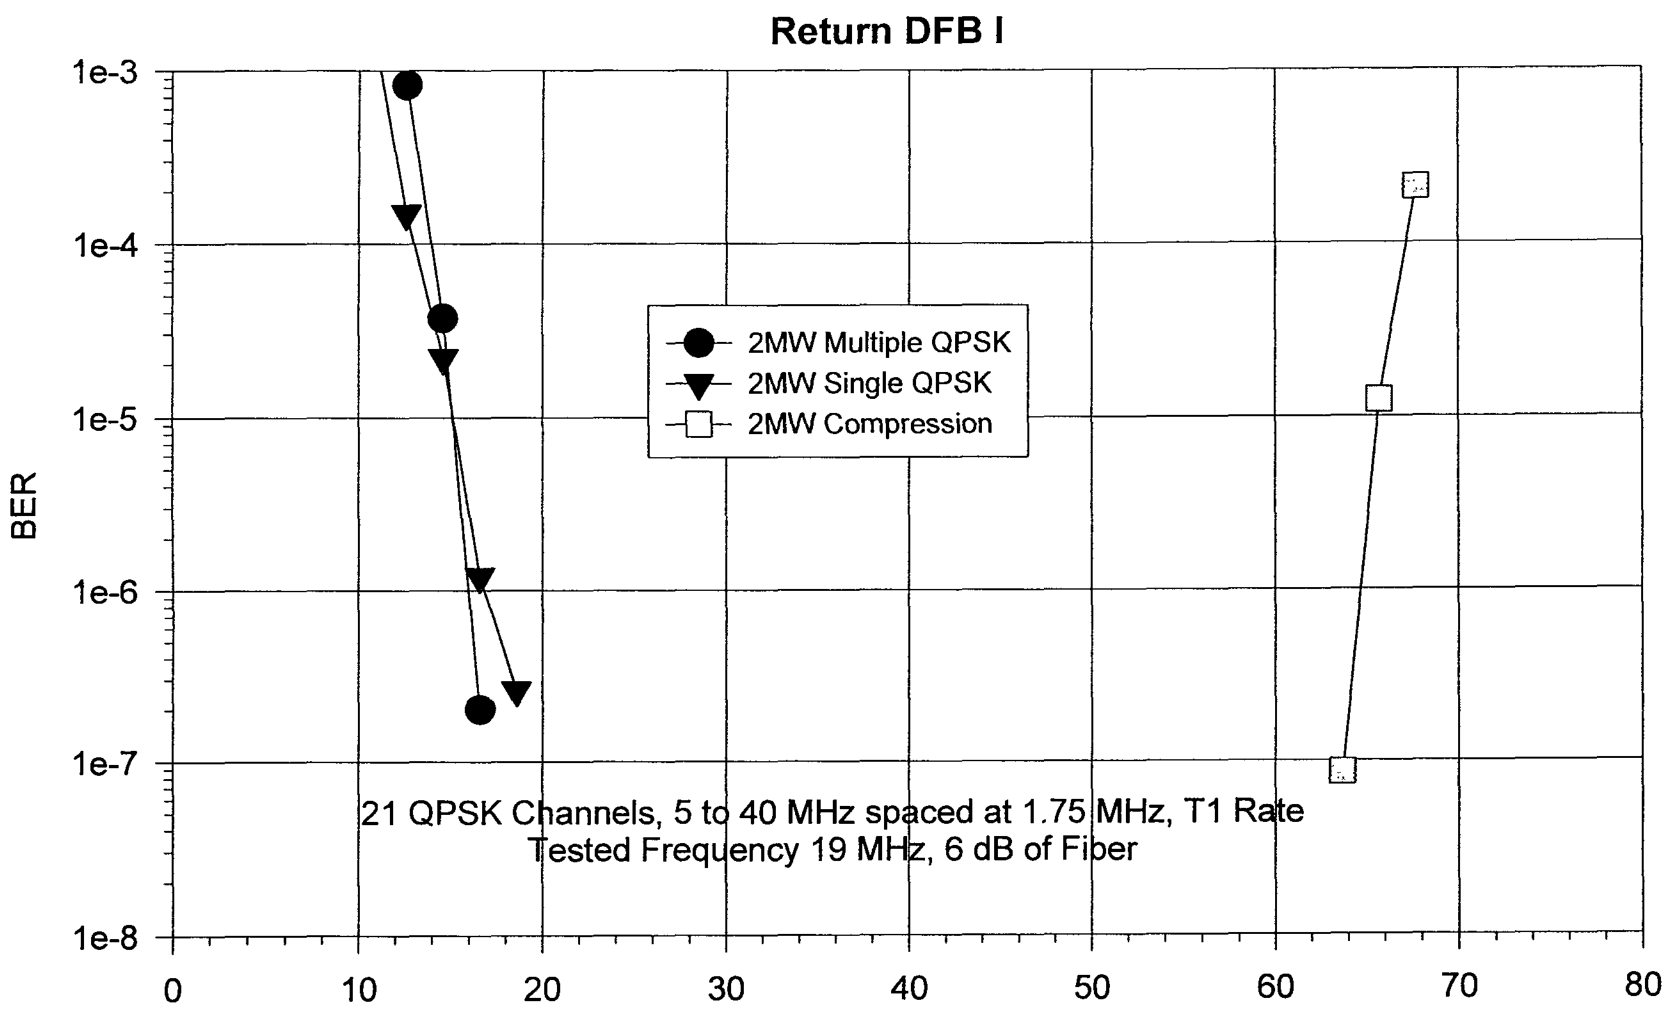
Return DFB I
1e-3
1e-4
1e-5
1e-6
1e-7
1e-8
BER
0
10
20
30
40
50
60
70
80
C/N
2MW Multiple QPSK
2MW Single QPSK
2MW Compression
21 QPSK Channels, 5 to 40 MHz spaced at 1.75 MHz, T1 Rate
Tested Frequency 19 MHz, 6 dB of Fiber

FIGURE 6

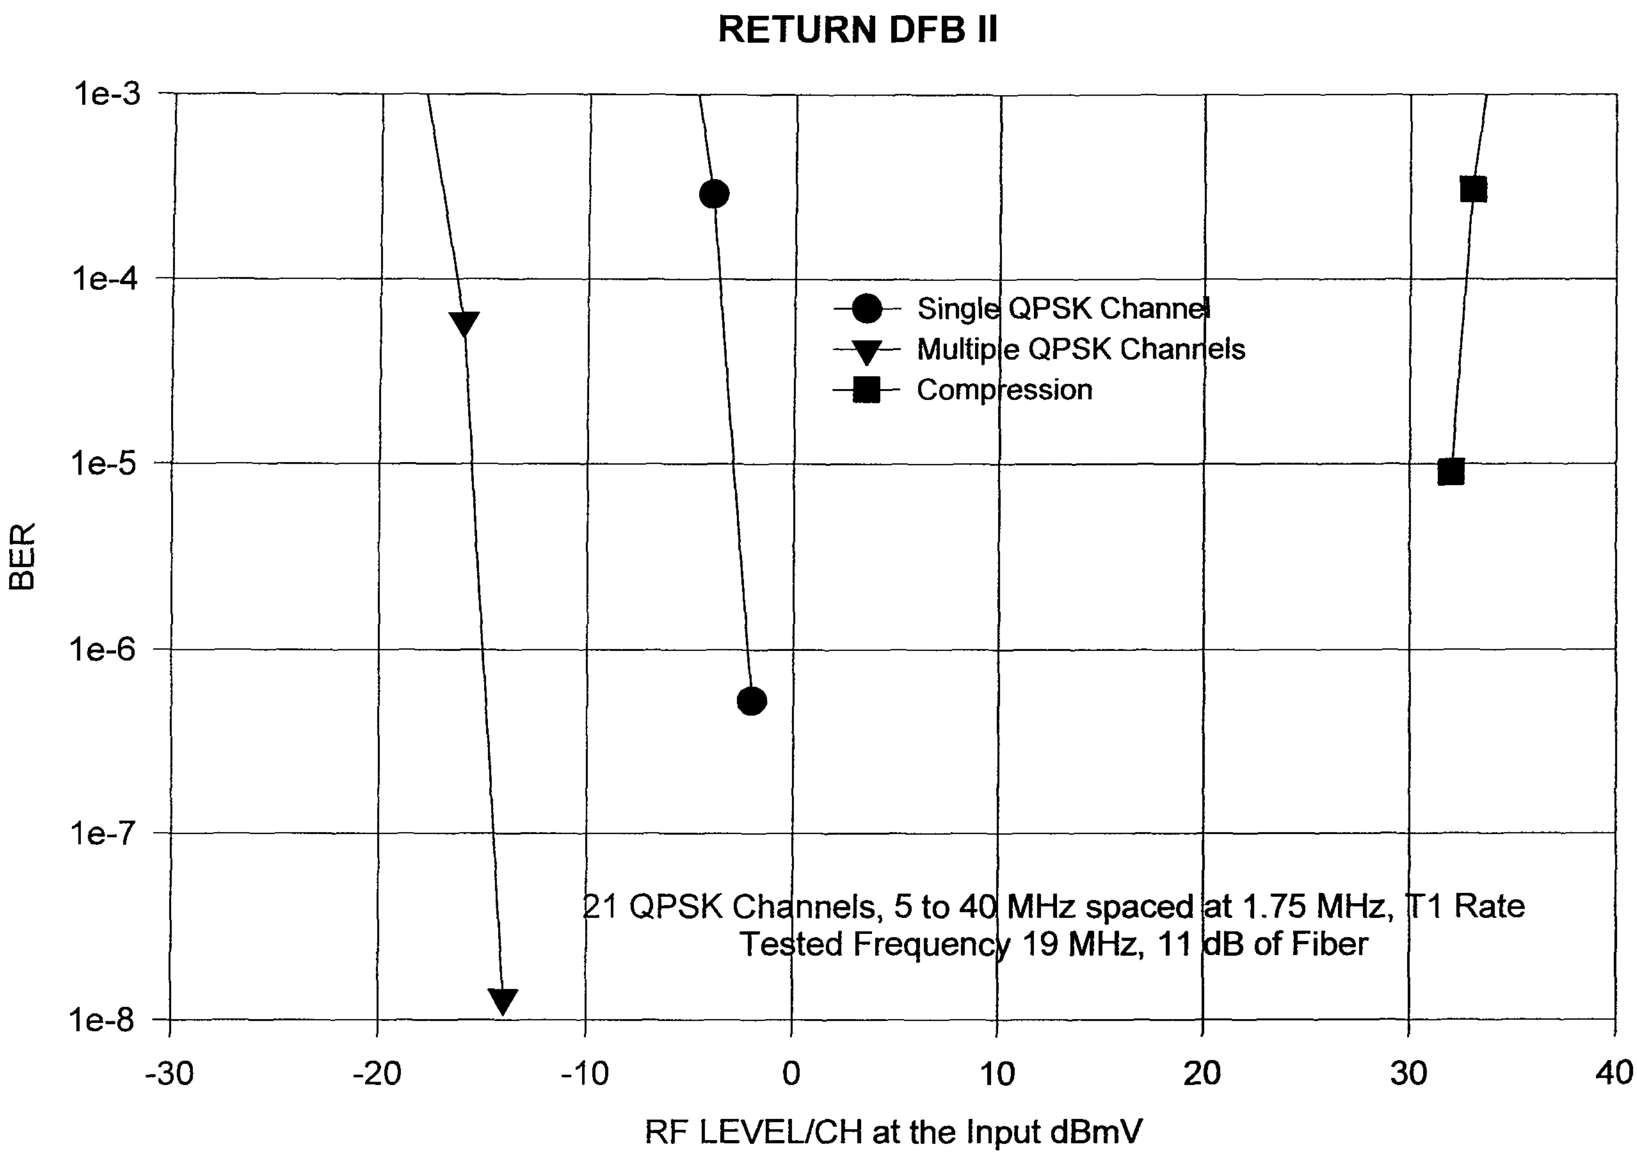
RETURN DFB II
1e-3
1e-4
1e-5
1e-6
1e-7
1e-8
BER
-30
-20
-10
0
10
20
30
40
RF LEVEL/CH at the Input dBmV
Single QPSK Channel
Multiple QPSK Channels
Compression
21 QPSK Channels, 5 to 40 MHz spaced at 1.75 MHz, T1 Rate
Tested Frequency 19 MHz, 11 dB of Fiber

FIGURE 7

System Diagram

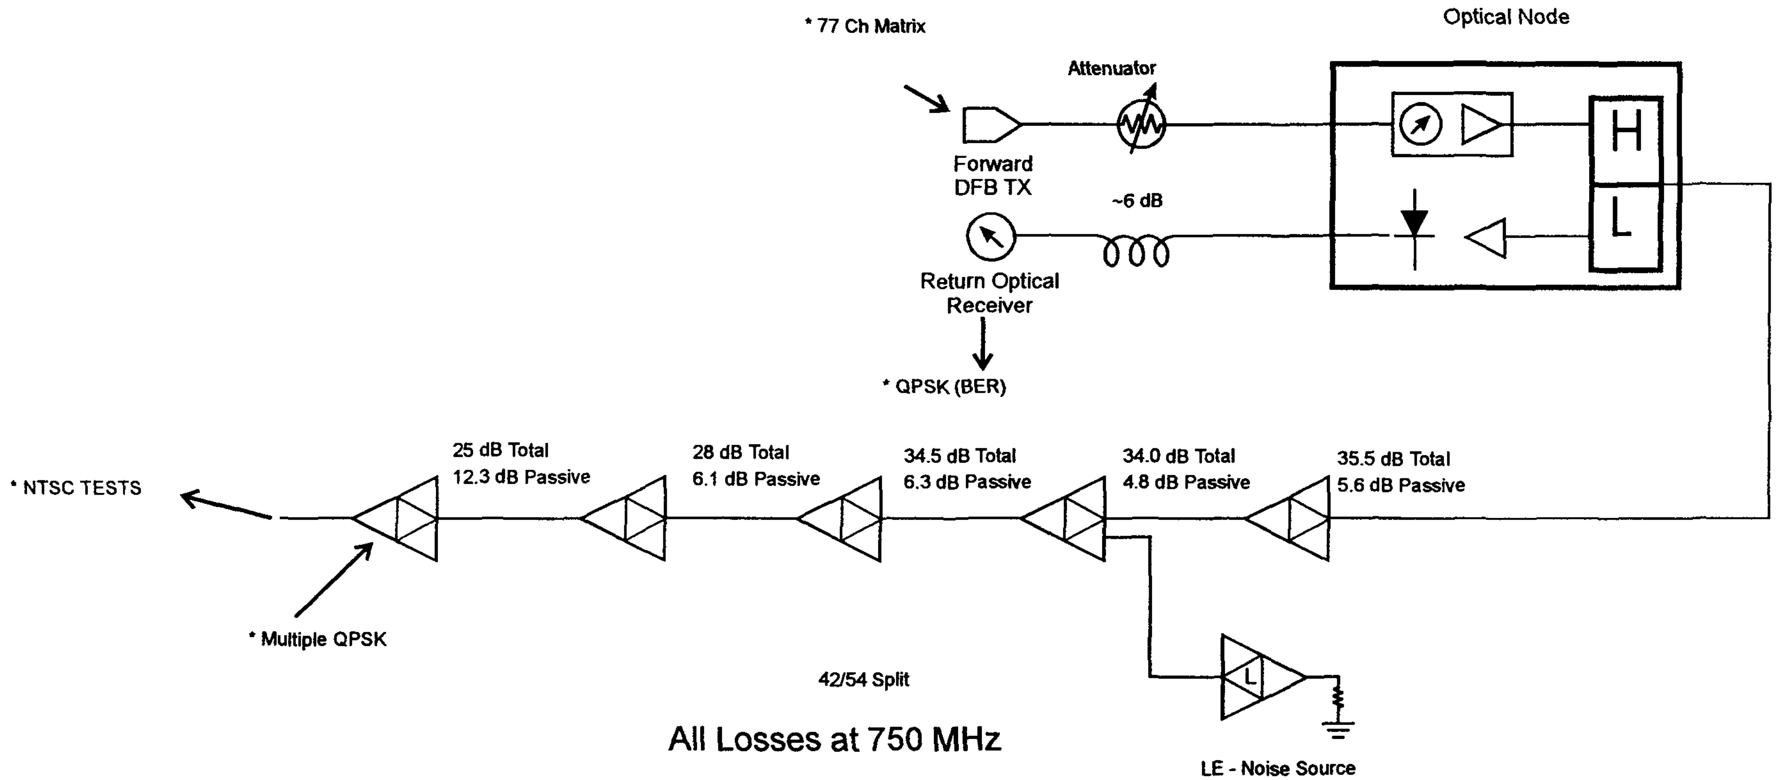

Figure 8

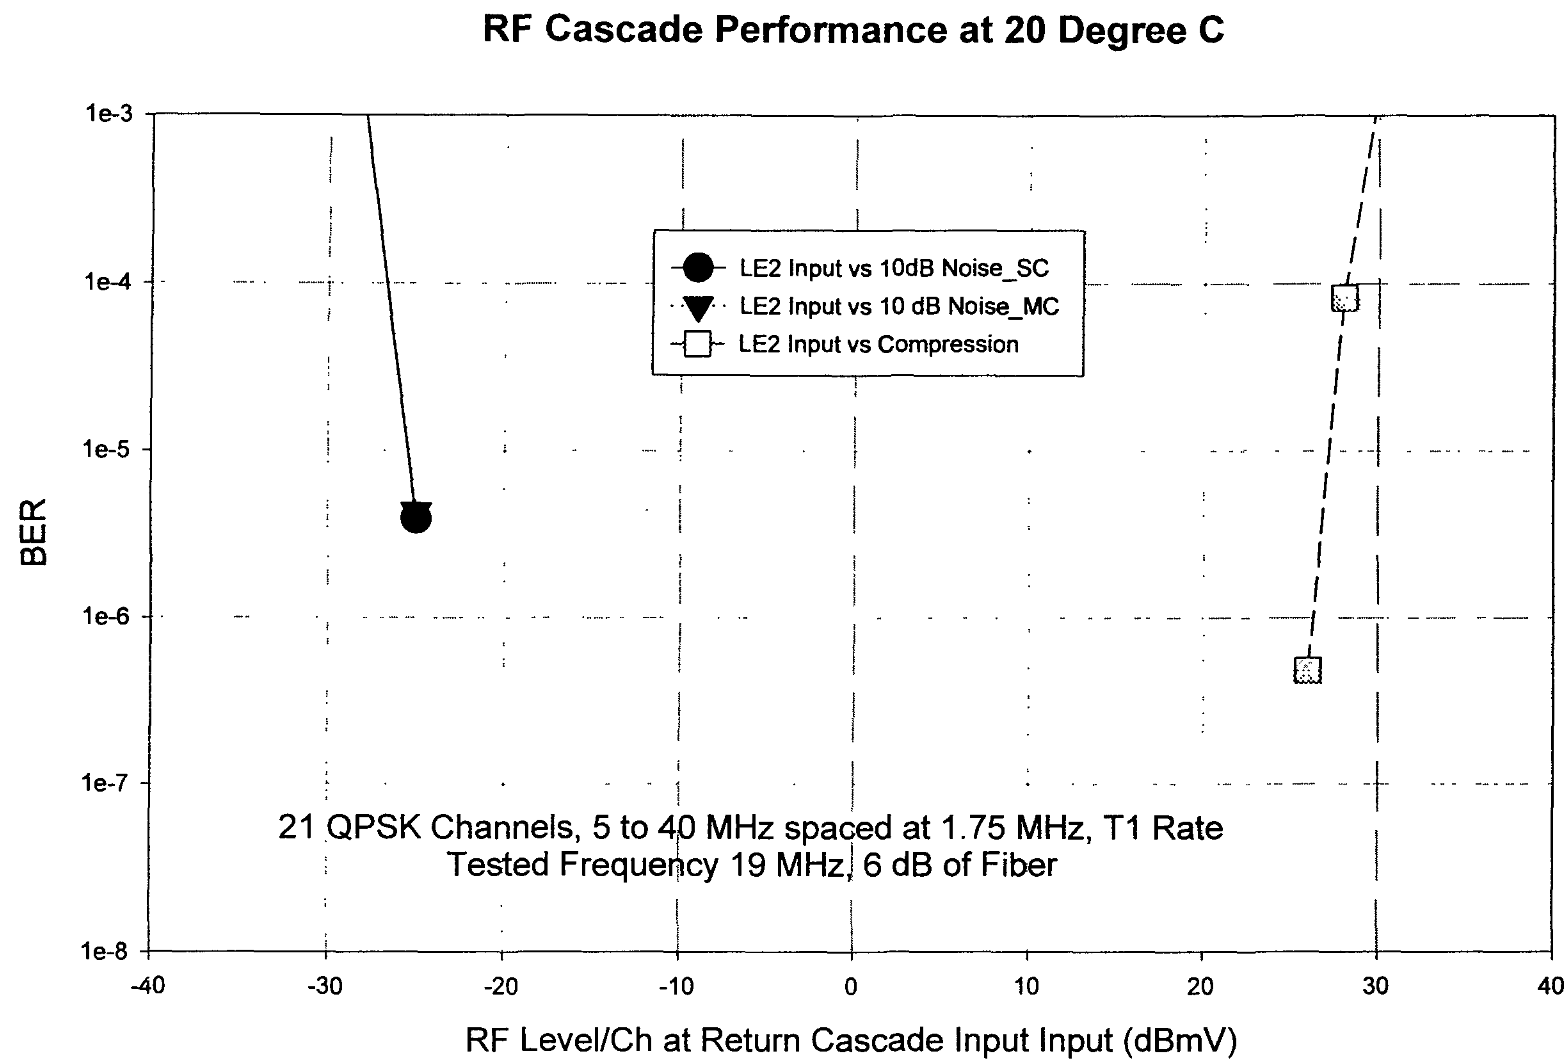
RF Cascade Performance at 20 Degree C
1e-3
1e-4
1e-5
1e-6
1e-7
1e-8
BER
LE2 Input vs 10dB Noise_SC
LE2 Input vs 10 dB Noise_MC
LE2 Input vs Compression
21 QPSK Channels, 5 to 40 MHz spaced at 1.75 MHz, T1 Rate
Tested Frequency 19 MHz, 6 dB of Fiber
-40
-30
-20
-10
0
10
20
30
40
RF Level/Ch at Return Cascade Input Input (dBmV)

FIGURE 9

Full System Performance at 20 Degree C

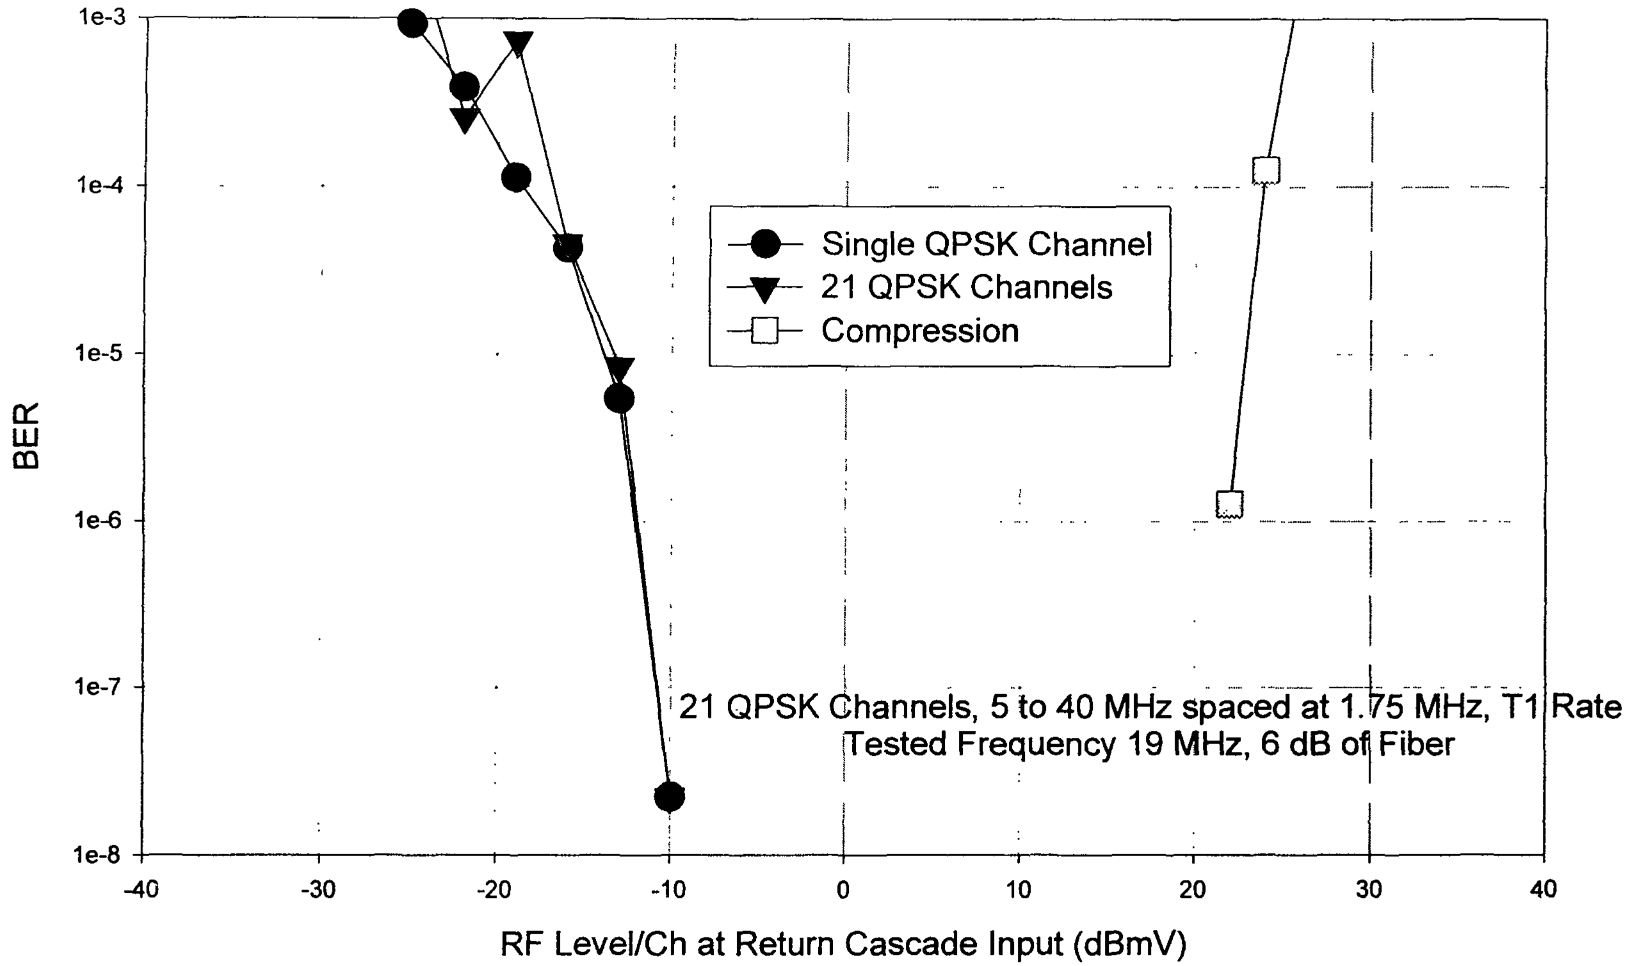

FIGURE 10

Full System Performance Over Temperature

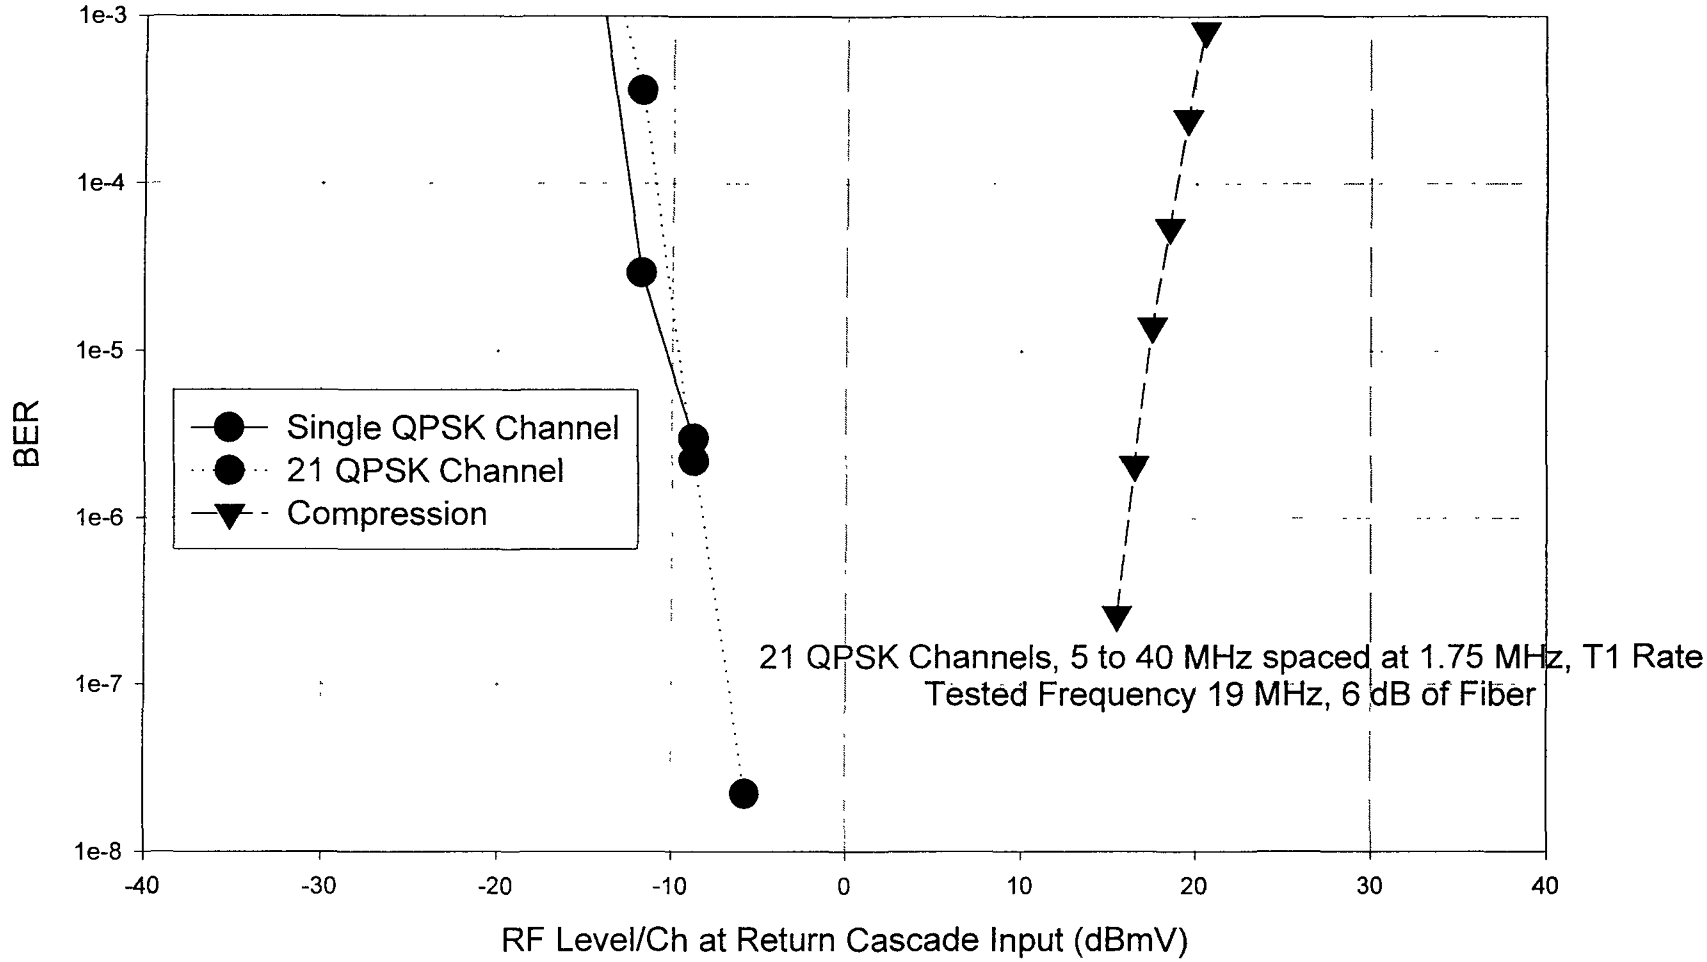

FIGURE 11

Ingressing Event

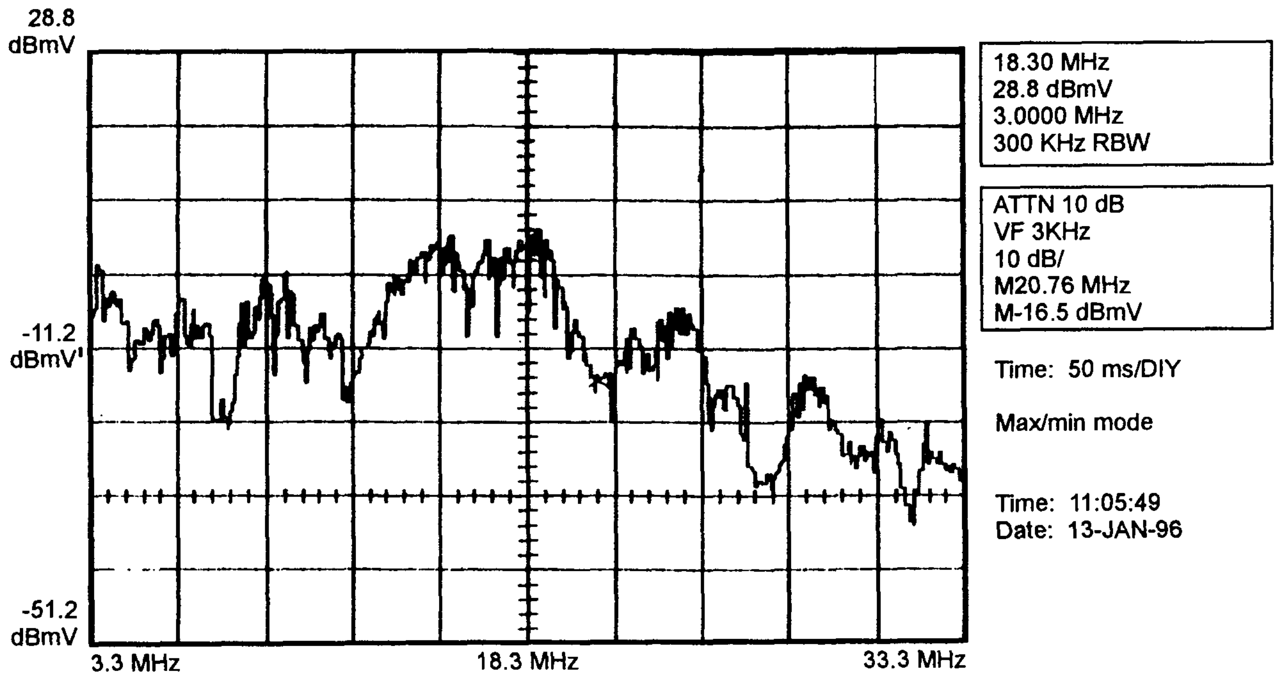

Figure 12