Cable Modems in the Home Environment

By
Joseph B. Glaab, Rob Howald,
Dean Stoneback, and Matt Waight
General Instrument

Abstract

This paper investigates the "in-home" environment when implementing cable modem systems in two-way HFC networks. Cable modems are hitting the market on a wide scale in 1998, with standardized performance and interfaces per the recently developed MCNS (Multimedia Cable Network System) consortium. Their ability to coexist with the existing downstream infrastructure is crucial to their acceptance as a broadband access tool. However, the modem may be required to operate at high transmit levels, potentially interfering with the downstream TV and/or settop converter unit. Also, the downstream receive equipment may introduce ingress into the return path, which can impair digital communications. Suggestions that address these serious issues are discussed. Measurement results are presented that justify these conclusions.

1.0 Problem Description

1.0.1 Cable Modem Interference on Forward Video

Addressable settop converters and TV receivers, both already deployed and currently being deployed in modern plants, have not necessarily been designed to be compatible with newer cable modems. A typical return path is considered to be 5-40 MHz, with allowances for performance degradation recognized at either band edge due to likely frequency response distortion. Also, the ingress problems at the low end are well documented. Modern cable modems, and particular those being designed to the specification written by the MCNS, will be able to operate up to 40 MHz. Actual return services and bands allocated are, of course, determined by individual MSO's. However, it is all but guaranteed that use of the complete available return spectrum will be expected in the near future to support the new services, which we be counted on as important revenue streams. A traditional converter's return transmitter typically operates in the low end of the return band. Thus, the forward-related equipment in those units was designed to reject high level interference in that band. The CFT-2200 settop, for example, a successfully received GI product which is widely deployed, was designed to be used with any return-path transmitter operating from 5 to 15 MHz at levels up to +57 dBmV with no recognizable video interference. Maximum cable modem levels will be roughly the same, but with wider spectral range. Thus, it becomes important to characterize the effect on video performance against cable modem levels and the wider upstream frequency range anticipated.

Related data has been measured, and is available in the report from Carl T. Jones Corporation, "Non-Video Interference Test Results report" (12/94). That data is important because it characterizes this phenomenon when it is associated with TV receivers and VCR's. A brief comparison of those results will be made with the results obtained in our measurements. Consumer electronics of this type are in the same situation as "old" converters - potentially seeing return band interference that they were not designed to handle.

1.0.2 Return Ingress Issues

In section 1.0.1, the return system's interference on the forward plant was introduced. A second issue is impairments imposed on the return communication link due to forward-related TV equipment. The fact that the home is a major source of ingress, and, in fact, the dominant source, is well known. Many external sources can be blamed for home-generated electrical pollution of the return band. Common examples include sources such as appliances, garage door openers, hair dryers, light dimmer switches, etc. These effects are aggravated by poor electrical practices, such as poor in-home connections and grounding, cheap, flea market quality splitters, and unterminated RF ports. Much effort has gone into the system design efforts for products such as MCNS-based cable modems to assure reliable communications through many of these impairments using advanced signaling schemes and powerful digital receivers in the Headend (HE). These receivers leverage the latest techniques in equalization, forward error correction (FEC), synchronization, and ingress avoidance to meet performance requirements and maintain a suitable percentage of error-free channel availability. However, a perhaps overlooked portion of ingress management is the effect of traditional settop converter units and cable-ready TV's themselves in contributing to the interference. While HE techniques to avoid and/or transmit signals through ingress don't care about where it is coming from, it is important to understand local sources of upstream impairments. In this way, future equipment can be modified and/or augmented properly for maximum return path availability. Also, there is an opposite extreme to very high modem transmit power. The range of output power required per MCNS specifications is 8 dBmV to 58 dBmV. Thus, the low end is quite susceptible to home-generated interference. In addition, for a properly aligned return plant that is heavily loaded, the transmitter should not be made arbitrarily high to achieve detection. For example, a transmitter that just increases its power until heard at the HE, only because high interference is corrupting the path, will not allow for comfortable coexistence with other users sharing the available power load.

It is important to note that both of these important issues have not been completely ignored in the past. In fact, committees exist that are now attempting to come to a consensus on return path emissions. Also, there has been some discussion with the release of the MCNS specification about whether the cable modem interference problem is in the court of the modem developers, or if it is the settop designer who must assure further protection of the video signal from return band frequencies that he does not use. Inspiration for this paper stems from the idea that not enough attention is being paid to the issue at this point, and now is past the time to do so.

2.0 Test Set-ups

2.0.1 Continuous Return Band Interference - Frequency Modulation

The first set of tests used FM to simulate a signal from the home (many traditional boxes implement FSK systems). The modem is simulated by a signal generator that is frequency modulated by an internal +/- 75 KHz tone. To eliminate harmonics from the test signal generator (much care must be taken for these types of measurements) the signal is low-pass filtered (see Figure 1). In fact, for ultimate rejection purposes, two filters are used in cascade, separated by a 3 dB pad. This signal is combined with a would-be cable-drop carrying a 77 channel forward band load. This combined signal is connected to the converter. Other test setups have been used, including one in which the modem simulator and settop converter are connected, as in a real application, to splitter outputs driven by a simulated drop from a tap carrying the downstream signals. However, side issues developed with this test set-up, including differences in splitter isolation, tap return loss, and signal harmonics from test amplifiers needed to boost the simulated modem signal to levels high enough to overcome the splitter port-to-port isolation. Thus, it was determined that the most reliable measure of interference that, in turn, could be used to back out required performance of the passives, was to combine the inputs *into* the converter. Thus, by measuring the level *into* the converter that causes degradation, rather than the levels *out of* the modem that cause degradation, the data is useable with any arrangement of passives associated with the drop and the home. Because it quickly became apparent that second harmonic distortion was the primary impairment, sensitive return band frequencies to test were quickly converged upon.

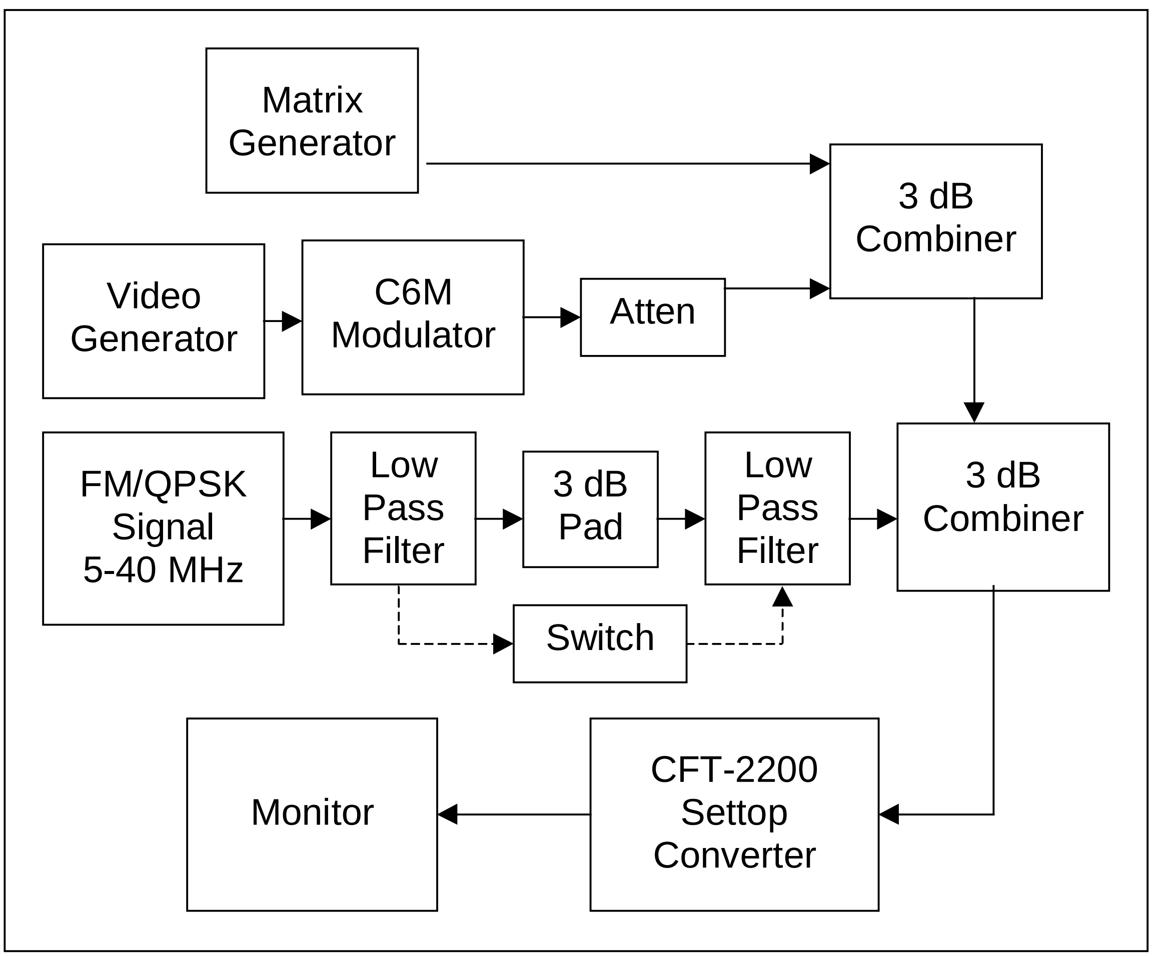

Figure 1 - Setup: Continuous FM or Burst QPSK Interference in Return Band

A 50 IRE test pattern was used, along with a forward matrix-generated load, with the test limit being the threshold of visibility (TOV). The desired signal and the matrix signals were set at 0 dBmV, the minimum specified video carrier level at the converter input. At higher levels, the AGC circuit in the tuner will reduce the return-path signal as well as the desired signal. It is important to note that there will be a correlation between homes that see the minimum video carrier level, and homes who's modems require the high transmit output levels. In this test, the four lowest channels were examined, before subsequently concentrating on the most sensitive for the situation being considered (Channel 5). Test results for this and all cases are given in section 3.0.

2.0.2 Continuous Return Band Interference - QPSK

The above test, using the setup of Figure 1, was repeated with a QPSK interferer for comparison. Two different QPSK data rates were used. One modem was a 2 MBPS unit (1 MSPS), while the second was a 256 KBPS unit (128 KSPS). This contrast of narrowband and wideband digital modulation, it was felt, would be informative. All other information described in the above setup remained the same. Again, care was taken to check the equipment harmonic performance, in this case of the modems, which were wideband and agile. TOV criteria was again the performance criteria (after giving our squinting eyes a break from the previous test).

Other important parameters of this test, and each of the subsequent QPSK tests, continuous or burst, include a forward band loading of 77 channels at 0 dBmV. In the Jones report, it was pointed out that a low end video level was more limiting with regard to interference performance, rather than excess distortion from a fully loaded high end, where all channels are at 15 dBmV. Also, the 50 IRE flat field was used still for TOV observations, and Channel 5 (77.25 MHz) was considered exclusively, for reasons that become more clear after discussing the first set of measurements. However, this is not meant to imply that a complete characterization across all channels should not ultimately be evaluated.

One other item of interest in the test setup is that the limiting distortion is known to be related to second harmonic energy. Thus, modem center frequencies were set close to one-half the Channel 5 frequency. Because of the band of occupancy of QPSK at 2 MBPS, the frequency of the modem in this case was shifted relative to the 256 KBPS case. In order to assure that the interfering energy of the QPSK signal (about 1.5 MHz of RF bandwidth) fell within the video bandwidth, it was offset so that the second harmonic fell at RF video carrier plus 750 kHz (39 MHz). The 256 KBPS unit was centered such that the second harmonic fell at carrier plus 250 kHz (38.75 MHz), a known location of poor TOV. For this narrowband signal, this made the most sense. For the 2 MBPS signal, it made the most sense to make sure the noise energy it represented was both near the sensitive region, and with all of the spectral energy contributing.

2.0.3 Continuous Interference Directly on Downstream - QPSK

In addition to this "overload" testing noted above, whereby tuner nonlinearity was the mechanism to create interference, TOV measurements were also done with direct interference. In other words, QPSK modulation was placed directly under the Channel 5 video signal by summing it in, along with the downstream load (see Figure 2). This has a couple of advantages. First, it allows a rough calibration of "eyeballs". That is, by putting the modems in CW mode, confidence in observations is high, because one can pull out the common chart on interference thresholds of CW interference and compare it to measurements. The second advantage, as previously described, is that it is very straightforward to go from what level of disturbance is actually at the converter that causes poor performance to useful specifications. Unlike the previous description, however, in this case not only is the measurement of tolerable level made into the converter, it is even made in the band of the channel being disturbed (Channel 5 in this case). Thus, it is an accurate gauge of allowable harmonic content in the forward band. For Channel 5, the QPSK modems were set directly at 77.5 MHz (256 KBPS) and 78 MHz (2 MBPS).

2.0.4 Burst Return Band Interference - QPSK

While it is certainly the case that interference all of the time represents a worst case scenario from the standpoint of the behavior of an in-home cable modem, it was considered possible that the burst nature of a transmission may, in fact, be as recognizable or more so. Because the cable modem operates in a TDMA (and FDMA) system, someone "surfing" while another is watching the tube will cause intermittent interference associated with the particular temporal characteristics of the upstream transmissions. This will vary due to overall return plant loading, physical location in the plant, selected modulation and symbol rate for the user, type of web surfer, etc., all which ultimately determine transmit levels, packet lengths, and interpacket periods.

For burst QPSK testing, a logic signal was used to toggle an RF switch. The input to the switch was the QPSK signal from the test modems. Data rates of 256 KBPS and 2 MBPS were again used. The modem burst parameters were 500 usec bursts of QPSK at a 200 Hz rate, roughly consistent with some preliminary studies of the temporal characteristics of upstream traffic. Obviously, these can vary widely. The QPSK signal again was set to interfere with the Channel 5 RF video carrier frequency near it most sensitive TOV point. The QPSK center frequencies were set as in 2.0.2 to land on this known most visible part of the video spectrum. The switch output is lowpass filtered, in an attempt to recognize that a cable modem output itself will be driven through an output lowpass filter, which thus significantly impacts the spectrum in a pulsing situation. It is also worthwhile to point out that, while MCNS does not consider the modem transmit level's ability to overwhelm the input of a converter, the specifications do carefully assure very low (intolerably low, if you ask modem vendors) harmonic distortion and spurious performance at the modem transmit output.

The same test criterion was used in these measurements - TOV. The return band burst test setup is also shown in Figure 1, with the dotted portion representing the addition of the path through the switch for burst testing.

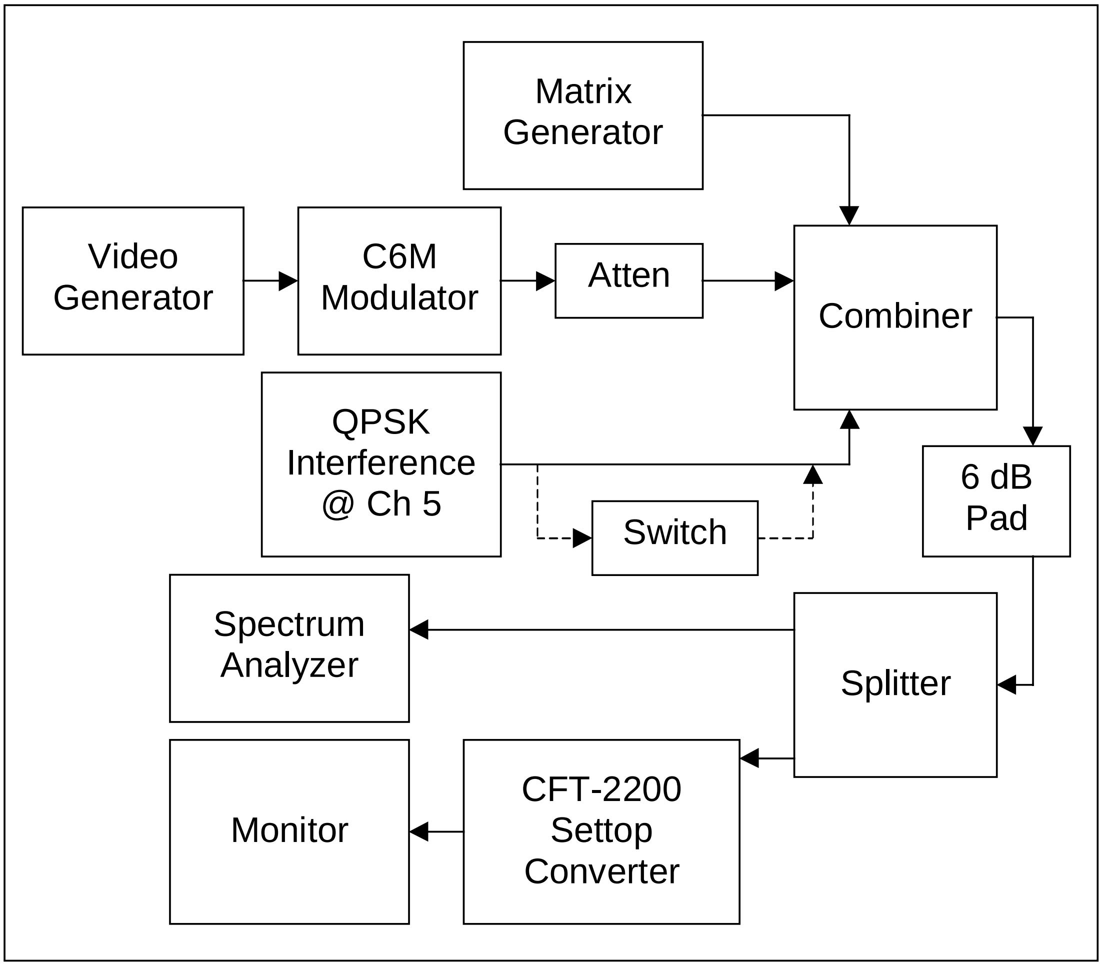

Figure 2 - Setup: Continuous or Burst QPSK Directly on Downstream

<u>2.0.5 Burst Interference Directly on Downstream - QPSK</u>

Finally, the bursty signal was placed underneath the downstream carrier directly, once again at the most sensitive point (Figure 2) Both data rates were implemented, and TOV measured, using the same QPSK center frequencies as 2.0.3.

3.0 Test Results

<u>3.0.1 Continuous Return Band Interference - Frequency Modulation</u>

The results of TOV testing with the continuous FM signal are shown in Figure 3. The primary problem is the high-level signal entering the converter and mixing with itself in the tuner, producing a second harmonic that falls in-band of the tuned channel. Even if the cable modem completely eliminates all harmonics, a fundamental signal that is above 27 MHz may interfere. The levels which cause a problem are consistent with the second order distortion performance of the converter’s tuner.

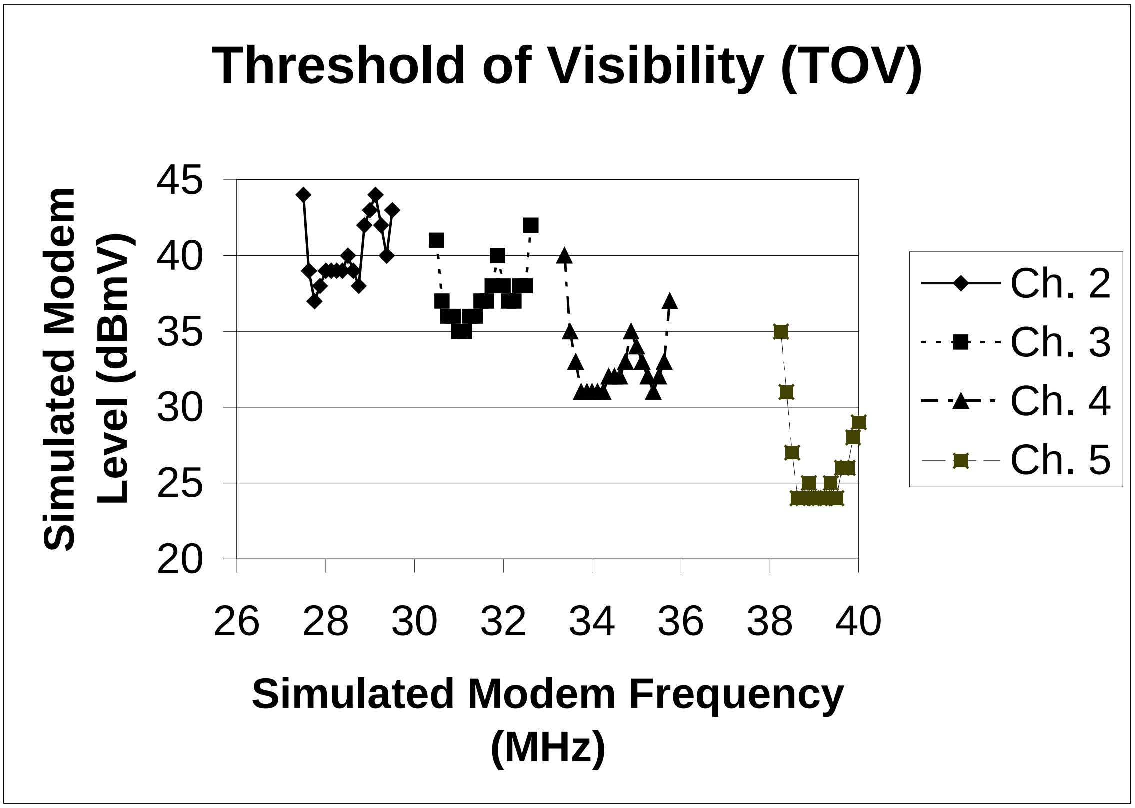

Figure 3 - Continuous FM Interference in Return Band vs. Channel

The highest frequency modem outputs (around 40 MHz) create the most degrading situation when second order distortion falls into Channel 5. Of course, the video TOV is dependent on where within the video channel the interference falls, and in the case at hand, this occurs the worst in the 39 MHz range. It is apparent that, for the measurements taken on Channel 5, the converter has an interference threshold of only 24 dBmV at its input. Thus, a splitter which has only 25 dB of isolation between converter and modem would result in video degradation of Channel 5 if the modem were to transmit above 49 dBmV. MCNS-based cable modems have a maximum requirement to transmit at 55 dBmV (QPSK) and 58 dBmV (16-QAM). And, link analysis will show that these types of levels can be expected in the plant. Both link analysis and the desire to overcome in-home ingress point to the need to transmit from the home as high as possible, while the converter on the other side of the splitter begs for consideration of more tolerable transmit levels to keep the low end video channels clean. For the lowest channel, Channel 2, which is degraded at its worst point by a modem frequency of about 28 MHz, an input level of 37 dBmV establishes the TOV. Here, clearly, a 25 dB splitter isolation would be adequate to avoid video interference (37+25 = 62 dBmV, above modem transmit maximums). Channel 4 becomes an in-between or marginal case, showing a TOV when a 31 dBmV signal is delivered at the converter input. In this case, a 25 dB isolation would just be enough (QPSK), just miss (16-QAM), or simply be very sensitive to the quality of the splitter and the return loss of the tap.

The modem signal level that sees its way to the converter is a combination of splitter isolation and the return loss of the tap. The reflected signal from the tap is reduced by splitter's loss two times, once traveling out to the tap, and once traveling back, for a total loss of somewhere between 6-9 dB on top of the tap return loss. The keys to minimizing problems are to use a good splitter having a port-to-port isolation much better than 30 dB in the high end of the return band, and using taps that have a return loss much better than 20 dB in the same region. Both of these numbers can be demanding relative to what is known about typical values in the field, and relative to some of the poorer commercial quality passives available for use in the home today. The example below illustrates the issues before discussing more about the performance of the RF passives.

Example:

A cable modem transmits at 48 dBmV (MCNS range required is 8 dBmV to 58 dBmV, modulation dependent) into a splitter having 30 dB of isolation and a loss of 3.5 dB. The tap has a return loss of 20 dB. The converter or TV receives the return signal directly from the splitter at (48 – 30) 18 dBmV and from the tap (ignoring return band drop loss) at (48 – 3.5 –20 –3.5) 21 dBmV. The combined signal level is 22.8 dBmV. The tap is the primary problem in this example, but both contribute, and a poor splitter could easily be the dominant source of interference. However, from the measurements taken, this level is below TOV for the worst case situation observed with the converter. Later burst measurements will add yet more insight into what levels to be concerned about.

Clearly, then, it is the case that there is a bit of "competition" between interfering paths, with either being capable of being dominant. It is as well the case that, if they are close to presenting the same level at the converter port, they can add. The exact adding relationship, being the same signal, would be dependent upon the relative phase shifts of the two paths. Certainly, in the worst case, there is the potential for voltage or nearly voltage addition that would create more disturbing peaks than either path alone.

Also, the above example is not representative of every situation. For example, very good splitters may have 40 dB of isolation over a portion of their band. This would be what is necessary to stay below TOV for a 58 dBmV signal in the example above. Fortunately, these better values typically exist nearer to 40 MHz, where the video interference problem is the worst, and they are poorer, for example 25 dB, at 5 MHz. A bigger problem is consumer TV home-improvement efforts. The well-infiltrated consumer quality brand can have isolation as poor as 10 dB in the worst part of its band, and perhaps averaging 15 dB across the band. Other splitters exist of quality between these two extremes. Clearly, this type of home environment will significantly aggravate the problem. A comparison of two-way splitter data of consumer quality and professional quality models is given in Tables 1 and 2, respectively.

	5 MHz - 40 MHz			50 MHz - 750 MHz		
	Worst Case			Worst Case		
	Level	Freq	Avg	Level	Freq	Avg
S11	-15	5 MHZ	-17	-13	600 MHz	-16
S22	-10	5 MHz	-15	-8.95	750 MHz	-14
Insertion Loss	-4	5 MHz	-4	-4.55	750 MHz	-4
Isolation 1-2	-12	5 MHZ	-18	-25.47	52 MHz	-27
Isolation 2-1	-10	5 MHZ	-15	-25.72	52 MHz	-27

Table 1 - Consumer Grade 3 dB Splitter Measurements

	5 MHz - 40 MHz			50 MHz - 750 MHz		
	Worst Case			Worst Case		
	Level	Freq	Avg	Level	Freq	Avg
S11	-24	5 MHZ	-32	-24	750 MHz	-27
S22	-21	5 MHz	-30	-28	750 MHz	-33
Insertion Loss	-3	5 MHz	-3	-4	750 MHz	-4
Isolation 1-2	-26	5 MHZ	-39	-33	631 MHz	-37
Isolation 2-1	-26	5 MHZ	-39	-33	631 MHZ	-37

Table 2 - Professional Grade 3 dB Splitter Measurements

In addition to splitter values, the tap return loss described in the example also can vary significantly. The one advantage here is that this is an item out of reach of the consumer. However, it is still the case that, in the return band, tap return loss values of about 18 dB may be more typical. Again, however, to our advantage, better return loss numbers in the upstream band occur closer to 40 MHz. These results indicate that typical the tap products made by the industry are not well suited to coexist with cable modems. Because of this, additional converter input filtering will probably be needed in the future.

Another tap-related concern, aside from the trouble that a modem in the home will do to a TV in the home, is what the same modem may do to his neighbor's TV. The level issues are the same with regard to what can be handled at the input to the neighbor's settop, but the specification of interest with regard to the passives becomes the port-to-port isolation of the taps. Here, a typical value is on the order of 25 dB. At lower frequencies of the return band, it may be as poor as 20 dB. It is important to recognize that, as with any RF passive, both return loss and isolation performance can be sensitive to the VSWR's seen at the other ports of the tap.

Ignoring any drop losses, the 25 dB isolation number results in interference still a few dB above TOV for the tested converter when a 58 dBmV modem output is assumed (58-25-7 = 26 dBmV). A solution that solved the return loss problem at the settop in a subscriber's home would correspondingly work, technically, for the isolation problem to his neighbor. However, the headache of possibly moving into multiple consumers' homes for adding one new subscriber is not a pleasant scenario. Note that use of return path equalizers would reduce the isolation problem and likely aid with VSWR as well, if implemented in the tap arms. This, however, is not the preferred implementation. But, this complexity must be weighed against the significantly complex effort of possibly needing major performance improvement in existing RF passives to support cable modem introduction. In addition, a settop from another manufacturer showed a TOV criterion of only 18 dBmV, another 6 dB more sensitive than the unit described above. Clearly, this settop is sensitive to cable modem levels that are output closer to the middle of their transmit power range.

<u>3.0.2 Continuous Return Band Interference - QPSK</u>

The results for the Channel 5 tests in each of the QPSK cases are summarized in Table 3. There, it is apparent that the QPSK interference tested was less visible than the FM in the previous tests. There are two likely reasons for this: the wider bandwidth (albeit slightly in the 256 KBPS case) and the noise-like modulation (as opposed to the FM, "more" periodic tone-modulated waveform). However, the differences are relatively small (3-4 dB), enough such that, along with the subjectivity of the test, they could be considered not very significant. There does seem to be a pattern established that the broader band modulation that spreads the energy away from the sensitive TOV, and across the video band, is less visible.

<u>3.0.3 Continuous Interference Directly on Downstream - QPSK</u>

The CW video interference charts have been around for a long time. These graphs show that the location of the interference is important to what level it can be to upset performance, and is the reason that specific interfering frequencies are setup as previously described. Most charts show the lowest visible level is about -57 dBc near the carrier offsets already mentioned. In other locations, it can be as high as -40 dBc. In the baseline done here, the CW threshold observed was -63 dBc, possibly due to higher quality monitors of today, and the close range, intensive observations. The modulated data effects on video observed for the tests here are best compared to the relative measured TOV of -63 dBc.

As in the previous section, the QPSK modulated signal is able to be higher than the CW carrier, as shown in Table 3. Again, the reason for this is likely related to bandwidth and periodicity, as the "snow" effect of digital noise fades more easily into the background than discrete lines. Both the 256 KBPS and 2 MBPS cases have this "snowy" interference characteristic. Also, as in the previous case, the differences are rather small between the two, but show an impressive 7 dB to 9 dB more forgiveness to interference when modulated, relative to the -63 dBc baseline. Based on the discussion so far, it would have been expected that modem at 2 MBPS would be less discernible, but, in fact, the opposite is true from the measurements taken. However, it is felt that, with the small differences and different days of measurement, this could easily be attributed to measurement error.

3.0.4 Burst Return Band Interference - QPSK

The concern that transient bursts of interference would be more troublesome was put to rest in this round of tests. As can be seen in Table 3, another 5 dB or 6 dB of transmit level, as measured during a burst, was tolerable. The likely reasons are perhaps, again, spreading of spectral energy due to the pulsing, or simply a lower average interference power at the detector. In fact, it is suspected that truly random interference, such as a cable modem user in a home would induce, would be less visible yet. The interference in these tests had line-like qualities, with snow "within" them, and these periodic characteristics would be removed in a real application. It is likely that the ability to recognize what disturbance was being looked for effected the ability to see it.

It is very important to note the absolute numbers measured for this part of the test. Consider this 5 or 6 dB improvement, and assume a 5 dB or 6 dB improvement in tap return loss over currently accepted typical industry performance. Assume a high, but achievable, splitter isolation. All of this could allow the prior example to remain below TOV, even with a modem output at 58 dBmV. In addition, the true randomness of transmissions will likely provide margin above that measured here. However, improved tap isolation is unlikely given that wider tap bandwidths are being accommodated, causing VSWR compromises. Return path pads may be the only solution. Also, it is unlikely that great strides will be made in consumer grade splitters.

These important results emphasize the need to examine more cases of picture and modulation under more traffic (burst) conditions. This will provide a better handle on the realistic magnitude of the problem.

3.0.5 Burst Interference Directly on Downstream - QPSK

The results of section 3.0.4 were supported by the measurements taken using a direct QPSK burst interferer. In this case, 6-7 dB of additional transmit level was tolerable over the continuous and direct interferer.

	TOV on Channel 5	
	Interference Characteristic	
	Continuous	***Burst***
Interference Type		
FM @39 MHz	24 dBmV	
QPSK (Return Band)		
256 KBPS @ 38.75 MHz	27 dBmv	32 dBmV
2 MBPS @ 39 MHz	28 dBmv	34 dBmV
QPSK (Direct C/I)		
256 KBPS @ 77.5 MHz	-54 dBc	-48 dBc
2 MBPS @ 78 MHz	-56 dBc	-49 dBc

Table 3 - Summary of TOV Data for Channel 5

3.0.6 Comparison of TOV Measurements with Jones Report

The test report "Non-Video Interference Test Results Report", prepared for the EIA by the Carl T. Jones Corporation in 1994, describes measurements taken on ten different TV models and ten different VCR models which, as of 1994, were of "recent manufacture". EIA compliance requires meeting specifications only up to 30 MHz, the traditional limit of return band signaling. In the Jones report, "non-video" means CW interference. The inspiration for the Jones report work was basically the same as this paper. That is, an increase in upstream services means more of the band will be used, and at high signal levels. Thus, the report closely augments nicely the work that has been described above for settops by concentrating on TV's and VCR's. Also, Jones defined the TOV numerically, as a 55 dB C/I, a number that some may question based on subjective perceptibility data. However, it is in the ballpark of most reasonable attempts to characterize this subjective phenomenon. Jones also used Channels 2-6, because of their low-end location, and their relative relationship to return band signal harmonics. The test uses CW interference up to 48 MHz, to characterize IF interference effects, which turn out to be the worst kind. That is not the topic of our experiment, however, and we will concentrate on Jones' results for second harmonic distortion, the most deleterious also in the measurements of consumer electronics.

Summarizing the relevant results, the report concludes that 90% of the "recent" TV's have TOV problems for around 34 dBmV, with 32 dBmV being about the bottom of all tested. For VCR's there is a bigger cause for concern, as 90% bottom out at around 22 dBmV, with 18 dBmV being the lowest tested. This later was also the lowest interfering level to cause a TOV problem among the settops measured. Thus, the conclusions that can be drawn do not significantly differ. In order to provide reliable return services that can comfortably co-exist with existing equipment, it appears some changes will have to be considered to the existing forward-related hardware, whether it is by CATV equipment providers, or consumer electronics equipment providers.

4.0 Ingress on the Return

Traditional return equipment was very limited in capability and in spectral occupancy. Egress from such sources, and from television sets, had no reason to be a major concern to MSO's, because there was such limited use of the spectrum, and signaling techniques, while not very bandwidth efficient, were very robust (i.e. FSK). However, the situation going forward is once again decidedly different than the one that existed during mass deployments of traditional set-top converters and televisions. It is generally expected that more and more plants will become two-way, and that more services covering more return bandwidth will be offered. Eventually, spectral real estate will be at a premium. While ingress will always be a problem due to uncontrollable sources in the home and surroundings, interferers that are imposed by the CATV equipment and consumer electronics gear that make return channels unavailable will not be tolerated from the manufacturers. It is likely that this problem will come to the forefront as older plants move into upgrade phases, in search of providing much more traffic on the returns. Some systems use blocking filters for all inactive returns. This certainly helps ingress management, but is inherently self-defeating, as adding revenue from the return requires adding as many new subscribers as possible. Thus, the filters ultimately will be mostly removed if the return band is used successfully.

4.0.1 TV's and VCR's

The egress associated with existing consumer equipment is designed only to meet FCC requirements. Unfortunately, consumer equipment will do just that - meet these "minimal" targets and do little else. Our region of interest for the return is, of course, the 5-40 MHz band. FCC cable ready consumer electronics specifications are detailed about what can emanate from the port above 54 MHz, but there is little information to go by for return band interference. About the only recognizable "rule-of-thumb" idea is that interference from the port does not interfere with services in the band (CB radio, for example, is in the 27 MHz range). However, this is normally associated, for example, with situations such as TV's receiving broadcast services, but with an unused cable port left open that may output RF interference

When actually using the cable network, the requirement becomes a more complex assortment of possible leakage regions, including cables, connector, and tap emissions. In essence, however, the specifications are about assuring no existing services, such as CB, are disturbed, and not related at all to disturbing any services that do or may eventually return on the cable. Recognizing this, one could peruse the FCC regulations for other "unintentional radiators", or other radiated emission violations, which are measured by field strength at various distances. A possible calculation could then back into what that emission specification would say about the level of interference on a CATV port of a settop, TV, or VCR. The main point, however, is that nothing in the FCC spec will relate at all to considering how interference on that port in the return band effects what services are in that band and on that cable. This is not the fault of the FCC, it is just that, until recently, there has never been a need for such consideration. But, now there is.

4.0.2 Settop Equipment

A second issue of concern is consumer-grade CATV equipment in the field, primarily the traditional set-top converter, and its potential for spewing signals onto the cable drop. For set-top boxes, such egress can be aggravated more so than with VCR's and TV's. This is because homes within a plant may employ various TV and VCR models from various vendors. This variable means that noise contributors are likely to vary, depending on where across the band the different equipment spews interference. However, for settop units, there is a strong likelihood that everyone in the plant will be using the same or a similar model unit. As such, the contribution to upstream interference will be very close to the same frequency for every box, with the difference being related to the reference crystal's drifting, and possible the channel being tuned to. This is simultaneously both good news and bad news. The good news is that the interfering frequencies are more predictable, and thus more able to be filtered out, at the expense of not using those bands, of course. The bad news is that, because of this and the noise funneling effect of the return system, there is more of a likelihood of interference increasing in a more destructive fashion up towards the node.

4.0.3 Addition of Spurious Signals in the Feeder System

The first step in calculating how the spurs add in the feeder system is determining the loss from each tap port to the nearest amplifier. The loss from each tap port can be used in conjunction with the number of homes connected to each tap to determine the total signal power at the amplifier. It is sufficient to calculate the loss to the nearest amplifier since, in a properly aligned plant, the net gain between amplifiers is zero. The loss between each tap port and the nearest amplifier was calculated for a sample of actual plants with low, medium, and high densities. The results are shown for densities of 24, 97, and 230 homes per mile (HPM) at 5 and 40 MHz in Figures 4 and 5. As expected, there are less high value taps in the rural area (24 HPM), since there is almost always some cable loss between the amplifier and the nearest tap. There are only minor differences between the distribution at 5 MHz and at 40 MHz.

The distribution of tap losses can be combined with the number of homes connected to each tap to determine the average loss from an average home to the amplifier port. This calculation was performed for all three densities at both 5 and 40 MHz and the results are shown in Tables 4 and 5. All losses are from the tap port to the nearest upstream (toward the node) amplifier station port. Table 4 shows what the total spurious level would be if the spurs add on a power basis. Table 5 shows the results if the spurs add in voltage. Measurements taken show a tendency to add higher than as power. However, in a statistical sense, over many plants, the power addition probably represents "most" plants, but the capability to add more destructively, the case that must be designed for, certainly exists. Since like crystals tend to behave like one another with regard to drift and aging, considering them as independent and random may be too lenient. Further, because they move slowly, any existing summation condition is likely to last for prolonged periods of time.

Example:

Consider a medium density (approx. 100 HPM) node with 1000 devices each emitting a signal at -35 dBmV at 5 MHz. If the signals add on a power basis, the total power at the node will be -35 + 11.63 = -23.37 dBmV. Note that Gain per Home (-18.37 dB) + 1000 homes (30 dB) = 11.63 dB. Thus, the number of homes is weighted by averaging it over the loss distribution. It is easy to show that, for large number of HPM, significant loading of the return due to interference alone can occur.

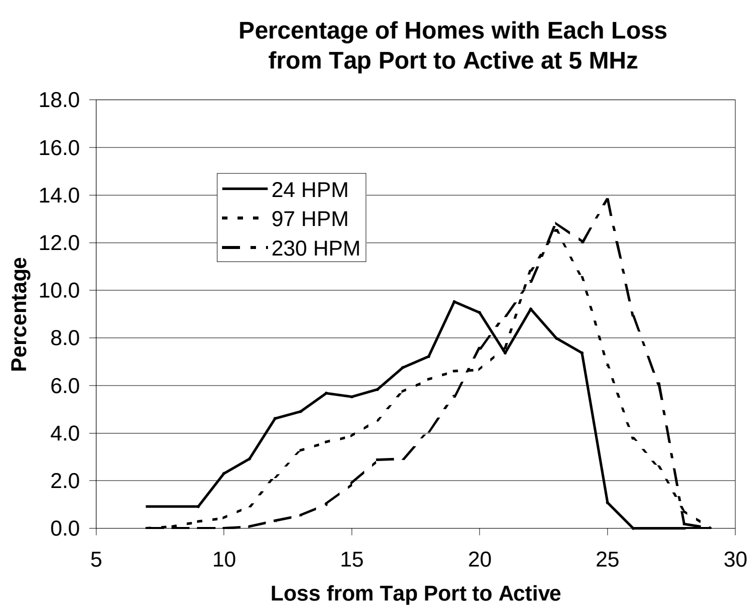

Figure 4 - Statistical Distribution of Return Path Loss @ 5 MHz

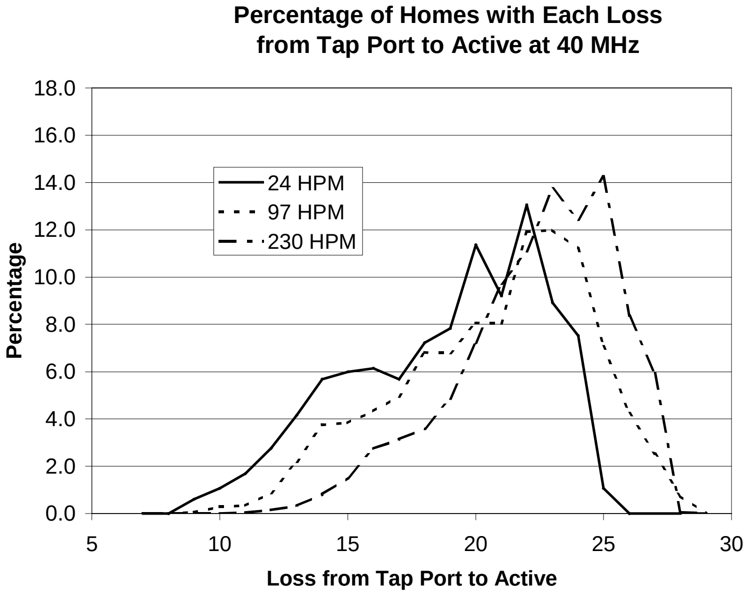

Figure 5 - Statistical Distribution of Return Path Loss @ 40 MHz

Frequency	5 MHz			40 MHz		
Homes per Mile (HPM)	24	97	230	24	97	230
Homes Passed in Sample Area	217	1306	1772	217	1306	1772
Total Gain (dB)	7.48	12.79	11.55	6.18	12.09	11.30
Average Gain per Home (dB)	-15.88	-18.37	-20.93	-17.19	-19.07	-21.19
Gain for 500 Homes (dB)	11.11	8.62	6.06	9.80	7.92	5.80
Gain for 1000 Homes (dB)	14.12	11.63	9.07	12.81	10.93	8.81
Gain for 1500 Homes (dB)	15.88	13.40	10.83	14.57	12.69	10.57
Gain for 2000 Homes (dB)	17.13	14.65	12.08	15.82	13.94	11.82

Table 4 - Power Addition of Spurs

Frequency	5 MHz			40 MHz		
Homes per Mile (HPM)	24	97	230	24	97	230
Homes Passed in Sample Area	217	1306	1772	217	1306	1772
Total Gain (dB)	29.75	42.92	43.33	28.69	42.39	43.15
Average Gain per Home (dB)	-16.98	-19.40	-21.64	-18.04	-19.93	-21.82
Gain for 500 Homes (dB)	37.00	34.58	32.34	35.94	34.05	32.16
Gain for 1000 Homes (dB)	43.02	40.60	38.36	41.96	40.07	38.18
Gain for 1500 Homes (dB)	46.54	44.12	41.88	45.48	43.59	41.70
Gain for 2000 Homes (dB)	49.04	46.62	44.38	47.98	46.09	44.20

Table 5 - Voltage Addition of Spurs

5.0 Conclusion

It has been demonstrated that the expansion of CATV networks to include broadband communications may have troublesome consequences if the pitfalls are not recognized. Specifically, return path transmitters have the potential to disturb forward services, a definite no-no. This can occur both for concurrent users in the same household, and possibly in a neighboring household. Aggravating the situation is the lack of quality off-the-shelf gear that finds its way into consumers' homes. A second important issue is the potential for ingress in the return band that is caused by the CATV equipment and consumer electronics. While upstream garbage is a well-known fact of life, any further obstacles in the upstream that are generated by this equipment should no longer be tolerated in new designs, in anticipation of the return path bottleneck on the way. Also, a strategy to deal with already deployed offending hardware must be thought through to make the most of the limited available return bandwidth.

Acknowledgments

The authors would like to thank Mr. Detlev Wundershock, Mr. Mike Short, and Mr. Dick Gresko for their invaluable contributions to the successful completion of this paper.

Mr. Glaab and Mr. Waight work in the Advanced Network Systems unit of General Instrument. Mr. Stoneback and Dr. Howald are with the Transmission Network Systems unit. GI's new address is 101 Tournament Drive, Horsham, Pa. 19044 (800-523-6678).