Increasing HFC Capacity: Design and Field Test of Return Path Frequency Stacking

Robert Howald, Michael Aviles, and Frank McMahon
General Instrument Corporation

Abstract

Thinking ahead during HFC deployment and upgrading is the key to assuring the network's long-term success. In the race to build the information superhighway, it is important to recognize that now may be the time to add the few extra lanes that could make the difference between smooth sailing and gridlock. For the case of HFC, this means making the most of the performance-constraining return path. One cost-effective way to expand return capacity is to allow each node port to be translated to its own 35 MHz of RF bandwidth using frequency conversion. This paper describes system analysis, key hardware, and summarizes the key results of the first known field trial of a complete node-based frequency stacking system.

Fattening the Pipe

For Hybrid Fiber Coax (HFC) infrastructures, there are a couple of ways to assure that the network is future-proof as the subscriber base for two-way communications and multimedia services continues to increase. HFC utilizes analog fiber optic transport between the Headend and the neighborhood node. Fiber optic nodes (optical to RF transducers) located throughout the community output the broadband downstream onto coaxial cables, through which it is then transported to subscribers as in an all-coaxial network. In the return path, signals travel this same coax, and return on an upstream fiber. A single neighborhood node branches into multiple coaxial outputs from the downstream fiber, serving from hundreds to as many as 2000 homes.

In the return band, all 5-40 MHz node upstreams share spectrum on a single fiber. As more services are deployed and subscribers lined up, the constrained bandwidth will become a sure bottleneck, unless the network is designed ahead of time to be ready for the onslaught. One solution is to add multiple fiber and lasers, such that there is a return laser and receiver for each RF port on the node. However, this is not a particularly low cost or low power solution, and is wasteful of the generous bandwidth available to implement a laser to transport only about 35 MHz of bandwidth. A more cost effective and efficient method of expanding capacity allows each subset of the subscriber community sharing a fiber optic node port to have their own 5-40 MHz return. Frequency division multiplexing the returns to the node using frequency stacking, also called block frequency conversion, can do this. Then, each port on the node accommodates a unique 35 MHz of bandwidth, providing N times the capacity for an N-port node, and also isolating port-to-port ingress. This composite signal can then modulate the single laser, and the upconversion function can be inverted at the Headend by downconversion. The frequency stacking system (FSS) concept is shown in Figure 1. This paper will outline the relevant issues involved in the design and implementation of a robust, high performance, FSS.

Communications Issues

The design goal of any hardware added in the middle of the pipe is to make sure that nobody on either end knows it is there. This means allocating specifications such that any degradation introduced goes unnoticed by any application. Ideally, any application's link budget will be negligibly effected. Obviously, this is an imposing goal. To fulfill it would require knowledge of every potential application, the modulation technique used, and more information on the HFC plants themselves. The latter two items can at least be quantified to generate numbers based on some assumptions, while the first item would require a visionary in marketing (always a challenge!). Today's fast-paced markets in telecommunications and wireless often place technology companies and their product developers back on their heels, and predicting services and take rates for two-way HFC is one such instance.

There are several primary enabling technologies that will be big players as HFC digital transport mechanisms. Among these transmission and modulation schemes are QAM, in conjunction with TDMA and FDMA, and possibly CDMA and OFDM. Designing FSS components to support complex modulations puts significant constraints on phase noise, amplitude ripple and flatness, group delay variation, among other important parameters. In order to maintain a quality uncorrected BER in the return path (our measure of performance), the converters are each allocated a portion of the tolerable amounts of the various degradations. Care must be taken in recognizing which impairments particular transmission schemes may be especially sensitive to.

Link Design Philosophy

It must be determined what effect FSS will have on the existing link budgets, which are the key to any successful communication system design. Consider the cable modem application. Degradation of 16-QAM performance can become quite substantial with relatively small amplitude and group delay distortion, since 16-QAM has information in its amplitude. The ability to correct minor disturbances with a simple equalizer is well known, and it is not cost effective in the design of the transmission network to impose difficult specifications that would be otherwise mitigated by proper modem design. Similarly, it is well known that the impulse noise problem is best handled by proper code and interleaver design. Optimum implementation strategies would utilize concurrent development of the system and the plants. Unfortunately, existing infrastructures are already in place.

Three primary link parameters of interest for FSS are thermal noise and signal-to-noise ratio (SNR), intermodulation distortion (IMD) and spurious, and phase noise.

FREQUENCY STACKING SYSTEM

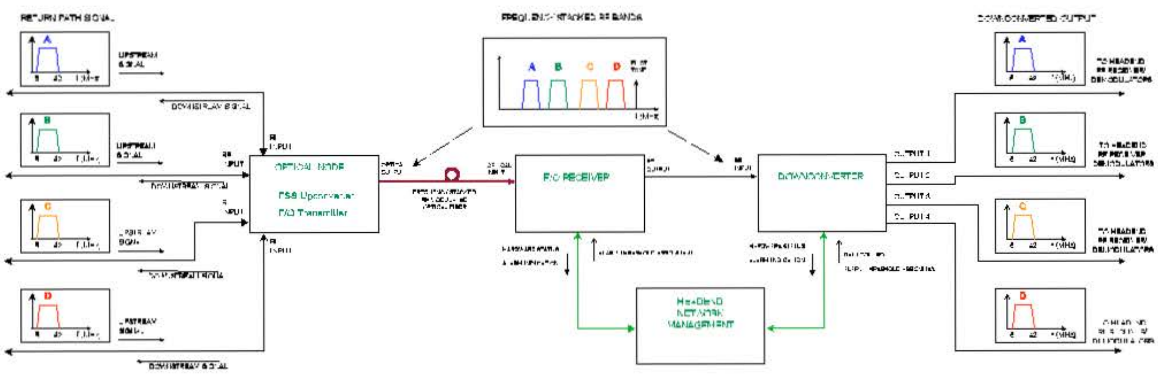

Figure 1 - A Four-Band Frequency Stacking System

Thermal Noise

Perhaps the most significant channel quality on the beneficial side for HFC networks is its inherently high SNR, where SNR implies the ratio of signal to thermal noise (AWGN). In the upstream, and even more so in the downstream, HFC networks are capable of quite high SNR's, which translates into high theoretical channel capacity. Exactly how much of this capacity can be taken advantage of is a function of how well the other impairments can be mitigated through proper modem design. For HFC networks, the primary contributor to the thermal floor is the fiber part of the system. Other noise contributors include the coaxial part of the plant, particularly in the noise funneling upstream, the noise figure associated with the upconverter at the head of the RF cascade, and, to a lesser extent because of location, the downconverter and demodulator.

The Fiber Optic Link

The fiber optic portion of the network, consisting of the laser transmitter, fiber optic cable and optical receiver, typically dominate link SNR capability. The RF/cable portion of the network, and the post-optical receiver electronic hardware, usually contribute in only a small way to overall SNR. Fiber optic limitations generally reveal themselves by two means: unacceptable minimum SNR and distortion effects, including those associated with excess loading, causing clipping. The initial setup and operation of a return link requires a different philosophy than that for the forward path, where a fixed number of signals are located within known video bandwidths at constant levels. The simplest assignment of return signal levels, although not ideal from a communication link perspective, is on a per bandwidth basis. The approach has important implementation advantages, such as its setup and testability. Also, it yields a constant SNR for all channels regardless of bandwidth, and allows operators to be prepared for eventual full channel loading without having to adjust signal levels. The total power allocated for return services is determined by the recommended composite signal level at the laser transmitter needed to maintain acceptable SNR yet avoid clipping effects.

Optical Parameters Effecting SNR

The main contributor in HFC links to SNR degradation is the laser diode used in the transmitter. Unless operating through the longer fiber networks, the transmitter diode's internal noise limits SNR. The noise is quantified as relative intensity noise (RIN). The RIN of a diode is expressed in dBc/Hz and depends on the type of laser used. The two types in use for HFC are Distributed Feedback, or DFB, and Fabry-Perot (FP). The RIN is typically anywhere from -110 to -160 dBc/Hz. DFB's normally have lower noise characteristics than FP's. Thus, for high quality digital communications, DFB's appear best suited to providing good SNR at the low signal levels desired to avoid laser overdrive and clipping in heavily loaded returns. Using RIN, together with the optical modulation index (OMI), the SNR associated with the laser diode section of the optical link can be found. Figure 2, CNR vs. Fiber Length (non-FSS), shows this value, identified as CNRtx, to be a straight line at 41 dB (carrier-to-noise, or CNR, is often utilized to be more consistent with analog video CNR, already familiar to the industry). CNRtx performance tracks on a dB to dB basis with RIN variation.

At the optical receiver, shot noise in the photodiode (typically PINs) contributes to SNR degradation. Shot noise limits are determined by the optical power and diode responsivity. Post detection RF amplifier circuitry also contributes to SNR degradation, its effect usually defined by an equivalent input noise current (EINC). Degradation due to shot noise tracks on a dB to dB basis with the optical input power level, and the variation due to EINC tracks on a 2 dB to dB basis with the optical power. In Figure 2, the two contributors are combined and the value is identified as CNRrx. The plot shows CNRrx to vary about 20 dB versus length. Overall link performance can be found by combining transmit and receive performance. As can be seen from the Total CNR, for fiber lengths out to 20 km, the laser transmitter is the dominant contributor to link degradation in this network. The receiver is not a major factor until distances of 25 km and longer are reached. From a BER perspective, the lowest SNR at the longest link is well above the minimum required to support the return services anticipated.

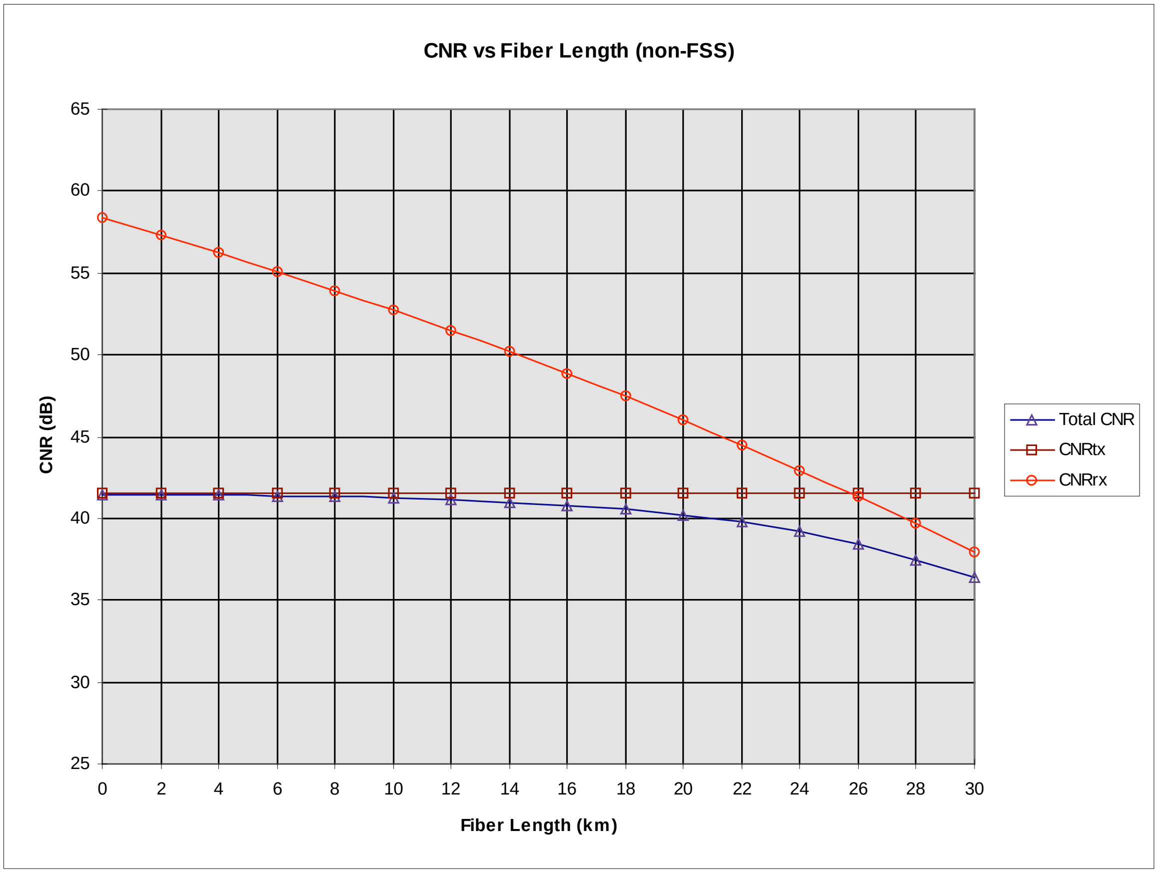

Figure 2 - Performance w/o FSS

Link Performance with Frequency Stacking

In Figure 2, without FSS, the link was shown capable of providing greater than 35 dB SNR always, and typically > 40 dB. Now consider the additional noise generated by the up and downconverter hardware. Figure 3, CNR vs. Fiber Length (FSS) shows what the addition of frequency conversion hardware does to the overall link performance.

At the shorter fiber lengths, for constant downconverter input power, optical receiver gain control attenuation settings are highest, degrading the subsequent CNRrx from the non-FSS system as shown in Figure 3. This has little effect on the link since the laser diode noise is still dominant. For these shorter links, the contribution of the FSS upconverter causes minimal degradation, by design, to the CNRtx under all conditions. At the longer links, the added noise associated with the receiver/downconverter hardware is masked even more by the equivalent noise degradation due to fiber losses. The result is even less difference in total SNR between FSS and non-FSS. Thus, incorporation of a properly designed FSS has minimal impact on an operator's ability to provide quality return path services while significantly increasing subscriber density per node.

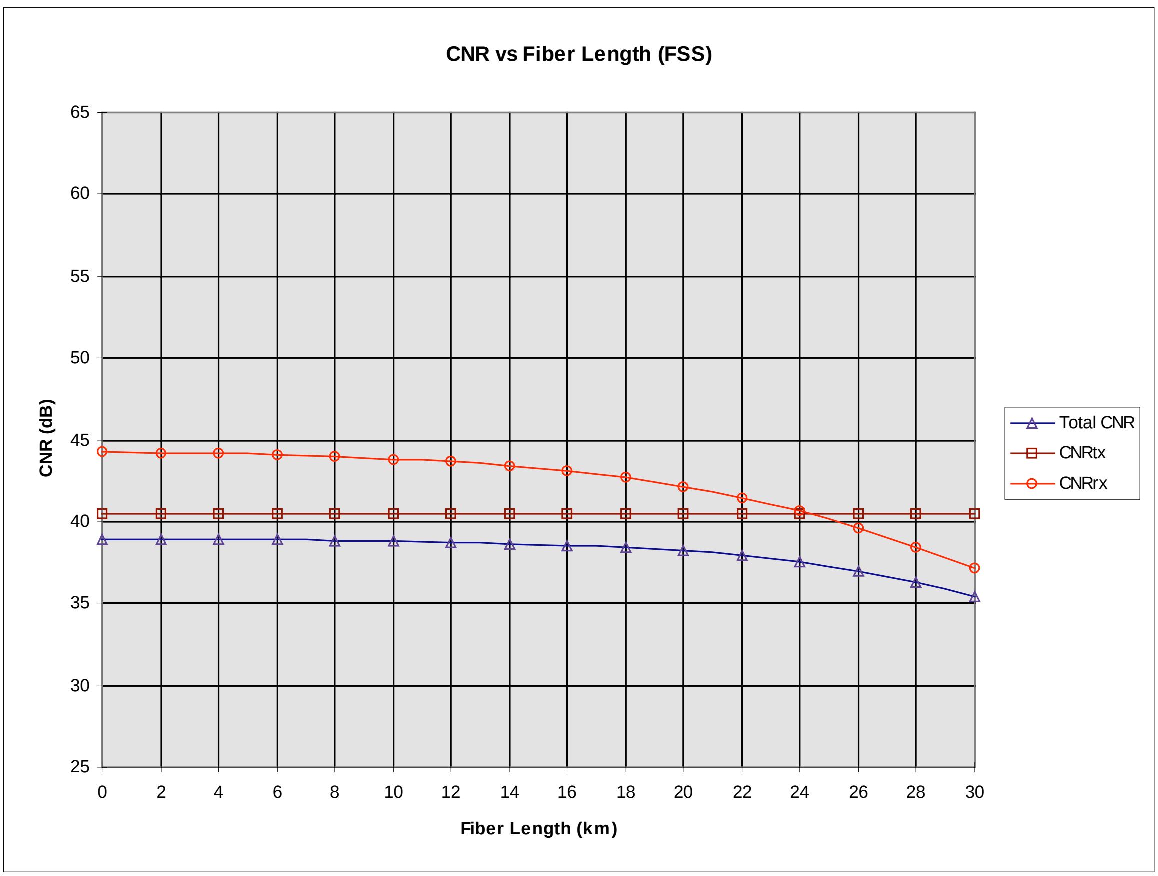

Figure 3 - Performance with FSS

Channel bandwidths above are not discussed, because the power-per-unit Hz allocation equalizes SNR in any channel. However, to discuss performance in terms of BER and data rate, SNR and bandwidth are needed. For video, this bandwidth is about 4 MHz. This bandwidth, at the SNR's calculated above, is adequate for high performance M-QAM with significant margin of between about 10-20 dB for QPSK through 64-QAM. For example, 16-QAM at 1e-8 symbol error rate requires about 22 dB of SNR. Naturally, QPSK needs less SNR, 64 QAM more, etc. The 4 MHz of RF bandwidth would represent at least 8 MBPS for 16-QAM.

Intermodulation Distortion (IMD) and Spurious

Use of frequency synthesis and conversion hardware results in the need to analyze and quantify intermodulation and spurious performance. Because of the broadband nature of the system, multiple intermodulation beats exist, and the number grows drastically as the number of signals increases. Of most interest are products that contribute to the degradation of digital communications performance by causing a significant signal-to-interference ratio (S/I). The products that dominate broadband performance can be either second order or third order beats, in contrast to a narrowband system which can often ignore second order products. Any part of the RF chain called on to process a broadband input and produce a broadband output must be concerned with second order products. Between these items, any filtering that reduces second order products will benefit second order performance. For HFC, the second order performance is typically laser dominated. Another advantage of frequency conversion is the freedom to design a frequency plan that helps mitigate second order interference.

The third order intercept (TOI) is typically used to characterize third order intermodulation characteristics of RF components. Third order products require consideration of the effects of multiple carriers, as this degrades the overall third order intercept of the cascade relative to the common two-tone reference. Unlike noise figure, cascaded intercept point is typically dominated by elements at the end of the chain.

In broadband systems, care must be taken in understanding the many possible sources of spurious frequencies. Unwanted signals can point to many areas: undesired higher order mixing products, frequency synthesizer related and local oscillator (LO) leakage spurs, intermodulation distortion in active components, DC power distribution, RF leakage, etc. Spurious contribute to S/I degradation, decreasing link margin. Locally generated spurious in the upconverter in the node can have the capability to be large relative to incoming signal levels. Proper design for adequate S/I is a combination of proper RF chain gain and intercept allocations at full load, and quality RF circuit board design.

Phase Noise

Traditional CATV frequency synthesis techniques result in typically phase noisy carriers, because analog video requirements are non-demanding in this regard. Thus, low cost, direct divide

PLL synthesis is dominant, and, combined with fine resolution of desired channel frequency outputs, results in high division ratios and the resulting noisy output. One of the key items to recognize in any digital communications system being implemented over HFC is that noisy synthesizers are generally not well suited to reliable digital communications, particularly of the high M-QAM variety. And since phase noise is a burst-type error mechanism, to mitigate via FEC requires interleaving, burst correcting codes, or both. Much can be gained by making relatively simple modifications to frequency generation. Once again, with the freedom to choose a frequency plan, designs that minimize phase noise are possible. Because of the myriad of applications to be served, the most useful specification of phase noise is to quantify its rms jitter performance over each decade of offset. This allows ease of identification of the portion of the phase jitter spectrum of significance to each application modem.

In addition to generating low noise local oscillator signals, implementation of frequency tracking benefits every application. The amount of frequency error that is acceptable varies by application, since there are so many different types of modems both existing and being developed. Lack of standardization for return path transmissions has resulted in the proliferation of various techniques, including modulations such as FSK, QPSK, QAM, and signaling strategies like CDMA, or OFDM and its wavelet-based cousins. A zero-frequency error approach requires a tracking PLL in the downconverter, and eliminates on both sides the need for very high stability references. In addition, the pilot recovery PLL serves to track out some of the upconverter's phase noise contribution. Zero frequency error means the FSS can be ignored in the complex analysis and allocation of requirements of synchronization, and, in particular, for sensitive burst modems, CDMA, and OFDM applications. Designs based on free-running crystal references contribute to frequency offset that must be handled by an application demodulator, which means that the FSS has become intrusive.

Upconversion in the Node

In terms of RF hardware, frequency conversion is a considerably mature technology. It is, however, unique in end-to-end design to nearly every application. For this case, upconverter design constraints include size, environment, power dissipation, induced phase jitter, spurious generation, and nearby image frequencies. What the design can take advantage of is that the signal levels are not high in this part of the plant, and that there is a zero gain requirement.

A dual conversion (see Figure 4) approach provides a good compromise of straightforward filtering of images and LO leakage, mixer product spurious management, and commonality of parts in getting from the multi-octave 5-40 MHz input to UHF outputs in four isolated bands. Converting all four bands eliminates multi-octave RF design in the node, and allows ease of laser implementation using ordinary forward-band units.

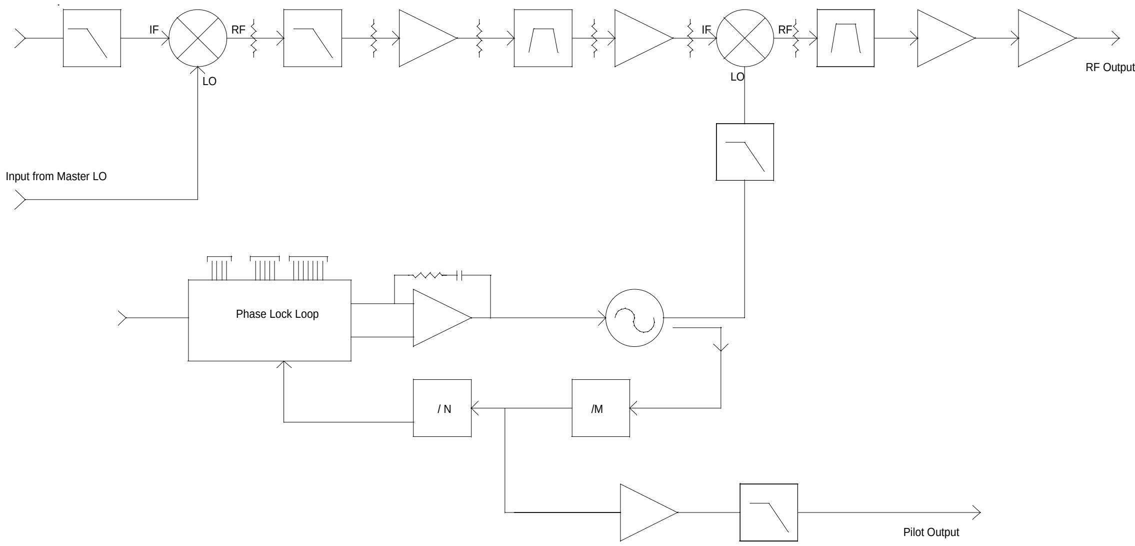

Figure 4 - Dual Conversion RF Design

Because of its location in the chain, upconverter noise figure can contribute to overall RF-related thermal noise. As such, it is important for cascaded gain blocks to be mixed into the front end and evenly dispersed. However, for IMD, filtering of wide out-of-band mixer products prior to the amplification is important, while simultaneously providing sufficient termination of the mixer ports so its performance does not degrade. Because of the multi-octave input, the first mixer and associated circuitry are critical to spurious performance, and therefore it is important that a quality mixer be used, and that the amplifier after the mixer have good dynamic range. These elements and isolation of circuitry drive the spectral purity of the design, provided proper matching around the high selectivity IF filter and output RF bandpass filters is maintained.

Other critical performance parameters in the RF path include the amplitude and group delay responses. The use of high selectivity filters can have consequences in both aspects, and can be troublesome should equalization of sensitive modulations be ignored. Another key RF parameter is the forward path isolation. Because the upconverter shares a compact, highly integrated RF environment in the node, isolation of return and forward signals, mostly analog video, is very important.

The frequency synthesis part of the design uses a common PLL synthesizer IC for all LO's, such that only varying of the divide ratios in the PLL are required. To be compatible with M-QAM, the design of this subsystem was based on maximizing phase comparison frequency for minimum divide ratio, and optimizing loop filter design values. In addition, because frequencies are not required to be programmable, low cost, low noise, narrowband discrete VCO's can be implemented. Integrated rms phase jitter on the order of less than one degree rms is typical over a 100 Hz to 1 MHz range.

Finally, zero frequency offset is achieved in the link by using a pilot tone, as previously described. With the flexibility of choosing a frequency plan, the signal can be placed well out of band of payload traffic.

Downconversion at the Headend

At the Headend, the purpose of the downconverter is to take the return path RF signal, consisting of four individual bands stacked in frequency, extract them, separate, and downconvert each to the original frequency bands. The downconverter is also typically required to interface with Headend network management equipment. As described, one important characteristic of block conversion is the frequency error introduced. For zero frequency offset at the unit's output, the downconverter implements the pilot tracking PLL, which is used to exactly re-derive the LO frequencies generated at the upconverter.

A critical design requirement for the downconverter is its spurious performance. Since the output signal band of the downconverter may exceed three octaves, special care must be taken in assuring highly linear amplifiers and mixers. Other important parameters, again, include amplitude flatness, phase linearity, induced phase jitter, noise figure, gain control, electrical isolation and power consumption.

Spurious signals can be generated in the downconverter through the mixing, nonlinearity in PIN diode and FET attenuators, and in switches. These distortion products need to be low enough not to interfere with desired transmissions. An important step in minimizing the levels of in-band spurious is in the analyzing of mixer products for the chosen frequency plan. Having specified the approach, functional block performance allocations are defined using the system level requirements for gain, noise figure, output signal level, second and third order intercept performance, gain control range, etc.

Determining the downconverter's output signal level is dependent on the number and types of return path services to be supported. The larger the number of different services accommodated, the larger the number of RF power splits, and correspondingly the higher the RF losses between the downconverter and application demodulators. The output signal level required is defined by the range of level requirements for the various demodulators, adding the splitting losses for present and future services, and then providing some level of margin. In order to provide sufficient output signal level in each output arm, the downconverter may have to deliver output signals on the order of 40 to 50 dBmV per converter channel (i.e. 35 MHz bandwidth). In order to accommodate various Headend configurations, it is desirable to provide some level of gain control within the module.

Performance in the Field

In the fall of 1997, GI's FSS was deployed in a field trial of the new SG 2000 node. The upconverter was installed in an existing four-port node configuration (the upconverter in the node was designed to be field replaceable as an identical form fit to an existing passive combiner RF board). Two RF ports on the node were connected to a functional plant, and the second two were wired to a motel room, where a QPSK modulator, taking in a pseudorandom bit stream, was located. At the Headend, located at the end of a relatively short fiber optic link, a QPSK demodulator followed the downconverter. The returns in this case implemented 5-42 MHz bands. Performance testing consisted of measuring the error rate statistics and average BER performance on one band for a period of time, and subsequently rotating through each band repeatedly. Because the network in question had no operating return prior to the testing, only ingress characterizations prior to running the error rate testing were available to gauge the nature of the return being used. Note that each band (i.e. each node port) had the QPSK upstream signal summed in, but only one band at a time was measured.

Unmaintained Plant

Before performing any plant upgrades in anticipation of employing return services, data was taken with the QPSK operating, providing a measure of the raw networks' readiness for digital communication. It is well known from return HFC characterizations, ongoing for about five years now, that major impairments that exist include noise funneling, narrowband ingress, frequency response distortion, impulsive noise, and 60 Hz interference coming in both hum and impulsive varieties. Because of these known problems, equipment being designed today is employing sophisticated equalization and error correction techniques to provide mitigation. Some equipment manufacturers are implementing advanced modulation and signaling approaches, such as the CDMA and OFDM, built specifically to mitigate these known impairments. For this test, we selected the simplest practical scheme anticipated for modern advanced services, QPSK, and did not implement any error correction. Thus, raw data availability parameters could be obtained, as well as important BER data, to help characterize the plant's capacity for digital communications. For the length of the test, the QPSK data rate was 2 MBPS, using about 1.5 MHz of RF bandwidth.

Single Signal Testing

On the raw plant, two weeks of QPSK data showed obvious impairments to uncorrected transmission. There was a very high degree of channel availability, as measured by the percentage of error-free seconds (EFS) and severely errored seconds (SES). In other words, there are very long periods with no errors (typically measuring 99.75% of the time), followed by an impulsive burst of errors (about .2% of time). The percentage of time without SES was therefore greater than 99.95%. This is important, because even rudimentary error correction can fix errors not related to severely errored seconds, because they tend to be more randomly distributed. The severely errored seconds, which occur in bursts, are more difficult to correct. This is the reason for the strong recognition of sophisticated error correction in the Multimedia Cable Network Systems (MCNS) specification (a standard for cable modems). The forward error correction (FEC) to be employed consists of a concatenated trellis and Reed-Solomon implementation, as well as interleaving. Interleaving is the technique by which the symbol sequence is transmitted such that some designed number of symbols, ideally associated with the expected burst statistics separates adjacent symbols. Thus, it is implemented specifically to aid in burst correction.

MCNS FEC specifications were built around the anticipated statistics of return burst-type interference, of which very recent studies have indicated a strong presence of the power-line related type. These findings indicate an important need to provide quality AC distribution and grounding in the plant, as well as the need to consider the nature of home-generated disturbances associated with major appliances. RF ingress levels associated with HFC returns, both steady-state and impulsive, have been accumulated by many sources for statistical analysis. The sources of ingress are quite well understood, and most will be present on every two-way HFC plant, no matter how clean. However, it is important to point out that ill-maintained plants, be it by poor grounding, poor in-home wiring, and/or poor connections, will significantly aggravate ingress levels at the Headend, potentially effecting the EFS performance and the ability of errors to be corrected.

Finally, looking at average BER during this test is also somewhat informative. For the length of the single signal testing, the QPSK power was set such that it represented the level as if the modem had to share a fully-loaded channel using a power-per-Hz allocation methodology. In other words, the input signal used was about 14 dB below (10log(1.5/37)), the maximum port input suggested. Looked at yet another way, the QPSK level was scaled to match its percentage of band occupancy in the 37 MHz return. Raw data indicates that, not counting SES periods, average BER's on the order of E-10 occurred during quiet times, on the order of E-8 and E-9 during nominal periods, and on the order of E-7 in the worst periods. Including the SES periods, there was variations on a per-day basis of days as good as E-9, nominal in E-5 and E-6 range, and E-4's at the poor extreme. Measurements were auto-recorded from the BER tester, and it is important to note that the lowest extreme (the zero-error limit) of the BERT is, in fact, 1E-10. Thus, periods of zero errors would correspond to this average BER as measured by the BERT.

Noise Loaded Testing

After about two weeks of gathering data on the unmaintained plant, the MSO implemented a well-coordinated effort to go through the plant methodically, tightening down all connections, and assuring good plant grounding and powering. Not all details were immediately available about every maintenance item addressed, as an outside contractor performed much of the work. Following the plant upgrade, the single modulated signal test was repeated, and it was immediately apparent that there were zero errors 100% of the time. This behavior continued on upstream band 1, until, after continuing to count zero errors, verification on the other three bands was done to make sure things were connected and ready to go to the next test step - noise loading.

To further demonstrate the capability of the FSS-2000 platform, the unit was subjected to a full loading of every band. The input of each port has the maximum suggested input level, uniformly spread across the full 37 MHz of bandwidth. The white noise occupied the entire 37 MHz band, except for a small portion. Using a notch filter, part of the spectrum is cleared out, and inserted in this open real estate is the QPSK signal. A plot of this "noise notch" transmit signal is shown in Figure 5. This loading configuration harshly tests the FSS and fiber link dynamic range capability, as Gaussian noise samples have a higher likelihood of producing large peaks capable of clipping the laser as well as driving RF amplifiers into saturation briefly. The fully loaded spectrum also fully stresses the ability to provide highly linear 5-40 MHz RF outputs from the downconverter at a high signal level.

Results of this test were also extremely encouraging. The same basic procedure was implemented, where the error rate measurements were accumulated on one band at a time. This test lasted for about one week, with data again recorded around the clock. On bands one and two, there were nearly identically zero errors during all measurement periods (it is reasonable to assume that each band had one fourth of a week of measurement time). Band one still actually had zero errors, and band two counted two bit errors in 40-some hours. On bands three and four, there were more errors, and logically so, since these two bands were connected to the operating plant. As such, they were exposed to the sources of upstream ingress. Even given that, over 40 hours of monitoring band three produced only about 350 bit errors, while band four showed only 651 bit errors. This was quite astonishing, but, given that the plant was relatively small (about 150 homes, 75 on each port), there were fewer sources of home-generated ingress. With ingress correspondingly reduced, excellent BER resulted, even without error correction.

While QPSK is a very robust modulation, this performance, without any error correction, was better than had been anticipated. The dynamic range of the FSS had been thoroughly and harshly tested, and performed admirably at the high end. Noise power ratio (NPR) tests show virtually identical performance (about 41 dB) between a system with FSS and without FSS using the suggested power loading.

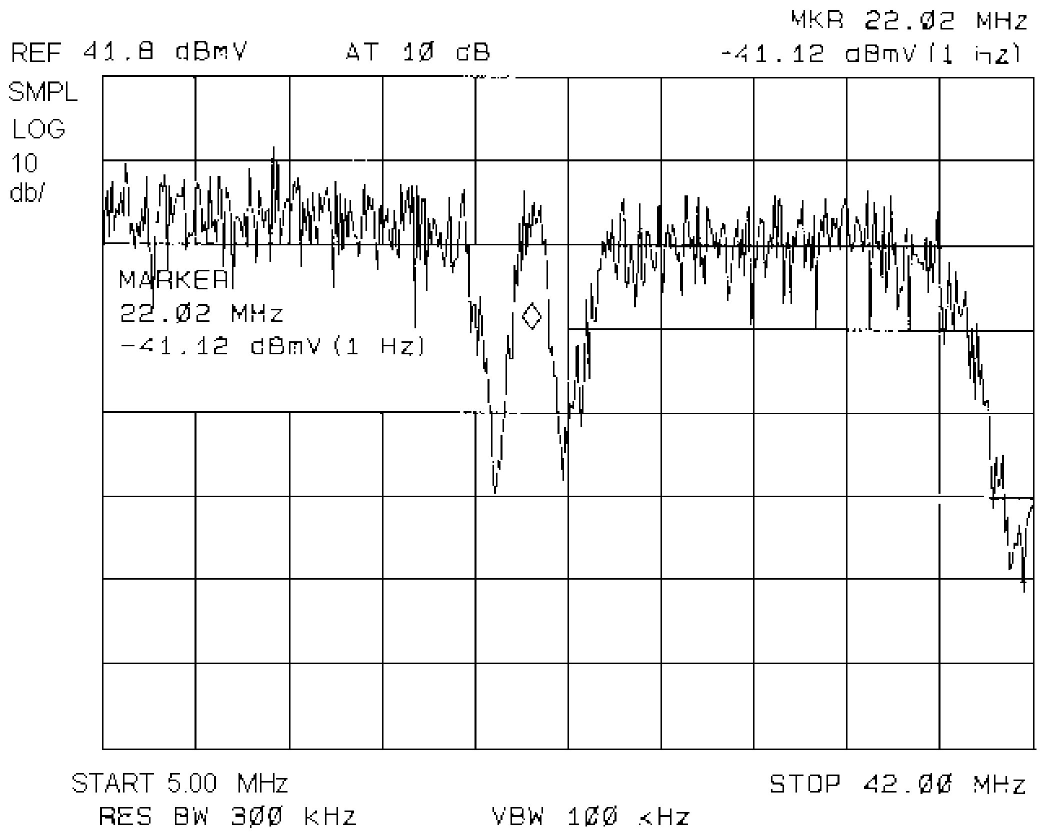

Figure 5 - Noise Load Testing

Conclusion

This piece has discussed a straightforward and cost-effective way to add capacity in the return path for HFC networks. More importantly, FSS has been proven through link testing, both in the lab and in the field. While telecommunications roll-out onto cable networks has been slow, if MSO's are committed to the growth of digital data transport in the return path as an important revenue stream, now is the time to provision that network for this growth. Technologies such as FSS effectively multiply the shared bandwidth, allowing the revenue stream to grow without sacrificing network quality.

Acknowledgements

The authors would like to recognize Mike Short and Gary Hitchner, whose efforts were critical to the completion of this paper.

Dr. Howald (215-323-2276) is the group leader of Systems Engineering in GI's Transmission Networks Systems unit. Mr. Aviles (215-323-2269) is the lead systems engineer for the FSS product line. GI is located at 101 Tournament Drive, Horsham, Pa., 19044.