Managing the Return Spectrum to Optimize Interactive Revenue Opportunities

James O. Farmer
Antec Corporation

ABSTRACT

Issues related to reverse transmission in cable plant are considered. A topology is described which removes the largest portion of the interference in the return band. A few of the issues related to protocols suitable to return data are mentioned, and the idea of spectrum management is introduced.

INTRODUCTION

For years the cable industry has had available technology that *almost* enabled the return band from 5 to 30 MHz. Some systems have employed the return band to transport video from local venues to the headend, and a very small number of systems have employed the return path to retrieve data from homes. However, little use has been made of the return spectrum until now. Cable operators have recognized that they will have to provide return for new services, however, and most systems now either have return path or can add it. In the days when systems were built according to tree and branch architectures (and there are a lot of such plants still in use), the return band was frequently considered impractical to use. However, with the modern fiber to the node architectures, the return band is usable.

A survey of several cable operators indicated expectation of revenue from services using reverse spectrum, ranging from $3.00 per month (average for 100% of homes served) to over $50.00 per month (in 25% of homes served). Revenue comes from such services as on line service (Prodigy, AOL, Internet, etc.), telephony (including long distance access) and interactive advertising.

ISSUES IN SELECTION OF A FREQUENCY PLAN FOR UPSTREAM USE

Any attempt to set a frequency plan for upstream traffic must first take into account existing services if any. Most of the existing services are either video return from remote venues, or impulse pay per view systems with RF return. While we expect future services to be controlled with some sort of intelligent system, existing systems will continue to be controlled manually. Future control systems are discussed below.

Interfering carriers will also render some frequencies inoperable for all except perhaps very low data rate services. A review of the literature suggests that frequencies below 10 MHz are very difficult to use, and that frequencies above 20 MHz are much less susceptible to interference. A point of concern with the use of frequencies above 30 MHz, is the potential for interference to TVs and other devices connected to the cable. TVs in North America use the band from 41-47 MHz as their intermediate frequency. A typical level for a return signal leaving the home, is about 50 dBmV according to current thinking. This is as much as 50 dB above the incoming signal the TV is receiving. If a significant part of the upstream energy strikes the TV (due to limited coupler isolation), then interference is possible.

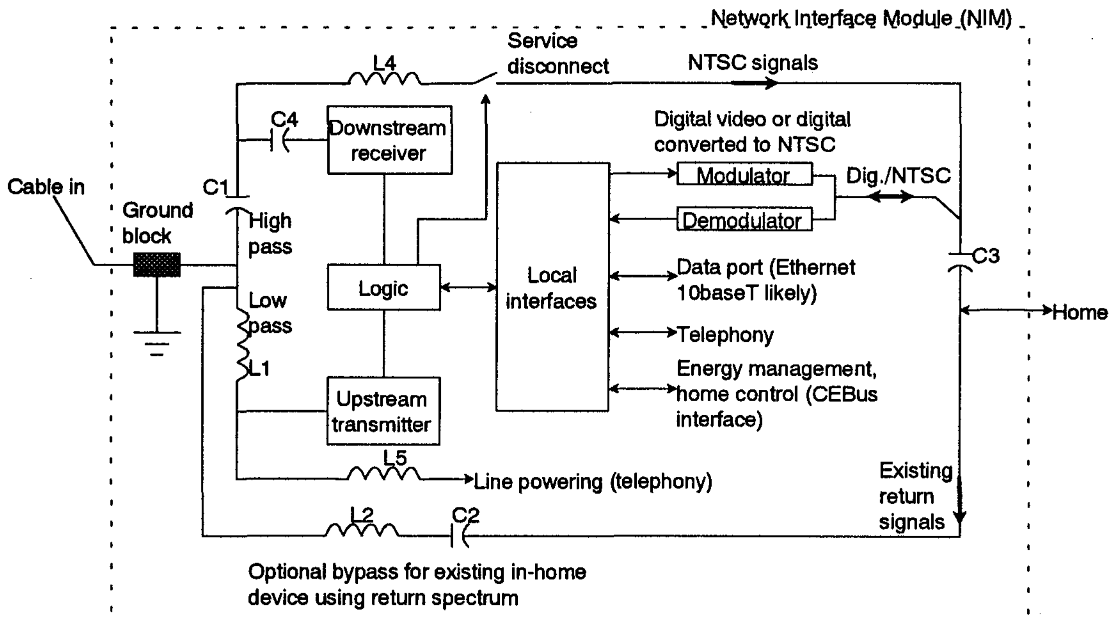

Figure 1. Configuration of a Fairly Complex NIM

INTERFERENCE IN THE UPSTREAM BANDWIDTH

Many studies have been conducted of the interference in the upstream path. One of the few conclusions that may be drawn from these studies is that the interference varies considerably from one system to another, and within a system, from time to time. Most investigators do agree, however, that the primary source of interference is the home. Drop cable is responsible for the next largest contribution, but is much less of a contributor than is the home wiring. Finally, a little interference may be picked up on the distribution plant, but this is by far the smallest source of interference.

The above considerations lead one to favor an approach to return architecture, in which reverse signals are added in a box on the side of the home (preferably inside, to reduce the cost of meeting the temperature range), located ahead of any subscriber devices. In this case, it is possible to provide for a high pass filter on the cable entering the home. The filter will reduce interference generated in the home, and will also protect the subscriber's equipment from the effects of reverse signals in the 30-40 MHz band. Such interface boxes have been called Network Interface Modules (NIMs) or Network Interface Devices (NIDs). Other names we have seen include Coaxial Termination Unit and Premises Interface Device.

Figure 1 illustrates one possible NIM configuration. This is a rather complex NIM, in that it is provisioned to handle not only downstream NTSC television, but also upstream signals from an RF impulse pay per view (RF IPPV) box of the type in existence today, and bi-directional signals from a data port, a cable telephone and a home control system. The likely configuration for such a device is modular, so that only the functionality needed will be plugged in.

The signal enters through a ground block, which may be included if the NIM is mounted outside the premises, but not if it is mounted inside. The C1 high pass section couples all the signals above 50 MHz, including downstream NTSC and digital signals.[i] The L1 low pass section couples upstream signals below 40 MHz, to the cable. In a certain band, signals from an existing set top box having an RF impulse pay per view module, are coupled from inside the house to the incoming cable, through bandpass filter L2-C2. In this band, any noise coming from the home is coupled to the cable, rendering this spectrum no more useful than it is already.

The upstream transmitter in the NIM is able to utilize any return spectrum not used by the in-home RF IPPV device. Presumably any new upstream equipment is frequency agile and managed by a management system as shown below. Notice that, with the exception of the band passed by L2-C2, no energy in the upstream band is coupled from the home to the distribution plant, so the major source of return interference is eliminated. Only systems using RF IPPV will need to provide the L2-C2 bypass.

As shown, NTSC downstream and possibly digital signals, pass through the C1 high pass section, where they are presented to the inside wiring. A service disconnect is shown to allow for disconnection of all services. We are not sure if this will be demanded of a box that normally sits in a homeowner's basement or garage: the chance of bypassing the box may be so great that the service disconnect is not considered appropriate.

Low pass section L4 is needed to attenuate the digital (550-750 MHz) band under some circumstances. A downstream receiver accepts data signals that the NIM is to process. These signals may, depending on the services ordered by the subscriber, include telephony, data, home control, and perhaps even digital video programming. Elsewhere in the NCTA Technical Papers for 1995[ii] is shown a system in which digital TV signals are extracted in the NIM and either decompressed and converted to analog in the NIM, or transmitted into the house to a lower cost digital set top converter.

Additional services may be provided at this point. If data services are provided to the subscriber, many operators are favoring providing a 10baseT (Ethernet, twisted pair, 10 Mbps) interface. This interface is becoming standard on many personal computers, and low cost boards are available for computers not so equipped. The 10baseT interface uses RJ45 connectors (big brothers to the RJ11 telephone connectors), and twisted pair cable. The cable is much easier to wire than is coaxial cable. We question the philosophy of putting computer interface connectors on set top converters, because most people don't watch TV and use a computer in the same place, and wiring from the TV to the computer is likely to be deemed undesirable by many subscribers.

A second service contemplated by many cable operators, is the provision of cable telephony services. This may be provided from the NIM by inserting a line card, which carries all of the interfaces needed to support telephone services. The line card will make the telephone "think" it is connected to a central office switch, and the logic will handle the details of getting signals to and from the headend. Downstream signals are received by the downstream receiver. (In some cases it may prove more economical to provide a separate receiver for telephony service due to signal routing complexities.) The line card may be replaced with an ISDN card for provision of ISDN services, a class of service that has not penetrated the residential market to date. ISDN promises faster data services and simultaneous voice traffic with a single telephone line.

Of serious concern in the providing of telephone services, is the lifeline problem of powering the telephone when local power at the home is not available. Subscribers have long expected that their telephones will work when the power is out. Some operators are talking about not providing first line service because of the powering problem. Others are considering home powering with battery backup. This would normally work, but we must be concerned with battery reliability under temperature extremes, and that batteries do create maintenance issues in the long run. Still other operators are planning to provide enough network power to operate the telephone interface without commercial power. This power is either provided on a separate twisted pair "siamese" cable, or on the center conductor of the coax. Taps are available now for the siamese cable, and will be available shortly for center conductor powering. In the figure, L5 extracts power for the telephone equipment, from the center conductor of the drop.

Another service of interest to some operators, is energy management and home control. Some utility companies have indicated a willingness to participate in the cost of providing the NIMs in order to gain control of the homeowner's major appliances such as air conditioning. The air conditioning and other loads, are cycled on and off during peak demand times. Alternatively the power company may increase its rates during peak hours. This service gives rise to other possible low speed data services such as security monitoring. Such services are similar to the data services shown above, but differ in the data rate (very low) and the priority of communications: if energy management commands take a minute or so to be delivered, nothing is lost. Of course, this is not true for security services. The most likely interface for these lower speed data services is the power line communications standard of the EIA's Consumer Electronics bus (CEBus), codified in the EIA IS-60[iii] voluntary standard. Communications is by power line, so no special wiring is needed (though in some cases an RF coupler is needed between the two line phases).

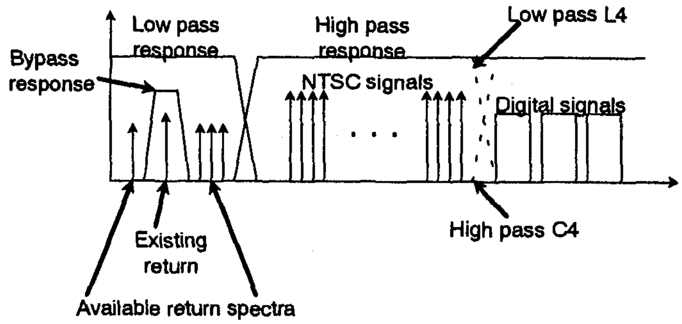

Figure 2. Spectrum Management in the NIM

Figure 2 shows the spectrum and how it is handled in the NIM. The 5-40 MHz reverse band is low pass filtered in L1 (figure 1), and contains all upstream signals. Most of the upstream signals likely will be generated in the upstream transmitter shown in figure 1. However, if an existing RF IPPV converter is in use, the bypass shown is provided by L2-C2. This spectrum bypassed would normally not be used for services other than RF IPPV.

The downstream spectrum is assumed to consist primarily of NTSC signals from 54 to 550 MHz, and digital signals above that. The NTSC signals are routed directly to the home as they are now, through L4, the service disconnect, and C3. The reason we show low pass filter L4 in figure 1, and in figure 2, is that some proposals involve reusing the 550-

750 spectrum in the home, for signals generated in the NIM. If this is done, it is desirable to attenuate the digital signals, which are not needed in the home for these scenarios.

HIGH RETURN BAND

Some investigators have considered using a high return band around 900 MHz. This has reportedly been used in one or two systems. The likelihood is that in a few years, demand for return spectrum will exceed that available between 5 and 40 MHz, even in systems built with small nodes. We believe that use can be made of the high return spectrum, but that the demand is a few years into the future. Practical filtering topologies have been successfully tested in the laboratory.

POSSIBLE DATA PROTOCOLS

Data services on cable are characterized by the almost universal need for two way communications. Some services such as telephony require symmetrical spectrum, with the same data requirements up and down stream. Others, such as most data services (on-line connectivity being a prime example), usually have need for more bandwidth downstream than upstream. Particularly in the upstream direction, the data flow is very bursty, with long periods of no communications, but with need for fairly fast response when a transmission is sent. Telephony requires low latency (that is, low delay), and arrival of packets in fixed time intervals (low jitter). Television signals require low jitter (or buffering, which can be expensive), but do not require low latency. Data services usually don't require either low latency or low jitter.

No one protocol is optimum for all types of data, but it is possible to combine elements of more than one protocol to produce a communications system that is reasonably good for all needs. In the downstream direction, data originates (or at least is codified) at the headend. Two methods of transmitting the data are frequency- and time-division multiplexing. Frequency division multiplexing (FDM) is used in NTSC transmission, in which every signal is modulated onto its own carrier at a different frequency. In time division multiplexing (TDM), different signals (or data streams) are combined on one carrier by transmitting one for a time, then transmitting another. This latter method can be more efficient in terms of the total bandwidth required. If more than one service is required, such as suggested in the NIM (figure 1), then TDM allows only one receiver to be used for multiple services, reducing cost.

In the upstream direction, FDM can be used, but is usually not chosen due to the relatively inefficient use of the upstream bandwidth. TDM is not used, as there is no point at which all upstream traffic is brought together to be formatted into a single data stream. Most of the protocols we have seen recently have involved in some way, reserved time division multiple access (R-TDMA. Each upstream talker is assigned a time slot during which he can transmit upstream. The time slots are assigned from the headend. In some cases, the time slots are fixed, with one talker being assigned exclusive use of a certain time slot. In other cases the time slot is assigned to those talkers who have indicated that they have something to say.

Other methods have been proposed to return signals upstream from multiple talkers. The most talked-about is probably code division multiple frequency modulation, a form of spread spectrum communications. Each talker is assigned the same wide frequency band and each talks at the same time. The difference is that different talkers use different spectrum spreading sequences. Knowing each talker's spreading sequence allows one to "take apart" the individual signals at the headend.

Fixed frequencies are usually not assigned for all time in the return band. At least two reasons exist for not assigning fixed frequencies. The nature of interference in the return spectrum is such that a frequency may be good at some times and not at others. Also, in order to efficiently use the return spectrum, it is necessary to allocate bandwidth only when a particular service needs it. For example, one would not like to dedicate bandwidth to a remote meter reading service that needed a few milliseconds each month to read meters in each home. Rather, one would like to use that bandwidth for another service when it is not needed for meter reading.

THE SPECTRUM MANAGER

These concerns may be addressed by employing the concept of dynamic spectrum management. This means that there exists some mechanism that knows what use is being made of the return spectrum (and downstream spectrum, too), and is able to assign different services to parts of the spectrum appropriate for their immediate needs and the condition of the return spectrum. This mechanism has been called a "spectrum manager," or a "spectrum management agent." Generally, it includes some sort of ability to assess the quality of the return spectrum, knowledge of the demands being placed on the spectrum, and the ability of different services to control their functionality. The spectrum manager is able to command different services to operate together to maximize efficiency in the return band.

The spectrum manager typically includes some sort of detector to assess the return spectrum quality by frequency, and is able to tell various services what frequency and power levels to use, and may be able to change modulation methods. A computer is assigned the task of the spectrum manager. It will allow parameter monitoring and automatic or manual adjustment. In more advanced systems, it will interface with other levels of a hierarchical system that monitors all aspects of a cable system's performance, and allows rapid pin pointing of faults. Such higher level systems are known as operational support systems (OSS).

LOGICAL ACCESS TO FREQUENCY MANAGEMENT

The following discussion describes how one logically controls the system using a frequency manager. This does not represent the physical data flow. The terminology is borrowed from the telephone industry, which has been using systems such as this for a long time.

Figure 3 illustrates the logical method by which monitoring and control of the system is achieved. A cable television system is subdivided into a number of **logical nodes.** In the limiting case the logical nodes are the same as the fiber nodes. In the interest of economy, it is likely that most systems will want to combine several nodes at the headend, until adequate revenue streams exist to support the additional equipment needed to serve each fiber node individually. Thus, for the present purposes we define a logical node that is equal to the fiber nodes that are supplied by one

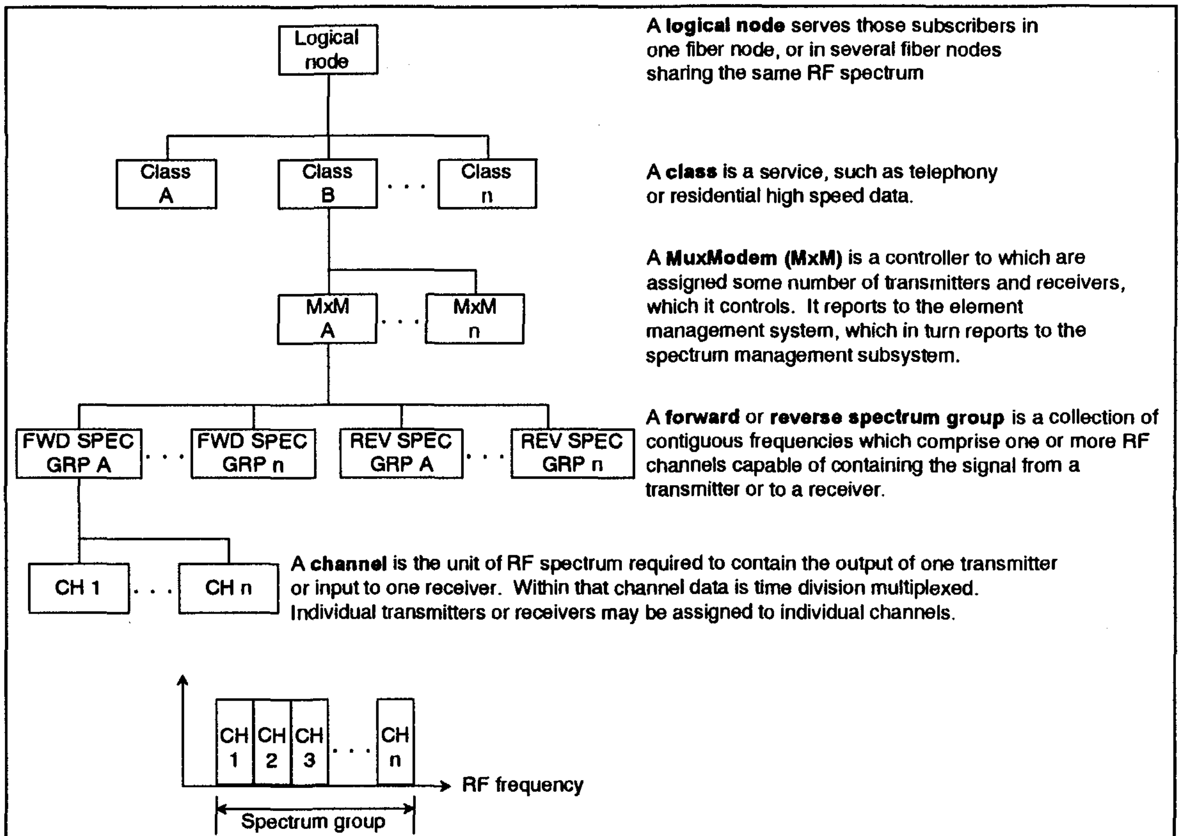

Figure 3. Logical Interfaces

fiber transmitter and which use one fiber receiver for the return path. This grouping is characterized by the need to coordinate all frequencies within the node, but no need to coordinate with other nodes.

Classes refer to services that need, in some way, differing characteristics in the data communication system. At the moment the following classes are defined.

1. Telephony.
2. Residential high speed data.
3. Switched digital services.
4. Utility communications services (including security monitoring).
5. Converter return data.
6. Integrated multimedia services.
7. Transponders (conventional cable status monitoring of distribution plant).
8. Invalid (a method to mark a certain RF bandwidth which has been identified as being unusable, for example due to ingress. The identification may be either automatic or manual.)
9. Unavailable (RF bandwidth reserved for future use, or occupied by services not covered by the system).

The **MuxModem** is the controller for all modems serving a common logical node. More than one MuxModem may serve the same logical node if adequate frequency spectrum is available. The MuxModem microscopically controls the assignment of modems to various spectrum groups and channels, within the limitations assigned to it by the element manager. In addition, it handles other aspects of modem operation. It communicates down to the individual transmitters and receivers, and up to the element manager.

The **forward** and **reverse spectrum groups** are a set of contiguous RF frequency bands which contain one or more RF channels. Each channel is capable of containing one RF path to or from a logical node. If a particular class is eligible to communicate on non contiguous bands, those bands must be set up as separate spectrum groups.

Each spectrum group is made up of one or more **RF channels**. One RF channel is wide enough to contain one transmission from the headend or from the home. An RF channel may not be subdivided within a class, as it is the minimum entity that can carry data. In different classes, RF channels are of different width (the width is expressed in KHz or MHz of RF spectrum). Thus, in one class, a spectrum group that is eight MHz wide may contain 4 RF channels, each two MHz wide. In another class, the same spectrum group may contain 8 RF channels, each one MHz wide. Within a spectrum group, all RF channels are of the same width.

Within each channel, a number of subscriber services are time division multiplexed, to enable service to many subscribers at once within that RF channel. However, since the spectrum cannot be broken down below the RF channel level, the spectrum manager doesn't care about how the channel is used.[iv] The only metric of use to the spectrum manager is the percent utilization of the RF channel. If all RF channels available to a particular MuxModem are approaching their capacity, then the spectrum manager may have to find additional RF channels in spectrum groups not presently available to the MuxModem.

UPSTREAM SIGNALS NEEDING TO BE CONTROLLED BY THE SPECTRUM MANAGER

The table on the following pages lists the services needing to be controlled by a spectrum control system, and the parameters that would need to be controlled. This is an early attempt to catalog the return services we expect to see in the next few years, but is not by any means a comprehensive list. It should serve to illustrate the subject, however. In addition to the management shown, we expect that downstream equipment will also be controlled, but in the interest of needed brevity, we omit it here. The table is discussed at its end.

Components needing control by Spectrum manager

Component	Description	Characteristics	Tuning Res.	Monitor	Control
1. Upstream data. (Includes transmitter at headend and receiver in field)	Existing upstream data for pay per view	FSK or BPSK, variable bit rates. Usually a combination of polling and ALOHA protocols	Varies by manufacturer.	Manual, cannot monitor	Must mask out manually.
2. Point to point modem, upstream receiver (headend)	Nailed up connections, continuous communications.	1.544-6.176 Mb/s above data link layer. QPSK, < 1 MHz to < 4 MHz occupied. Symmetrical data rate. Low jitter.	0.25 MHz to mesh with STD, HRC, IRC assignments.	-Receive frequency. -RSSI. -Eye closure.	-Receive frequency. -Expected signal level.
2a. Point to point modem, upstream transmitter (in field)	see 2	see 2	see 2	-Receive frequency. -Transmit frequency. -Output attenuator.	-Receive frequency. -Transmit frequency. -Output attenuator.
3. Point to multipoint upstream receiver for constant rate services (headend)	Switched connections, telephony and related services	1.544 Mb/s above data link layer. QPSK, ~1 MHz occupied. Symmetrical data rate. Low jitter.	0.25 MHz to mesh with STD, HRC, IRC assignments.	-Receive frequency. -Expected signal level. -RSSI. -Eye closure. -Utilization.	-Receive frequency. -Expected signal level.
3a. Point to multipoint modems for constant rate services, field unit	see 3	see 3	see 3	-Transmit frequency. -Receive frequency. -Output attenuator. -RSSI. -Eye closure. -Unsuccessful attempts.	-Transmit frequency. -Receive frequency. -Output attenuator.

Component	Description	Characteristics	Tuning Res.	Monitor	Control
4. Point to multipoint upstream receiver for variable rate services (headend).	Most data services (computer communications). Data delivery may include considerable jitter.	~1.544 Mb/s above data link layer (may want to reduce), QPSK, ~2 MHz occupied.	0.25 MHz to mesh with STD, HRC, IRC assignments.	-Receive frequency. -Expected signal level. -RSSI. -Eye closure. -Utilization	-Receive frequency. -Expected signal level.
4a. Point to multipoint modem for variable rate services (field unit).	see 4	see 4	see 4	-Transmit frequency. -Receive frequency. -Output attenuator. -RSSI. -Eye closure. -Unsuccessful attempts.	-Transmit frequency. -Receive frequency. -Output attenuator.
5. Telemetry receiver (headend)	Low bit rate data. For remote meter reading (poling), security (slotted ALOHA?), etc.	~2400 b/s FSK (may be faster but stay within channel)	~50 KHz occupied bandwidth.	-Receive frequency. -Expected signal level. -RSSI. -Eye closure. -Utilization	-Receive frequency. -Expected signal level.
5a. Telemetry modem (field unit)	see 5	see 5	see 5	-Transmit frequency. -Receive frequency. -Output attenuator. -RSSI. -Eye closure. -Unsuccessful attempts.	-Transmit frequency. -Receive frequency. -Output attenuator.

General notes to the table:

1. RSSI = received signal strength indicator. This is a voltage (or word) proportional to received signal level.
2. Output attenuator on a transmitter is an indirect indication of output level. We assume that in the majority of cases transmitter manufacturers will supply this information as an output level indicator, but we must have more information to derive output power. Since output power is important, we assume that manufacturers will supply a tabulation for each piece of equipment, relating power level to attenuator setting. The accuracy of the table will be only moderate, however, and a deviation of several decibels can be expected. This is important primarily in the interpretation of data.
3. "Unsuccessful attempts" is a measure of the number of times an attempt was made to access the system without success. The criteria and reporting format are not determined yet.
4. Eye closure is taken as a general measure of received signal quality. Other metrics could be used, but this seems to be a good choice.
5. "Expected signal level" in headend receivers refers to the need in such receivers to tell them what the proper signal level to which they should control the upstream transmitters. This will vary with headend configuration. A manual, or better, automatic method of setting this must be developed.
6. Other monitoring and control functions may apply to some pieces of equipment. Only those which seem related to spectrum management are shown here.

Discussion of elements of the table

1. The first entry is represented by RF IPPV systems in use today. These systems are not capable of being controlled by a higher entity, and therefore would be omitted from the spectrum manager (class 9 above).

2. Point to point modems are those intended for use between two points. For example, we are seeing new interest in T1 modems for use between business locations. These modems operate at 1.544 Mbps and can carry 24 voice channels or other combinations of voice and data in 64 Kb increments. For example, a medium size business may lease a T1 line to get all of its telephone traffic to a telephone central office or to the point of presence of a long distance carrier. This traffic is carried by the telephone company today, but is capable of being carried by cable, possibly at very competitive rates.

The service is characterized by continuous communications between two points. Communications that is always carried out between the same points is referred to as a "nailed up circuit," as opposed to a switched circuit. We refer here to a field modem and a headend modem, which are identical except for complementary frequency plans. The headend receiver is expected to report to the spectrum manager the following items.

a. Receive frequency - confirmation of the frequency to which it is tuned.
b. RSSI - received signal strength indication is a measure of the received signal strength. This information is used by the spectrum manager to control the field modem transmitter. Because the system gain from the field location to the headend is not precisely known, it is necessary to measure the signal level received at the headend, and to remotely control the filed transmit power to make the received signal level correct.
c. Eye closure - a measure of the quality of the received signal. This data tells the spectrum manager whether anything is interfering with the received signal, when the interference is low enough so as not to corrupt the data. The spectrum manager may then make a decision to raise signal levels above the nominal target, or to move upstream transmission to a new frequency.

In turn, the spectrum manager controls the received frequency and the expected signal level. The received frequency must be changed if the manager determines that the interference on the present frequency is too high for reliable communications at the highest signal level permitted, so that a frequency change is indicated. The headend receiver must be informed of the expected signal level because the loss in the headend is not generally constant. System design will dictate the target signal level at the headend from the field transmitter, based on reverse amplifier and transmitter signal levels. In general, from the input to the headend to the headend receiver, some gain or loss is expected, based on individual headend design. During system set-up, the expected signal level is computed, and the receiver so informed.

2a. The point to point upstream modem communicates with the receiver in 2 above. It reports its receive and transmit frequencies (for confirmation) and its output attenuator setting. The output attenuator setting is a measure of transmitter power output, which we choose not to call "output power" for technical reasons. It is a measure of how much power the transmitter is putting out, a metric useful for diagnosing upstream transmission problems. The spectrum manager controls the transmit and received frequencies (so it can change them based on current conditions) and the transmit attenuator (power).

3. Point to multipoint communications for constant rate services are particularly suited for telephony and related services requiring constant symmetrical data rates. The monitoring and control strategy is much the same as shown above, with possible the addition of a report on utilization and unsuccessful attempts. These additional items are particularly germane if contention is a part of the system. That is, if the number of possible users exceeds the number capable of simultaneous service. The utilization tells the spectrum manager how much excess capacity exists on the frequency, in case more subscribers need to be placed there. The number of unsuccessful attempts ("busy signals") tells the manager when more capacity is needed.

4. Point to multipoint communications for variable data rate services represent most computer communications which might be carried on cable. For example, cable is in a good position to offer communications for the commercial on-line services such as Prodigy, America On Line and Compuserve, and for Internet access. These services are characterized by very bursty traffic, with long pauses in between. Furthermore, the services are highly asymmetrical, with more information being sent downstream than upstream. For illustration, we assume an upstream data rate of 1.544 Mbps. The table from which this was extracted showed a downstream rate of 10 Mbps. The monitoring and control issues are the same as for the constant rate services shown above.

5. Telemetry services include energy management and home control. They have many of the same characteristics as point to multipoint communications, with the exception of lower bit rate and possibly slightly more symmetrical data requirements. In fact, these services could be combined with the point to multipoint, except that less efficient but lower cost modulation formats may be used.

CONCLUSION

The above examples show some of the monitoring and control tasks that are likely to be accomplished by a spectrum manager, related to services involving upstream communications. We have concentrated on upstream communications because this is likely to be the portion of the spectrum in which we find the most congestion in the next few years, so it is the place to put most of the effort to improve the utilization. One way to improve utilization and service availability is to provide a spectrum management system, which automatically monitors usage of the spectrum and which can move services around as needed to provide the most efficient utilization and the maximum availability.

ACKNOWLEDGMENTS

The author is grateful for the education provided by Jack Terry of Northern Telecom, Tom Engdahl of Digital Video (division of Antec), Michael Pritz of Antec, and Jim Stratigos of Wireless Communications Consultants. Shellie Rosser

of Antec surveyed operators concerning their expectations of revenue from various services. Valuable information and insight were provided by Tom Williams of CableLabs.

---

[i] Real diplex filters are, of course, much more complex than the simple representations shown here.

[ii] Terry, J. B., "Challenges and Solutions in the Introduction of Digital Services in Home Coax Wiring," *NCTA Technical Papers*, 1995.

[iii] Available from the Electronic Industries Association Consumer Electronics Group, 2500 Wilson Blvd, Arlington, VA., (703) 907-9600.

[iv] A system employing time division multiplexing or some other sharing arrangement within a channel would arbitrate within itself, the multiplexing within that system. The spectrum management would only come into play when the system was running out of capacity and needed more spectrum.