Network Availability Consumer Expectations, Plant Requirements and Capabilities of HFC Networks

Tony E. Werner
Vice President of Engineering
TCI Communications, Inc.

Pete M. Gatseos, Ph.D.
Vice President of Strategic Research
TCI Communications, Inc.

Abstract

Within a consumer expectations framework, this paper investigates the value equation of cost versus network availability.

Clearly there are certain low levels of network availability that will be deemed unacceptable by customers. Equally as clear is that you can continue to increase network availability at a greater cost. Of course, this embedded capital cost in turn translates into higher service costs, ultimately becoming a price hurdle to the customer. Figure 1 provides a graphical view of this equation.

The network provider's objective is to satisfy customer expectations for a given service while providing a reasonable network cost — in other words, a reasonable investment that can be reasonably recovered through the price of the service or product.

This paper takes a wholistic view of key cost elements involved in network availability and reliability vis-à-vis customer expectations as they relate to cable TV services.

Figure 1
"Network Availability
The Cost-Value Relationship"

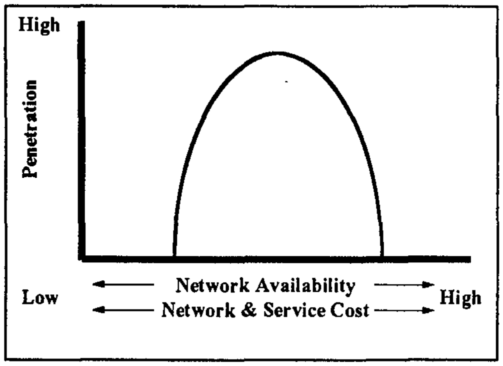

Introduction

Let's begin by defining the expectations of cable TV consumers with respect to network availability, *i.e.*, the ability of the cable network to perform at a given instant in time.

Very briefly, consumer expectation levels are set by their experiences with all other service providers that they have encountered. When a variety of services are typically delivered at a given level by most service providers, this level becomes the "norm", or in other words, the "expected". Of course, when expectations are met, consumers are satisfied that they have received delivery of services purchased. When expectations are not met, consumers become dissatisfied.

Further, in the free marketplace where competitors seek a means to generate greater revenues, service performance is continually improved (to gain a competitive advantage), thereby raising service level expectations over time. Indeed, as can be seen in Figure 2, holding all other variables constant and when services are at normative levels, (*i.e.*, customer expectations are being met), this results in parity with the competition. More importantly, when service levels exceed customer expectations, a competitive advantage is achieved.

Here in this discussion of cable network availability, consumers have electronic equipment such as their television sets, stereo systems, and telephones as a basis for comparisons. More specifically, because their TVs, stereos, and telephones perform error-free in nearly every instance that they are used, consumers have come to expect very high levels of service reliability from products/services that appear to be similar in nature, *i.e.*, "electronic" products and services. Therefore, cable services are intuitively measured/assessed by consumers against the reliability levels delivered by today's TV/stereo sets and telephone services.

It is also important to recognize that consumers have a higher probability of detecting a cable TV network failure. Indeed, according to Bellcore, cable customers are ten times more likely to experience a cable outage than phone outage. This is due to the fact that television sets are in use about 7 hours per day, while phones are in use only about 30 minutes per day on average. (See Table 1.)

Figure 2: Service Delivery and Consumer Expectations *vis a vis* Competitive Standing

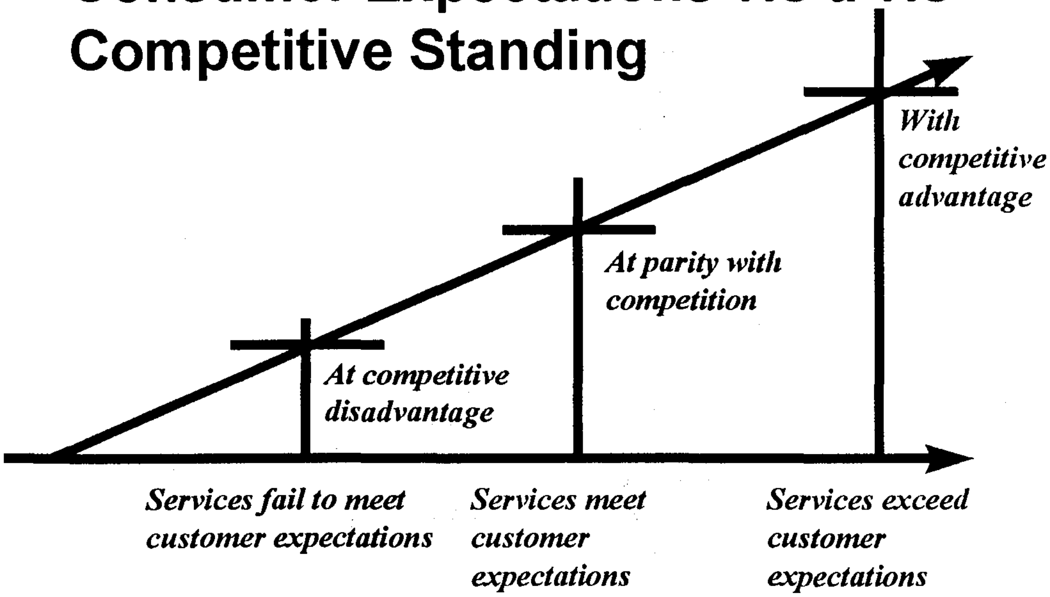

Table 1
Phone vs. TV Outage Perceptions[i]

	Usage per day Averages
Message Telephone Service (MTS)	30 minutes
Television viewing	***7 hours***

Definitions

By definition, network reliability is the probability that the system will not fail in a defined period of time. The frequency of failures over a defined period of time reflects the probability that the system will fail in a defined period of time, representing system unreliability. The well known term, Mean Time Between Failures (MTBF), reflects the inverted value of the frequency of failures and is usually used as a measure of reliability.

While network reliability is important to a network provider, a more important measure is the network availability. This measure is defined as the ability of a network or a unit to perform a required function at a given instant in time.

It can also be expressed as the percentage of time, within a given time interval, during which the network is capable of providing the service. This assumes that the external resources, if required, are provided.[ii]

In mathematical terms, the same can be expressed as MTBF/(MTBF+MTTR), where MTTR is the Mean Time To Repair with given external resources. This formula demonstrates that the network availability can be improved by increasing the MTBF, which represents reliability. Another way to improve network availability is to decrease repair time (MTTR) by providing more maintenance resources. Design, although not indicated directly in the formula above, is also a critical factor, since network availability can be greatly increased by limiting the number of cascaded devices.

Elements of Network Availability

Network availability is a function of three basic elements, which are:

- MTTR
- MTBF
- Architecture

Mean Time To Repair is often overlooked in network availability, but can provide significant improvements to network availability at little or no cost.

Figure 3 — HRC

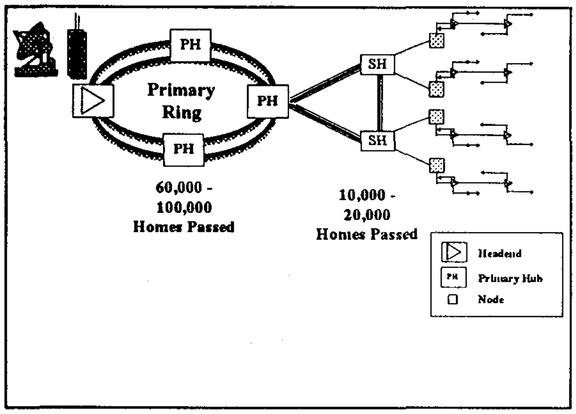

Using the Hypothetical Reference Circuit (HRC) in Figure 3 and empirically derived failure rates outlines in Table 2, we will evaluate repair time sensitivities.

Table 2
Failure Profile[iii]

	AFR (%)	MTBF	MTTR
FWD FO Tx	2.3	43[iv]	150
RET FO Rx	1.4	71[v]	150
FO Cable	0.44	227[vi]	270[vii]
FWD FO Rx	1.4	71	150
RET FO Tx	2.3	43	150
Cable	0.23	443	152
Connector	0.01	7,110	91
Power Supply	2.21	45	51
3rd Party Damage	3.56	28	99
Other	3.71	27	225
Passive	0.04	2,241	79
Fuse	0.66	152	44
Amplifier	1.75	57	63

Based upon the figures in Table 2, the reference architecture would provide annual network downtime[1] of 39.01 minutes, based upon a statistical model. The Pareto of the downtime contributions are contained in Table 3.

Table 3
Downtime Analysis[viii]

	Redun-dancy	NDT (min.)
Primary Ring	Total	0
Secondary Ring	Total	0
Fiber to Node	None	15.94[ix]
Express Coax 1	None	7.64
Express Coax 2	None	6.46
Tapped Coax 1	None	3.35
Tapped Coax 2	None	3.40
Tapped Coax 3	None	2.23
Total		39.01

Figure 4 illustrates the impact of increased MTTR on network down time. The reference level at the beginning of the graph utilizes mean repair times as outlined in Table 2.

Figure 4
MTTR Impact on NDT

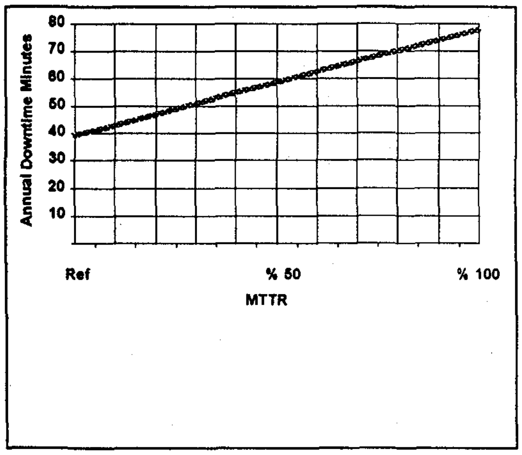

---

[1] Excludes simultaneous failures of standby and commercial power systems.

Table 3: Real Time to Repair Comparisons

(excluding power failures, drops, head end equipment, and network interface units [NIUs].)

	System A with aggressively managed repair times	*System B with little emphasis on repair times*	*System C with little emphasis on repair times*
	With minimal fiber options (Minutes)	*With adequate fiber options (Minutes)*	*Very good plant (Minutes)*
Mean time to repair	62	135	124
Median repair time	48	105	95
Mode repair time	35	120	60
Average network availability	99.987	99.950	99.986
Average network downtime	66 minutes	243 minutes	72 minutes

Source: Computed from "Network Availability Requirements and Capabilities of HFC Networks" by Tony Werner and Oleh Sniezko, November 1995.

MTTR and Its Impact

Greater network availability begins with management, and the agenda and standards it sets for its employees.

In Table 3, three systems are compared along with several key measures. While System A's plant architecture was older, with minimal fiber upgrades, management's focus on maximizing network availability paid great dividends. System A, using any measure, out-performed the less well managed System B — despite System B's superior (but not atypical) plant. Note also how these two Systems compare with System C, which enjoys superior plant technology.

Very clearly, when management has employees focused on minimizing impacts of all the factors surrounding outages, great strides can be taken in meeting the demands placed on cable TV network availability by consumer expectations.

(This data should also be interpreted to mean that, when employees perceive network availability from "end to end" to be critical to meeting customer expectations, all manner of remedies will be pursued vigorously to minimize disruptions to the customer TV viewing experiences.)

To illustrate further, the repair times in Table 1 are actual repair times detected in a well maintained Midwest system. In reviewing several other systems, we have found several that exhibit repair times close to this system. We have also found several that have significantly longer repair times — some extending to over three times as long.

By having an MTTR of 150% over the reference, which several systems do, you will cause an annual network down time (NDT) of over 150 minutes versus the 39.01

minutes, based upon the reference values. These two different availability outcomes are achieved with identical system performance for MTBF and architectural design. Managing outage response programs can significantly impact system availability at minimal or no cost!

The average repair time, often referred to as MTTR, is comprised of three major components affecting the total outage duration. They are:

- Identification;
- Response; and
- The actual repair.

Figure 5 diagrams the markup of a typical Hybrid Fiber Coax (HFC) outage.

Figure 5
Repair Time Components

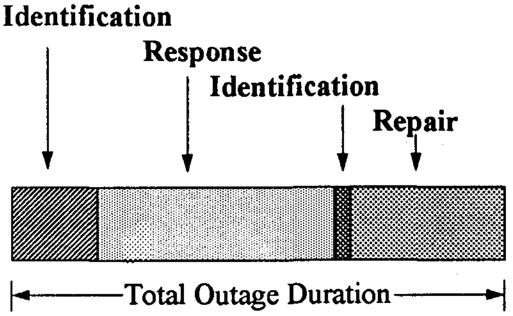

In a typical HFC system, these contributions break down at: 3-10% for identification, 50-60% for response time, and 30-40% for the actual repair. These exact distributions vary, based upon several factors, but usually the major portion of the repair time is not the actual repair, but rather the identification and response time. This is because of the modular nature and short repair times associated with a bus network.

MTTR Opportunities

Taken from the consumer's perspective, even the most sophisticated plant performance tracking databases monitor only a small set of total sources of outages to the customer's TV set. As illustrated in Table 5, among five areas that consumers believe the network (*i.e.*, from "end to end") failed, only two are tracked by one of the cable industry's best plant monitoring systems.

Table 5
An Illustration of Customer Perception of Sources of Outages

Customer reported source of outage	*Actual plan performance reports*
Weather	Not reported
"Cable network"	Reported
Power company failure	Reported
Subscriber drop	Not reported
Converter	Not reported

Taken in this light, the conventional industry wisdom that consumers overestimate the actual number of plant outages is contradicted. From the consumer's "end-to-end" delivery perspective, plant monitoring measures are too limited in scope. Cable industry measurements, therefore, need to be more thoroughly articulated, standardized, and continually monitored.

Now let's look at a few select techniques to improve identification and reduce identification times. A tool that gets a lot of attention in this area is status monitoring. Historically, status monitoring in a CATV architecture was not significantly beneficial. With high penetrations of customers and trained phone staff, outages could be determined with 5 to 10 minutes of the occurrence via phone calls. Exact location identification was performed from continued phone calls while a technician was en route. The overall savings from status monitoring was typically the first 5 to 10 minutes that it took to determine an outage from phone calls. This would be traded off against the capital and operating cost of status monitoring systems, and in most cases, more value could be provided from investing that money in some other portion of the network.

It should be noted that, in the past, return networks would have to be activated and maintained purely for the status monitor, adding this capital and operating cost directly to the cost of status monitoring.

Today, status monitoring is going through a rebirth in HFC networks under the name of *network management*. While there are several reasons for the re-entry of network monitoring and management today, be cautious not to see this as a panacea. While we like to call it network management, it is still largely monitoring today. Some sophisticated networks may use it to enable backup fiber routes or even backup equipment, but mainly it is a monitoring tool. In the future, as systems offer significant levels of telecom services, the same upper level management systems may also provide provisioning, and then it will be graduating into network management.

With this said, there are situations in which monitoring is more practical today than in the past. First, computers have grown in power and user friendliness significantly in the last 5 to 10 years, allowing for a much more practical, useable implementation. Secondly, many operators are activating the return path of the network for new services, so this cost does not have to solely burden the cost of monitoring. The third reason is that several new services will go in at low penetrations, and due to this and their less-than-full-time usage pattern, return path failures may go unidentified for several hours without status monitoring.

Let's look at some of the other methods of reducing MTTR by focusing on repair times. Several operators enforce zone maintenance practices. This procedure has a particular maintenance technician responsible for a geographical area of plan, and as such is required under most conditions to be in that area.

This means that his or her drive time (response time) will be minimal. Another technique is to operate two shifts of maintenance technicians. The first shift works from 7:00 a.m. to 4:00 p.m., and a second shift works from 10:00 a.m. to 7:00 p.m. This ensures that staff is in the system for more hours of the day and ultimately insures a shorter response time for system failures during this portion of the day. The wider the spread of the work shifts, the greater the system coverage, but also the greater the cost for supervision and shift premiums (if applicable).

The final important point is to manage the outages in general. This means monthly reviews of every outage, looking for the opportunity to eliminate or minimize the occurrence in the future.

These opportunities include: training, truck inventory levels, pre-made optical cables, and individual employee performance levels.

The MTBF Opportunity

Now let's focus on the second component of network availability: MTBF. This is another area that can offer significant improvements in network availability at minimal cost. It should also be pointed out that, while it is important to minimize repair times, thus improving the overall network availability, it is better to have no interruption at all. This objective is largely affected by MTBF.

There are several basics for improving MTBF, but the main area is again system-specific analysis. Of all the systems that we have analyzed, they all have significant low cost opportunities for improvement. Whether it is poorly performing equipment, improper grounding or improper fusing policies, the latter is simple to correct, and in several of the systems reviewed, it can reduce the system outages by up to 30%. Figure 6 identifies a seasonal outage pattern that is characteristic of most of the systems analyzed.

Figure 6
Outage Seasonally

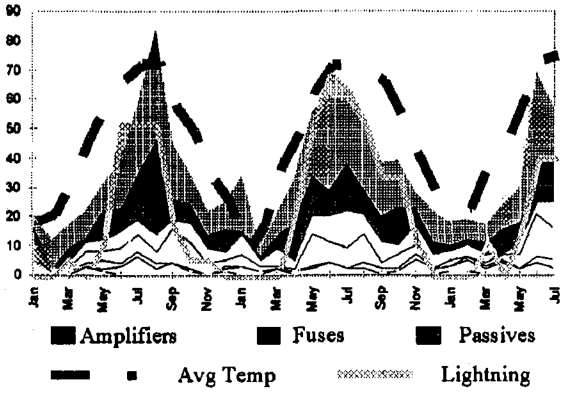

The area graph shows outages from January 1993 through July 1995. The dashed line indicates average temperature, and the light solid line thunderstorms. After identifying the seasonal nature of outage activity, we obtained and overlaid weather^{x} information for the same period. While somewhat intuitive, both temperature and thunderstorm activity obviously influence the system reliability.

To understand this better, we analyzed day-by-day activity for May-July of 1995. These results are graphed in Figure 7.

Figure 7
Day by Day Correlation

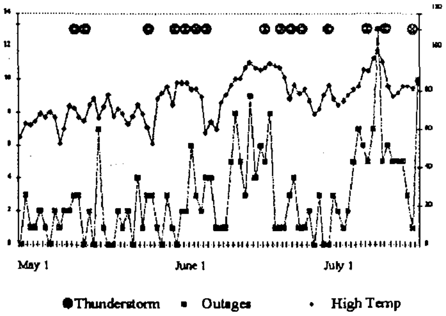

While there are times when the thunderstorms caused outages directly, the majority of the increased outages coincide with high temperature. This was demonstrated on July 13, 1995, when the high temperature hit 101 degrees Fahrenheit. This provided for the highest device failure of the entire year. Thunderstorms occurred on July 12 and 14, but not on July 13. As you review Figure 6, you can see several other points when thunderstorms occurred without significant increases in network failures. A close examination of the failures on those days often revealed water damage, not electrical damage. July 21, 1995, the last data point on the chart, revealed high electronic failures resulting from the electrical storm. Therefore, it is clearly a combination of effects, but the predominance in this system is temperature, not thunderstorms.

The device failures during the high temperature days were predominantly components burned by excessive amperage or shorted devices themselves. It is unclear whether these failures are the result of thermal instability of certain of the electronic devices or if the commercial power system becomes more unstable because of the increased loads for air conditioning.

The next step needs to be a careful power analysis to determine the exact cause or causes of the heat-related failure mechanisms. In either case it is likely that significant improvements can be made at minimal expenditure through selective power conditioning. If even the top 30% of these peaks can be eliminated, major improvements can be achieved in network availability. Furthermore, the savings in staff for technicians and phone personnel would likely go a long way toward paying for these improvements.

Network Architecture Opportunities

Now let's review architecture and its impact on network availability. This is often the element most focused upon for achieving high network availability. As demonstrated earlier in Table B, while architecture is an important element to network availability, even the best architecture will not achieve the desired results if not managed properly. Likewise, even classic cable systems can achieve remarkably good network availability with diligent management.

This is not to say that architecture isn't important, but rather that it is not the only factor in network availability. It can also be one of the higher cost methods of achieving high reliability, and if not carefully managed, can push prices beyond the ability of most consumers to pay (as illustrated in Figure 1).

If you are building a new system or upgrading an old system, network availability should be a primary factor in architectural decisions. Several architectural designs affect reliability. These include the following:

- Self-healing rings;
- Number of serial devices; and
- Redundant components.

All of these architectural designs help mitigate perhaps the most strategically important area of concern, losing customers to competitive TV delivery systems because consumers perceive competitors to provide higher levels of TV signal reliability.

For example, recurring outages frequently are confined to specific geographic areas in which network failures have proven elusive or resource limitations have precluded implementing a more permanent solution. In these situations, the probabilities of customers experiencing an outage greatly increase. Even more ominous, the groups of customers impacted will be repeatedly denied what they have already paid to receive.

Taking this state of affairs one step further, let's assume that 40% of the customers who watch TV daily are impacted by one outage and are angered by it (because they were very engaged in the program). If there is one outage per day, the laws of probability tell us that 16% will not only experience two outages, but be angered by them on *both* occasions. If there are three outages, 6% of our customers will have been moved to anger on all three occasions! Of course, the number of customers impacted can be quite great should these outages occur during prime time, *i.e.*, peak TV usage "dayparts" or times of day.

To help prevent recurring outage situations, architectural design alternatives now will be considered.

First, self-healing rings, which are typically used in fiber optic networks and offer significant improvements in network availability by providing virtually 100% network availability[2] to a portion of the network that has a higher MTTR and affects a significant number of customers. The question is how deep to run fiber rings.

---

[2] Redundant rings do not in theory provide 100% availability; however, with limited numbers of high MTBF, low MTTR, parallel devices, the reliability is practically 100%.

It is clear that fiber rings to hubs serving 100,000 to 200,000 customers can be implemented at a low cost per passing and provide significant improvements in network availability. It is also clear that fiber rings to the home with an alternate power network can provide virtually 100% network availability, but will also carry a ridiculous price. Determining the best value is again system-dependent. Figure 8 provides a cost versus reliability curve for extending fiber rings to various depths of the network.

Figure 8
Reliability versus Cost
(to close the ring)

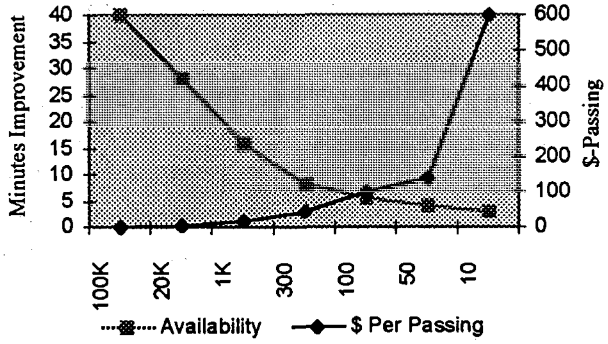

The right-hand axis is cost per passing associated diamonds and the upward trend. The left-hand axis is annual minutes of improvement from ring implementation. The X axis is the size of homes passed grouping that is ring backed up. The exact relationship for cost versus availability improvement varies by system and exact implementation. Figure 8 provides a generic representation illustrating the diminishing returns as you extend the self-healing rings deeper into the network. At the large primary hubs, rings can be installed for as little as $3.00 per home passed[3] and provide availability improvements of nearly 40 minutes. Significant availability improvements are still maintained at the 20,000 homes and at the 1,000 homes passed levels. These hub or node sizes can be ringed without significant cost penalties. The exact point of maximum value is dependent upon the specific system (aerial/ underground/ density/ etc.) and the technology implementation. The further you move out on the curve, the smaller the improvement in availability and the higher the per passing cost.

Serial device or cascade units also affect network reliability. For each amplifier that you shorten your cascade, you improve network availability by between 3 and 7 minutes, based upon Table 2[xi]. Cost modeling based upon actual network design indicates that upgrade costs increase by $20 to $40 per home passed to go from 4-amplifier cascades to 3-amplifier cascades.

Limiting other devices such as passives also provides improvements, but they are quite minor in nature. The improvements that you obtain from limiting these devices is likely more pronounced in frequency response and other technical parameters than in network reliability.

A final architectural tool is to employ redundant devices such as lasers and receivers in serving area fiber systems. Based upon Table 2, this could reduce the 15.94 minutes to 4.75 minutes. This comes at a significant, although you now require redundant lasers and receivers for up-and-down stream operation at every node. Depending on node size, this can add over $10 per home passed to provide a theoretical improvement of 10.69 minutes.

We refer to it as a theoretical improvement based upon historic experience with fail-safe switching at low levels of the architecture. While redundant switching at primary and secondary hubs is quite effective, this technology has been less effectively implemented at lower levels of the architecture. Several reasons probably

---

[3] This is the incremental cost above a non-ring installation.

account for this. The first is that, in order to have economics prove in at low levels of the architecture, less reliable switching and detection circuits have been employed. In addition, just because of the sheer numbers, management and maintenance has been difficult. This is not to say that node redundancy cannot be implemented reliably, but it is just cautioning based upon previous experience with fail-safe switching in RF amplifiers. In this experience, undesired switches caused almost as many problems as they solved.

Strategic Implications

We began this paper by noting that service providers who exceed consumer expectations will enjoy a competitive advantage over their rivals. We also noted how cable networks can be compared to service reliability levels achieved in telephony. It is therefore logical to try to determine the service availability levels in delivering telephone services. This can be taken one step further as well — by determining the TV signal availability of direct to home satellite TV providers, *e.g.*, DBS/DSS competitors.

In this manner, benchmarks for reaching service parity or service superiority can become the target levels for cable operators — to meet or exceed — as technological advances as well as managerial/financial resources permit.

For telephony standards, Bellcore has published a telephony service availability standard of 99.99% or 53 minutes per year downtime maximum.[xii]

In regard to DBS/DSS, signal availability (excluding environmental factors such as "rain fade") averaged 99.90%.[xiii]

Unfortunately, overall cable industry norms are unknown — as such data is not nationally or systematically collected. At this point in time, it remains up to each individual cable system to assess their signal availability performance relative to current DBS/DSS competitors as well as to consumer expectations (created when making comparisons to comparable products and service at a given cost). It is sufficient here to restate what was said at the beginning of this paper: It is only when service levels exceed that of the norm and that of competitors that a competitive advantage may be achieved.

Minimally, it is clear that you must reach a certain level of network availability to not be perceived inferior to your competitors. The counterpoint is, of course, that you can use superior network availability as a competitive edge.

Caution should be applied, however, once you reach an acceptable level of network availability, because you can continue to increase network uptime, but at a higher capital cost. If this increased capital cost is translated into higher service cost, you can actually hurt your sales penetration levels and competitive position. The objective is to find the optimum intersection of capital cost and network availability.

---

[i] Bellcore, TA-NWT-000909, "Generic Requirements for FITL Systems Availability and Reliability Requirements", December 1993

[ii] TR-NWT-000418

[iii] Werner/Sniezko SCTE Emerging Technologies 1966, "Network Availability Requirements and Capabilities of HFC Networks"

[iv] Based upon manufacturer's AFR and validated with TCI 1994 field results. The number is thought to be conservative.

[v] Werner/Sniezko SCTE Emerging Technologies 1966, "Network Availability Requirements and Capabilities of HFC Networks"

[vi] Victor T. Hou (Bellcore), "Update on Interim Results of Fiber Optic System Field Failure Analysis"

[vii] Bruce Corrigan (Rogers Cable TV), "Regional Fiber Outage Report, July 1990 to June 1993)

[viii] Werner/Sniezko SCTE Emerging Technologies 1966, "Network Availability Requirements and Capabilities of HFC Networks"

[ix] Include up- and down-stream components for full duplex availability.

[x] National Oceanic and Atmospheric Administration (NOAA)

[xi] Werner/Sniezko SCTE Emerging Technologies 1966, "Network Availability Requirements and Capabilities of HFC Networks"

[xii] Bellcore, TA-NWT-000909, Issue 2, "Generic Requirements for FITZ Systems Availability Requirements" December 1993.

[xiii] *Multichannel News*, July 25, 1994, p. 103.