USER-PERCEIVED AVAILABILITY OF HYBRID FIBER/COAX NETWORKS

David J. Large
Principal
Media Connections Group

ABSTRACT

Much attention has been focused recently on network reliability in both the telephone and cable television industries. CableLabs' Outage Reduction Task Force has formalized the modeling of distributions systems and suggested approximate component failure rates. Rogers Cable, in an extension of that work, analyzed several architectures and compared their performance.

This paper extends that work by predicting actual and perceived service availability to individual users, including the effects of drop and terminal equipment. The availability of a reference architecture is calculated, along with attainable improvements. Using the modifications suggested, the user-perceived availability of a video service is calculated to be nearly 0.9999, including the effects of headend, distribution system, drop cable and converter.

INTRODUCTION

The hybrid fiber/coax (HFC) distribution system has become the network of choice for both cable companies seeking to upgrade their facilities and for telephone companies seeking wider bandwidth and lower maintenance costs. While it is intuitively apparent that the reliability of HFC networks is much greater than long cascades of coaxial amplifiers in conventional trunk-feeder plants, it is important to quantify that improvement, as it must compare favorably with other alternatives for provision of both entertainment and non-entertainment services.

The Federal Communications Commission was so concerned about this issue (as well as the performance of other transmission alternatives), that it asked the Network Reliability Council to reconstitute itself for the primary purpose of studying the effect of new distribution technologies on the reliability of the nation's telephone system.[1] While wireless and satellite distribution will also be considered, HFC networks will be a primary area of emphasis because of its near universal planned deployment by multiple industries.

HOW GOOD IS GOOD ENOUGH?

Before analyzing how reliable the new networks are, it is worthwhile setting realistic goals.

Some Benchmarks

The widely quoted availability[2] of local telephone service is 99.99%, corresponding to a yearly outage time of about 53 minutes. That figure, however, does not include subscriber's telephones or in-house wiring (which is no longer the responsibility of the telephone company) and may not include the physical cable between homes and central switching offices.

A recent paper has proposed considerably more modest requirements for long-haul data circuits: 99.84% availability, with a predicted 99.96% availability of the local exchange carrier circuits on each end.[3]

Within the cable television industry, one of the few studies of viewer tolerance of failures was undertaken by CableLabs' Outage Reduction Task Force. Its principal report, *Outage Reduction*[4], was published in 1992 and included sections on customer expectations, methodology for tracking outages, reliability modeling methods, techniques for improving the reliability of network powering and system restoration techniques.

Based on a number of studies of customer reactions, CableLabs found that there was a sharp knee in customer perceptions of reliability at about 0.6 outages per month. In other words, customers who experienced less than about two outages every three months found the service to be acceptably reliable, while those experiencing more outages had a very negative opinion of the cable system reliability. Based on a mean time to repair (MTTR) of four hours, CableLabs translated 0.6 outages/month into a minimum acceptable availability of 99.7%.

Even given the different nature of the services provided over the public switched telephone network (PSTN) vs cable television systems, the difference in acceptable availability is very large. User perception of failures should be even greater.

User Perceptions of Availability

Standard telephone circuits are affected by many more outages than users perceive, both because of the low percentage of time the telephone is in use and because of how customers react to problems. For instance, if a call is interrupted because of an outage, the parties simply stop talking and one of them re-dials to re-establish the connection. Since the vast portion of the average telephone call is handled over shared facilities, there is a high probability that an alternate path will automatically be established by the network around any failed element. While the callers may be irritated by the call interruption, they are unlikely to identify the cause as an outage.

Since both parties are immediately aware of the failed connection, no loss of information occurs. Furthermore, since the average telephone user has the phone "off-hook" for less than a half a hour per day, outages occurring the other 23.5 hours do not affect perceived availability. Of course, as usage of the network for extended-duration data connections increases, subscriber are exposed to many more failures.

The situation for classical cable systems is radically different, both from network design and subscriber usage. In existing cable television systems, little if any of the network is redundant, so that circuits interrupted by failure cannot be re-established until the failed element is replaced. Furthermore, the desired programming that is not viewed generally represents irretrievably lost data (they don't stop the Superbowl just because your local cable system had an amplifier fail!). Finally, the average household watches television about five hours per day, so the exposure to outages is much higher than for voice telephone.

Taking these factors together, it is surprising that cable customers are willing to accept an availability as low as 99.7%. Clearly, this will not represent acceptable performance for the provision of switched voice, data circuits or PCS base station interconnect.

THE SHORT HISTORY OF CABLE AVAILABILITY MODELLING

CableLabs' Outage Reduction Task Force

CableLabs' member companies wishing to enter not only the telephone business but the switched video market were concerned about the lack of perceived reliability of cable systems. They needed to quantify current network performance and develop tools for modeling the availability of systems in a systematic way so that different architectures could be compared quantitatively.

Given the historic lack of communication links among regional cable systems, the task force limited its studies to headends and local distribution networks. A key decision was to study only outages affecting two or more subscribers, and thus to eliminate the effects of individual drop cables and converters. Given that about half the individual subscriber outages in a typical cable system arise from drop and converter problems,[5] Cablelabs' results do not accurately reflect customers' perceptions of availability. On the other hand, the analysis tools developed are very valuable in comparative analysis of various distribution architectures.

CableLabs Availability Modeling

CableLabs' method was to gather data from participating companies on actual outages and their causes. From this data, they calculated average failure rates of various components involved. These failures rates were applied to a reference coaxial system architecture and classical reliability analysis techniques used to predict performance as a function of such parameters as the number of amplifiers and the reliability of various component types.

As a check on the model's applicability, the results were compared with actual recorded failure rates for systems with similar characteristics and the results compared with what the participants felt was adequate accuracy.

Outage Reduction was distributed to Cablelabs' member companies and highlighted at the Cable Tec Expo, among other technical gatherings. The computer model, in a generalized form, was distributed as a spreadsheet along with the study so that systems could apply it to their own situations.

Despite the tremendous effort put in by members of the Outage Reduction Task Force, the accuracy of the results are limited by the original data. In most, if not all, cable systems, outages are manually logged and accurate failure analysis and documentation is second in importance to restoring service. Thus, outage durations, the number of affected customers and the cause analysis are all of limited accuracy.

A final factor that must be mentioned is that the modelling and field data were taken on all-coaxial systems, so that the effect of fiber-optics on overall reliability was not determined.

A key finding of the group was that system powering problems dominated all other outage causes in most systems unless adequate standby powering was employed. Given that, the task force extensively analyzed utility power systems and techniques for minimizing utility outages and damaging transients. Many cable systems who are less than rigorous in deploying or maintaining standby power supplies might want to study this part of the document.

Rogers Cable's HFC Extensions

Nick Hamilton-Piercy and Robb Balsdon of Rogers Cable, in July, 1994,[6] used CableLabs' techniques to evaluate the theoretical availability of several HFC architectures. They found that various configurations had theoretical availabilities ranging from 0.9994 to 0.99985.

REFERENCE HFC MODEL

Architecture

For purposes of calculating the affects of various reliability-enhancing techniques, the reference system architecture shown in Figure 1 was be used. This is typical of many modern cable systems where a central headend is connected by fiber to independent nodes, each containing small, non-interconnected coaxial distribution systems.

This is a scaleable model whose variables may include the number of cascaded amplifiers and other components, differing numbers of homes passed and different fiber interconnection options with the headend.

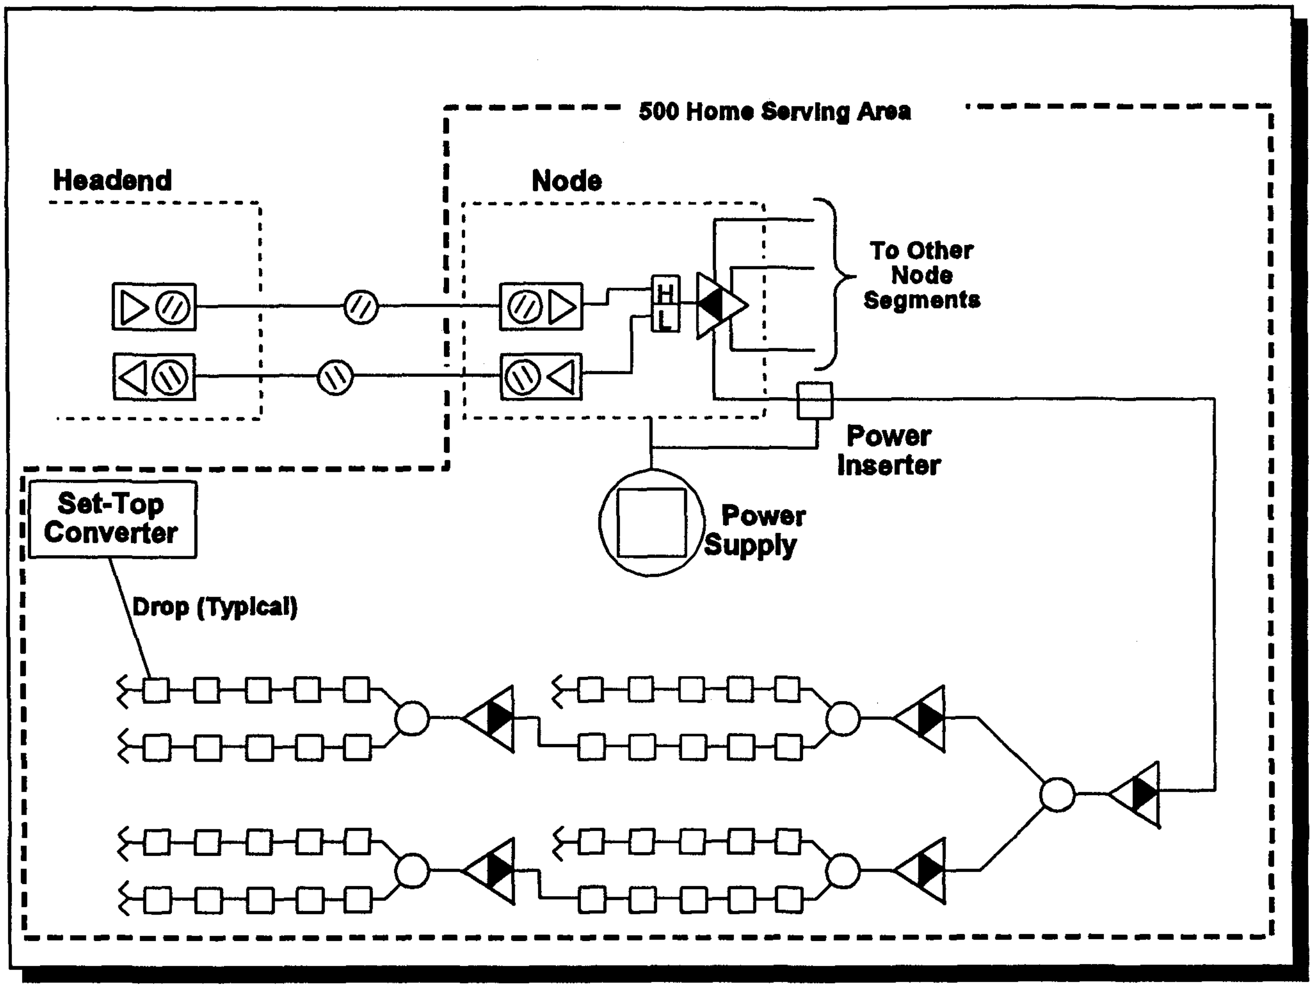

Figure 1: Reference Architecture for Availability Analysis

For modeling purposes, the following parameters were assumed:

- 500 homes per node.
- Five strand-miles of fiber between HE and node.
- One power supply per node serving area.
- Four coaxial trunks from each node.
- 160 four-port subscriber taps per node (78% tap efficiency).
- Three amplifier cascade past the node.
- 100 homes per mile density.

Repair Time Assumptions

Unlike the CableLabs and Rogers models, availability calculations included headend, plant, drop and converter failures and, therefore, true subscriber-perceived network availability. The algorithm also allowed for different average repair times ("mean time to repair" or MTTR) in the headend *vs* plant, which more closely resembles typical field situations.

For the reference model, it was assumed that the MTTR for plant failures was four hours (based on CableLabs data). It should be noted that this is much less than the seven hour MTTR found in an NRC study of telephone system fiber cuts, however the telephone cables cut were much larger, on average, than typical CATV cables, so the difference was expected.

One hour MTTR was assumed for headends, based on their proximity and accessibility to repair personnel. Actual headend MTTR will vary widely, of course. Large urban system headends may be manned part or all of the day, while small rural headends may be unmanned and remote from on-call personnel. This is one of the factors that is leading to consolidation of small headends into larger regional centers connected by fiber to large distribution areas.

Component Failure Rate Assumptions

Headend

An important component of headend reliability is its effect on viewers, *e.g.* if a subscriber is not watching a channel, he is unaffected by its failure. For modeling purposed, it was assumed that the average viewer would be affected by ten channels in one viewing session. In this respect, the model used was more liberal than that used by CableLabs, which counted any single channel failure as an outage.

In the author's model, headend failure rate was predicted by counting the number of pieces of equipment required to process ten channels, then multiplying that by average equipment failure rates. The average yearly failures rate for equipment was assumed to be 5%, the same rate used by CableLabs in their analysis.

The headend equipment configuration is not shown in figure 1, however, it was estimated that generating ten channels would require three satellite antennas, ten microwave receivers, five satellite descramblers and ten RF modulators. The mathematical model also includes three series-connected headend amplifiers (required to provide sufficient isolation between node-specific signals fed to individual F/O transmitters), each with a 3% annual failure rate. Although CableLabs assumed a 7% failure rate for trunk amplifiers and 5% for line extenders, it was felt that the lower rate was reasonable for an indoor mounted unit.

This combination results in a failure rate per year of 149%. With a one hour MTTR, the resultant availability of the headend is 0.99983. This failure rate correlates reasonably well with actual recorded outage data from a recently built headend where the average failure rate among 189 pieces of equipment was 6.3%.year.

Distribution Plant

The distribution plant includes everything from the input of the headend optical transmitter to the subscriber tap port, but no drop components. The yearly component failure rates assumed for the initial analysis are as follows:

Component	Yearly Failure Rate
F/O Transmitter	7%
F/O Cable	0.06%/mile
F/O Node	7%
Power Supply	5%
Passives	1%
Taps	0.5%
Connectors	0.5%
Amplifiers	5%
Coax Cable	0.06%/mile

These numbers generally follow the CableLabs recommendations with some exceptions:

- F/O transmitters and nodes were each assumed to have a reliability, based on complexity and heat generated, similar to trunk amplifiers.
- Rogers' data was extrapolated to predict F/O cable failure rate. Then, on the assumption that accidental damage represents the highest failure exposure

and similarly affects all types of cable, the same rate was applied to coax lines.

- Based on the author's experience, CableLabs' predicted 2% failure rate for standby power supplies, including batteries, was increased to 5%.

- The failure rate for trunk passives was taken from CableLabs' indication that this is a typical manufacturer's specification, despite their recommendation to reduce to 1/10 of that value.

- Tap failure rates were estimated at half the failure rate of trunk passives based on the lower average currents carried.

- Connector failures were only estimated in aggregate by CableLabs, with no scaling for quantity. The author's model uses an estimate of 0.5%/connector/year based on field experience.

- Since HFC systems typically use simplified amplifiers with fewer components than trunk stations, CableLabs' recommended failure rate for line extenders was applied to all post-node amplifiers.

It was assumed that the power supply is a typical cable television standby unit with approximately 2 hours of batteries capacity and no monitoring to alert operators that commercial power has failed. Given that, the model assumes a 30% probability of a power outage that will exceed the standby supply by 4 hours each year. Since there is a wide variance in the quality of local power grids, this parameter will, in practice, vary widely from location to location.

The estimated failure rate of the distribution system shown in Figure 1, using these component failure rates is 89.8%/year, as measured at the tap port feeding the most distant customer from the node. With a four hour MTTR, this results in an availability of 0.99959.

Drop

Although the details of the drop wiring are not shown in the figure, each drop is assumed to include the series connection of cable, four drop connectors and a ground block or splitter. Since CableLabs did not analyze single customer outages, their model does not suggest failure rates for drop components. Based on the author's experience, the following failure rates were assumed:

Component	Yearly Failure Rate
Drop Cable	3%/drop
F-Connectors	5%
Splitter	5%
Ground Block	5%

Given the above failure rates, a drop will experience annual composite failure rates of 28%, which translates into an availability of 0.99987.

Converter

Set top converters have historically been a frequent cause of single subscriber service outages. Although design and construction has improved in recent years, the newest converters are also much more complex. For modeling purposes, an 11% annual failure rate was assumed. Thus the converter alone will have an annual availability of 0.99995.

REFERENCE MODEL AVAILABILITY CALCULATION

Using standard reliability/availability calculations, the series connection of headend-plant-drop-converter was predicted have an

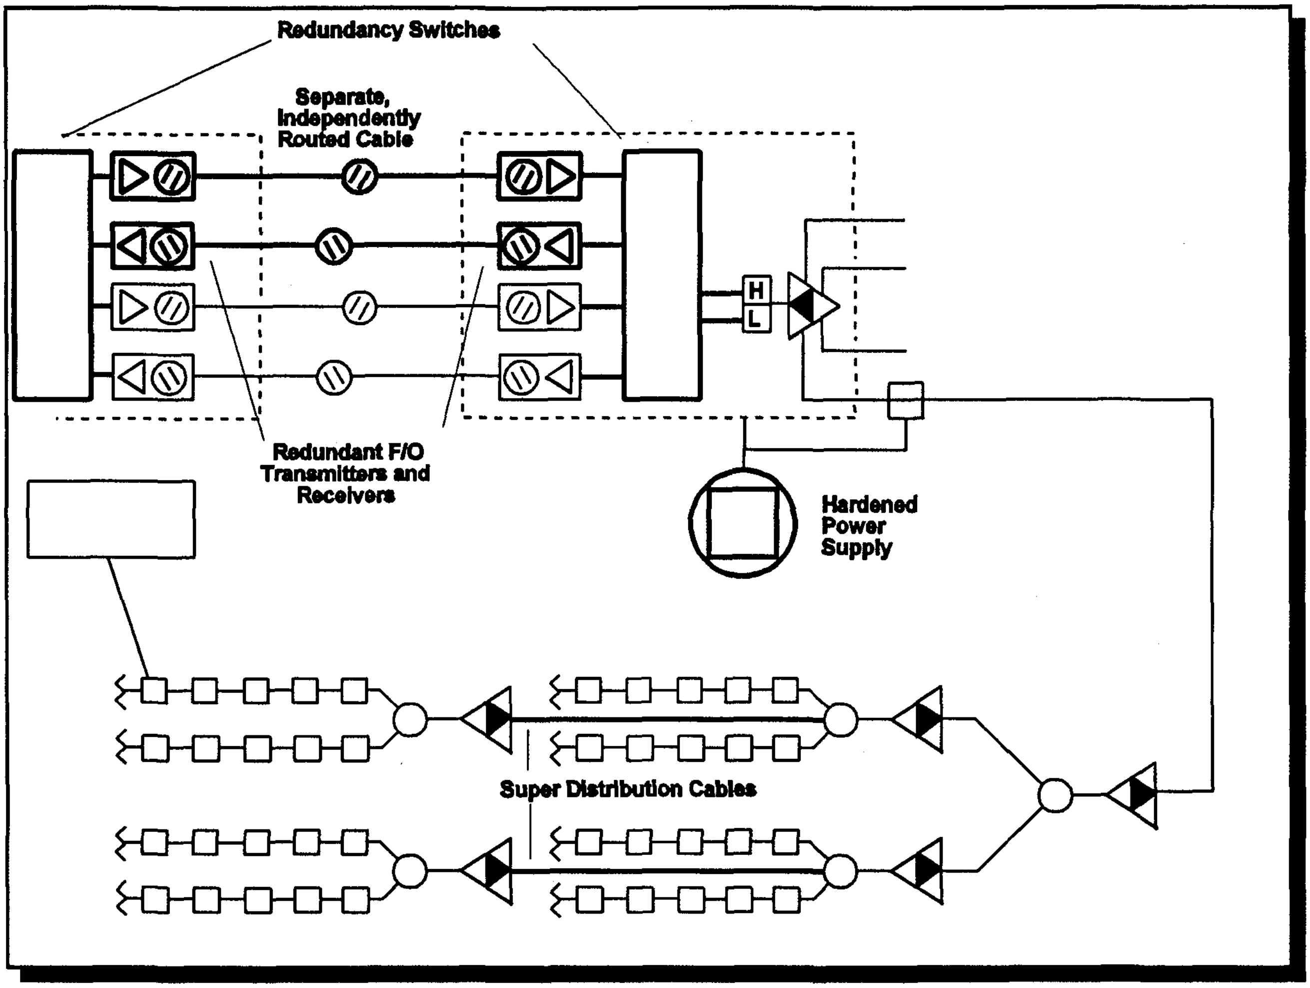

Figure 2: Improvements in Reference Architecture

absolute availability of 0.99924 equivalent to 398 minutes per year of outage.

A typical television subscriber, however, will only experience those outages which occur while viewing is occurring and, assuming a random distribution of failures, will perceive an availability of 0.99975. This represents only one twelfth of the outage rate that CableLabs found to be the critical viewer "irritation threshold".

It can be argued that, unless subscriber expectations change radically, this is sufficient reliability for video entertainment services. It may well exceed the reliability of today's telephone system as well, when plant and terminal equipment failures are added to known switch reliability. Nevertheless, it is worthwhile understanding what improvements could be made, if required.

AN IMPORTANT CAVEAT

The above analysis theoretically applies to communications in either direction between subscriber and headend, as the component cascade is the same. In fact, however, communication in the reverse direction will be degraded due to two factors:

- Ingress signals which may cause communications failure, even if all components are properly working.

- Failures of terminal equipment or upstream transmitters which cause continuous transmission and thus block upstream communications from other locations.

Analysis of these "soft" failures, as opposed to component failures, is beyond the scope of this paper, but must be considered when predicting overall system communications reliability.

IMPROVEMENTS IN THE REFERENCE MODEL

Several modifications were analyzed to ascertain the degree of resultant improvement in the reference model's reliability. Those affecting the distribution plant are shown in Figure 2.

Hardened Node Power Supply

Among the network sections, the lowest availability occurs in the distribution plant, whose predicted failures represent over half the total. About 1/3 of those are the result of commercial power outages in excess of the standby capacity of the node power supply. These may be eliminated or reduced by increasing supply capacity to greater than the duration of most outages, by status monitoring so that crews can be dispatched before the batteries fail, or both.

The regional Bell operating companies are generally opting for either eight hour battery capacity or inclusion of natural gas powered generators in each supply and, in either case, are deploying status monitoring.

Use of Super-Distribution Architecture

Super-distribution[7] (also known as "express feeders") can be used to reduce the number of series-connected taps and connectors.

Redundant Fiber Feeder Cable

A final increment of plant improvement can be gained by paralleling the fiber-optic transmitters, cable and receivers with an automatic changeover switch at the node. Given that cable reliability is primarily a function of accidental dig-ups and car-pole accidents, it is essential that the redundant cable not have any common routing with the primary cable.

These three improvements, taken together, improve the predicted plant availability from 0.99959 to 0.99983, equivalent to 88 minutes per year of outage. The overall system availability improves from 0.99924 to 0.99948.

Reduced Headend MTTR

After the distribution plant, the next highest contributor to unavailability is the headend. Although its effect on individual customers is comparable to the plant, outages affect the entire customer base and so contribute to a disproportionate number of customer-hours of outage.

Possible tools for reducing down time include hot-standby redundancy with automatic changeover and/or 24 hour manning of the headend.

Even if the basic equipment reliability cannot be improved, reducing the MTTR from 1 hour to 15 minutes improves the headend availability from 0.99983 to 0.99996 which is comparable to many estimates of telephone switch performance.

Improved Drop Connectors

From a reliability standpoint, F-connectors are inherently poor. Since there is no wiping outer conductor contact, good electrical performance, as well as mechanical strength, depends on having an appropriately tightened locking nut. Technician skills, temperature cycling and corrosion combine to create a low-reliability situation.

Although improved connectors have been available for a number of years, the basic failure mechanism remains. It is exacerbated by use of the connection as the primary means of turning service to homes on and off. When un-mated connectors are left exposed to the weather, the resultant corrosion makes it even harder to assure a reliable connection.

There are several alternatives available and recently proposed. For instance, the German DeutcheBundesPost (DBP) uses a permanent, connector-less attachment of the drop at the tap and does the electrical connect/disconnect function in a side-of-building weatherproof housing[8]. An alternate approach has been proposed by a U.S. connector manufacturer which has suggested building a sealed switch into taps, so that the drop-to-tap connection can be made permanent[9].

Reducing the connector failure rate from 5% to 2% per year would increase the availability of individual drops from 0.99987 to 0.99993.

Improved In-Home Descrambling Hardware

With the planned increase in complexity in set-top converters necessary to support on-screen graphics, digital decompression, etc., there is little reason to believe that the reliability of these components will increase significantly without a disproportionate increase in cost.

If the FCC, as part of its rulemaking implementing the consumer interface section of the Cable Act of 1992, mandates the availability of set-back decoders, they should theoretically be more reliable because tuners will not be required. While such improvement was not considered in this availability estimation, system reliability may benefit from it in the future.[10]

Achievable Availability

Taken together, the above measure allow overall absolute availability to be increased to 0.99967, equivalent to 175 minutes per year of outage. The perceived availability to a typical video subscriber is 0.99989 or 58 minutes per year, roughly one thirtieth of the outage time found by CableLabs to be the threshold of

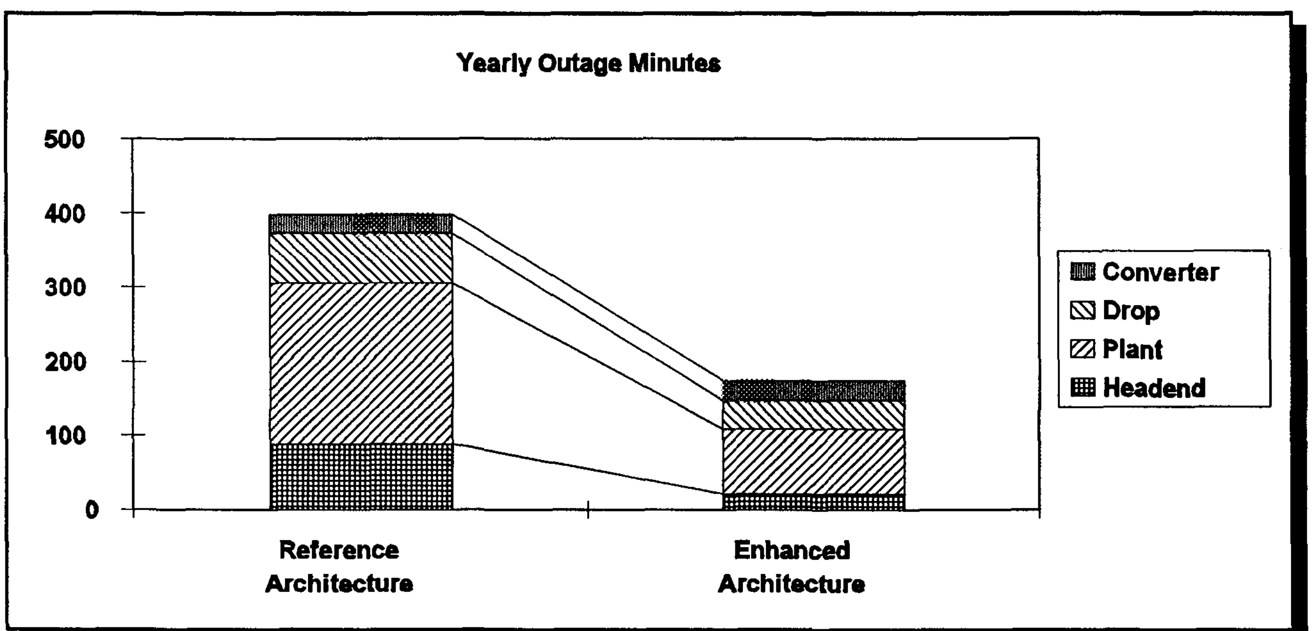

Figure 3: Predicted Outage Improvement with Enhancements

serious subscriber dissatisfaction.

Figure 3 shows graphically the predicted minutes of outage for the reference model and that attainable with the improvements discussed.

WHAT'S WRONG HERE

The problem with this analysis is that it is based on failure rate data that is of marginal accuracy, at best. Until comprehensive failure information is gathered or use of automatic monitoring becomes common, field data will be suspect. Under typical field conditions today, many small outages go unreported and/or not analyzed. Individual customer outages are not included in failure analysis by most systems.

Compounding this, cable operators have not generally required MTBF data from their suppliers, so that, in many cases, not even manufacturers' predicted data is available.

As a result, while the analysis tools developed by CableLabs are invaluable for evaluating relative performance of various architectures, meaningful predictions of actual network performance will only be possible when much-improved component data is available.

CONCLUSIONS

Based on the estimated component failure rates listed earlier, this analysis predicts that typical small-serving-area hybrid fiber/coax networks, coupled with reasonable repair times, can achieve outage times sufficiently low to satisfy customer expectations for current video services.

Various measures, particularly hardened power supplies and improved headend availability, can be undertaken to considerably increase availability. Using all of the measures described, a perceived availability of 0.9999 is predicted.

ADDENDA

Reliability/Availability Terms

Failure Rate: The percentage of devices which fail in a specific period of time (specified as %/year in this paper) and expressed with the symbol λ.

Reliability: The probability of failure in a specific time period, mathematically:

$$R(t)=e^{-\lambda t}$$

For a series connected system, the net reliability is the product of the reliability of the individual components:

$$R_S(t)=R_1(t)\ R_2(t)\ \ldots\ R_n(t)$$

MTBF: Mean time between failures, mathematically

$$MTBF=\frac{1}{\lambda}$$

MTTR: The mean time required to restore operation after a failure.

Availability: The decimal amount of time that a mechanism (in this case a cable system) is available to the user. Mathematically:

$$A=\frac{1}{1+\frac{MTTR}{MTBF}}$$

Outage Time: Time the system is unavailable in a specified time period. For instance:

$$Outage\ (\mathrm{min}/yr)=minutes\ /year\ (1-A)=525{,}600\ (1-A)$$

References

1. The FCC established the NRC in 1992 in response to several large failures of the public

switched telephone network (starting in 1988) which affected thousands of telephone circuits. Its scope was limited, however, primarily to failures of switches and cuts of major inter-office fiber cables. The re-constituted council (with cable television representation) is just beginning to deal with the issues raised by new local distribution means.

2. See the Glossary for an explanation of the relationship between reliability, repair time and availability.

3. To, Michael and Neusy, Philippe, "Unavailability Analysis of Long-Haul Networks", *IEEE Journal on Selected Areas in Communications*, vol 12, no 1, January, 1994.

4. *Outage Reduction*, an internal publication of Cablelabs, Boulder Colorado, distributed to all CableLabs member companies (who together serve 85-90% of the nation's cable subscribers). The first chapters of this manual and study were distributed in September, 1992.

5. Based on the author's experience and analysis of trouble call resolution data from a large multiple system operator.

6. "Network Availability and Reliability", *Communications Technology*, July 1994.

7. Super Distribution has the added advantages of allowing increased physical spacing between amplifiers and of decreasing the effects of multiple signal reflections among the imperfectly matched coaxial components.

8. *Breitbandverteilnetze de DeutschenBundespost, 2.,überarbeitete und erweiterte Auflage*, Dipl.-Ing Hans Steckle, published by R.v.Decker, Heidelberg, 1988 (German). Like the U.S. telephone network, the DBP cable system terminates at the building entrance, where all provisioning takes place. Cables between the dwelling and tap are generally of semi-rigid construction with solid copper jacket and center conductor and are hermetically sealed at the connection point with the tap.

9. "Approaching the 'Seamless Drop'", Brian Bauer, Raychem Corporation, *Communications Technology*, June, 1994.

10. *First Report and Order In the Matter of Implementation of Section 17 of the Cable Television Consumer Protection and Competition Act of 1992*, Federal Communications Commission, released May 4, 1994, paragraph 34.