Transmission plans for the traditional long distance telephone network

Douglas A. Kerr

Issue 1
September 7, 2025

INTRODUCTION

As initial efforts were made, in the late 1800s, to provide "long distance" service to U.S. telephone subscribers, the biggest challenge was the great loss (attenuation) of the transmission lines, which initially were wholly passive. The subsequent invention of the electron tube made it practical to add amplification to the lines to overcame the loss. We might at first think that we could now choose to nullify all the loss, giving a very desirable end-to-end transmission result.

But that idyllic vision was spoiled by the impact of the echo that occurred in the telephone connections. Lower loss for a connection (a subjective performance improvement) increased the relative level of echo (a subjective performance degradation). The way these traded off resulted in there being a "sweet spot" for the loss of a connection that optimized the overall subjective performance experience to the "average" user.

This then, over the years, as the long distance network evolved, led to a number of strategies for determining the "optimal" overall loss for a connection, and in turn, a number of plans for actually implementing that strategy.

The emergence of digital transmission systems brought a rethinking of those plans. Part of this was a very "creative" outlook on the attribution of the losses in a transmission path to individual trunks.

In this article, the underlying concepts of this area are discussed, followed by descriptions of the resulting plans and the way they were realized.

1 BACKGROUND

1.1 "Long distance" and "toll"

What are called "long distance calls" to the using pubic were called, inside the telephone industry, "toll calls". The tem comes from the fact that for a long distance call there was a charge (a "toll") that usually depended on the duration of the connection and the distance between the two stations involved. This was as contrasted with local calls, which were either "free" (any number could be made under the "rate plan" the subscriber had) or were charged for with a fixed amount for each call after basic monthly ration that was available under the rate plan.

The result was that almost anything associated with the provision of long distance service had the "toll" moniker (such as toll trunks, toll switchboards, toll operators, toll cables, toll central offices, etc.).

1.2 Decibel notation

1.2.1 *General*

The reader is almost certainly familiar with decibel (dB) notation. But later in this article, we will encounter it in a rather tricky way, so I thought it best to review here the fundamentals.

Decibel notation is a logarithmic system to describing the ratio between two powers (in our case, usually two signal powers). The advantage of such a system is that we can reckon the measure of the ratio between the output and input power of a chain of transmission "blocks" just by adding the measures of the ratio of output power to input power of all the blocks.

1.2.2 *Definition*

The decibel measure of the ratio between two powers is defined by:

$$R = 10\log_{10}\frac{P_1}{P_2} \quad \text{(dB)} \qquad (1)$$

where R is the decibel measure of the ratio and P_1 and P_2 are the two powers involved. When used for transmission measurement, we would typically apply that this way:

$$T = 10\log_{10}\frac{P_{out}}{P_{in}} \quad \text{(dB)} \qquad (2)$$

where T is the decibel measure of the transmission through a block and P_{out} and P_{in} are its output power and input power, respectively. We note that if $P_{out} > P_{in}$, T will be positive, and its absolute value is spoken of as the *gain* of the block. If $P_{out} < P_{in}$, T will be negative, and its absolute value is spoken of as the *loss* of the block. If $P_{out} < P_{in}$, T will be zero (**no** gain nor loss).

1.2.3 *Etymology and usage*

It might seem that the unit "decibel" (abbreviated "dB") should be a unit one tenth the size of the unit "bel", and in fact it is. That basic unit was named in honor of Alexander Graham Bell. But its size is too large for convenient use in many cases. So the use of the smaller unit became common. And, in fact, by custom we almost never use the unit "bel", nor do we ever use the smaller units such as "centibel", "millibel", and the like.

1.2.4 *Use for describing absolute power*

I emphasized above that decibel measure pertains to the **ratio** between two powers. Thus we cannot express a power itself (such as what we might describe in ordinary terms as "6 mW") in decibel notation.

But there is an indirect way to do that. We can select some reference power, and then say, of a certain power, "It is +4.5 dB with respect to <some reference power>.

In the field of basic telephone transmission, a reference power of 1 mW is almost always used.[1] Thus we might say of a power of 5.5 mW, "It is approximately +7.4 dB with respect to a reference power of 1 mW."

But there is a widely used shorthand for that. With it, we would say of a 5.5 mW power, "It is +7.4 dBm". There, the "m" is an arbitrary abbreviation for "with respect to a reference power of 1 mW".

2 EARLY TELEPHONE TRANSMISSION

2.1 Loss in telephone lines

Early telephone lines were for the most part carried by pairs of wire supported on glass insulators on cross arms on wooden poles. While these "open wire" lines were soon replaced in urban areas by underground cables, the lines between cities remained largely implemented in open wire form for many years.

At first, these lines were wholly passive. The lines of course had a certain "loss" (attenuation) per mile of length. The accumulated loss in a connection of course resulted in a substantial decline in the signal power at the distant end. This was manifested as a decline in the delivered acoustic signal, making it hard for the listener to hear what was being said.

The preponderance of the loss came from "I^2R" loss occurring in the series resistance of the conductors. If we increase the diameter of the conductors, and thus the cross sectional area, the resistance per unit length decreases, and with it the loss. But of course larger diameter conductors use more material (in open wire lines, generally hard-drawn copper), increasing the raw material cost. And the larger-diameter

---

[1] Other reference powers are widely used in such field as radio transmission, and even in such aspects of telephone transmission as characterization of noise.

conductors were of course heavier per unit length, increasing the load on the supporting structures, and potentially their cost.

The largest diameter conductors used on any regular basis in open-wire long-distance lines had a diameter of 0.165 inch (very near that of 6 AWG[2] conductors). The largest diameter conductors used on any regular basis in cable pair long-distance lines were 10 AWG. For comparison, note that in normal house wiring, the most common conductor gauge is 14 AWG, with 20 A circuits using 12 AWG conductors.

An open wire pair using the 0.165 inch diameter conductor has a loss at 1000 Hz (the frequency at which transmission losses in telephone circuits are usually stated) on the order of 0.04 dB per mile. If we contemplate a circuit from New York to San Francisco (a route distance of perhaps 3350 miles). its loss would be on the order of 140 dB, and would hardly be suitable for communication.

2.2 An important early development

A major achievement in the development of the long distance service was the completion of a line running from New York to San Francisco, a route mentioned hypothetically above. The route length of the line was about 3350 miles.

The open wire line used hard drawn copper conductors of 0.165 inch diameter[3]. There were four conductors (two pairs), operated in a "phantom" configuration to yield three circuits. The circuits operated on a "2-wire" basis, meaning that speech signals traveled in both directions over the same conductors (just as in a modern "wireline" copper subscriber line).

The "raw" end-to-end loss of such a circuit[4] was about 140 dB.

After a couple of stages of evolution, 12 "repeaters" (amplifying units) were added at various point along the line, with an average gain of

---

[2] American Wire Gauge, the system generally in use today in the U.S. for describing the cross-section of electrical conductors. The cross sectional area **decreases** by a factor of two for each **increase** of three in the AWG number.

[3] Sometimes said to be "British Wire Gauge (BWG) No. 8". The accurate term for that system is "Standard Wire Gauge" (SWG), but that term does not so well hint as to which standard it is! In any case, the diameter of an SWG No. 8 solid conductor is 0.160 inch. The conductor used is also very close to AWG No. 6 (which, precisely, has a diameter of 0.162 inch).

[4] In the phantom setup, the two "side circuits" have a greater loss than the derived "phantom" circuit, and I'm not sure to which these values pertain. If I had to guess, I would guess that the values apply to the side circuits, which each operated over a physical pair of conductors.

10.7 dB each. The overall net loss of a circuit was now about 11.6 dB, making the circuits practical for use in coast-to-coast long distance connections.

But, having those repeaters, could we not have made the net loss of a circuit even less, leading to even better transmission between the two stations? The answer to that question is a centerpiece of this article.

3 REPEATERS

3.1 Introduction

The invention of the electron tube made possible the additional of gain to telephone circuits to overcome some of the loss. The amplifier-based systems used for this were known as “repeaters” for a fascinating historical reason.

3.2 The bidirectional nature of the basic telephone line

Basic telephone lines used a single pair of wires to carry speech signals in both direction (in just the way we are familiar with in our traditional “landline” telephone station line).

3.3 A complication in adding amplification

This immediately raises a complication as we endeavor to add amplification to a telephone line. We might at first naïvely think of just doing what is shown in Figure 1 (shown in single-line block diagram form).

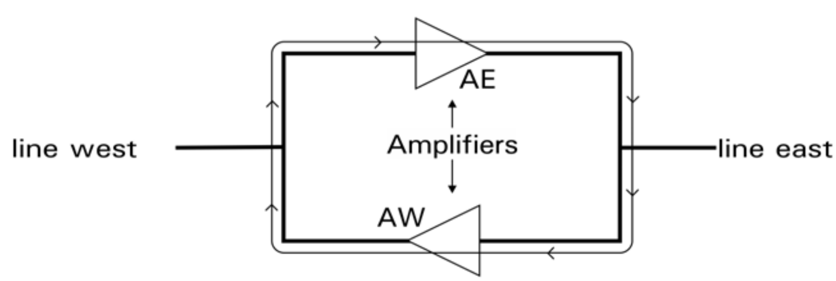

Figure 1. Amplifiers in a two-way telephone line?

Clearly this won’t work. There is a path through both amplifiers (shown by the light line) which would turn this rig into a serious oscillator.

What we would like to have is shown in Figure 2.

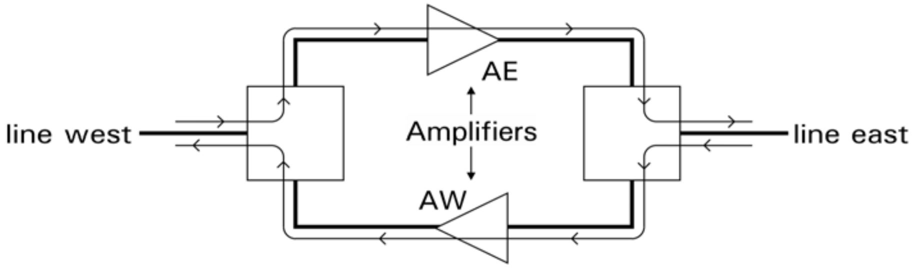

Figure 2. Amplifiers in a two-way telephone line?

Here, "on the blackboard", our problem has been solved by using some sort of "magic box" to couple the two amplifiers' inputs and outputs to the two lines.

But what might be in those magic boxes? Well, they are what are called *hybrid coil* [5] circuits. But they are very simplified in that figure. Let me unpack them one layer in Figure 3.

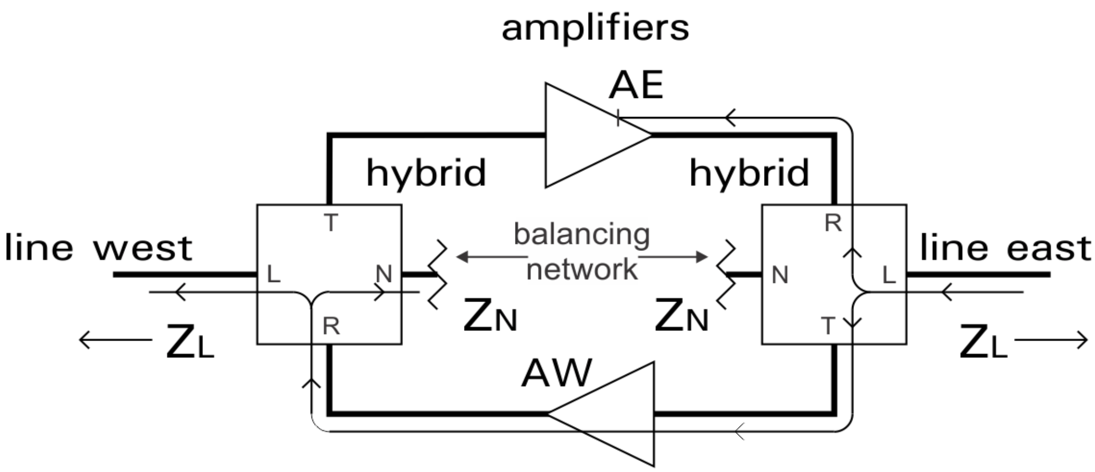

Figure 3. The hybrid coil circuit in place

The square boxes are the hybrid coil circuits proper. But we see a critical accessory: the *balancing network*. It is sort of a single ended artificial line. Its job is (ideally)to exhibit, at each frequency over the band of interest, the same impedance (Z_N) as is seen "looking into" the actual line at that end of the repeater (Z_L).

If each of the two balancing networks fulfils that duty exactly, the signal flow is as shown by the light lines.

Consider first a westbound signal entering the repeater from the line east. Assuming a certain other condition is achieved, half of that

---

[5] Often just "hybrid" for short. "Coil" here is short for "repeat coil" (or even earlier, "repeating coil", the name used in telephone technology for an audio transformer directly associated with a telephone line. By the way, we don't ever say "audio" either, but rather "voice frequency".

signal power flows into the input of the westbound amplifier (AW); sounds useful. The other half of the signal power flows into the output of the eastbound amplifier, where it of course accomplishes nothing. Why would we want half of the arriving signal power to take such a futile route? Well, we don't really want that, but it is a (nearly harmless) artifact of how this system works.

The signal power that enters the input of AW is amplified and goes to the west side hybrid circuit. There half of that power goes out over the line west (sounds useful), and half of it goes into the balancing network on that side of the repeater. Again we don't really want that, but it comes along with the deal.

We note that there is now no path for signal to "circulate" and turn this repeater into an oscillator. And we do want that situation.

Now with half of the input signal power not reaching the input of AW, and with half of its output not reaching line west, if we want the insertion gain of this repeater to be, for example, 10 dB, we need to make the actual gain of amplifier AW 16 dB. But we can do that. And we need to kick in a bit more to account for losses in the hybrid coil circuit itself.

Of course, the very same story applies to transmission in the opposite direction.

What if the impedance of the balancing network on one side of the repeater does not exactly match, at some or all frequencies, the impendence looking into the line on that side? The result is seen in Figure 4.

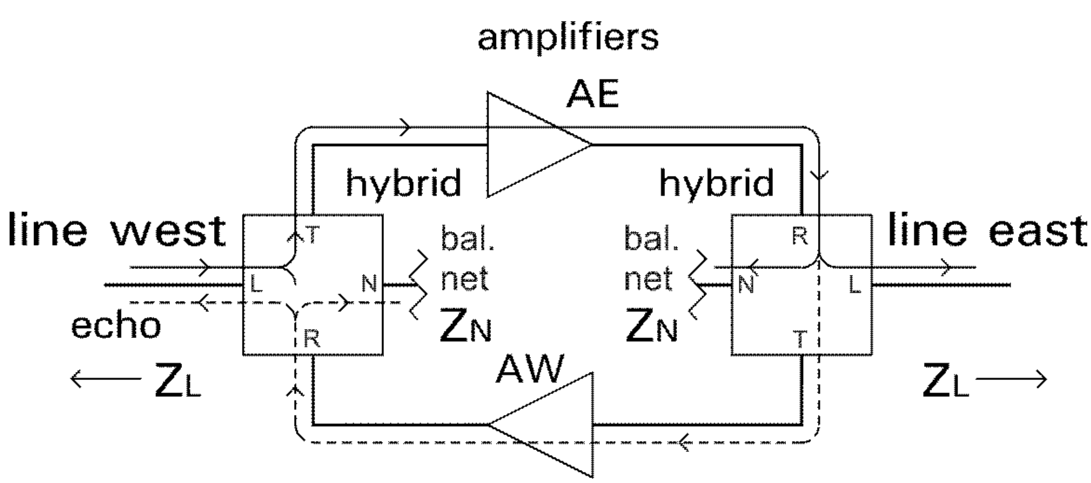

Figure 4. Imperfect balance of the hybrid coil circuit

Here we assume that Z_N on the east side does not exactly match Z_L on that side. Thus, some fraction of the output of amplifier AE goes across the hybrid (follow the dashed path) and thence into the input of amplifier AW. There it is amplified and sent to the west hybrid, where half of it goes on to the line west. And to the talker (out west someplace) who generated this speech signal, this returned signal is an echo of his own voice. Hold that thought.

4 2-WIRE AND 4-WIRE CIRCUITS

In Figure 5 we see a telephone circuit in the long distance network that has three repeaters "enroute".

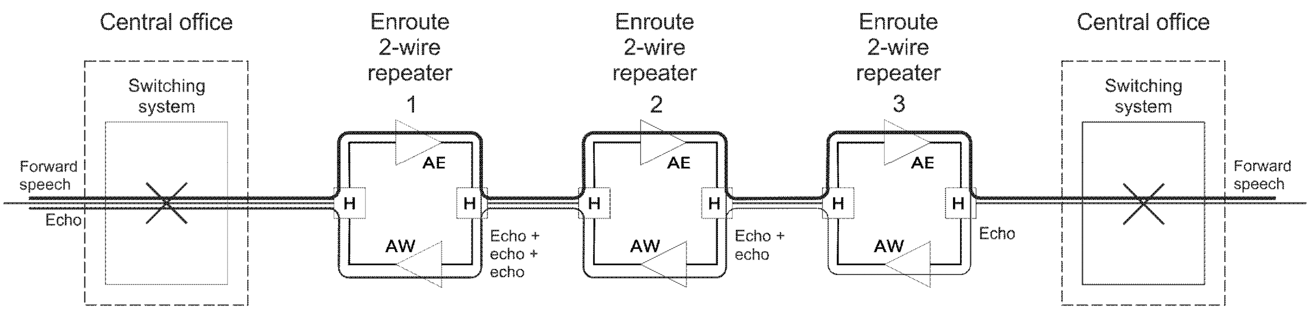

Figure 5. Repeatered circuit

We will assume that none of the hybrid coil circuits enjoy "perfect" balance; that is, the impedances of their balancing networks do not exactly correspond (at all frequencies of interest) to the impedances seen looking into the adjacent lines (the reality of this situation).

We see the "echo" implications of this at repeater 3 when the subscriber at the west end of the scenario is speaking. Part of the signal that is amplified by amplifier AE passes across the east hybrid into amplifier AW, where is it amplified some more by amplifier AW and passes (half of it, that is) through the west hybrid on to the west.

But at repeater 2, the same thing happens, and so the echo returned from repeater 3 is joined by more echo at repeater 2. It joins the echo from repeater 3, which is also is amplified there.

And at repeater 1 we have the same. So the caller gets the "benefit" of hearing his own voice returned from three different places (with some "time stagger" owing to the different transmission delays).

The reality is that, considering that it is difficult on an ongoing basis to maintain "nearly perfect" balance at all the repeaters, this whole notion doesn't work out as well as we might wish.

At one point the question was asked: "Suppose we dedicated two pairs to each circuit, and use one for each direction of transmission?"

If we did, there would be no hybrids at the repeaters. The repeaters would be simpler, and setting them up, and maintaining them, would be simpler, and no echo would be introduced at the repeaters. That sounds really attractive.

Of course, now the cost of the actual conductors (a major part of the cost of the lines, not surprisingly) would be about twice as much. But a very thorough cost-benefit analysis showed that, overall, this would be a good move.

AT&T now took a bold step of adopting this new paradigm for general use in toll transmission facilities.

Not surprisingly, this new paradigm was called "4-wire operation".[6] And now that there was another kind, the original paradigm (which had no name, being "just how telephone lines worked"), became called "2-wire operation".

But there was a fly in the ointment. The switching systems of the network did not have two paths through them, one for each direction of transmission. That is, they operated on a "2-wire" basis. We see the implications of this in Figure 4.

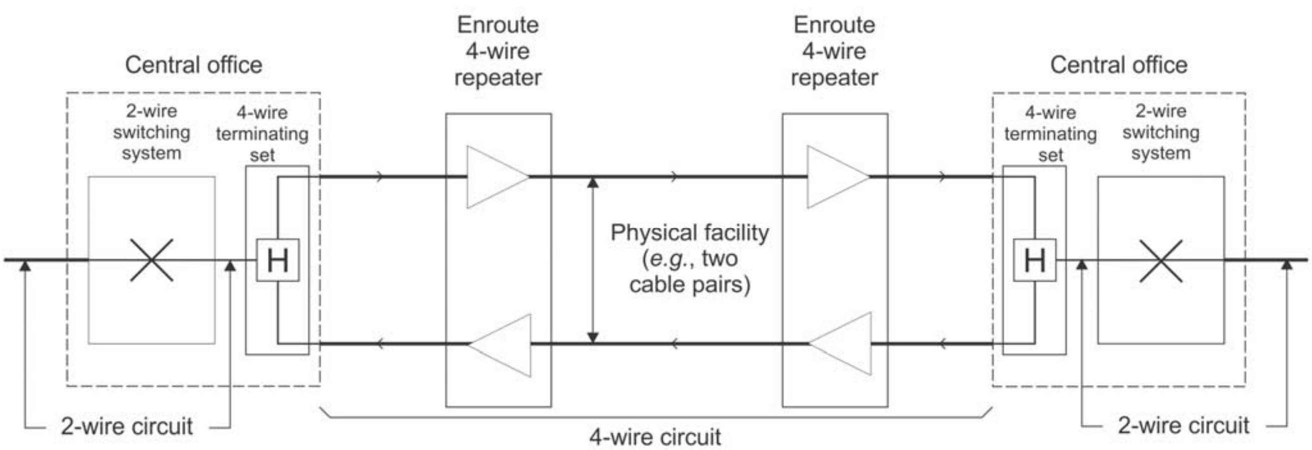

Figure 6. 4-wire circuit between 2-wire offices

We see the circuit proper, with two repeaters enroute, operating on a 4-wire basis. But at the ends of that circuit, where it interfaces with switching systems that operate on a 2-wire basis, we must have a hybrid to mediate between the 4-wire circuit between the offices and the 2-wire path through the switching systems themselves. I don't show the balancing networks; you can just imagine them.

These hybrids are usually actually part of (perhaps the most important part of) an auxiliary unit called a "4-wire terminating set".

[6] And in fact, even when, later, the circuits were not carried by individual physical conductors (but perhaps by channels of a frequency-division multiplex system), the term "equivalent 4-wire" was used. Eventually, just "4-wire-" came to mean any situation where the two directions of transmission were wholly separate.

So is there still a potential echo problem at the 4-wire terminating sets? Yep. We see this in Figure 7.

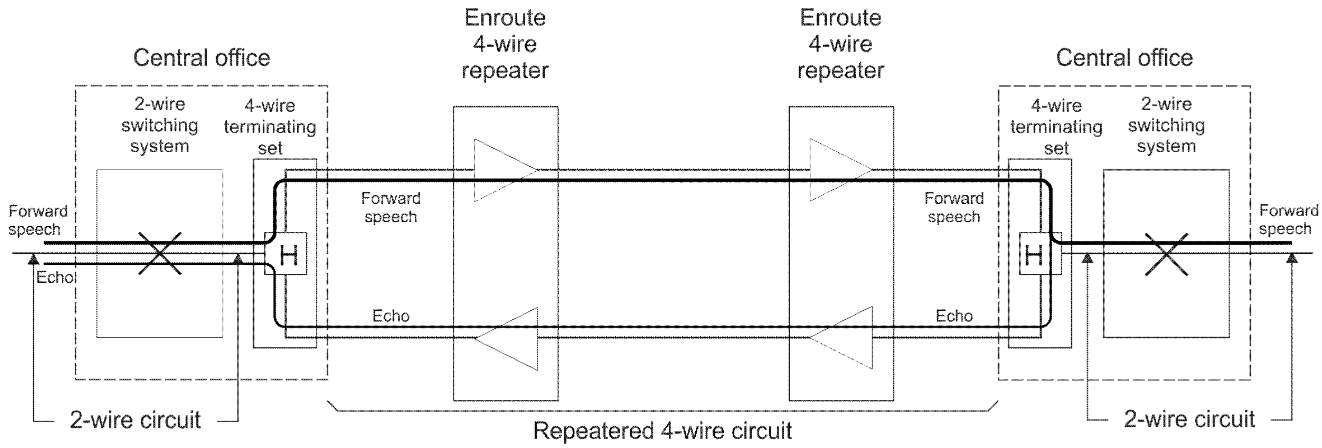

Figure 7. 4-wire circuit between 3-wire offices, showing echo

By now, what is going on here should be self-explanatory.

So, is the echo situation in this configuration less bothersome, and less troublesome, than with 2-wire transmission? Absolutely.

So can we now completely ignore the matter of echo? Not at all.

5 MITIGATING ECHO

5.1 Active intervention

One way to quash echo essentially completely is to use a device called an *echo suppressor*. It watches for speech to emerge in one direction and when it does the circuit inserts a large attenuation in the transmission path in the opposite direction. But this was not even close to being a panacea. For one thing, with the technology that was available at the time, their performance had any number of undesirable side effects. And they were very costly, and required delicate adjustment on an ongoing basis.

So echo suppressors were applied on a very selective basis.

5.2 Echo and net loss

Elsewhere in the network, we took advantage of an important aspect of user reaction to echo. Not surprisingly, the adverse impact of echo is affected both by the relative level of the echo (with respect to the level of the outgoing speech that spawned it) and the delay in its arrival. For any given delay, echo at a higher level is more noticeable and disturbing; for any given relative level, echo at a greater delay is more noticeable and disturbing.

We might at first think that this delay would usually be very short. But not necessarily so. For example, when a technique called *inductive loading* is used to decrease the loss of cable circuits, the velocity of propagation of the signal waves along the circuit might be as low as 0.3 of the speed of light. A round-trip delay of 50 ms can easily occur in longer connections.

A brute force tool we have to reduce the relative level of the echo is to intentionally increase the loss of the circuit beyond what we might otherwise use. There is a nice leverage available here. If we increase the end-to-end loss of the connection by 3 dB, the result is a 3 dB decrease in the signal loudness as heard by the distant party (a negative factor from an overall performance standpoint), and a **6 dB** decrease in the relative echo level as perceived by the talker (a positive factor). This is because the echo (at least in the 4-wire mode) must traverse the end-to-end loss of the connection twice, once going and once coming.

So, if we have a metric for overall "satisfaction" of performance that recognizes both the negative impact of lower delivered speech volume and the positive impact of lower relative echo level, we can imagine that for any given overall round trip delay there would be an optimal end-to-end loss of the connection.

And it is in fact this matter that leads to the transmission plan for the long distance network as it existed in the era of perhaps 1930-1980.

6 THE LONG DISTANCE TRANSMISSION PLAN

6.1 Introduction

The architectural and operational nature of the long distance network continually evolved over many decades. And in parallel, the development of a conceptual transmission plan based on the principle I mentioned just above, and the way it was to be implemented, also involved. The result is a very complicated story.

But I will jump to a point in time where all these things had somewhat settled down: the time (perhaps 1951) that Direct Distance Dialing (DDD) (the dialing of long distance calls by the calling subscriber, no human operator being involved) was just being introduced (and over the prior some years, the network had progressively evolved to be ready for that).

6.2 The conceptual loss plan

After a great deal of subjective testing, a vision emerged of a mathematical function that would tell us the "ideal" end-to-end loss of a toll connection as a function of the round trip delay (of course predicated on certain assumptions as to the degree of reflection to be expected at the ends of the connection). That formulation was predicated on the statistical objective that 99% of the users (whose individual perceptions and preferences of course differ) would be "satisfied" by the balance between echo and end-to-end loss attained in 99% of the connections.

But implementing that "ideal" function would be hard to do with what was seen as practical technology at the time. Fortunately, this was not a situation in which a certain "precise" result was needed (given that the whole matter was highly subjective anyway). So the next step was to approximate that ideal function with an function that could be implemented with a practical scheme. That function was:

$$L = 0.1t + 5 \quad \text{(dB)} \qquad (3)$$

where L is the target end-to-end loss in dB and t is the round trip delay of the entire connection, in milliseconds.

A further reality is that the actual loss of each circuit was not necessarily the one "assigned" to it by the engineering department. There could be small errors in setting the loss, and the loss might drift slightly over time.

We might think that the best thing to do would to just live with that; likely the overall discrepancy in actual loss of a connection would average zero.

But the reaction of the user to a "departure from ideal" in the loss of the connection was not symmetrical. A 1 dB greater loss than "ideal" would provide a small impairment from a 1 dB decrease in received volume, while a 1 db less loss than "ideal" would give a 2 dB increase in echo, a greater impairment. Thus, if the actual losses differed uniformly above and below the "target" values, the average user echo experience would be worse than the bogey.

Thus, it was decided to push the target loss up, from the "ideal" value, to provide a "cushion" in this regard. The distribution of the departure of the total loss (whether up or down) would be greater for a larger number of links (since each link would be subject to a similar distribution of actual loss about its "assigned" value).

Of course, recognizing this rigorously (following the laws of statistics) would lead to a rather complicated function, but the watchword here was to "simplify, simplify". So dealing with this issue was also done simply, by just kicking in an additional 0.4 dB for each line in the connection, thus:

$$L = 0.1t + 0.4n + 5 \quad \text{(dB)} \qquad (4)$$

where again L is the target end-to-end loss in dB and t is the round trip delay of the entire connection, in milliseconds, but now n is the number of links in the connection.

6.3 "Net loss"

To be precise, the losses of trunks I will be discussing here are known as the "net losses" of the trunks. That term is to recognize that this is made up of several losses and/or gains in cascade (some of them small losses in interface equipment and the like).[7]

For conciseness, I will generally here say just "loss" when I really mean "net loss". But the term "net loss" will pop up later in a place where we can't avoid it!

6.4 The toll switching plan

The way in which the plan described above is actually executed is based the structure of the long distance network. What we need to know about it is shown (as of perhaps 1955) in Figure 8.

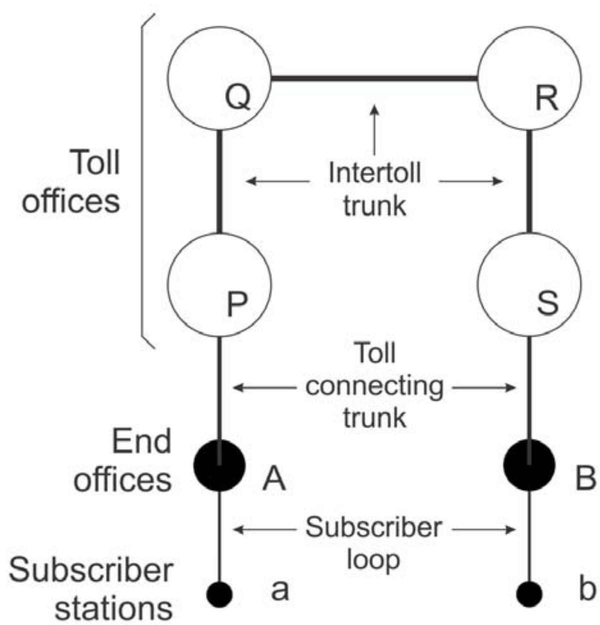

Figure 8. Toll connection—generalized

The central office of interest are divided into two major types, end office and toll offices. End offices serve subscribers. Toll offices are the switching nodes of the toll network proper.

[7] Actually, in very rigorous writing, the term "inserted connection loss" (ICL) is used.

The central offices of the toll network proper are structured in a hierarchy, which has 4 levels, called “classes”, but that has no influence on the story here.

Toll connecting trunks are the links between the end offices and the toll network. In most cases, a given end office has toll connecting trunks to only one toll office.

Each toll office will have “direct” groups of intertoll trunks to other toll offices where the degree of traffic between the two offices warrants. But each toll office always has a group of trunks to the office of the next higher rank on which it “homes”. If all the trunks in the group to the “destination” toll office are busy, or if in fact there are no “direct” trunks, an alternate route is followed.

But that has no influence on the working of the toll transmission plan. Thus, in Figure 8, we do now know, nor care, what are the classes of the toll office that ended up involved.

The “minimal” kind of route is shown in Figure 9.

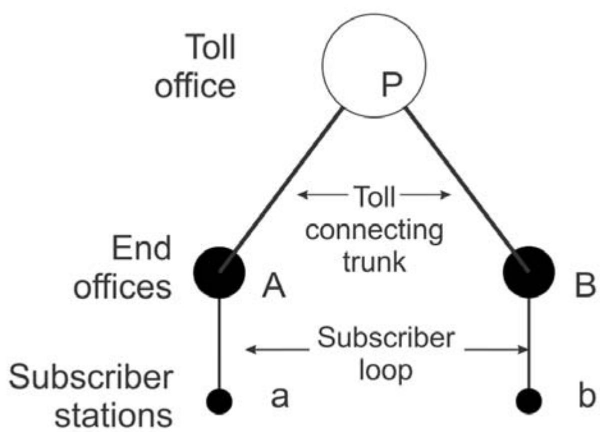

Figure 9. Toll connection—minimal

Here, the same toll office serves both the origin and destination end offices, and no intertoll trunk is required at all.

The punch line of this story is that, in any toll connection:

- There will be zero, one, or several intertoll trunks

- There will (always) be two toll connecting trunks

Figure 10 is a special case that does not follow those rules (but we will see later how it fits fine into the transmission plan).

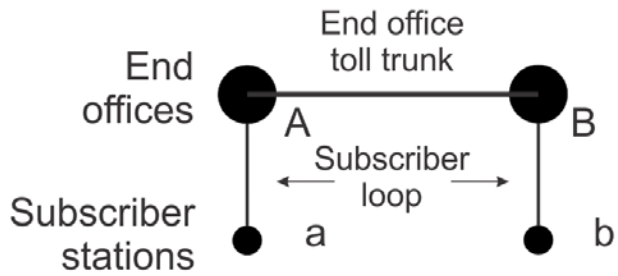

Figure 10. Toll connection—sub-minimal

Here, the "interior" of the toll network has shrunken into nothingness. In this case, the two serving end offices are perhaps fairly close geographically (but a call from subscriber a to subscriber b is still a toll call).

6.5 Loss between what points

The target loss defined as in equation 4 is measured between the two end offices.

We note that the overall loss between the two **stations** is rather variable owing to the varied losses of the different subscriber loops. That is just taken as a "fact of life" with regard to overall transmission system planning.

6.6 Implementation of the overall plan

Now suppose we:

- Assign to each intertoll trunk this loss:
  $0.1t + 0.4$ (dB) (5)

- Assign to each toll connecting trunk this loss:
  $0.1t + 2.9$ (dB) (6)

Then we would find that for any possible toll connection (with *n* links altogether, 2 of them always being toll connecting trunks) the sum of the target losses would be:

$$0.1t + 0.4\,n + 5 \qquad (7)$$

which of course is just what the plan calls for.

6.7 The special case of the end office toll trunk

In the case of the situation shown in Figure 10, the target loss of the end office toll trunk is defined as just:

$$0.1t + 6.4 \qquad (8)$$

This is what we would have for the "generalized" situation with a toll connecting trunk at each end and a single intertoll trunk with a loss of 1.0 dB in the middle (but with only one 0.4 dB allowance for "cushion").

6.8 Special rules for toll connecting trunks

Many toll connecting trunks were provided over passive facilities (often 19-gauge copper pairs with *inductive loading*), where the is no simple way to adjust the loss to that dictated by equation 6, and many of them have losses (under an older set of guidelines) of slightly greater than that. Accordingly, the actual rules for toll connecting trunk target losses were:

- For passive trunks, a loss in the range 2.0-4.0 dB.
- For trunks with repeaters or operated over multiplex channels (such that the loss can be readily set), a loss of 3.0 dB. This of course is essentially what equation 6 would call for if we artificially set t to 1 ms.

6.9 About "via net loss" (VNL).

In an earlier form of the toll network, the trunks between the toll network and the local network (what are, in the plan described above, characterized as *toll connecting trunks*) were not really part of the toll connection loss scheme. They were then called *toll switching trunks*.

In a given connection, a given intertoll trunk might either be the last one at one end of the intertoll connection or not the last one. **In a given connection**, a trunk that was the last one at one end of the connection was said to be in the "terminal" situation; one that was not the last one at either end was said to be in the "via" condition.

Now, in order to fulfill the transmission plan discussed above, a trunk that was (in a given connection) in the *via* situation needed to have a loss of $0.1t+0.4$ dB. That was said to be its "via net loss" (VNL). But when (in some other connection) it was in its terminal situation, it needed to have a loss of 0.1t + 2.9 dB.[8] That was said to be its "terminal net loss" (TNL). And, numerically, that could be said to be "VNL + 2.5".

Now how could a single trunk exhibit both of those net losses? Well, at each end of any trunk that might find itself in either the "via" or "terminal" situation, there was a 2.5 dB attenuator ("pad") that could be switched into the transmission path with a relay.

---

[8] Actually, sometimes it was. 0.1t + 2.9 dB and other times 0.1t + 3.4 dB; to make this story simpler, I spilt the difference.

And the loss of the trunk itself was set up to be "VNL".

A clever circuit recognized when one end of an intertoll trunk was not connected to another intertoll trunk, and then switched that pad into place. Thus the trunk took on a loss of VNL + 2.5 dB—that is to say, a loss of its TNL value.

Now, moving to the later plan I described earlier, we recognize that all intertoll trunks always operate at VNL. And toll connecting trunks essentially always operated at TNL, but it was not said that way; rather, it was said that they operated at VNL + 2.5.

And in fact because the intertoll trunks always operated at a loss of VNL, this plan was often spoken of as the "VNL loss plan".

7 THE SWITCHED DIGITAL NETWORK

7.1 Introduction

The discussions above are predicated on the use of "analog" transmission facilities. Either the speech signals were carried directly by physical pairs, or by channels of a frequency-division multiplex system. In either case, the speech signal was carried as a waveform (in the case of the frequency-division multiplex signals, by waveforms that were translated in the frequency spectrum).

But starting in about 1960, digital multiplex systems came into use. There, the instantaneous voltage of the speech waveform was sampled at periodic intervals (in this case, typically 8000 times per second), and the value of each sample represented in a digital form (as an 8-bit word). The resulting data streams from a number of channels (in the initial system, 24 channels) are interleaved into a single bit stream (in the initial system, at the rate of 1.544 Mb/s). That stream was sent as a stream of pulses over a physical pair.

But, the analog waveform represented by that digital stream was reconstructed at the distant end, and so each channel, end to end, was an "analog channel". And the trunks these channels conveyed had their losses set in accordance with the plan discussed above, and they connected to switching systems just as if they were analog trunks.

7.2 Emergence of digital switching

Digital switching came into the toll network in 1976 with the 4ESS ("No. 4 electronic switching system"), intended primarily for use in the interior of the toll switching network. In 1982, the 5ESS, intended primarily for end office use, but also usable as a toll office, completed the tools for a completely digital telephone network.

These systems carried the speech signals through their switching "fabric" in digital form, essentially the same as in the digital transmission systems. When two trunks carried by digital transmission systems were connected together, in most cases, the stream of 8-bit words describing the speech waveform was carried verbatim through the switching fabric; the speech was never decoded into analog form.

This then had an influence on the toll transmission plan as it pertained to connections that were wholly switched on a digital basis.

7.3 The switched digital network (SDN) transmission plan.

Of course, the conversion of all trunks to use digital transmission was a lengthy progressing process (and was never fully attained). During the period in which analog and digital trunks might both be a part of any given toll connection, a cleaver scheme was used for the transmission plan.

But to set the stage for that, I need to first jump ahead to the plan that was predicated on all trunks being carried over digital channels, with all switching done on a digital basis. Figure 11 is a map of this battle zone, in a little bit of detail.

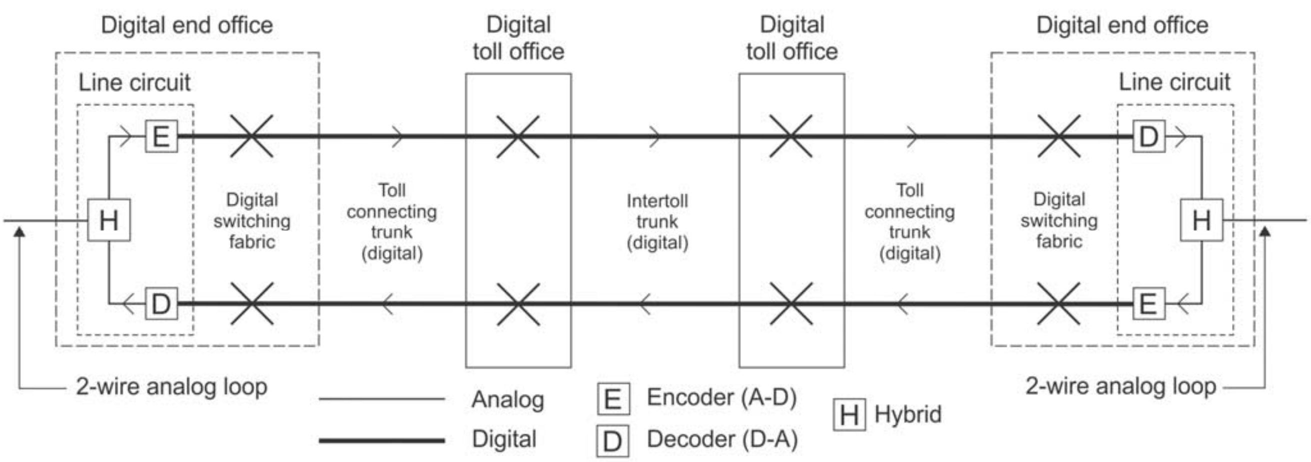

Figure 11. Wholly digital connection

To describe it as "wholly digital" is in fact a little imprecise; the subscriber loops are, of course, analog. But here the end offices switch on a digital basis (as in the 5ESS switching system). As a result, the *encoder* (analog to digital converter) and *decoder* (digital to analog converter), and the hybrid, are located in the individual line circuit.

In addition, since the transmission in the digital transmission systems is of course *4-wire* (in the sense that the two directions of transmission are wholly independent), but transmission in the subscriber loop is *2-wire* (in the sense that the two directions of transmission are conducted over the same medium, which of course also does consist of two conductors), a hybrid is needed, here also located in the individual line circuit.

Now, in this contest, if we wanted to implement the "VNL" loss plan, just how do we introduce some (small) loss in the digital trunks? Well, we could do that by, at each digital toll office, decoding the speech to analog, attenuating it as needed, and the recoding it into digital form. Today we could do that by digital signal processing, no conversion to and then from analog being required.

But in either case, this is not attractive, at the very least, the impact of quantizing error in distorting the reconstructed analog signal is exacerbated.

However, by the time digital switching came into use, some subtle things had changed in the context of the toll transmission plan. For one thing, the digital transmission channels were almost universally borne by "high velocity" media, and so the maximum amount of round trip delay was less. And various improvements meant that the degree of reflection at the hybrids was typically less.

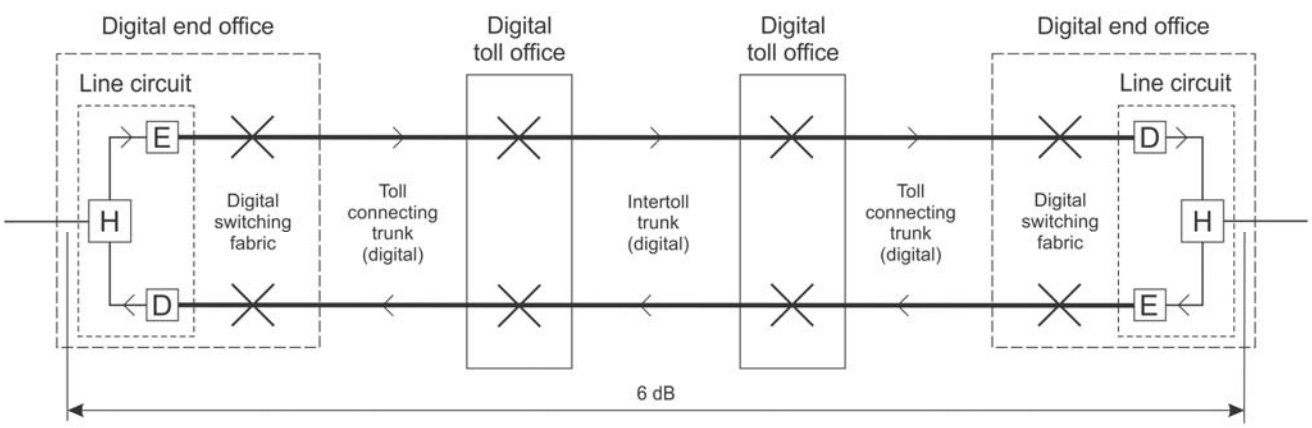

Figure 12. Wholly digital connection—target loss

Taking all that into account, AT&T took a deep gulp and decided that the target overall end-to-end loss of a fully-digital toll connection (measured between the two end offices) should be 6 dB. Period.

That would be consistent with the end-to-end loss prescribed by Equation 3 is we assumed t to be 10 ms.

Figure 12 shows this situation.

And if all connections were suddenly like this one, that would be the whole story. But, as I mentioned earlier, there was a very long period in which a toll connection might include some analog trunks and some digital trunks, and in some cases the switching at the toll offices would be digital and in some cases analog. And to deal with this, we must have some way to attribute that 6 dB loss over the trunks involved (three in the example). And this required some very clever mathematical sleight of hand.

This new plan was called the “switched digital network” (SDN) plan.

8 TRANSMISSION LEVEL POINT AND SUCH

8.1 Introduction

Before I proceed, I need to discuss some further concepts of “dB things” we will be encountering in the discussion of the Switched Digital Network Plan

8.2 Transmission level point (TLP) designation

8.2.1 *The transmission level (TL) of a point*

In traditional telephone transmission work, various points of interest in an overall transmission chain are assigned a *transmission level point* (TLP) value. We can consider it to be defined thus:

> A point in a transmission system where power of a signal passing through the chain is x dB different from the power of that signal at the end office (which serves the subscriber line) where that signal originates is said to be an “x TLP”

If the value is positive, the plus sign is shown (as, “+7 TLP”).[9]

Although the reckoning is in terms of dB, that “unit” does not appear in the TLP designation (but see Section 8.2.2 for alternative forms).

The “end office” is considered to be a 0 TLP.

An important corollary is:

> The difference in the TLP values between two points in a transmission system is the gain (or loss) between those two points.

The TLP system can be thought of as a way to specify a “scaling” for the signals at various points in the overall transmission chain.

This designation is useful in setting up a trunk so that the range of powers of the speech signal at the sending port of the multiplex channel is appropriate. It is also useful in making the trunk setup give the desired end-to-end loss for the trunk.

---

[9] And it was the custom in the telephone industry to speak the *minus sign* as “neg” (for “negative”), the premise probably being that “minus” would refer to that sign as an indicator of subtraction, not an algebraic sign as it is here.

8.2.2 *Alternative forms for the TLP designation*

The form "-16 TLP" was seemingly the most widely used in Bell Telephone System documents.

But in various places, TLP values are shown as "-16 dB TLP" or even as "-16 dBm TLP".

8.2.3 *Normalized signal power–dBm0 notation*

This actually doesn't directly figure in the story here, but I include it just for completeness of this area.

A signal with an actual power of

-19 dBm at a -16 TLP

is "as potent" as a signal with an actual power of

4.0 dBm at a +7 TLP

or "as potent" as a signal with an actual power of

-3 dBm at a 0 TLP

(where everything is expected to be 16 dB "hotter" than at a -16 TLP.)

We describe the "potency" of a signal ("normalized power" is the more rigorous term), at any point, in terms of the power it would have at a 0 TLP. For the signal I just spoke of (wherever we might observe it), we would write that as:

-3 dBm0

That's a zero at the end, but is customarily spoken as "oh".

In a digital transmission system

Inside a digital transmission system, there is no actual "power" for a speech signal—it is just a series of digital words. But we do declare a certain standard digital representation to be logically equivalent to a 0 dBm0 signal. For example, in many digital transmission systems, that digital signal represents a sine wave whose "power" was exactly half the "power" represented by the largest signal that the encoding system could represent (the "clipping level" signal).

That digital signal (a standardized series of sample values) is called the *digital reference signal* (DRS).[10]

[10] It is sometimes called the *digital milliwatt*, by parallel with the "milliwatt" (0 dBm) test signal used in te analog portions of the network.

An actual digital signal (perhaps originating as an analog test signal at some earlier point) can have its potency described in terms of what I call dBD (dB with respect to the potency of the DRS).[11] Conversely, the DRS has a "potency" of 0 dBD[12]. Thus an actual digital signal that was 5 dB "less potent" that the DRS would be described as a "-5 dBD signal".

This notation is somewhat the parallel (in a digital transmission environment) of the "dBm0" notation used in an analog environment.

8.3 Standard test tone

It was early adopted in the telephone industry that the standard tone we would send into a connection to measure loss and such would be at a power of 1 mW (0 dBm) at a nominal frequency of 1000 Hz (which is essentially at the "geometric center" of what we consider the telephone voice transmission band). Thus that test tone can be considered, wherever we might encounter it, as a "0 dBm0" signal

That is not to suggest that a power of 1 mW is typical of even the loudest speech signal from a telephone set; that is typically substantially lower.

Rather, that test tone power was chosen with a "handy" reference power and large enough to make usable "signal power meters" that were passive (just like your old non-electronic multimeter) for the most common measurement situations.

In modern times, in order to eliminate certain artifacts when transmission is over a digital channel, the frequency of the standard test signal was changed to 1004 Hz, and even more recently, to 1013 Hz. The rationale for this is beyond the scope of this article.

It turns out that sending a test toe at a level of 0 dBm0, especially when it might be done on many circuits at the same time, for any extended period was not desirable. In th case of circuits conveyed by physical pairs, this could result in objectionable "crosstalk" into circuits carried by adjacent pairs. It could also cause "overloading" situations in multiplex systems.

So when a test toe is to be applied to a circuit for any appreciable time, or to multiple circuits, or in modern times in all situations, a test tome level of perhaps -10 dBm0 is used.

---

[11] This is not found in the formal literature. There is no "official" terminology for this.

[12] But not every 0 dB EPL digital signal is a DRS; that term only applies to a specific sequence of sample code words. There is one defined for a 1004 Hz signal, and one for a 1013 Hz signal.

But the 0 dBm0 test tone remains, if only on our "blackboards", the mythical hero of this field.

8.4 Encoder and decoder level points

At a certain analog input to a digital transmission system, what happens almost immediately is that the analog signal is "encoded" into digital form. If the "scaling" of the encoder is such that the analog signal power we would need to present to generate a digital signal at 0 dBD were -7.5 dBm, then we would describe that point as a "-7.5 ELP" (encoder level point).

Similarly, if at a certain analog output interface of a digital transmission signal, where the digital signal has just been decoded into analog form, the digital signal is at 0 dBD, and the analog output signal decoded from that has a power of +2.5 dBm, then we would describe that point as a "+2.6 DLP" (decoder level point).

9 BACK TO THE SWITCHED DIGITAL NETWORK PLAN

9.1 Local connections

To lay the groundwork for a point to be made later, we need to look into the transmission plan for interoffice local connections in an all-digital context. We see this in Figure 13.

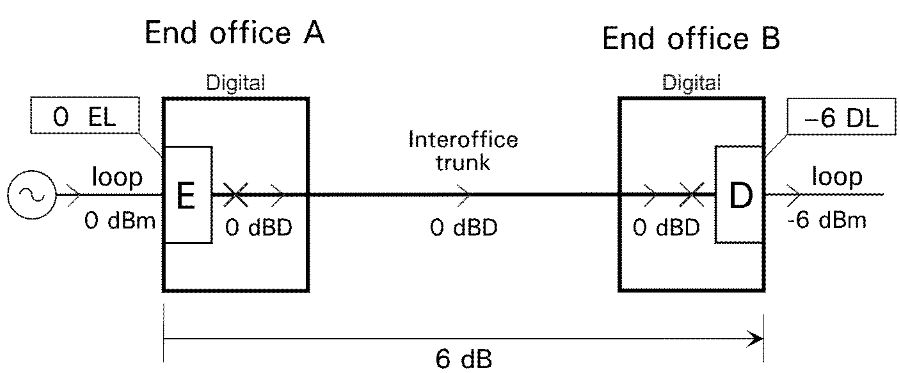

Figure 13. Loss in a local digital interoffice connection

By now, this should be self explanatory. The loss plan for fully-digital local connections calls for an end-to-end loss of 6 dB, which is attributed to the interoffice trunk.

We note that this loss (which is an analog metric) results from the difference between the ELP at office "A" and the DLP at office "B". We also recall that in a digital end office, the encoder and decoder are part of the individual line's line circuit. Note that in particular, this plan requires the DL at an end office subscriber line circuit of -6. Hold that thought.

9.2 Attribution of trunk losses

In section 7.3, I said that in the Switched Digital Network plan “fully in effect” (all transmission, except for the subscriber loops, and all switching, being digital), the entire plan was “end-to-end loss for a toll connection of 6.0 dB.” Period. That should be very straightforward to implement.

But for many years, we had both digital and analog trunks, and both digit and analog switching, with each kind possibly appearing in a toll connection. And to provide continuity with the VNL transmission plan used in a the analog context, we had to be able to attribute a loss to an individual digital trunk. The way that would work (as seen on a wholly digital connection) is shown conceptually in Figure 14.

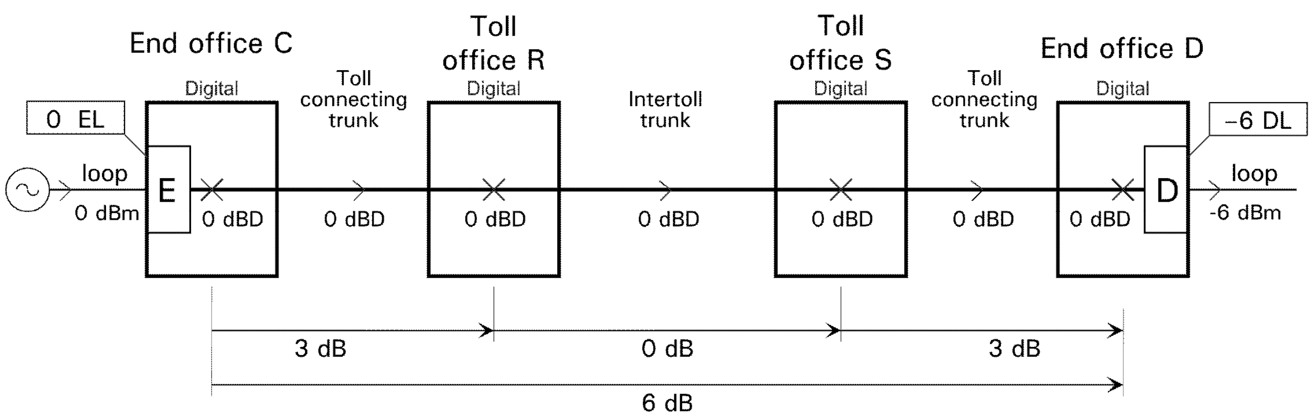

Figure 14. Attribution of trunk loss in a fully digital connection

Here, as in several following figures, I have ignored such things as hybrids and the fact that transmission may be “2-wire” in some places and “4-wire” in others. The figure only follows one direction of transmission, and from the standpoint of trunk losses between offices. And I have arbitrarily shown the losses of trunks as being defined between the switching fabric of the various central offices. As before, E represents an encoder; D represents a decoder.

In this outlook, we consider that the loss of a toll connecting trunk is 3 dB, and the loss of an intertoll trunk is 0 dB.[13] The former is “hard to believe”, since we know that there really isn’t any “loss” to a digital transmission channel. We note, in fact, that the “potency” of the digital signal (stated in terms of dBD) does not vary at all along the entire digital portion of the connection (not surprising given that the digital word stream passes verbatim through the switching systems).

[13] Note that this exactly parallels the “VNL” (analog) plan, if we assume that the round trip delay is 0 and if no safety margin is included for “variation”.

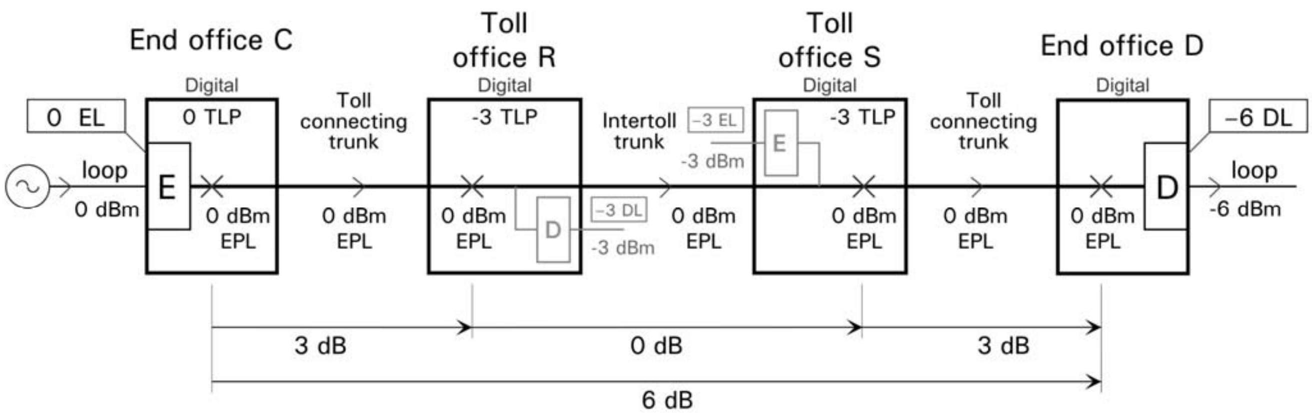

Figure 15. Attribution of trunk loss in a fully digital connection (2)

The justification for this outlook can be seen in Figure 15. It's a little tricky

We have noted here that the TLP of the end offices is defined as 0, and the TLP of the toll offices is defined as -3.

That means that if, at toll office R, we had an arriving digital signal at 0 dBm ELP, and decoded it to analog form (as suggested fancifully by the decoder shown in gray), the DL of that decoder would properly be -3 (that is, the TLP of the office). Thus, the delivered analog signal would have a power of -3 dBm.

That is 3 dB less than the power of the analog test signal introduced at end office C. Thus we can consider that the loss of the C-R toll connecting trunk is 3 dB—as prescribed by the digital transmission plan.

Now consider toll connecting trunk S-D, for which the transmission plan also prescribes a loss of 3 dB. Imagine that we actually introduce our test signal into an analog interface at toll office S. We would of course have to have an encoder for that, and again we show a fanciful one, in gray. Its EL would be -3, in keeping with the office TLP of -3. So properly we would introduce a test tone at a power of -3 dBm. Based on the -3 EL, it would be encoded as a 0 dBD digital signal.

When that arrives at end office D, it is decoded by the decoder in the called line's line circuit (which has a DL of -6) into a -6 dBm analog signal. This is 3 dB less than the power of the analog test signal we introduced at toll office S. Thus we can say that the S-D toll connecting trunk in fact has a loss of 3 dB.

These are both consistent with the outlook seen on Figure 14.

So why would the actual decoder at end office B have a -6 DLP? To make this story work.

Now will that directly fit in with, for example, the digital version of the loss plan for local connections (as discussed in Section 23 and Figure 13)?[14] No. here, the DLP at the line circuit has to be -3. But on a local call, the DLP at the line circuit has to be -6. How can these be reconciled?

The answer the end office digital switching system (typically a 5ESS) is very smart. It knows what kind of call has led to the connection to this line, and it commands the line's decoder to take on the proper DLP.

9.3 It might seem . . .

It might seem that in this setup we should consider the loss of the toll connecting trunks to be 0 dB in the direction out of the end office and 6 dB in the direction toward the end office. But of course a trunk with different losses in the two directions would not fit into the traditional notions of transmission design.

By defining the TLP of the toll office as -3, the mathematical sleight of hand leads us to describe the toll connecting trunk as having a loss of 3 dB in each direction.

9.4 Toll connection with both analog and digital trunks

I said earlier that the only reason we even need to consider the individual trunks in a wholly-digital toll connection as having individual losses was to deal the case in which both digital and analog trunks are used in a connection. Now, let's see how this works out. We see in Figure 16 one illustrative such connection.

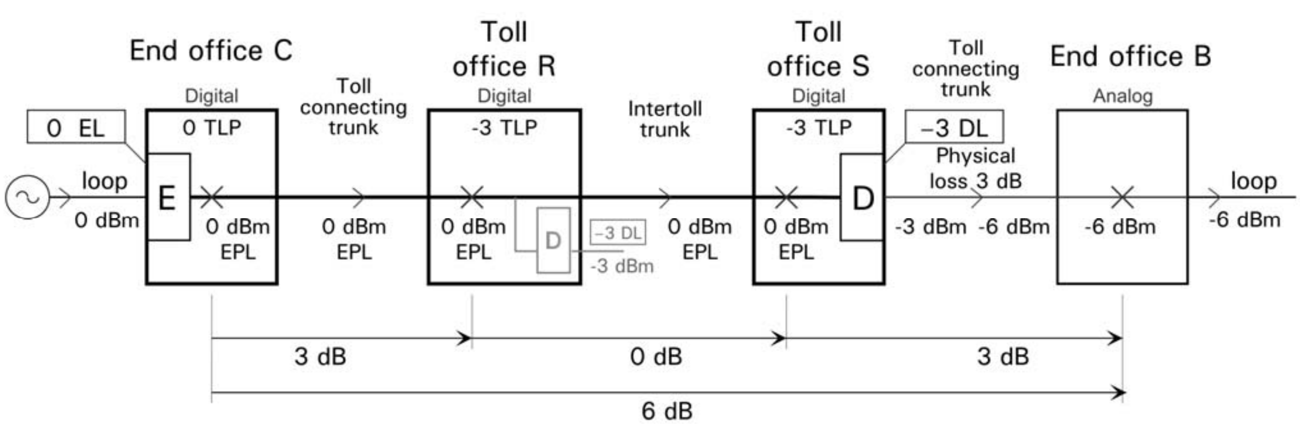

Figure 16. Connection with digital and analog trunks

---

[14] The target end-to-end loss for a fully digital local connection is 3 dB. Why is that not consistent with the 6 dB loss for a digital toll connection? Essentially this was done so as to have continuity ("no surprises to the user") with the prior local and toll plans, which in turn originally grew from the realties of implementation in each case.

Here, the toll connecting trunk from end office A to toll office R (both of them digital) is digital, and (as previously described) we ascribe to it the loss 3 dB, and we ascribe a loss of 0 dB to the R-S intertoll trunk. But the toll connecting trunk from toll office S to end office B is analog.

Under the VNL plan (which governs analog trunks), analog toll connecting trunk S-B should have a target loss of 3 dB, and we assume that here that trunk is implemented over a passive physical pair with a physical loss of 3 dB.[15] Thus we have no difficulty believing that its loss should be considered to be "3 dB".

Now, at digital toll office S, the interface to that trunk has (for this direction of transmission) a decoded and, keeping with the -3 TLP, it would properly have a DL of -3.

Thus our test signal would emerge from the decoder at a power of -3 dBm, fulfilling the idea that the accumulated loss "to" office S should be 3 dB. And the 3 dB loss of the S-B toll connecting trunk adds to this handily, for an overall loss for the connection of 6 dB. That is, the overall loss is the sum of the loses prescribed for all the trunks in the connection, per the VNL plan for analog trunks and the SDN plan for digital trunks.

In Figure 17, I have extended this situation to one in which the second toll office, the second intertoll trunk, and the second toll connecting trunk are analog.

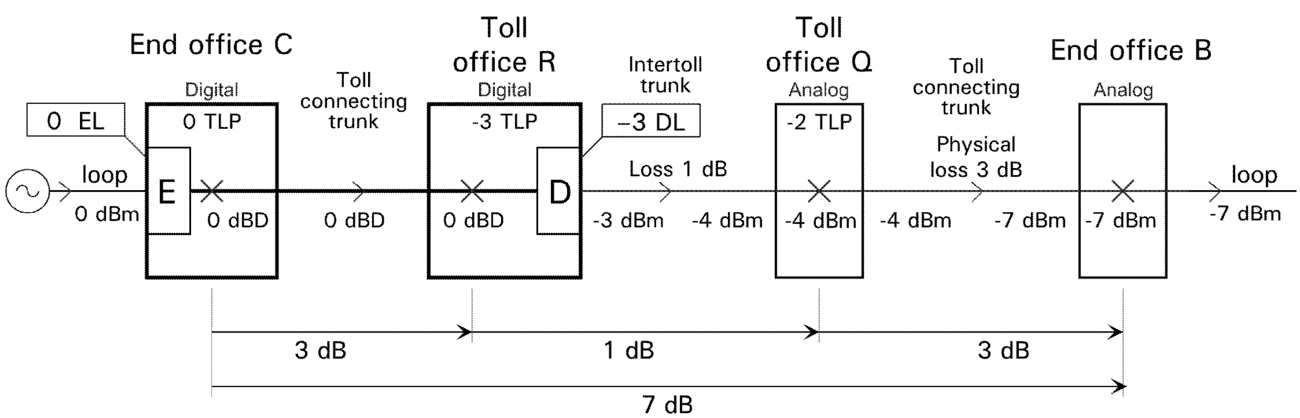

Figure 17. Connection with digital and analog trunks (2)

For completeness, I have shown the TLP of the analog toll office as -2, as is the norm (differing from the norm, +3, for digital toll offices), but this doesn't figure into the story being discussed here.

[15] The loss of the physical facility (usually a passive cable pair) may be greater or less than that, but the actual DL of the decoder that feeds it can be "diddled" to being the effective loss to exactly 3 dB.

By this point, the reader should be able to follow the action here. The result is an overall loss for the connection of 7 dB, just as should result from using. for each trunk, its target loss per the VNL or SDN plan (for analog and digital trunks, respectively).

By the way, here the physical loss of the facility for the Q-B toll connecting trunk may actually be less or greater than 3 dB (just as we discussed for the VNL plan in section 6.8), and in this case we have no adjacent decoder whose gain can be "diddled" to bring the effective loss to exactly 3 dB. So the effective loss of this trunk may not be exactly 3 dB.

In "fairly modern" times, an analog intertoll trunk would likely not be implemented over an individual conductor pair (or a pair of pairs), but rather over a channel of a multiplex system. In that case the power levels I show at both ends of the circuit for that trunk may not actually be meaningful in practice, but serve in this simplified illustration to illuminate the principle of interest.

Incidentally, trunks that terminate on an analog switching system at one end and a digital switching system at the other are called *combination trunks*.

10 (MIS)UNDERSTANDINGS IN PRACTICE

The reader should easily recognize that in all this there was, at any stage of the evolution of the toll transmission plan, great opportunity for bafflement and misunderstanding. Fortunately, much of the work was done under detailed standard practices which "cookbook-ized" the work and mitigated the need to truly understand the underlying doctrines and rationales. This was of course the blessing and the curse of the telephone industry.

11 SUMMARY

We can see that, as the long distance network evolved over the years, the transmission plan evolved to follow, with clever schemes devised to ensure continuity.

-#-

Appendix A
The transmission level point (TLP) of an analog toll office

A.1 INTRODUCTION

In section 9.4, I glancingly mentioned that the transmission level of an analog toll office is declared to be -2. But this did not play any role in the story I was telling. so I did not give it any further attention there.

And in fact that matter has hardly any effect on the whole topic of this article. But it is closely related, so in this Appendix I will discuss the matter.

What does it mean that the TLP of an analog toll office is defined as -2, and how was that decided?

A.2 MULTIPLEX TRANSMISSION SYSTEMS

The matter is most clearly illuminated if we imagine an intertoll trunk that is carried by a channel in a multiplex transmission system, as was often the case for longer intertoll trunks during the era of interest. The kind of multiplex system most widely used in this context (most particularly, a *frequency-division* multiplex system) carries 12 bidirectional voice *channels* over two cable pairs (one used for each direction of transmission). Each of these channels would become a *circuit*, which in turn could carry an intertoll trunk.

For any channel, the voice frequency signal amplitude modulates a carrier frequency, a different one for each channel. The 12 modulated signals (each occupying a different frequency slot) are combined and launched over the cable pair. At the distant end, the 12 modulated signals are separated by filters and each is demodulated to recover the corresponding voice frequency signal.[16]

Although the actual power in the modulated signal for any channel when, for example, the channel is carrying a test tone, is of considerable importance to the system designers (and in fact, to

[16] When this concept was first introduced, especially in the “popular technology” press, much was made over the parallel between this approach and radio broadcasting. This brought great attention to the role of the carrier frequency signals, which led to this approach often being described as “carrier telephony”. The inevitable result was that the early multiple systems were formally designated “carrier (transmission) systems”, and that terminology persisted. In fact. since “multiplex” systems were usually called “carrier systems”, when digital multiplex systems emerged, which did not at all involve the “carrier” concept at all, they were still formal designated “carrier systems”.

technicians working with it), we do not ordinarily speak of it in the context of transmission planning. Rather, we just consider the multiplex channel, as accessed at a standard interface at each end (as configured for the particular situation), to be a "black box", with a certain loss.

A.3 OUR BATTLE ZONE

In Figure 18 we see a familiar-looking figure, representing a hypothetical toll connection with analog switching and transmission.

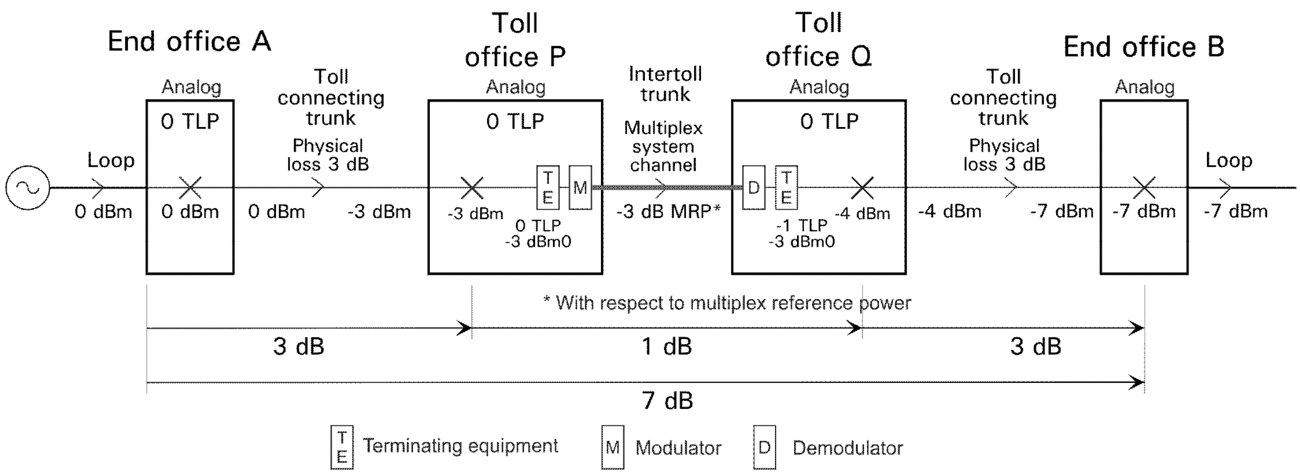

Figure 18. Toll connection with toll office 0 TLP

This in fact shows the P-Q intertoll trunk as being carried by an analog multiplex channel. At the "beginning" end (at toll office P), we see the modulator that creates the online signal for this channel, and some vague "terminating equipment", which embraces various details of the interface between the modulator and the switching fabric. At the "far end" (at toll office Q), we see the complementary modules, in this case involving the demodulator that recreates the voice frequency signal.

In this figure, we assume that the TLP of the toll office had been declared to be 0 (which is what we might think would be the "obvious" thing to do.

The 0 dBm test tone we introduced at end office A arrives at the switching fabric at toll office P at a level of -3 dBm (by virtue of the loss in the A-P toll connecting trunk being 3 dB).

The interface into the multiplex channel is defined as 0 TLP (in keeping with the definition for the office). That means that if we were to apply a 0 dBm test tone into that interface, the "online" signal it would generate is what I will call 0 dB MRP (multiplex reference power[17]).

[17] My term, not found in the literature.

We can think of that as essentially "the greatest power that the modulator-transmission-demodulator chain can safely handle".[18]

But in this scenario, the test tone is fed into the channel interface at a level of -3 dBm. As a result the online signal is at a power we can describe as -3 MRP.

This of course does not "screw up" the loss of the intertoll trunk. At the "far end" (toll office Q), the TLP of the interface out of the multiplex channel is -1, and so the test tone (at -3 dB MRP) will come out at -4 dBm. Thus the intertoll trunk will, as we need, exhibit a loss of 1 dB.

Why is the TLP of that output interface -3 ? It is made that way so the loss will come out to -1 dB, which we have been assuming is what is called for by the VNL transmission plan.

But this situation is not ideal. With our test tone being carried, the power in the online signal is -3 dB MRP. And, when this connection carries actual speech, the speech signals are 3 dB lower than "normal". And this means that the signal-to-noise ratio (SNR) on the speech is 3 dB "worse" than it might have been.

A.4 WITH THE TOLL OFFICE TLP DEFINED AS -2

In Figure 19, we see the same setup, but now with the TLP of the two analog toll offices defined as -2.

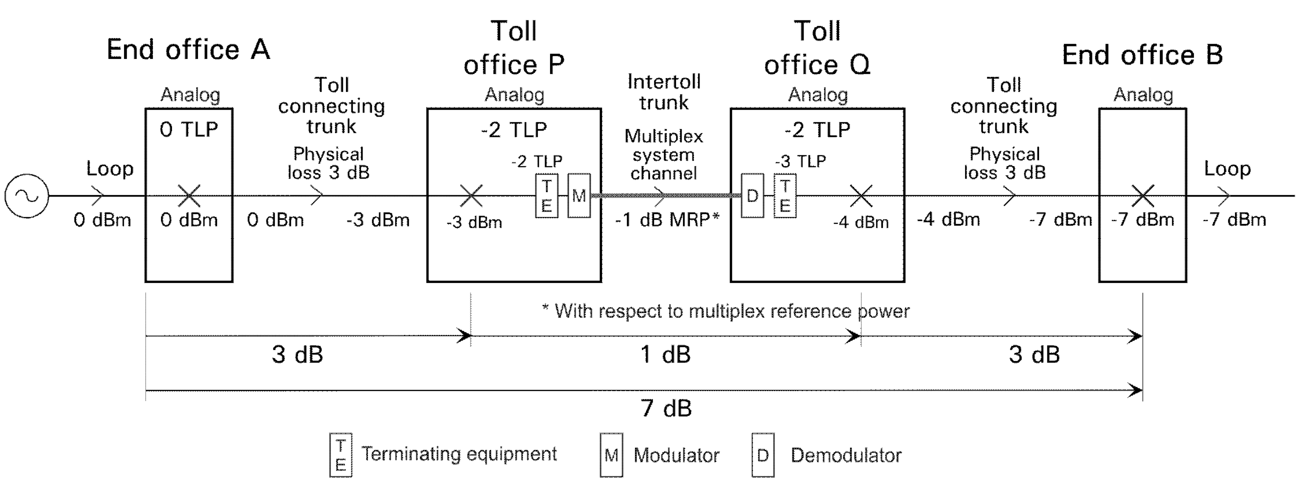

Figure 19. Toll connection with toll office -2 TLP

Now, with the interface into the multiplex channel at toll office P at -2 TLP (in keeping with the designation of the office), we find that in

---

[18] There may actually be as much as 3 dB of "headroom" above this before "clipping" occurs. But we should not "go there", even in testing.

our scenario the test signal “online” is at -1 dB MRP, 2 dB “hotter” than in the prior case. Again, the loss of the intertoll trunk is still 1 dB by virtue of the TLP coming out of the multiplex channel at toll office Q being set at -3.

But given that the “online” signal is 2 dB “hotter” than in the prior case, the SNR ratio will be 2 dB better than in the prior case. And that’s a good thing. So that is how we do it.

A.5 WHY NOT +3 TLP

Now, why did we not make the TLP of the toll offices -3? Then the SNR would have been 1 dB better yet.

The reason is that the loss of the toll connection trunk might not actually be as much as 3 dB. We remember that, in the case that the toll connecting trunk is implemented with a “passive” cable pair, it is “allowed” to have a loss as low as 2 dB (see section 6.8).

Then, if the toll office had a TLP of -3 (and thus the input to the multiplex channel would also have that TLP), the line signals in the multiplex channel would be 1 dB “hotter” than normal. And that is a no-no.

So the -2 TLP designation is a “non-greedy” compromise.

-#-