Guide to 10-Mbps Ethernet Configuration

Charles Spurgeon

Networking Services
University of Texas at Austin

Document version 2.5

NOTICE: The examples and other information in this guide are intended solely as teaching aids and should not be applied to any particular network without independent verification. Independent verification is especially important in any application in which an incorrect network design could result in loss of data or time. For these reasons, while every effort was made to provide accurate information, there is no warranty expressed or implied that the examples, specifications, or other information shown here are free of error, or that they will meet the requirements of any particular network application.

1.0 Introduction

This is a guide to the configuration rules for Ethernet. The rules are provided in two models described in the IEEE 802.3 standard for multi-segment 10 Megabit per second baseband Ethernets. The two models provide two different approaches to Ethernet configuration. Note that these configuration rules do not cover the 100 Megabit per second Ethernet standard.

The first configuration model is a set of rule-based configuration guidelines, and the second model is a calculation method that can be applied to more complex Ethernet systems. These IEEE multi-segment configuration guidelines cover baseband media only, including 10BASE5 thick Ethernet, 10BASE2 thin Ethernet, 10BASE-T twisted-pair

Ethernet, and 10BASE-F fiber optic media systems. The 10BROAD36 broadband system is not covered by the IEEE multi-segment baseband guidelines.

This guide does not provide any descriptions of the IEEE acronyms and various media types, nor does it explain how the Ethernet system works. That information is contained in the **Guide to Ethernet**, which you are strongly encouraged to read if you are not familiar with the IEEE terminology or the range of media types and standards.

2.0 Scope of the Configuration Guidelines

Designing an Ethernet based on a single segment of a given media variety is pretty straightforward. To stay within the specifications when using thick Ethernet, for example, you must make sure that the thick Ethernet cable segment does not exceed 500 meters in length and does not contain more than 100 MAU connections to the cable. As long as the cable that you use to build the segment also meets the media specifications and is properly installed, the system will work correctly.

However, when you build a multi-segment network things get more complex. To help you maintain compliance with the Ethernet specifications when combining segments from a variety of Ethernet media types, the IEEE developed two models for verifying the operation of a multi-segment Ethernet. Transmission System Model 1 provides a set of standard configuration rules, and Transmission System Model 2 provides a set of calculation aids so that you can do the calculations yourself.

2.1 Collision Domain

The scope of the multi-segment configuration guidelines is limited to a single Ethernet, or “collision domain.” A collision domain is formally defined as a single CSMA/CD network in which there will be a collision if two computers attached to the system both transmit at the same time. To combine multiple segments into a single Ethernet system you use repeaters, MAUs, AUI cables, and the segments themselves to create a network that functions as a single collision domain.

Each computer, or DTE, contains an Ethernet interface that implements the medium access control (MAC) rules for Ethernet. If two DTEs with their associated Ethernet interfaces and MACs are attached to segments that are connected by repeaters then they are within the same collision domain, since repeaters are designed to propagate collisions onto all segments to which they are connected.

If two DTEs are instead separated by a packet switch such as a bridge or router, then they are in separate collision domains, since packet switches do not propagate collisions. Instead, bridges and routers contain multiple Ethernet interfaces and are designed to receive a packet on one Ethernet and transmit the data onto another Ethernet in a new packet.

Instead of propagating collision signals between Ethernets, packet switches interrupt the collision domain and provide separation for the operation of the Ethernets they link. Therefore, you can use packet switches to build larger network systems by connecting individual Ethernet systems. The point is that the configuration guidelines apply to a single collision domain only, and have nothing to say about combining multiple Ethernets with packet switches.

3.0 IEEE Guidelines for Baseband Multi-segment Networks[1]

The guidelines and calculations used for verifying an Ethernet configuration are based on maintaining two important features of an Ethernet system: the maximum round-trip signal propagation delay and the minimum inter-packet gap.

3.1 Round-Trip Signal Propagation Delay

The limit on the maximum round-trip signal propagation delay is related to the operation of Ethernet, which is based on Carrier Sense Multiple Access with Collision Detection (CSMA/CD). For the CSMA/CD system to operate properly, all devices attached to an Ethernet must be able to hear each other's transmissions and to respond to collisions within the correct time limits.

Each segment you use in an Ethernet provides a certain amount of signal delay. The important thing is that when you build a multi-segment Ethernet, the total round-trip signal propagation delay must not exceed the limits in the specifications. By following the configuration guidelines you can ensure that your system meets the limits set forth in the specifications.

3.2 Inter-Packet Gap

The standard for the inter-packet gap establishes a minimum spacing between packets sent over an Ethernet. Ethernet interfaces are designed so that they can receive packets that are sent “back-to-back” with only the minimum inter-packet gap between them. As long as the inter-packet gap is correct, an Ethernet interface chip can keep up with the full traffic rate and not miss any packets.

The signal reconstruction circuits in an Ethernet repeater, combined with other signal variations that can occur in the transmission path of a packet, can result in a shrinkage in the inter-packet gap between two packets travelling across the network. Therefore, the requirement to maintain a minimum gap between packets means that the amount of equipment that may be in the path between any two DTEs must be limited. By follow-

1. Portions of the information contained herein are copyrighted information of the IEEE extracted from IEEE Std 802.3j-1993, IEEE Standard for Local and Metropolitan Area Networks – Fiber Optic Active and Passive Star-Based Elements, Type 10BASE-F, copyright © 1993 by the Institute of Electrical and Electronics Engineers, Inc. The IEEE takes no responsibility for damages resulting from the reader’s misinterpretation of said information resulting from the placement and context in this publication. Information is reproduced with the permission of the IEEE.

ing the multi-segment configuration guidelines you can ensure that your network system is capable of maintaining the correct inter-packet gap and prevent the loss of packets.

4.0 Transmission System Model 1

In this model, a set of multi-segment configuration rules are provided based on conservative calculations for the components involved. Even though these configuration rules are based on conservative calculations, don't let this convince you that you can bend the rules and always get away with it. There isn't a lot of engineering margin left in maximum-sized Ethernets, despite the allowances made in the standards for manufacturing tolerances and equipment variances.

Therefore, even though there is some margin built into the rule-based configuration guidelines, this does not mean that an Ethernet system that is outside the guidelines will be sure to work. A system that violates the guidelines may work, and it may not. Worse yet, it may appear to work acceptably for a while but then proceed to fail when the network grows in size or the traffic load and the population of attached stations exceeds some level. If you want guaranteed performance and reliability, then you need to stick to the published guidelines.

The multi-segment configuration rules are as follows. (Bold face type indicates text taken directly from the IEEE standard.)

1. **Repeater sets are required for all segment interconnection.** The repeaters used must comply with all IEEE specifications in section 9 of the 802.3 standard, and do signal retiming and reshaping, preamble regeneration, etc. If you do not use true IEEE 802.3 repeaters for all segment interconnections, then your Ethernet system cannot be verified using either configuration model.
2. **MAUs that are part of repeater sets count toward the maximum number of MAUs on a segment.** Thick Ethernet repeaters typically use an outboard MAU to connect to the thick Ethernet coax. Thin coax and twisted-pair repeater hubs use internal MAUs located on each repeater port.
3. **The transmission path permitted between any two DTEs may consist of up to five segments, four repeater sets (including optional AUIs), two MAUs, and two AUIs**. The repeater sets are assumed to have their own MAUs, which are not counted in this rule.
4. **AUI cables for 10BASE-FP and 10BASE-FL shall not exceed 25 m. (Since two MAUs per segment are required, 25 m per MAU results in a total AUI cable length of 50 m per segment).**
5. **When a transmission path consists of four repeaters and five segments, up to three of the segments may be mixing and the remainder must be link segments. When five segments are present, each fiber optic link segment (FOIRL, 10BASE-FB, or 10BASE-FL) shall not exceed 500 m, and each 10BASE-FP segment shall not exceed 300 m.**
6. **When a transmission path consists of three repeater sets and four segments, the following restrictions apply**:

a. The maximum allowable length of any inter-repeater fiber segment shall not exceed 1000 m for FOIRL, 10BASE-FB, and 10BASE-FL segments and shall not exceed 700 m for 10BASE-FP segments.

b. The maximum allowable length of any repeater to DTE fiber segment shall not exceed 400 m for 10BASE-FL segments and shall not exceed 300 m for 10BASE-FP segments and 400 m for segments terminated in a 10BASE-FL MAU.

c. There is no restriction on the number of mixing segments in this case. In other words, when using three repeater sets and four segments, all segments may be mixing segments if desired.

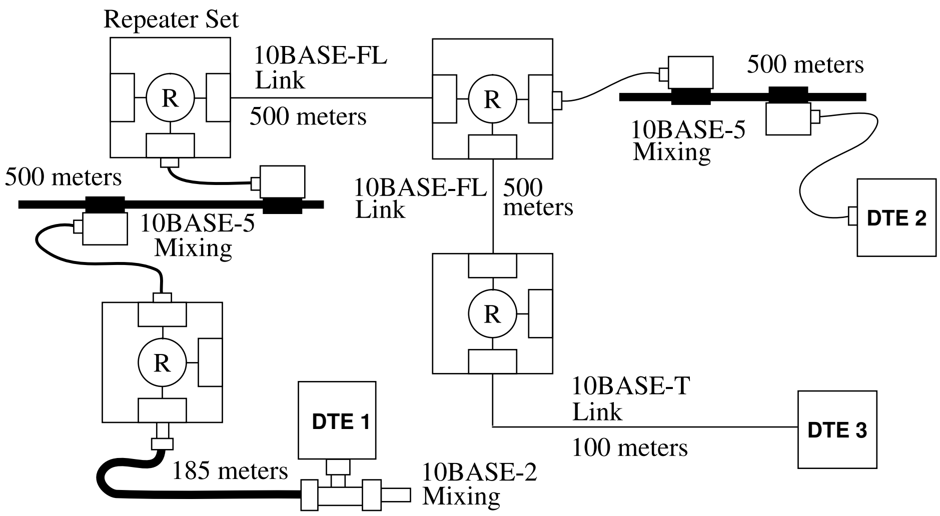

FIGURE 1. One possible configuration using four repeaters.

Figure 1 shows an example of one possible configuration using a mixture of segment types and a total of four repeaters. The maximum transmission path is determined by inspection; you look at the drawing of the network and determine which path contains the most segments and repeaters. The maximum path in Figure 1 occurs between DTE 1 and DTE 3, since this is the signal path with the most repeaters and segments in it. This path includes five segments linked with four repeaters.

While the configuration guidelines emphasize the maximum limits of the system, you should beware of stretching things as far as they can go. Ethernets, like many other systems, work best when they are not being pushed to their limits. As the performance of computers continues to increase, many network managers are finding that they need to limit the number of computers they attach to a given Ethernet system to prevent network overload.

The exact number of computers a given Ethernet can support is a local issue, and is determined by your mix of hardware and software and the network load the machines produce. However, you should be aware that network loads are steadily increasing at most sites and adjust your network designs accordingly.

5.0 Transmission System Model 2

The second model provided by the IEEE standard for qualifying an Ethernet configuration provides a set of calculation aids, so that you can verify the operation of more complex Ethernet systems than those covered by the rule-based model. In the calculation model there are two sets of calculations that must be performed for each Ethernet system that you wish to verify.

The first set of calculations is based on ensuring that the round trip signal propagation delay is within the correct limits required to allow the essential collision detection mechanism to work properly. The second set of calculations verifies that the amount of interfame gap shrinkage is within the correct limits. Let's begin with the calculations for signal delay.

5.1 Path Delay Calculation

In a valid network it must be possible for any two DTEs to fairly contend for access to the shared Ethernet channel if they happen to transmit at the same time. Each DTE attempting to transmit must be notified of channel contention (collision) by receiving a collision signal within the correct collision timing window. Also, the frame fragment created as a result of a collision must be less than 511 bits long.

These requirements place limits on the physical diameter, or maximum distance between DTEs, of a network. To verify that your Ethernet system meets these limits you need to calculate the path delay of the maximum packet transmission path in your system.

There are a number of engineering considerations that affect the total path delay for each media type including collision detect time and the rate of signal propagation. The calculation model is designed to simplify the process of verifying a network by providing delay values for each segment type that incorporate the complete set of engineering considerations.

The operation of the calculation model is based on a simplified network topology called a generalized transmission path model. This model includes a “left end,” one or more “middle segments,” and a “right end.” To perform the calculations you determine the worst-case path for your system and then you list the segment types that are used in that path. Next, you calculate the worst-case path delay based on the transmission path model and using the delay values provided in the standard.

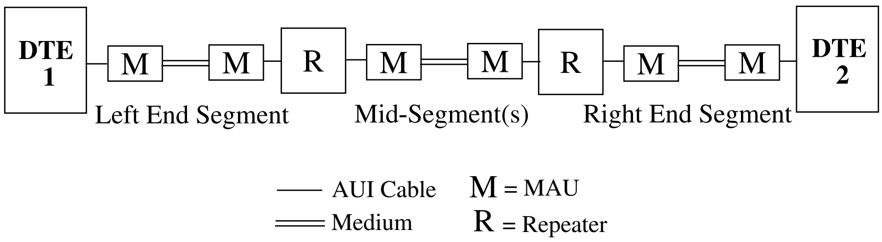

FIGURE 2. Generalized transmission path model from section 13 of the 802.3 standard.

The transmission path model provides a way to structure the process of evaluating your network segments to find the worst-case path delay.

The model is based on the following operational scenario: DTE1 transmits. DTE1's transmission propagates to DTE2. DTE2 begins transmitting at the last possible time before hearing the transmission from DTE1, causing a collision and transmitting a total of 96 bits. DTE2's transmission propagates back to DTE1. DTE1 detects the collision, sends out a collision enforcement jam signal, and stops transmitting.

Based on this scenario, the calculations for verifying the operation of an Ethernet must ensure that the following conditions are met as described in this excerpt from Appendix A of the standard.

> (1) DTE1 must detect collision before having transmitted the 512th bit, including the preamble and SFD bits.
>
> (2) DTE1 must stop transmitting before having transmitted a minimum length frame, 576 bits. (512 bits after SFD).
>
> (3) The overlap between DTE1's transmission and DTE2's transmission must be less than 575 bits (511 bits after the SFD transmitted by DTE1).
>
> For all existing segment types, the last condition is the limiting factor; if it is met, then the other two conditions are also met.
>
> The maximum time between the first bit and the last bit of the overlapping transmissions of the two DTEs colliding across a path will be called the Path Delay Value (PDV). Many factors contribute to this delay. Simplification of the delay calculation, ... can be achieved by using a set of base numbers, Segment Delay Values (SDV), for each segment type that combines the factors that contribute to round trip delay associated with that segment. The PDV (Path Delay Value) is the sum the SDVs that comprise the path. [1]

1. *IEEE Std 802.3j-1993, IEEE Standard for Local and Metropolitan Area Networks – Fiber Optic Active and Passive Star-Based Elements, Type 10BASE-F.* (New York: Institute of Electrical and Electronics Engineers, October 13, 1993), p. 34.

You begin the process of calculating the worst-case PDV by determining which path in your network is the maximum delay path. If you have a complete and up-to-date map of your network you can use that map to find the maximum path between two DTEs.

If your system is not well documented then you will have to investigate and map the network yourself. You need to know what kinds of segments are in use, how long they are, and how the system is laid out. Once you have this information, then you can determine what the maximum path is and what kinds of segments are used in the maximum path.

For each segment type used in the maximum-length path of your network, one of three segment delay values will be used in your calculation depending on whether the segment is a left end, middle, or right end segment.

5.2 Calculating the Worst-case Path Delay Value

To see how this all works, let's walk through a sample path delay calculation. To do this, refer back to Figure 1, which shows one possible maximum-length system using four repeaters and five segments.

As we've seen, the rule-based configuration method says that this system is OK. To check that, we'll evaluate the same system using the calculation method. The table shown next provides the segment delay values used for making the worst-case path delay calculation.

TABLE 1. Segment Round-Trip Delay Values in Bit Times from IEEE 802.3 Standard, Section 13

Segment Type	Max Length	Left Base	End Max	Mid-Base	Segment Max	Right Base	End Max	RT Delay/ meter
10BASE5 Coax	500	11.75	55.05	46.5	89.8	169.5	212.8	0.0866
10BASE2 Coax	185	11.75	30.731	46.5	65.48	169.5	188.48	0.1026
FOIRL	1000	7.75	107.75	29	129	152	252	0.1
10BASE-T	100[a]	15.25	26.55	42	53.3	165	176.3	0.113
10BASE-FP	1000	11.25	111.25	61	161	183.5	284	0.1
10BASE-FB	2000	N/A[b]	N/A[b]	24	224	N/A[b]	N/A[b]	0.1
10BASE-FL	2000	12.25	212.25	33.5	233.5	156.5	356.5	0.1
Excess AUI	48	0	4.88	0	4.88	0	4.88	0.1026

a. Actual maximum segment length depends on cable characteristics.

b. N/A: Not Applicable, 10BASE-FB does not support end connections.

To begin with, we need to determine which path is the maximum delay path in the system. By examination, you can see that the path in Figure 1 between DTE1 and DTE3 is the maximum delay path since it contains the largest number of segments and repeaters in the path between any two DTEs.

To perform the calculation for each segment in our system we need to compute the total segment delay value (SDV). The total SDV for a given segment is a combination of the base SDV and the delay represented by the length of the segment in question.

SDV = Base + (Length * (Round Trip Delay/meter))

The SDV value for each segment type used in a network is found by adding the base value for that segment type to the product of the segment length and round trip delay per meter (RT Delay/meter).

For a mixing segment like thin coaxial Ethernet the length of the segment is the length between the repeater connection and the farthest end of the segment for an end segment, or the length between two repeater connections for mid segments.

If you are not sure what length a segment is you can use the maximum value shown, which has been set equal to the maximum media segment length allowed in the specifications. In that case: **SDV = Max.** This assumes that your segment is not longer than the maximum value allowed in the media specifications for that media type.

5.3 Four Steps for Calculating Worst-Case Path Delay

Here are the four steps that will develop a total worst-case path delay value (PDV) for an Ethernet system.

1. Determine the SDV for each segment in the maximum delay path.
2. If a candidate for worst-case path has end segments of different types, perform the calculations twice using first one end segment as the left end, then the other. The maximum value obtained should be used as the worst-case path delay value.
3. From the table, determine the SDV for the sum of all AUI cables in excess of two meters, except the AUI cable associated with the left end DTE, which does not contribute to the total path delay value.
4. Sum all SDVs from steps 1 through 3, plus a margin of up to five bit times to form the total path delay value. The margin may be from zero to five bits, but five bit times is recommended. If the total path delay value is less than or equal to 575 bit times, then the path is qualified in terms of worst-case delay.

5.4 Calculating the Worst-Case Delay for the Sample Network

To start with, we set out the segments in our maximum delay path as left end, middle, and right end segments. Let’s begin by assuming that the thin Ethernet segment that DTE1 is attached to is the left end segment. That leaves us with three middle segments composed of a 10BASE5 segment and two link segments, and a right end segment which is a 10BASE-T link segment.

Next, we begin the calculations with the left end, using the 185 meter 10BASE2 thin Ethernet segment. We could calculate the segment delay value by adding the left end base (11.75) to the product of the round trip delay times the length (185 * 0.1026 = 18.981) to come up with a value of 30.731. Since 185 meters is the maximum segment length allowed for 10BASE2 segments, we could be lazy and simply look up the max left hand segment value in the table, which, not surprisingly, is also 30.731.

The segment delay values provided in the table include allowances for an AUI cable of up to two meters length at each end of the segment, except for 10BASE-FB segments which connect directly to special repeater hubs and never use AUI cables.

In many thick Ethernet systems, however, or in other systems with outboard MAUs attached to segments, the AUI is often longer than two meters. In that case, you can determine how long your AUI cables are and use that length times the RT Delay/meter for excess length AUIs to develop an extra delay value which is added to your total path delay value calculations.

If you’re not sure how long the AUI cables are you may wish to use the maximum delay shown for an AUI cable, which is 4.88 for all segment locations, left end, middle, or right end.

Since the left and right end segments in Figure 1 are different media types, we need to do two calculations to meet the requirement in step 2. To meet step 2 we first calculate the total path delay using the 10BASE2 segment as the left end segment and the 10BASE-T segment as the right end.

Then we swap their places and make the calculation again, using the 10BASE-T segment as the left end segment this time, and the 10BASE2 segment as the right end segment. The worst-case value that results from the two calculations is the one we use for our system.

Let’s continue by making the calculations for the middle segments. In Figure 1 there are three mid-segments composed of a maximum length 10BASE5 segment, and two 500 meter long 10base-FL segments.

By looking in the table under mid-segments we find that the 10BASE5 segment has a max delay value of 89.8. Let’s add in the delay for two AUI cables to allow for two maximum-length AUI cables in the segment, one at each connection to a repeater. That gives us an AUI cable delay of 9.76 to add to the total path delay.

We can calculate the segment delay value for the 10BASE-FL mid-segments by multiplying the 500 meter length of each segment times the RT Delay/meter, which is 0.1, which gives us a result of 50. We then add 50 to the mid-segment base value for a 10BASE-FL segment, which is 33.5, for a total of 83.5.

While we’re at it, let’s assume that we followed item 4 of the rule-based guidelines and used two AUI cables of 25 meters length, for a total of 50 meters of AUI cable on the segment. We can represent that length by adding 4.88 extra bit times to the total path delay.

Finally, we have a right end segment of 10BASE-T twisted-pair Ethernet, which is the maximum length of 100 meters. The max value for 10BASE-T right end segment is 176.3. Adding all the bit time values together we get the following:

Left End	10BASE2	30.731
Mid-segment	10BASE5	89.8
Mid-segment	10BASE-FL	83.5
Mid-segment	10BASE-FL	83.5
Right End	10BASE-T	176.3
Excess Length AUI	Quan. Three	14.64
	Total PDV =	478.471

To complete the process we need to perform a second set of calculations with the left and right segments swapped to fulfill the requirement in step 2. In this case. the left end becomes a maximum length 10BASE-T segment, with a value of 26.55, and the right end becomes a maximum length 10BASE-2 segment with a value of 188.48. Adding the bit time values again, we get the following:

Left End	10BASE-T	26.55
Mid-segment	10BASE5	89.8
Mid-segment	10BASE-FL	83.5
Mid-segment	10BASE-FL	83.5
Right End	10BASE2	188.48
Excess Length AUI	Quan. Three	14.64
	Total PDV =	486.47

Since the second set of calculations produced a larger value, this is the value we must use for the worst-case path delay for this Ethernet.

Finally, the standard recommends adding a margin of up to five bit times to form the total path delay value. We are allowed to add anywhere from zero to five bits margin, but five bit times is recommended.

Adding five bit times for margin brings us up to 491.47, which is less than the maximum of 575 bit times. Therefore, our sample network is qualified in terms of the worst-case delay. All shorter paths will have smaller delay values, so all paths in the Ethernet system shown in Figure 1 meet the requirements of the standard as far as path delay value is concerned.

Calculating the worst-case path delay value is not sufficient to qualify an Ethernet under Transmission System Model 2. To complete the process you must also calculate the worst-case inter-packet gap shrinkage that results for the set of equipment used in your Ethernet.

5.5 Calculating Worst-Case Inter-Packet Gap Shrinkage

The inter-packet gap (IPG) may shrink on an Ethernet system when two successive packets encounter differing bit loss on the same path. This occurs because repeaters are required to fully regenerate the preamble of each packet that passes through them.

If the first of two successive packets has experienced greater bit loss in its preamble than the second, then the process of regenerating the preamble in the repeater will result in a shrinkage in the inter-packet gap between the two packets.

The IPG shrinkage is calculated using a value for each segment type, called the segment variability value (SVV). For each segment type along the worst-case path in your system, one of two segment variability values is used depending on the position of the segment: whether it is the transmitting end or the mid segment.

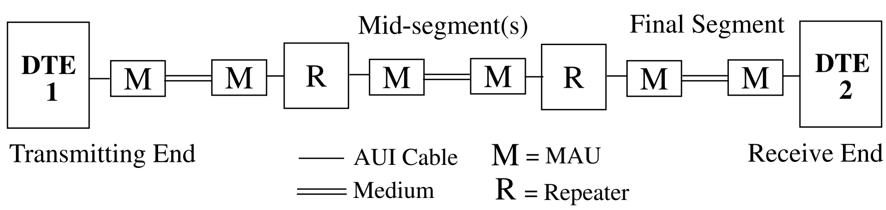

FIGURE 3. Variability model from Section 13 of the IEEE 802.3 standard.

The scenario for the operation of the variability model notes that the transmitting end segment and the mid segment variability values include the variability that may occur from the transmitting MAU through the associated repeater unit.

Since IPG shrinkage only occurs when a repeater restores lost preamble bits, the final segment connected to the receiving DTE does not contribute any variability and is not included in the calculations. To make the calculation for path variability, therefore, you do not count the final segment attached to the DTE on the receive end.

In a network where the receive and transmit end segments are not the same media type, however, you should use the end segment with the worst variability as the “transmitting end” for the purposes of this calculation. This will provide the worst-case value for inter-packet gap shrinkage.

The following table provides the values used in the segment variability calculations.

TABLE 2. Segment Variability Values in Bit Times

Segment Type	Transmitting End	Mid-segment
Coax	16	11
Link except 10BASE-FB	10.5	8
10BASE-FB	N/A[a]	2
10BASE-FP	11	8

a. N/A: Not Applicable, 10BASE-FB does not support end connections.

To make the calculation of worst-case path variability value (PVV) perform the following steps.

1. Determining the segment variability value for each of the segments in the worst-case path, excluding the end segment with the lower SVV.
2. Sum all of the SVVs from step 1 to form the total path variability value. If the PVV is less than or equal to 49, the path meets the specifications in terms of worst-case variability.

Let's finish the evaluation of the network shown in Figure 1 by calculating the worst-case path variability value for that network. This is done by evaluating the same worst-case path that we used in the path delay calculations. However, for the purposes of calculating segment variability we only evaluate the transmitting and mid-segments.

According to step 1, the transmitting segment should be set equal to the end segment in the network path that has the worst variability value. As shown in the table, the coax segment has the worst-case value, so we will assume that the 10BASE2 segment is the transmitting end. The mid segments consist of one coax and two link segments. That leaves a 10BASE-T receive end segment which is simply ignored. The totals are:

Transmitting End Coax	16
Mid-segment Coax	11
Mid-segment Link	8
Mid-segment Link	8
Total PVV =	43

As you can see, the total path variability value for our sample network equals 43, which is less than the 49 bit time maximum that is allowed by the standard. With a worst-case path delay value of 491.47 bit times and a worst-case path variability value of 43, the Ethernet system shown in Figure 1 meets the specifications and is qualified under the calculation method of Transmission System Model 2.