INCOMPLIANCE
ELECTRONIC DESIGN, TESTING & STANDARDS

Designing Ethernet Cable Ports to Withstand Lightning Surges

The subject of lightning damage to Ethernet ports has received increased attention in recent years. For reasons that are not fully understood, more manufacturers are experiencing field failures caused by visible surge damage to Ethernet ports, such as that shown in Figure 1. Other suspected surge failures do not show visible damage, but the timing of the failures correlates with local thunderstorm activity. It is not clear whether this recent increase in failures is simply due to having more Ethernet-connected equipment now in service, or because of changes in the types of equipment that are being interconnected via Ethernet cables.

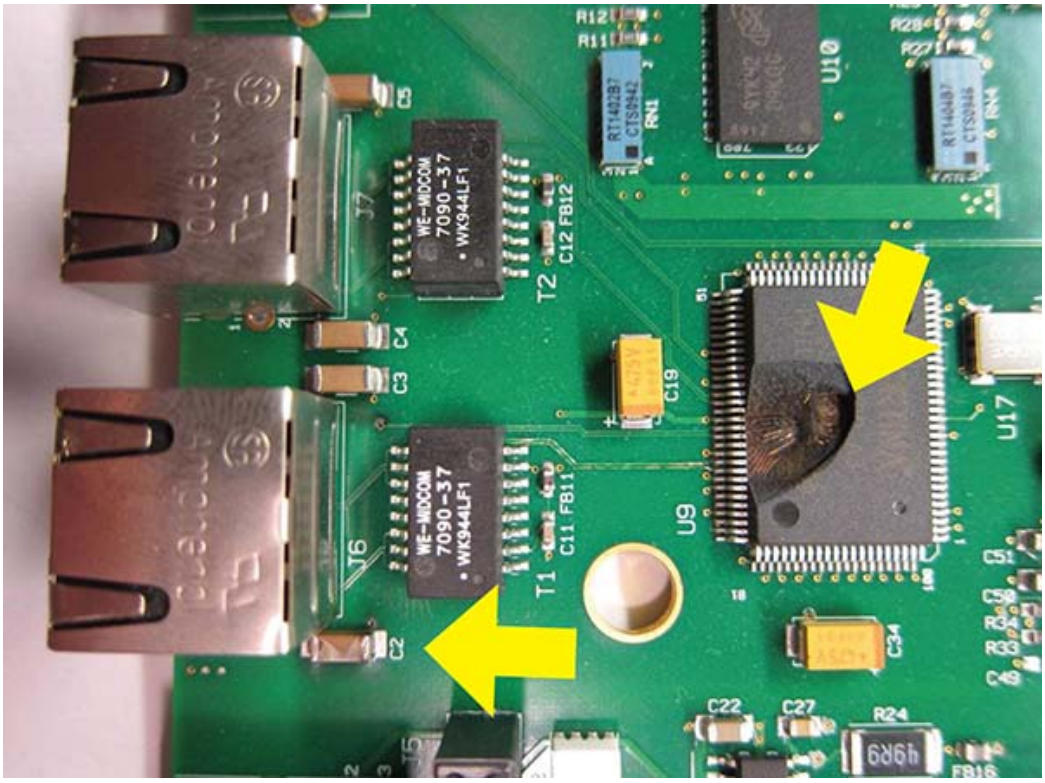

Figure 1: Lightning Surge Damage to an Ethernet Port

Due to increasing industry awareness of surge damage to Ethernet ports, considerable attention has been given to understanding the surge coupling mechanisms involved. Various theories have been put forward,

all of which are based on the established physics of how lightning-induced surges can appear on communication cables. At this time there is no clear consensus on which of these known mechanisms is the dominant cause of the observed damage to Ethernet ports.

This article will briefly review the known mechanisms that couple lightning surges onto communication cables, and will describe how these mechanisms apply for the specific case of Ethernet cabling. A basic understanding of these principles is useful for understanding the emerging industry requirements for surge tolerance of Ethernet ports.

Once the various types of Ethernet surge threats are explained, some specific design strategies will be suggested for obtaining various levels of surge tolerance. It turns out that useful improvements can be obtained at low cost by applying some simple design rules. High levels of surge tolerance will typically add some cost, but the same basic protection strategies still apply.

Background

Historically, many people assumed that Ethernet ports are not subject to lightning surges. After all, most Ethernet cables are routed entirely within a given building, and would not appear to have any exposure to lightning surges. In some cases, Ethernet cables are routed outside a building for applications such as security cameras that use Ethernet. Many people consider outdoor applications to be the only case where lightning surges need to be considered, but experience has shown that this is not the case.

As the incidence of Ethernet port surge damage became more widely recognized, various vendors and industry groups began to focus on understanding why this surge damage was occurring and how it could be prevented. Several large companies have developed their own internal requirements for surge tolerance of their Ethernet ports, and representatives of these companies have been working to influence emerging industry standards.

Standards Activity

The first public standard to address the increased incidence of Ethernet port surge damage was the 2011 edition (Issue 6) of the Telcordia NEBS GR-1089 standard that is used by most of the major telecom carriers in the USA. GR-1089 is not a regulatory requirement, but most telecom carriers in the USA use it as part of their purchase specifications for telecom network equipment. Telecom network equipment is a large market, so the changes in Issue 6 of GR-1089 generated a lot of attention. In Issue 6 of GR-1089, the surge tolerance requirements for Ethernet ports were greatly expanded compared to the previous edition.

Ethernet surge tolerance has also received attention at the international level. During the period of 2012-2015, the International Telecommunications Union standards K.20, K.21, K.44, and K.45 were all updated to include specific lightning surge tests for Ethernet ports.

Most recently, in 2016, the Alliance for Telecommunications Industry Solutions (ATIS) issued ATIS-0600036, “Electrical Protection for Ethernet Systems.” This document provides a comprehensive review of various Ethernet surge hazards and the methods that can be used to protect against these hazards. There is also a section containing recommended test requirements for specific types of Ethernet deployment scenarios.

In summary, there is growing recognition that Ethernet ports are being damaged by lightning surges, and a growing number of companies and standards organizations are taking steps to include Ethernet surge immunity as one of the requirements they impose for field reliability.

This article will not discuss the specific requirements in these emerging standards. Rather, the Ethernet surge problem will be discussed in general terms to familiarize the reader with the types of surges that the standards are trying to address.

Surge Coupling Mechanisms

It is important to understand that even for conventional outside wiring such as telephone lines, lightning almost never strikes the cable directly. Rather, lightning strikes an object nearby, and various coupling mechanisms induce a surge on the cable. In the very rare circumstances where lightning strikes the cable directly, the result is melted copper and extensive physical damage. Fortunately, direct strikes to a cable are extremely rare.

So, how can a lightning surge get coupled onto a cable if the lightning does not directly strike the cable itself? There are three basic coupling mechanisms:

1. Far-field electromagnetic coupling
2. Near-field electromagnetic coupling
3. Ground potential rise

Far-Field Electromagnetic Coupling

Figure 2 shows lightning striking the ground near an aerial cable. This creates an electromagnetic pulse that couples into the cable. The induced surge current I_S is only a small fraction of the actual strike current I_L, but even a small degree of coupling can pose a problem, since the actual strike current averages about 30 kA and can sometimes exceed 100 kA.

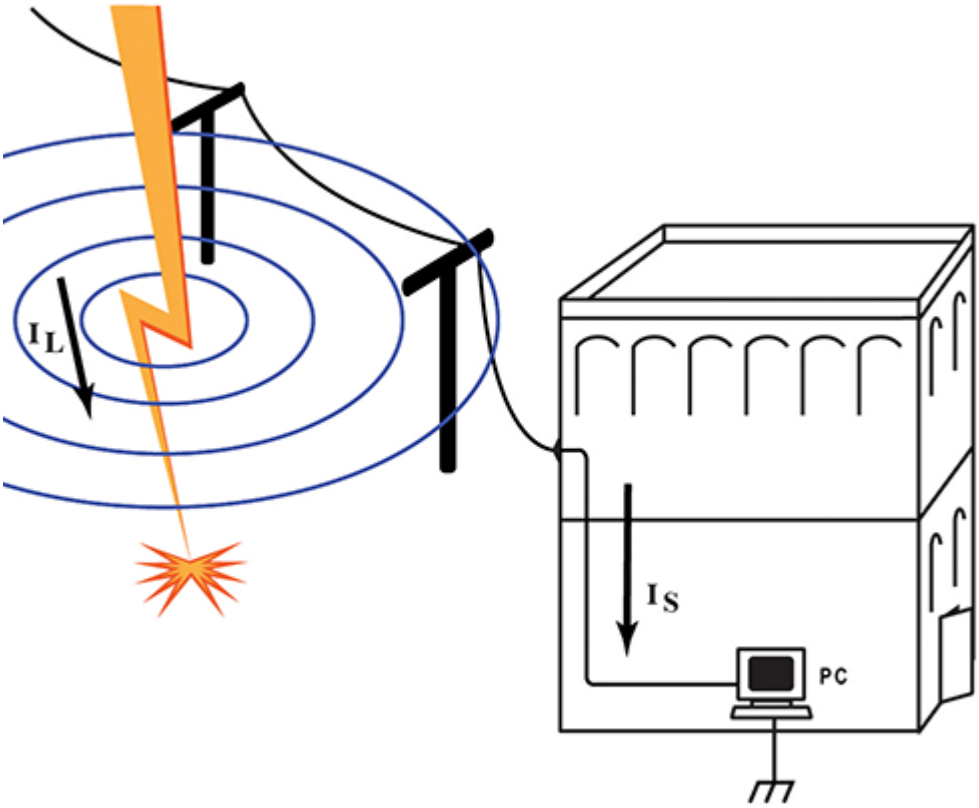

Figure 2: Far-Field Electromagnetic Coupling

Interestingly, having the cable strung inside a building rather than outside often has little effect on the degree of coupling. Materials such as wood, brick, and concrete provide almost no attenuation to the electromagnetic pulse, although the reinforcing steel used in some concrete construction can provide several dB of attenuation.

Near-Field Electromagnetic Coupling

Figure 3 shows lightning striking a building and being conducted to ground by either a down-conductor of the building's lightning protection system, or by the building's steel frame. If the cable is routed in close proximity to this path for some distance, a surge can be coupled into the cable.

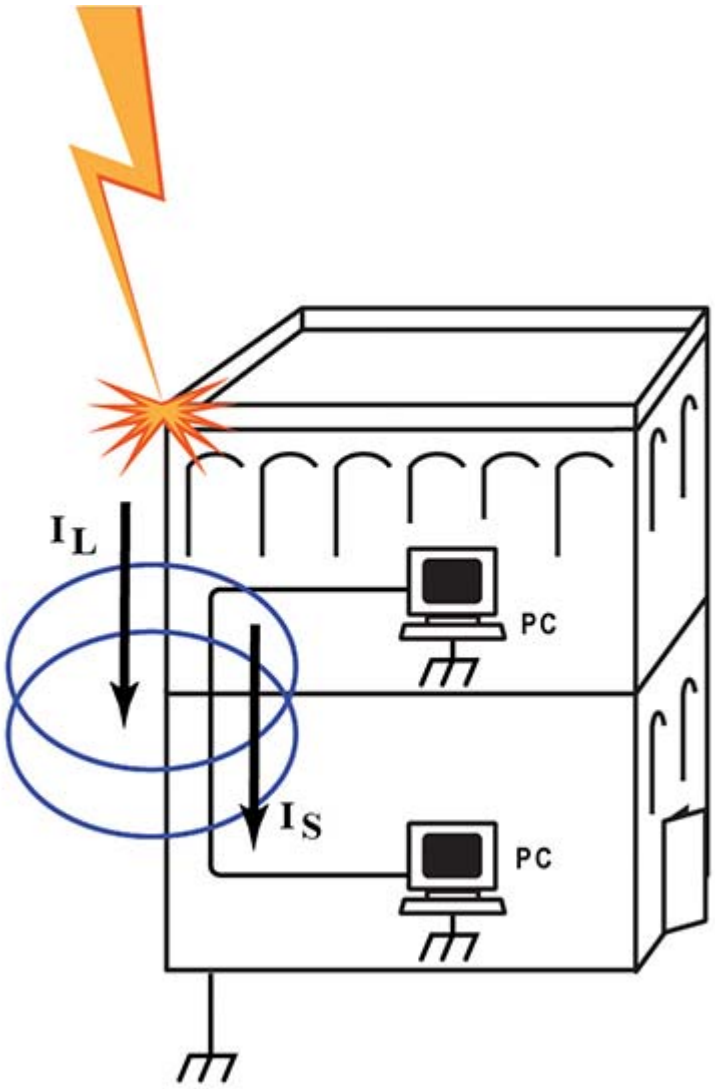

Figure 3: Near-Field Electromagnetic Coupling

Ground Potential Rise

Figure 4 shows a mechanism called ground potential rise (GPR). GPR is perhaps the most complicated mechanism to explain, but it is often the root cause of surge failures on communication cables.

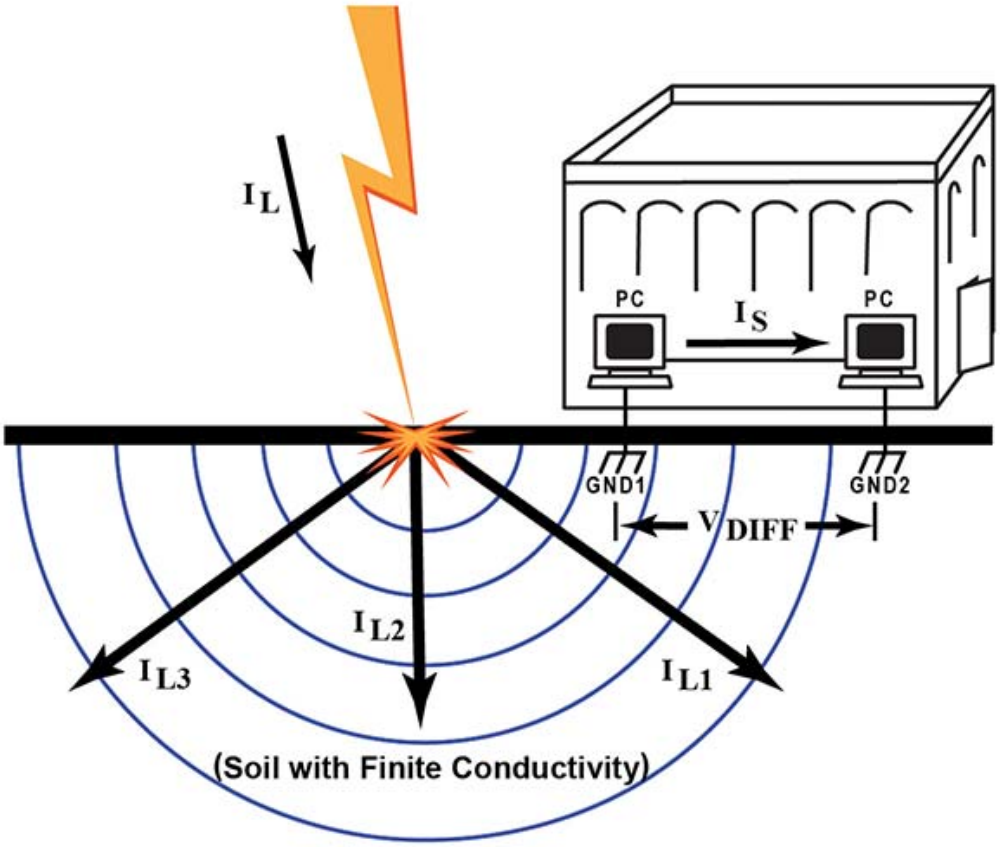

Figure 4: Ground Potential Rise

When lightning strikes the ground, currents up to 100 kA flow through the soil toward the center of the earth. However, since the soil has a finite resistance, the lightning current I_L spreads out in many directions, represented by I_{L1}, I_{L2}, and I_{L3} in Figure 4. This creates a voltage gradient that can be represented as a set of equipotential half-spheres, indicated by the blue lines in Figure 4.

If there are two devices nearby that each have a separate reference to ground, the surge current I_{L1} flowing through the soil creates a momentary difference V_{DIFF} in the potential at the two separate ground points GND1 and GND2. This voltage difference can be several thousand volts if the lightning strike is nearby. Reference [1] contains a more detailed description of the mechanisms that lead to GPR.

Applying the Three Surge Coupling Mechanisms to Ethernet

Now that we have identified three mechanisms by which lightning surges can couple into a cable, we can look more closely at the special case of Ethernet and draw some conclusions.

Since Ethernet cables are inherently limited to about 100 meters long, there is little opportunity to generate large voltage surges through the far-field coupling mechanism of Figure 1. Generally, a cable needs to be more than 300 meters long to induce surges over 1 kV. This suggests that far-field coupling onto Ethernet cables is not a major cause of damage to Ethernet ports.

The near-field coupling mechanism shown in Figure 2 can induce a large surge into a 100 meter Ethernet cable if the cable is routed for most of its length in a way that closely couples it to the lightning conductor. However, the coupling has to be very close, such as when the Ethernet cable is physically tied to a conductive element that carries high surge currents to ground. While such situations do occur (and should be avoided), most Ethernet cables are routed in ways that would not provide sufficient coupling to generate a large surge.

This leaves us with just GPR as a suspect for most of the observed surge damage to Ethernet ports. It turns out that for widely distributed networks that contain multiple terminal devices and routers, the opportunity for GPR-induced surges increases considerably. However, we should keep in mind two requirements that must be met for GPR to cause surge damage:

1. Lightning must strike the ground (or a grounded object) within about 100 meters of the affected equipment.
2. There must be different ground references at each end of the Ethernet cable.

Analysis of field installations that experienced Ethernet surge damage has shown that in many cases, the cause was most likely a GPR event. However, some failures have occurred in installations where it did not appear that GPR or either of the other two mechanisms described above could have caused the damage. It seems that some other mechanism is involved for at least some of the observed field failures.

Could Ethernet Surges Be Coming from the AC Mains?

Many people who are studying Ethernet surge failures have begun to suspect that somehow, surges that originate on the AC mains are getting onto Ethernet cables. Since the power distribution network uses very long cables in highly exposed environments, surges on the AC mains can be quite large. Surges in the range of 5 kV to 10 kV are somewhat common.

Various theories have been put forward to describe how surges on the AC mains could get onto an Ethernet cable. For example, most devices with an Ethernet port also have an AC mains port. In theory, a very large surge on the AC mains could break down the isolation barrier in the AC power supply and then break down the isolation barrier in the Ethernet port. Another theory involves capacitive coupling through the AC mains power supply onto the Ethernet port. A third theory involves unintended side effects of user-installed external surge protectors that combine both AC mains protection and Ethernet protection in the same protection device.

These theories are more fully discussed in Reference [2]. The evidence for some of these theories is intriguing but not conclusive. For the present discussion, it is sufficient to say that surges from the AC mains continue to be a key suspect for many Ethernet failures.

Surges on Multi-Conductor Cables

Common Mode Surges

All of the surge coupling mechanisms described above typically induce what is called a common mode surge (also referred to as a longitudinal surge). In a common mode surge, all of the conductors in the cable develop the same instantaneous voltage with respect to earth ground. There is no voltage difference between any two conductors in the cable. The majority of surges that affect communication cables are common mode surges.

Figure 5 shows a test arrangement that generates a common mode surge. Note that all the cable conductors are at the same potential with

respect to ground, so there is no reason for current to flow from one conductor to another. Rather, the only available path to ground is through the equipment.

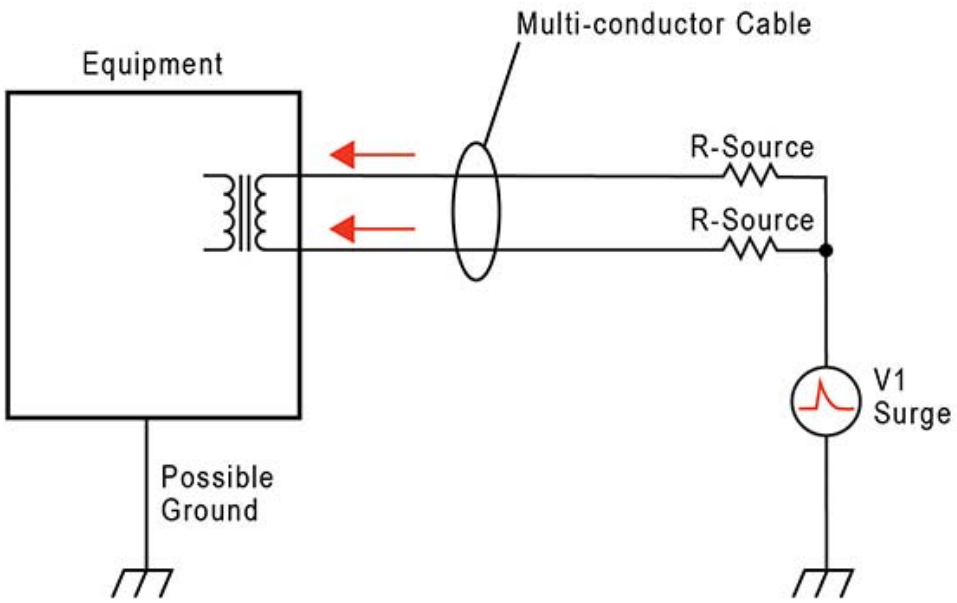

Figure 5: Common Mode Surge

If the equipment has a connection to ground, surge current will attempt to return to ground through that path. For example, if the equipment has an explicit chassis ground connection, surge current will usually return to ground through the chassis. Even if the equipment has no chassis ground, any other cable connected to the equipment presents another possible path to ground. So, sometimes the surge current enters one port on the equipment and exits via another port.

When considering common mode surges, it is important to understand that if there is no path for surge current to exit the equipment, no surge current can flow. Similarly, even if there is a path but the path contains an isolation barrier that is stronger than the applied surge voltage, no current can flow. As we will see, this particular strategy works very well for Ethernet ports.

Differential Mode Surges

Under certain circumstances, a surge voltage can appear between two individual conductors in a multi-conductor cable. This is referred to as a differential surge (also known as a metallic surge). Figure 6 shows a test

arrangement that generates a differential surge. Note that in this case, surge current will attempt to enter the equipment on one of the cable conductors, and exit the equipment on another conductor in the same cable. This is very different from the common mode surge described above.

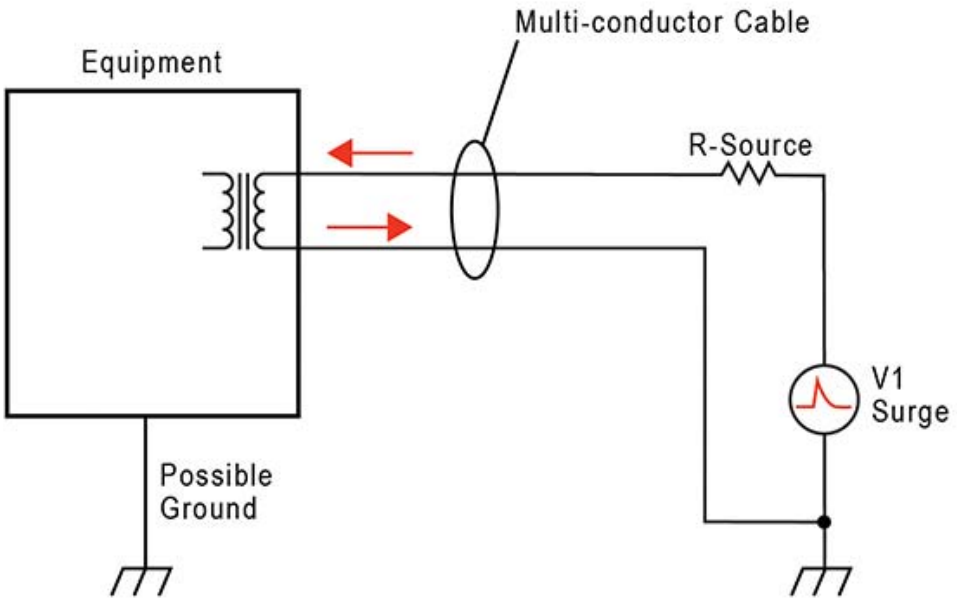

Figure 6: Differential Surge

Common-Mode-to-Differential Conversion

Interestingly, the physics of lightning induction onto multi-conductor cables almost always create a common mode surge. The appearance of differential surges between individual conductors is usually due to some characteristic of the cable terminations.

The most frequent cause of differential surges is due to the installation of a surge protection device on the cable. If protection devices have been installed between the cable conductors and ground, either externally or within the equipment itself, the protection devices create the opportunity for what is called a common-mode-to-differential-conversion.

Figure 7 shows a situation where individual gas-discharge tubes (GDTs) have been connected from each conductor to ground. The idea is that when the common mode surge is applied, gas tubes GDT-1 and GDT-2 will simultaneously trigger and conduct the surge current to ground, keeping the surge current from entering the equipment.

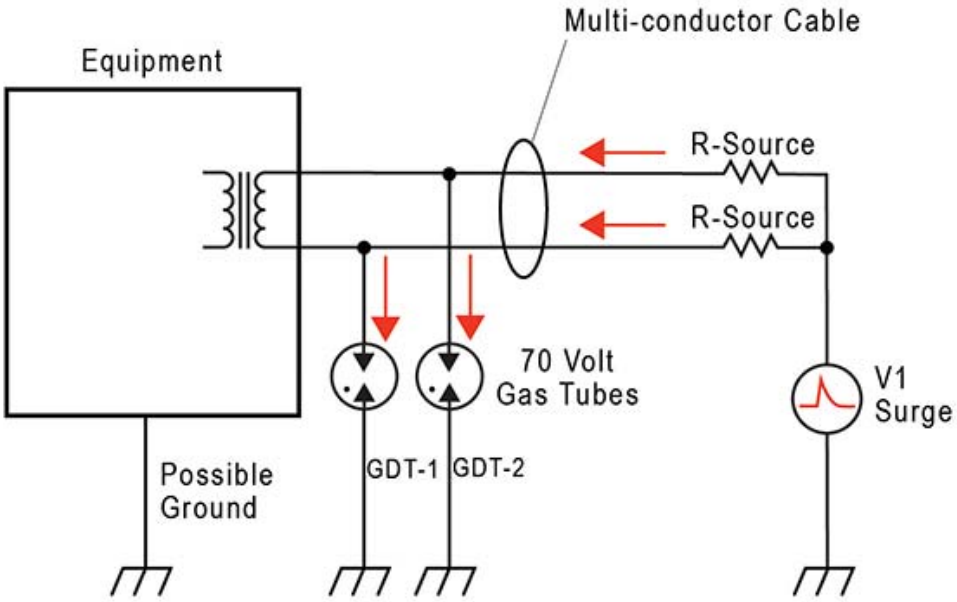

Figure 7: Common-Mode-to-Differential Conversion

Unfortunately, the individual GDTs never trigger at exactly the same time. Whichever GDT triggers first will effectively ground that particular conductor. The result is that a differential surge voltage with a very fast rise time is created between that particular conductor and all of the other conductors.

External surge protectors are the most common mechanism for creating a common-mode-to-differential conversion. The resulting differential surge can cause current to flow into the port on one conductor in the cable, and back out of the port on another conductor in the same cable. For example, if the differential surge occurs across one of the pairs in an Ethernet cable, current will flow though one of the primary windings of the Ethernet transformer. Through transformer coupling, this will induce a surge on the secondary side of the transformer, potentially damaging the Ethernet transceiver chip (often referred to as the PHY).

The irony of this is that if the protection devices had not been installed, there would have been no mechanism to cause a differential surge to appear on the port. For Ethernet ports, this is an important concept to understand. Adding protection devices between the cable and ground can actually create more problems than it solves.

Representative Test Surge

In order to discuss design strategies for protecting Ethernet ports from surge damage, it will be necessary to define some standard test surges to use for the circuit analysis. Laboratory test surges are typically described by four parameters:

1. Rise time
2. Decay time to half-crest
3. Peak open-circuit voltage
4. Peak short-circuit current

The industry has several different standard test surges that are called out for various tests. For some surge generators, the rise time and decay time can be different depending on whether we are referring to the open-circuit voltage waveform or the short-circuit current waveform. In addition, the short-circuit current waveform can be further affected by internal resistance that sets the peak short-circuit current.

To keep the present discussion simple, we will use just two representative test surges. The first will be a common mode surge, and the second will be a differential surge. Both test surges will use a 2/10 us waveform, meaning that they have a rise time of 2 us (to 90% of peak value) and a decay time of 10 us (to 50% of peak value). Our test waveform will be the same shape for open-circuit voltage and short-circuit current.

Figure 8 shows the basic shape of the 2/10 us waveform. This particular waveform has a fast rise time and short decay time, which is representative of the type of surges expected on short cables such as Ethernet cables.

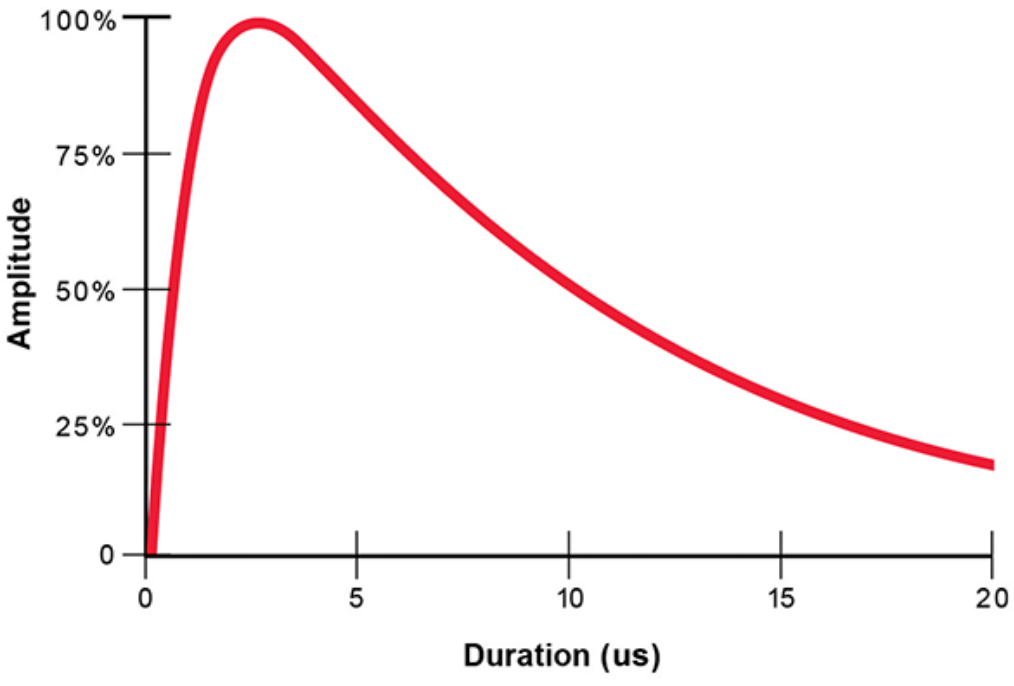

Figure 8: Surge Testing Waveform

For the purposes of the present analysis, our common mode test surge will have a peak open-circuit voltage of 6 kV, and a peak short-circuit current of 100 amps. The 6 kV value is based on analysis of actual field failures, and represents an estimate of the maximum surge voltages encountered in the field. The 6 kV peak value appears in some the emerging standards mentioned above.

Our differential surge will have a peak open-circuit voltage of 1 kV, and a short-circuit current of 100 amps. The reason for using a lower voltage for the differential surge is that we assume that surge protection has been connected between ground and each of the cable conductors, and that the individual surge protectors break down at a maximum of 1 kV.

In summary, our two laboratory test surges are:

1. Common mode surge: 2/10 us, 6 kV open-circuit voltage, 100 amp short-circuit current
2. Differential surge: 2/10 us, 1 kV open-circuit voltage, 100 amp short-circuit current

These are fairly severe surges, but there is considerable evidence that surges near these values actually do occur in the field. In addition, Ethernet ports that can survive these test surges have been shown to hold up very well in actual field conditions. As we will see, it is not difficult or expensive to design an Ethernet port that will survive these test surges.

Protection Strategies

When discussing surge damage to equipment, it is important to understand that voltage alone is not what causes the damage. Only current flow can cause damage. If the equipment can successfully stand off the induced surge voltage without allowing current to flow, no damage will occur. On the other hand, if the equipment cannot prevent current from flowing, steps must be taken to direct the current to a known, safe path. In summary, there are only two basic strategies to protect equipment from surge damage:

1. Block the surge current
2. Direct the surge current to a known, safe path

It is generally preferable to block the surge current, but this is not always possible. If the protection scheme relies on directing the current to a known, safe path, careful attention must be given to every portion of the current-carrying path to ensure that it can handle the expected current.

The Transformer is the First Line of Defense

A fundamental circuit element of all Ethernet interfaces that use twisted-pair cable is a coupling transformer. The IEEE standard 802.3 for Ethernet specifies a 1500 VRMS isolation barrier between the cable and the Ethernet PHY chip, and most commercially available Ethernet transformers are rated at a minimum of 1500 VRMS primary-to-secondary isolation to comply with the 802.3 standard.

If this 1500 VRMS isolation barrier is properly implemented in the circuit board, the entire port will be inherently immune to common mode surges up to the peak value of a 1500 VRMS waveform, which is about 2.1 kV peak. It turns out that most Ethernet transformers, while only rated for 2.1 kV peak, will actually stand off surge voltages in the range of 4 kV to 8 kV peak without breakdown of their internal isolation barrier. This level of surge tolerance is a very useful feature that can be exploited to improve the common mode surge immunity of an Ethernet port. Specially designed high-isolation Ethernet transformers are available with guaranteed surge tolerance that exceeds 8 kV peak.

Other Design Elements that Affect Common Mode Surge Immunity

Unfortunately, many designers inadvertently compromise the transformer's isolation by placing weaker isolation barriers in parallel with it. Some common mistakes are:

1. Insufficient spacings in the circuit board layout
2. Installing EMC filter capacitors that have inadequate surge tolerance
3. Placing overvoltage protection devices from the cable conductors to circuit ground

Figure 9 shows a representative circuit for just one pair in an Ethernet cable. Resistor R1 and capacitor C1 comprise the so-called Smith termination that is commonly used to reduce high frequency common mode noise on the cable. Typical values are R1 = 75 ohms and C1 = 1000 pF.

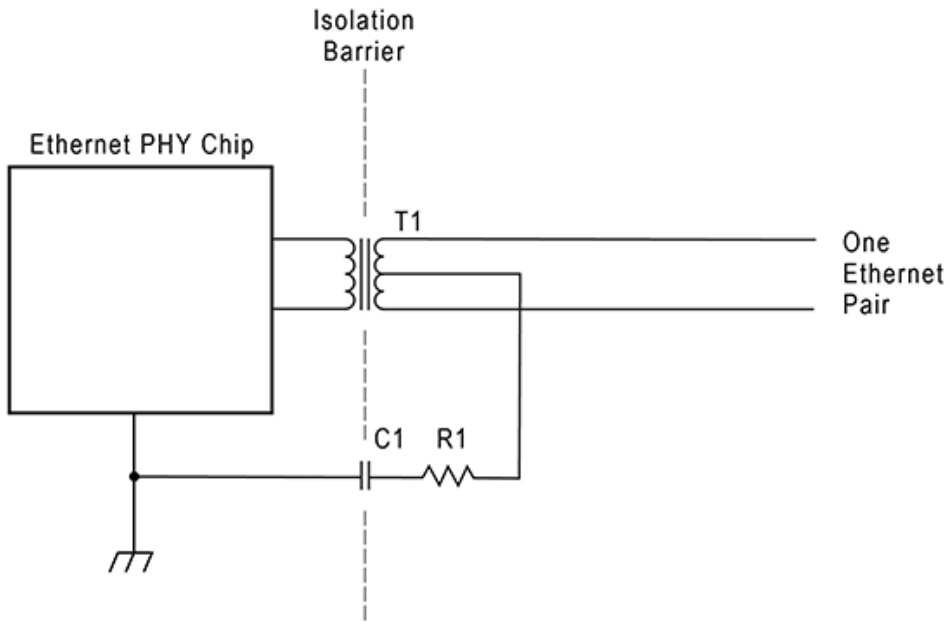

Figure 9: Ethernet Interface for One Pair

Note that since C1 is placed across the isolation barrier, it is effectively in parallel with the transformer barrier and can easily become the weak point in the overall barrier. Examination of a variety of surge-damaged Ethernet ports has revealed that breakdown failures are far more common in the Smith caps than in the transformer.

For example, the Smith capacitor in the lower left corner of the photo in Figure 1 was the failure point in that case. While it is difficult to see in the photo, there are signs of arcing across the capacitor. Once the breakdown occurred, enough energy got through the barrier to literally blow the top off of the Ethernet PHY.

It should be noted that the purpose of the Smith termination is related to EMC compliance and is not required by the 802.3 standard for Ethernet. Sometimes, careful EMC testing will reveal that the Smith termination is not required for a particular design, or that smaller capacitor values can be used.

Guidelines for Achieving Immunity to 6 kV Common Mode Surges

The best strategy for immunity to common mode surges is to have a strong isolation barrier that simply blocks the surge current. It is important that all elements of the isolation barrier be considered together as a system. Following are some general guidelines:

1. Do not place overvoltage protection components across the isolation barrier, or from the cable conductors to ground. Let the transformer itself serve as the first line of defense.
2. Maintain air gap spacings of at least 4 mm in the board layout of the isolation barrier. The goal is to keep the barrier's air breakdown voltage higher than 6 kV, so that the strength of the barrier is controlled by the transformer and Smith caps, not by the air gaps. If the Ethernet jack includes shielding or built-in indicator LEDs, careful attention must be given to maintaining the 4 mm air gap within the jack itself.
3. Use Smith caps that can handle a 6 kV surge. For multilayer ceramic capacitors, sometimes the limiting factor is the construction of the dielectric, and sometimes it is simply the physical spacing of the electrodes at each end. Note that by placing two equal-value capacitors in series, it is possible to double the effective breakdown voltage, although the effective capacitance is cut in half. Often this approach can be less expensive than using a single capacitor that is physically large and rated at a very high voltage.
4. Lastly, make sure that the transformer itself will withstand a 6 kV surge across its isolation barrier.

As noted above, most off-the-shelf Ethernet transformers will actually stand off surge voltages in the range of 4 kV to 8 kV, but there is considerable variation among different vendors, and there is also considerable unit-to-unit variation among parts from the same vendor.

If your goal is to have each and every unit of production withstand 6 kV common mode surges, you should purchase a transformer that is explicitly specified by the manufacturer to withstand 6 kV surges. However, if your product is extremely cost sensitive and all you want to do is get the best possible surge tolerance at minimal cost, focus your initial efforts on items 1) to 3) listed above. Failures due to deficiencies in these areas are far more common than transformer failures.

Guidelines for Achieving Immunity to 100 Amp Differential Surges

Differential surges create a very different kind of design challenge than common mode surges. For differential surges, the primary winding of the transformer presents a very low-resistance path, typically about one to two ohms. This comes very close to presenting a short-circuit load to the surge generator, so the parameter of interest will be the short-circuit current of the surge, rather than the open-circuit voltage.

Remarkably, the windings of most Ethernet transformers actually handle a 100 amp, 2/10 us surge without fusing open. While the current is very high, the duration of the surge is very short. In general, failures of the wire in the primary winding are uncommon.

Rather, the danger from a differential surge is the energy that couples through the transformer to the secondary winding. The secondary winding is usually connected directly to pins on the Ethernet PHY. Most PHY chips have only modest protection on these pins, typically targeted at static discharge rather than lightning surges. As a result, the PHY chip can be damaged by a differential lightning surge that couples through the transformer.

As with common mode surges, the first line of defense for differential surges is the transformer itself. It turns out that as the surge current in the primary winding rises, the transformer rapidly saturates and stops coupling current from the primary winding to the secondary winding. When the transformer core saturates, it becomes similar to an air-core transformer with a very poor coupling coefficient.

The result is that only a fraction of the surge current in the primary winding is coupled to the secondary winding. Reference [3] provides a detailed analysis of the coupling mechanism. The specific shape and magnitude of the let-through pulse on the secondary winding are affected by the choice of transformer, the electrical characteristics of the PHY side termination, and the choice of primary surge test waveform.

However, the general characteristic of getting a dramatic attenuation of the surge energy is very consistent. For our test waveform of 100 amps, 2/10 us, the peak current coupled to the secondary winding is typically less than 25 amps, and the duration of this pulse is typically just two or three microseconds. This represents a significant attenuation of the applied surge.

Figure 10 shows the relationship between the primary current and the induced secondary current. Note the large reduction in the surge due to the poor coupling through the transformer. In general, it is better to add surge protection components on the secondary side to handle the secondary surge, rather than trying to deal with the much larger surge on the primary side.

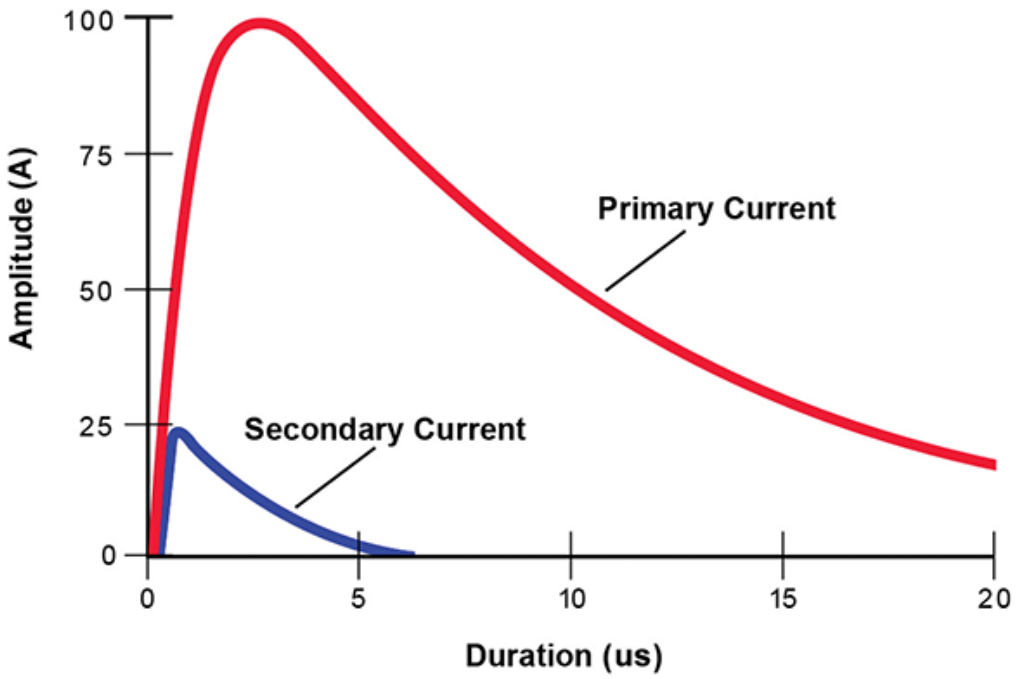

Figure 10: Secondary Current for a 100 A, 2/10 us Primary Surge

The actual reduction in the destructive power of the surge is even greater than the visual comparison in Figure 10 would suggest. This is because the energy content of a surge is related to the square of the current. To evaluate the amount of heat-causing energy in a complex surge waveform, engineers use a term called the action integral, or I^2t. In Figure 10, the I^2t energy of the blue secondary waveform is less than 2% of the I^2t of the red primary waveform. So, we get a 98% reduction in the surge energy at no additional cost, just by taking advantage of the saturation characteristics of the transformer.

Let's assume that the let-through surge in our case is 25 amps peak, with a duration of 3 us. This is still more than most Ethernet PHY chips can survive. To protect the PHY chip, many designers will place low-capacitance surge protection diodes on the secondary winding. This is a helpful step, but the external diodes are effectively in parallel with the weaker internal static discharge protection diodes in the PHY chip. So, we end up with two protection diodes in parallel, with the current being shared between the two paths.

To help direct most of the surge current through the more robust external diode, it is very helpful to place some series resistance between the external diodes and the PHY's internal diodes. As surge current tries to pass through the series resistance, a voltage drop develops across the resistance, which tends to force more of the surge current through the external diodes.

This series resistance will have some effect on the AC impedance presented by the Ethernet port, and on the shape of the pulses in the waveform. For this reason, the resistance should generally be limited to 5 ohms or less. If series resistors are used, some performance testing should be performed to ensure that the effect of the resistors is acceptable.

An alternative to using a series resistor is to use an active current limiter such as the Bourns transient current suppressor (TCS™). In its normal state, the TCS presents only a small series resistance of about two ohms. However, if the current through the TCS tries to exceed a specified threshold of say, 250 mA, the series resistance of the TCS increases and attempts to limit the current to 250 mA. The switching process is very fast and helps to force most of the surge current through the external diode.

The final choice of resistor versus active current limiter for the series impedance will depend on the specific transformer, external diode, and Ethernet PHY chip. Most PHY chips can be adequately protected by using simple resistors for the series impedance, but if the PHY chip is particularly sensitive to differential surges, an active current limiter provides a much higher degree of current limiting.

Figure 11 shows a basic circuit topology that will protect a typical Ethernet pair from both common mode and differential surges.

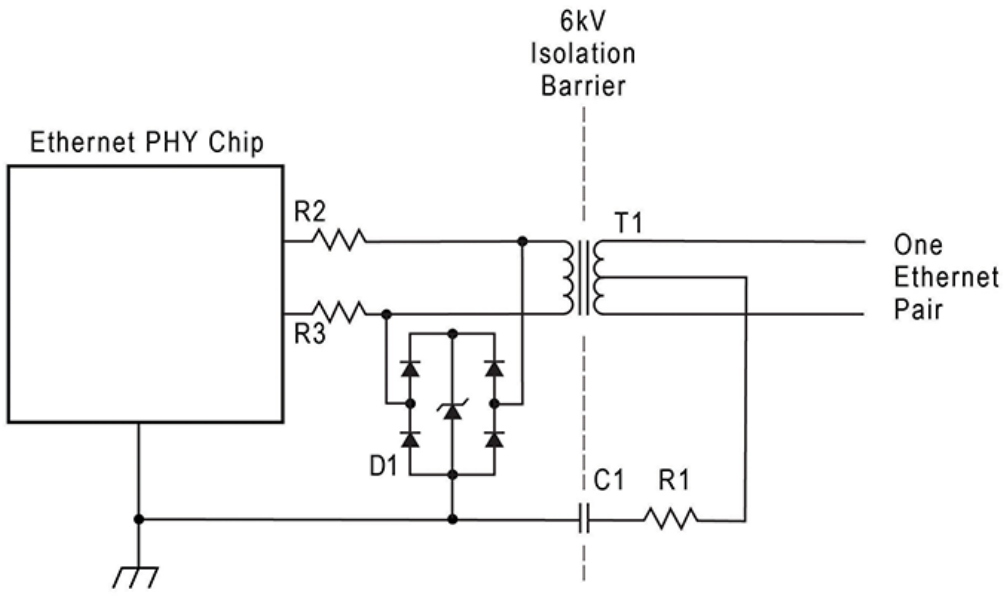

Figure 11: Ethernet Surge Protection for One Pair

Additional Considerations for Power Over Ethernet (PoE)

The IEEE 802.3 standard defines a method for delivering DC power over an Ethernet cable. Current versions of PoE use two of the four pairs in the cable to deliver DC current to the far end. One pair serves as the source, while the other pair serves as the return. The standard defines two choices for delivering the DC power. Mode A uses pairs 1 and 2, while Mode B uses pairs 3 and 4.

The DC voltage is typically in the range of 50 volts, with the current limited to approximately 270 mA. This provides the ability to deliver up to 13 watts of power to a device connected to the far end of the Ethernet cable. For remote devices such as telephones and security cameras, the main benefit of PoE is that it eliminates the need for a separate power supply to power the remote device.

Figure 12 shows a simplified diagram of DC power being delivered over two pairs that are simultaneously carrying standard Ethernet data. The engineering trick that makes this possible is that IEEE 802.3 simplexes the DC current onto the center taps of the two pairs, so that there is no net DC current flux to saturate the transformer cores.

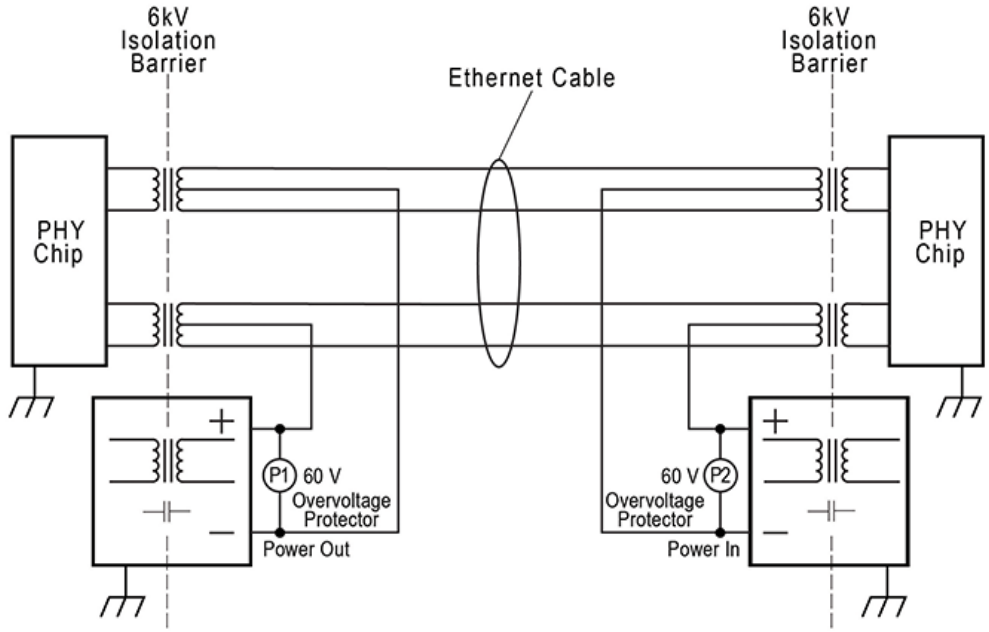

Figure 12: PoE Protection (PHY Protection Not Shown)

References [4], [5], and [6] provide additional detail on reported PoE lightning failures and industry efforts to deal with the problem.

Common Mode Surge Protection for PoE

Recall that the 802.3 standard requires a 1500 VRMS isolation barrier between the Ethernet cable and ground. This same requirement still applies to PoE. This means that the power source and the load must each contain an isolation barrier of their own. This additional barrier is placed electrically in parallel with the isolation barrier of the associated Ethernet data transformer in the equipment.

Since the previously recommended strategy for common mode surges was to rely on the isolation barrier for protection, this same strategy is easily extended to PoE. Typically, the PoE power supply circuit and the PoE load circuit are transformer-coupled DC/DC converters. So, the transformer in each of these DC/DC converters becomes part of the overall isolation barrier of the associated Ethernet port.

Most PoE DC/DC converters have additional components that bridge the isolation barrier, such as EMC filter caps and feedback opto-isolators.

When establishing the common mode strength of the isolation barrier, all components that bridge the barrier must be capable of standing off the desired voltage. In addition, care must be taken to ensure that any capacitance placed across the PoE barriers supply barriers does not allow damaging surge current to simply couple through the capacitors.

In summary, the basic approach for common mode surge protection of PoE circuits is the same approach used for non-PoE Ethernet. Namely, the isolation barrier required by the 802.3 standard is simply strengthened to stand off the peak voltage of the common mode surge. The only thing that changes with PoE is that there are more components that bridge the isolation barrier, and more circuits where the board layout spacings require careful attention.

Differential Mode Surge Protection for PoE

For each individual pair in a PoE interface, the differential surge considerations are identical to the previously discussed protection strategy for non-PoE ports.

However, the DC power circuits in PoE introduce a new surge vulnerability that is not present with non-PoE Ethernet. If a pair-to-pair differential surge occurs across the specific pairs used to provide DC power in a PoE port, the associated DC/DC converters will be subjected to a surge. For the power-sourcing converter, the surge will appear across the nominal 50 VDC output of the converter. For the load converter, the surge will appear across the 50 VDC input to the converter. Both circuits are vulnerable to damage from high-current differential surges.

When considering how to protect these circuits, it is tempting to simply install a metal oxide varistor (MOV) or a transient voltage suppressor (TVS) diode at locations P1 and P2 in Figure 12. In theory, an MOV or TVS diode that turns on at 60 volts could be connected across the 50 VDC supply.

Unfortunately, at currents above turn-on, the voltage across MOVs and TVS diodes increases with current. So, a device that measures only 60 volts when conducting 1 mA may rise to 100 volts at 5 amps and 200 volts at 20 amps. There are two problems with this situation:

1. For high-current surges, the peak voltage across the MOV or TVS diode can greatly exceed the nominal turn-on voltage. So, the degree of overvoltage protection that is provided may be less than it first appears.
2. The instantaneous power dissipation in the MOV or TVS diode is equal to voltage times current. So, a device that has 100 volts across it when conducting 10 amps is momentarily dissipating (100 x 10) = 1000 watts. All MOVs and TVS diodes have ratings for maximum pulse power.

In some cases, both of the above problems can be adequately addressed with careful attention to the device data sheets to ensure that the behavior under surge conditions is within expectations. In general, this leads to using parts that are physically large, and still requires that the peak voltage under surge conditions be allowed to exceed 100 volts.

An alternate type of overvoltage protection component is an SCR-type device such as the Littelfuse SIDACtor®. These devices exhibit what is sometimes called a crowbar behavior when triggered. After turn-on, the voltage across the device drops to about one volt and remains there until the current drops below a specified threshold, such as one amp.

The advantages of a crowbar device are that the surge voltage is strictly limited to the turn-on voltage, and the instantaneous power dissipation in the device is fairly low. A disadvantage is that the device remains triggered until the current drops below a specified threshold.

If a crowbar device is connected across the PoE DC supply, it will effectively short-circuit the output of the DC supply when the crowbar device triggers. If the PoE supply is capable of delivering sufficient current, the crowbar device will remain latched in the on-state and will keep the power supply output shorted.

Latch-up can be prevented by keeping the short-circuit current of the PoE supply below the holding current of the crowbar device. Alternatively, the PoE supply can be designed to shut down momentarily when it detects a shorted output. This momentary break in the output current gives the crowbar device a chance to reset to the off-state.

Summary

In recent years, many manufacturers of equipment using twisted-pair Ethernet ports have noted an increase in lightning surge damage to the Ethernet ports. The reason for this apparent increase may be simply due to the larger number of Ethernet ports that are being deployed. It is also possible that the types of different products now being connected via Ethernet has somehow increased the surge vulnerability of some equipment installations. There is some evidence that surges on the AC mains are being coupled onto Ethernet ports.

The known mechanisms by which lightning surges can couple onto Ethernet cables have been described. While further work is needed to clarify which of these mechanisms are the dominant causes of the Ethernet port surge damage, there is evidence that low field failure rates can be achieved by designing Ethernet ports to withstand a 6 kV, 2/10 us, 100 amp common mode surge and a 1 kV, 2/10 us, 100 amp differential surge.

Some simple design strategies have been provided for designing Ethernet ports that can withstand the above surges. A key aspect of these strategies is to rely on the Ethernet transformer itself as the first line of defense for both common mode and differential surges.

References

1. A. Martin, “Lightning Damage to Equipment Without a Metallic Connection to an External Communications Service,” InCompliance Magazine, September 2011.
2. J. Randolph, “Lightning Surge Damage to Ethernet and POTS Ports Connected to Inside Wiring,” 2014 IEEE Symposium on Product Compliance Engineering.
3. M. J. Maytum, “Differential Surge Stress Reduction by Ethernet Magnetics,” Essays in Information and Communications Technology (ICT), Surge Protection.
4. J. Wiese, “Power over Ethernet, Lightning Field Failure Analysis,” 2015 ATIS-PEG conference, Alliance for Telecom Industry Solutions.
5. T. Ardley, “Protecting PoE PSE Against Lightning and Power Fault in Internal and OSP Type Environments,” 2015 ATIS-PEG conference, Alliance for Telecom Industry Solutions.
6. T. Ardley, “Protecting PoE PSE and Ethernet to the Latest International OSP Standards,” 2016 ATIS-PEG conference, Alliance for Telecom Industry Solutions.