**ANALOG
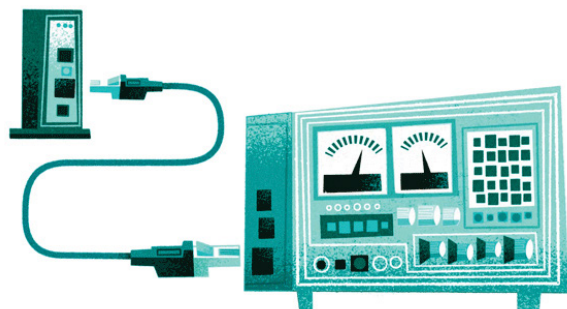DEVICES**
AHEAD OF WHAT'S POSSIBLE™

*Analog*Dialogue

RAQ Issue 217: Best Methods for Protecting Against Surge Events in Ethernet and Industrial Applications

James Niemann, Field Applications Engineer

Question

Is there an easy way to protect Ethernet from lightning damage?

Answer

With a deep understanding of magnetics and circuit theory along with good grounding and shielding techniques, there is a solution.

Abstract

Lightning-induced damage to Ethernet-connected devices can be prevented if the proper precautions are taken. Traditional treatments using protection components may not be completely effective except when augmented with a second deterrent inspired by a deep understanding of the underlying mechanisms of the transfer of energy from the lightning to the Ethernet cable and connected devices, all of which are described in this article.

Introduction

Failures due to surge events, with lightning taking the primary responsibility, are always on the minds of network administrators or anyone with responsibilities related to wired Ethernet. This topic does not limit itself to Ethernet either. It applies to any electronic or power system that is physically large. Other examples include electrical measurements brought back from a distance, power delivery, and industrial automation applications where the sensor is nowhere nearby. Traditional solutions operate on the principle of absorbing or clamping the event energy around a region to protect the physical layer components. The issue with this method is that the energy is not eliminated, and neither is the current flow that results from it. This transient current flow in an inductive path will always generate large voltages that can do damage. Therefore, when considering conventional approaches, we must ask what levels of protection are required, and how much time, effort, and resources are needed for implementation? Additionally, the protection method deployed must not only withstand but also survive the surge. A lightning strike can wield an energy force that is orders of magnitudes greater than you might think. Achieving reliable and safe operation requires robust protective measures that can meet this challenge.

Magnitudes of Energy

Ethernet systems need to protect against surge events that can differ by significant orders of magnitude. A lightning strike that is thousands of meters away might cause an energy surge that is five orders of magnitude less potent than one just outside your door. The size of an Ethernet system can also affect how well it handles differing magnitudes of energy. Even the orientation of a loop might add an additional three orders of magnitude to the surge resistance of a system.

Energy from Lightning

Surge events cause damage based on the energy of the event, where the energy surge happens, and how much energy a system can store when it is struck. Understanding these factors will help find a solution to protecting against this damage.

The energy from that lightning strike is stored in the area surrounding the lightning strike (for the discussion, let's exclude the possibility of a direct lightning strike). The primary concern with lighting strikes is that the energy is stored in the near field with the magnetic field the most significant for this low impedance source. The total energy in the magnetic field can be approximated by using the length of the lightning strike to find the total inductance and the familiar energy equation $E = 1/2Li^2$. The strike current does vary but can be as high as 50,000 A. Beyond this distance, in the far field, there will be little energy to be concerned with unless you are in the business of building radio receivers.

Our sun generates 3.846×10^{26} W of power per second. 93 million miles away on Earth, one square meter of space will see 1000 W of that total. There will always be 3.846×10^{26} W of radiated power if we integrate over the entire sphere surrounding our sun no matter what the distance from its surface, and that 1 square meter is a very small percentage of that total surface area 93 million miles away! Now let's talk in terms of energy instead of power. To see 1000 J of energy will require 1 s of exposure (watts have the units of J/s). This is a volume of energy equal to $1M^2$ times the distance light travels in one second or 3×10^8 M and, in this case, the total volume is also 3×10^8 M^3.

In order to understand the remainder of this article, the concept that both radiated and static energy (magnetic energy BxH and electrostatic energy ExD) are stored in space must be accepted. Poynting's theorem describes the motion, movement, or transfer of energy. Any transfer of energy always involves both the magnetic and electric fields. There can be no appreciable electric field inside a conductor so there cannot be any energy stored there either. Both near and far (radiated) energy is stored in the space surrounding the lightning event, plain and simple. This concept (that energy is stored in space) suggests the following solution to the issue of surge. Eliminating access to this energy will eliminate the surge problem.

To have access to that energy, a conductor geometry (the Ethernet cable) needs access to the space that contains the movement of energy. Time is involved, just as for our radiation example, even though we are in the near field. Since the Ethernet cable differentially does not have any appreciable loop area, it is not likely to couple any appreciable energy from this surrounding space. This is not true for the area between the Ethernet cable and the grounding system.

Surge is a high frequency loop current that involves the chassis grounding system. The chassis grounding system is really a part of every circuit ever built. In this context, it is only important for large circuits. See Figure 1 for an example of how the chassis grounding system is always present but becomes more important for larger systems and why Earth ground does not specifically have anything to do with the problem, any parasitic conductor will do. The two most common sources of surge current are described in the next sections.

Ground Loop Energy

Ground loops arise due to the fact that the ground potential is not a constant at any two locations. Figure 1 showed that every schematic has a second circuit, the parasitic ground loop. This ground loop also goes by the name of common impedance coupling[1] due to the fact that both the ground loop and your designed circuit can share a common conductor. Figures 1 and 2 show a more detailed example. Often this second chassis ground circuit is not so large but it's always present. Generally, the more distance the electronic system covers, the greater the potential difference between these grounds and the more inductance and resistance they will have between them.
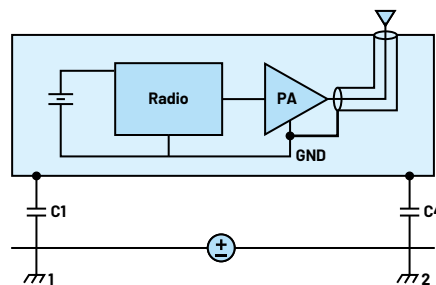

Figure 1. Even a system as small as a handheld device could technically be affected by the outside world. In this example, the ground loop is very small, and any interference current will flow in the shield rather than the radio ground.
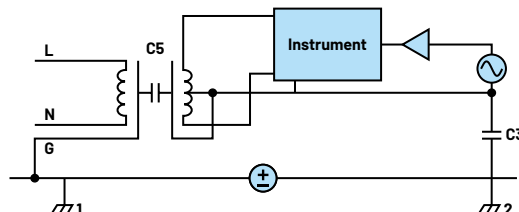

Figure 2. A line-powered instrument is shown above with a ground loop voltage between Chassis Ground 1 and Chassis Ground 2. This loop is also large enough for magnetically coupled interference to be significant. Also, note that the interference loop shares a common conductor with the instrument ground.

When lightning strikes the ground, current spreads out in all directions. This current causes a significant voltage to drop in the resistance and inductance of the ground through which the current travels. For some wired Ethernet installations, this potential difference can span the entire Ethernet cable (from end to end) and can cause large currents to flow. This effect is correctly categorized as a ground loop. Ground loops are also caused by currents originating in instrumentation and electrical machinery as well. A properly grounded building will reference a single ground conductor at the utility entrance. This leads to the conclusion that lightning-induced ground loops are not the largest contributor to equipment damage within a single building. This is obviously not the case for Ethernet runs outside or between buildings.

Whatever the source, the voltage in the ground can cause currents in the Ethernet cable even in the absence of great lengths or loop areas. All that matters is the difference in potential of the two grounds, the rise time, and the inductance of the chassis system between the two points.

Magnetic Field from Lightning

Lightning can also generate magnetically coupled voltages in any loop area by Faraday's law. This is probably the most concerning issue since it will affect Ethernet runs within buildings.

Lightning surge events due to a ground loop are different from events caused by magnetic coupling (Faraday's law). The next few sections will address possible solutions for each. For reference, Figure 3 represents an Ethernet connection without any of this article's solutions. Here, current pathways (due to ground loops or Faraday's law) through the circuitry and the ground reference plane (which is part of the circuit too) are the only path surge currents could pass. Conventional solutions attempt to shunt this current away from components, but hazardous $V = Ldi/dt$ events could develop in the current path.
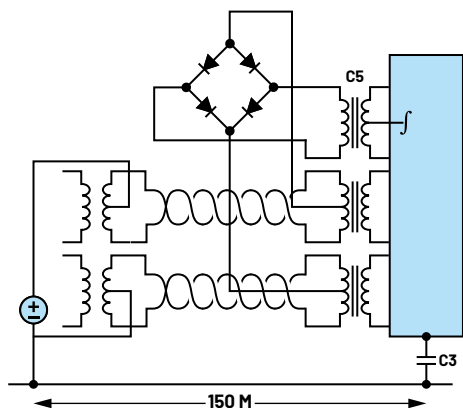

Figure 3. An Ethernet example that is susceptible to surge damage.

A Textbook Solution

A textbook solution to solving both the ground loop and magnetic energy problem requires guarding. Guarding is provided by surrounding the entire application with a shield. This guard has the effect of minimizing the capacitance from the application to anything except the guard itself. Considering Figure 4, a non-Ethernet simplified example, it should be clear that ALL ground loop or magnetically induced current will follow the guard metal and cross the isolation barrier at C5. It is not possible for the ground loop current to enter either application area encircled by the guard. In this case, the interference fields are entirely external to the application components. This textbook solution covers both sources of interference in addition to eliminating any electrostatic coupled noise. The textbook solution is truly remarkable in that it works even if C5 is minimized. A shorted turn is not strictly required.
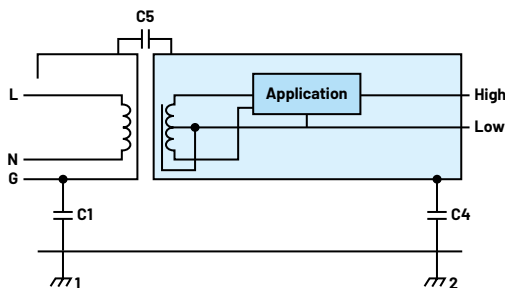

Figure 4. An instrumentation example showing the use of guards to eliminate energy resulting in less surge current in application circuits.

This is the only solution that works for both ground loops and magnetically coupled field energy. It is also usually more than required for this Ethernet application so some simplifications below will result in a solution for Ethernet that we can actually build.

The Shorted Turn

The energy causing the damage comes from the field generated by the lightning. To eliminate the energy for the Ethernet run, we need to eliminate the field, and we are going to do just that by designing a shorted turn into this transformer where the lightning bolt represents the primary and the Ethernet ground loop area is the secondary. If we can build an isolated shorted turn using the guard inside the Ethernet cable and planes in the application circuit with the ground providing the final conductor to close the shorted turn, we should be able to eliminate the energy. In practice, after implementing this shorted turn, the process of adding external shunting protection components will be much easier.

The simplification can be made by sacrificing the complete encirclement of the left and right halves of the system. This is shown in Figure 5 (the Ethernet configuration is shown in Figure 6). This simpler configuration can be effective, provided that the guard circuit can act as the shorted turn and the ratio of C3/C4 is minimized. This simplified method for eliminating surge only works if we can build a shorted turn compared to the isolated path.
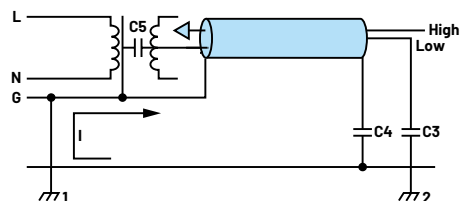

Figure 5. A simplified instrumentation example using shields to guide surge energy away from application circuits.
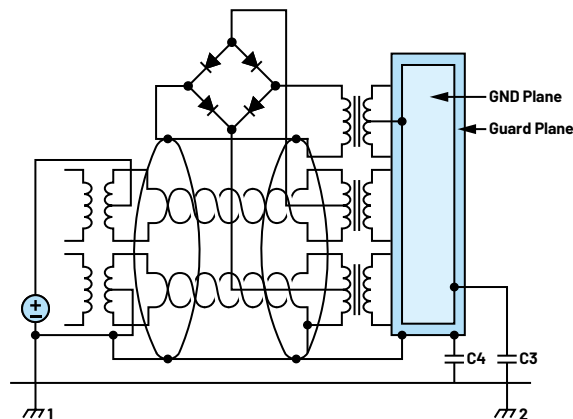

Figure 6. An Ethernet example using shields to guide surge energy away from application circuits, with C3 < C4.

Exactly how does this shorted turn eliminate the energy from the point of view of the Ethernet loop? To unravel this, we need to understand the transformer analogy at a deeper level. A real transformer is designed to move energy rather than store energy. This is true whether we are talking about an air core transformer or a transformer with a magnetic core. To achieve little stored energy in an air core transformer the windings must be wound directly on top of one another with very little space between them to store energy. Transformers built with magnetic cores

move energy (with hysteresis and eddy current losses) from winding to winding even if the windings are not directly on top of one another, but there must be little space between the windings and the core so little energy will be stored. When a core is used, you get the added advantage of lower magnetization current due to the higher inductance, which is a direct result of μr being large. Core or not a voltage applied to the primary will result in a current described by the familiar $V = Ldi/dt$ relationship, which will in turn result in a voltage developed on the secondary given by $V = (loop\ area)dB/dt$. The presence of magnetic material does not change the primary Ldi/dt or the secondary dB/dt. In other words, it does not change the transformer voltages. In the primary, the permeability μr is a constant that does increase inductance (by μr) but also lowers the di/dt to compensate. For the secondary, the large μr slows the dB/dt (since the primary di/dt is lower), but it also increases B by the same constant. The high permeability really just lowers magnetizing current by raising the primary inductance.

Since energy is not stored in a transformer, when the secondary is heavily loaded, more current will be demanded from the low impedance voltage source driven primary, and the primary current will increase to provide the energy.

In contrast, a lightning strike stores a lot of energy in a very large space. Nature will always arrange itself to a configuration that stores the least amount of energy possible. And this is exactly what our transformer is doing on the inside, and at the interface to the secondary winding where the current runs opposite to the primary current. These opposing currents guarantee that there will be no net external magnetic field (stored energy). At the highest level, this is called the principle of least action but in this context, it's referred to as Lenz's law. This is what is happening in the space around the Ethernet cable and the chassis ground return. The Ethernet loop (or our shorted turn, your choice) is providing the means to either move or dissipate this energy because either offers a way to store less energy. Just like the transformer example above, the secondary voltage developed is still $V = (loop\ area)dB/dt$, but without the tight coupling between the primary (lightning) and the secondary (your Ethernet loop). This poor coupling prevents the region from accessing an unlimited source of energy. The shorted turn will generate a current canceling/dissipating the energy that the lightning stored in this space. If the inductance of the primary could be measured with this shorted turn in place, it would be a lower value pointing to less stored energy, some of this missing energy being dissipated in the shorted turn. In other words, the magnetic field resulting from the secondary load will cancel the field generated by the lightning, leaving less energy stored in the Ethernet loop.

Incidentally, this is exactly what happens in a transformer when you short one of the secondaries. There is an important difference though. In the case of a real transformer, the shorted turn will dissipate all the energy available in the primary due to the tight coupling. For lightning, only the energy that was in the Ethernet loop space will be dissipated.

Let's look at an example. The H field generated from a lightning strike is I/2πR. Assuming the strike is 1 mile away (1600 M) from our Ethernet cable, and the strike current is 50,000 A, the magnetic field intensity will be 4.97 A/M.

The B field is then B = μH = $(4 π × 10E-7)(4.97)$ = 6.25E-6 Tesla,

The Ethernet loop area (one mile distant) is: 1 M × 150 M = 150 M^2

A lightning strike current rise time can be as fast as 1 μs, and its fall time is ~ 100 μs so the voltage developed in this loop can be approximated as: V = A (loop area × dB/dt) = 150(6.25E-6)/1 μs = 937 V

Let's look at a simulation to give us the exact value. Figure 7 shows a 50 kA lightning strike with a 1 μs rise time and a 100 μs fall time.
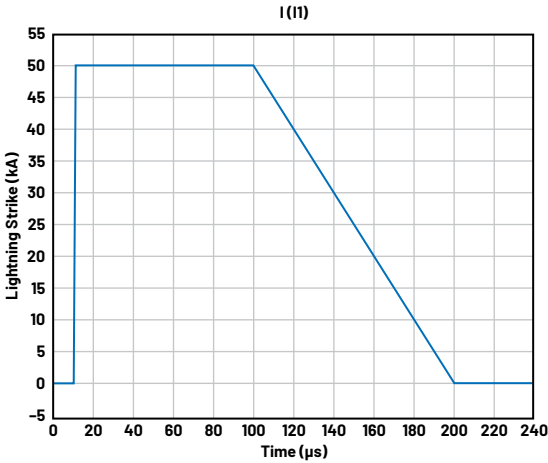

Figure 7. A 50 kA lightning strike with a 1 μs rise time and a 100 μs fall time.

Due to Faraday's law, this current will generate the voltage V1 as shown in Figure 8. E1 represents the surge voltage within an unprotected Ethernet loop. The 459 μH is the inductance of the Ethernet loop area with chassis, the
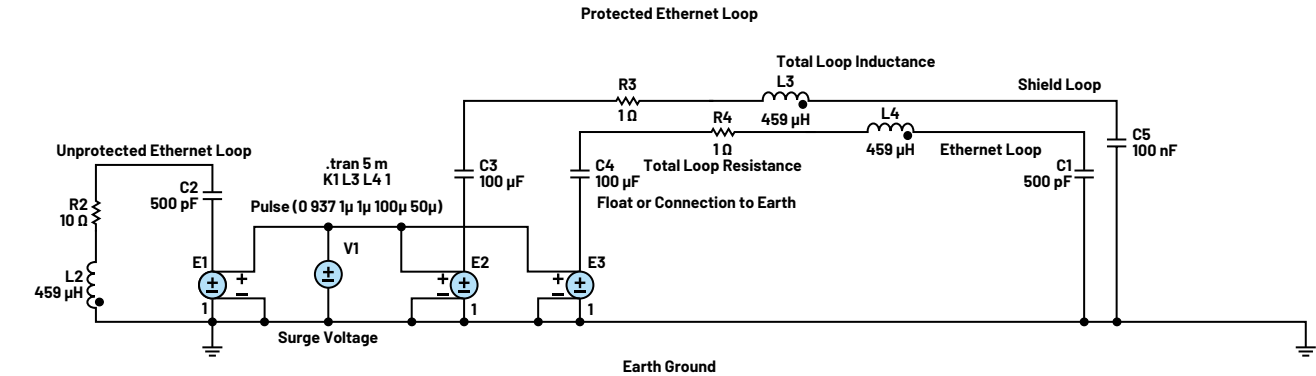

Figure 8. A SPICE simulation model illustrating how surge current can be reduced by using a second shorted turn tightly coupled to the Ethernet loop.
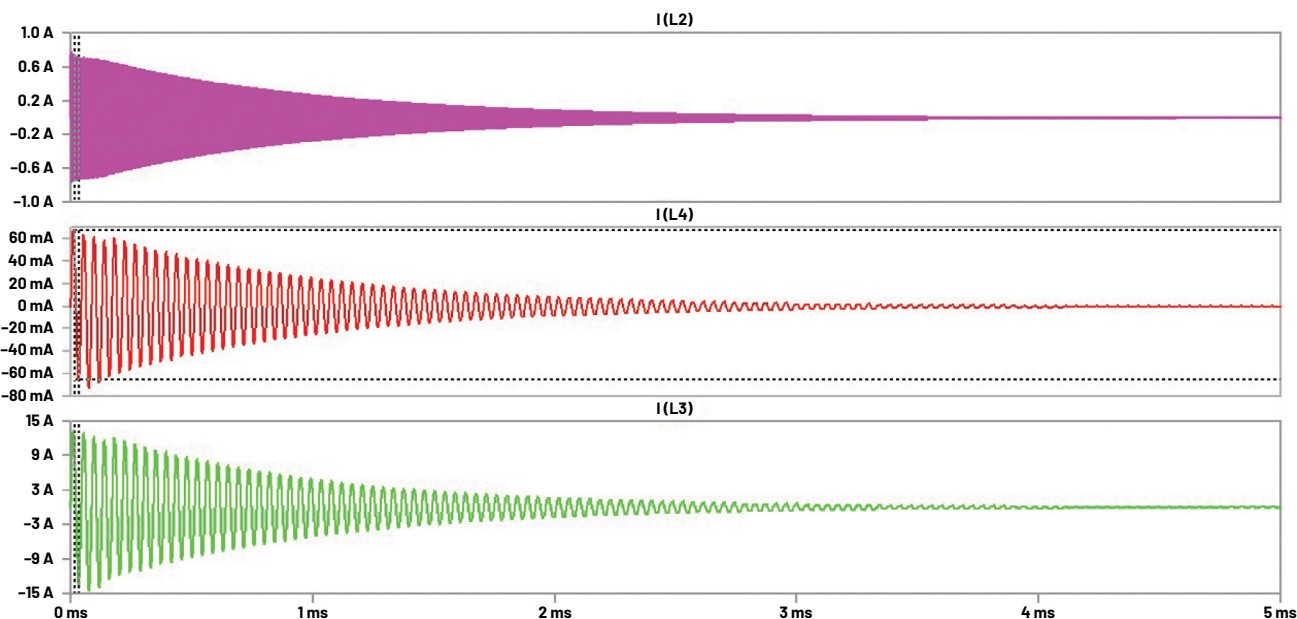
Figure 9. Surge currents for the example simulation in Figure 8.

500 pF represents the net series capacitance to ground of both the PSE and the PD sides of the Ethernet connection, and the resistor is the series resistance of the circuit. In the simulation, the value of R2 does not really change the peak values of current, but rather causes the envelope of the waveform to decay at a faster rate. This more favorable L/R time constant will dissipate the surge energy more quickly as heat throughout this distributed resistor.

The plot showing the resulting surge current, I(L2) is shown in Figure 9. The plot indicates that even if the lightning strike is 1 mile away, an unprotected loop will see 1.6 A of peak-peak surge current. Imagine how much more loop current would result if the lightning strike was much closer. Even this is enough current to cause damage.

Now let's consider the surge current in a protected Ethernet loop shown on the right half of the schematic—here, the inner Ethernet loop. This surge current can be reduced even more if the shield loop impedance (increasing C3 and C4) is lowered while maintaining good magnetic coupling to the Ethernet loop.

Isolation

There is one other way to eliminate the surge current and that is to isolate one or both ends of the cable. Ideally to isolate an application in this way requires an open circuit at all frequencies. This is typically done in the isolation transformer and, for Ethernet, this includes both the data as well as the power transformers (POE applications). Transformers do a great job at stopping DC; however, their primary to secondary capacitance shorts at higher frequencies, supporting high frequency surge currents. If vanishingly low capacitance transformers were available, we would not have a surge problem to begin with, so this is not the answer. Nevertheless, lowering the isolation capacitance will lower the currents caused by lightning strikes. However, the proposed solution provides a better isolated system at higher frequencies even though there is more capacitance across the isolation barrier. Why would the capacitance matter at all if it does not see any dv/dt?

What's the Catch?

The catch is that we will never be able to build a perfect guard around our circuits or eliminate all that magnetic field with our shorted turn or build a transformer without capacitance. In this case, what else can be done? To augment these solutions, we may need to also add protection components designed to divert any remaining surge current. Currents can be high in the shorted turn, but this is of very little concern since we have only used copper and capacitors to build it. The last possible improvement we can make would be to add a ferrite around the entire Ethernet link as shown in Figure 10.
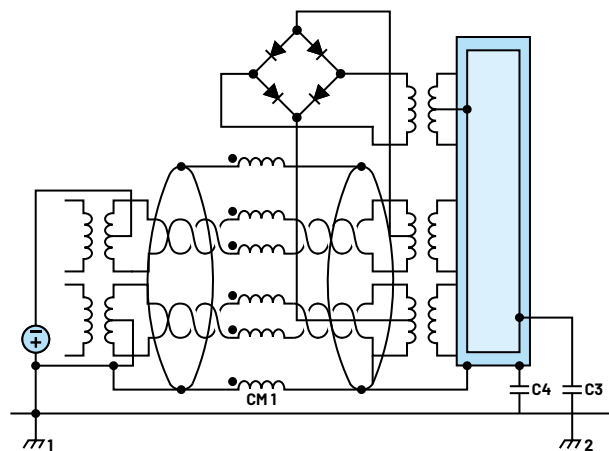

Figure 10. The common-mode choke, CH1, provides a low impedance to differential mode currents and an increasing impedance to common mode currents.

In the absence of our new shorted turn, this ferrite still does good. It provides an open circuit for high frequency currents that complement the open from the isolation transformers at DC and at lower frequencies. If we use the ferrite along with our shorted turn, we get some very amazing results. In this case, the ferrite further minimizes the C3/C4 ratio by providing an open for current around the ground loop.

Conclusion

Any application that requires long cables may be susceptible to damage caused by lightning. The source of this damage may be from the voltage drops in the ground impedance due to the high currents of the lightning strike (ground loops) in addition to the voltage developed due to Faraday's law (magnetic coupling). In some applications, the use of protection components to direct this damaging current may not solve the problem. In this case, the addition of a low impedance shorted turn placed directly along the Ethernet cable and circuitry (well coupled) can be used to lower the surge current dramatically. This method uses only copper and capacitors so we need not be concerned about the high currents that may result from this shorted turn. The addition of a common-mode choke on the Ethernet cable can also be used to safely lower the surge current.

References

[1]Alan Rich. "Shielding and Guarding How to Exclude Interference-Type Noise What to Do and Why to Do It—A Rational Approach." Analog Devices, Inc., 1983.

Niemann, Karl-Heinz. "Engineering Guideline Ethernet-APL," Version 1.14 19. September, 2022.

Feynman, Richard P., Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics, Vol. II: The New Millennium Edition: Mainly Electromagnetism and Matter.* Basic Books, 2011.

Morrison, Ralph. *Grounding and Shielding Techniques,* Fourth Edition. John Wiley & Sons Publications, 1998.

About the Author

James Niemann joined Analog Devices in March 2020 and is currently a field applications engineer in Cleveland, Ohio. James has 35 years of combined experience designing test and measurement equipment and working as an FAE at ADI. James has 14 patents.

VISIT ANALOG.COM

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™