THE THREAT

Just plugging in may be the most dangerous thing you do to your computers.

BY JAMES J. STANISLAWSKI

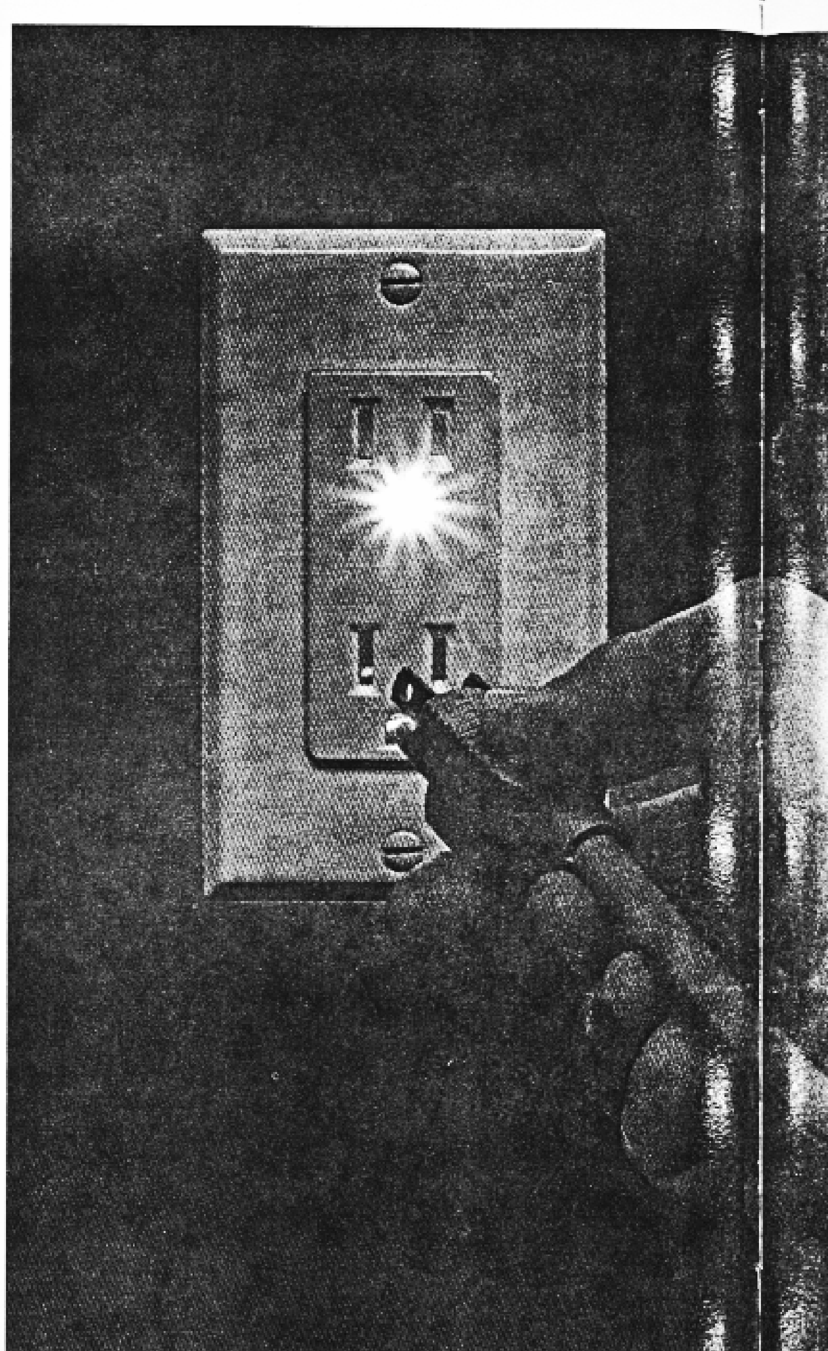

Electromagnetic events on AC power lines have been blamed for billions of dollars per year in lost revenues, but hard data on real probability of power deviation has been limited until now. A major new point-of-use study reveals that thousands of potentially dangerous power events can hit a user site in a single month. The power behind your outlets may be the gravest threat your systems face.

The four-year National Power Laboratory (NPL) study recorded 160,000 power disturbances during 1,057 site-months of monitoring at 130 randomly selected North American locations (see sidebar, "How NPL Did It"). Single-phase, line-to-neutral event data was collected at the standard wall receptacle, where it is normally used by computers and other electronic equipment.

Seven different types of problems were monitored on premises power lines: overvoltages, undervoltages, harmonic distortion, high-frequency noise, lightning, transients and outages. NPL also identified causes and protection modes to help reduce the dangers from each event.

This data was filtered to extract those power disturbances that probably would damage system processors or destroy data—an average of nearly 2,000 per site, per year. And of these, an average of 289 fell beyond the Computer and Business Equipment Manufacturers' Association (CBEMA) safety guidelines. This includes an average of 19 transients, 164 swell/overvoltage conditions, 90 sag/undervoltage conditions and 16 interruptions. The "cleanest" locations experienced no abnormal power events, while the worst locations experienced 164 transients, 6,714 swells/overvoltages, 7,121 sags/undervoltages and 146 interruptions. Some of these events lasted as long as 71 hours.

TRANSIENTS. Transients, also called spikes or impulses, are electrical peaks of more than 100 volts "root-mean-square" (RMS) that last between .5 and 2,048 microseconds. High-frequency transients can cause data errors, while high-voltage transients are more likely to destroy electronic components. NPL found that transients are caused by conditions both inside and outside the building.

With more than 1,000 hits per year at the average study site, transients constitute more than half of all the power problems detected. Surprisingly, however, only 18 recorded transients had peaks greater than 500 volts. One way of looking at this result is to say that an equipment-damaging high-voltage transient will probably hit an average site once every five years. Another interpretation is that one out of every six corporate sites will experience a damaging transient during any nine-month period. This transient data actually may be understated, since some NPL study sites may have had transient protection at the building service entrance or somewhere upstream from the monitor.

High-energy events, most commonly associated with utility capacitor switching, can damage electronic components, microchips and even networked equipment thousands of feet from the event source. One solution is to use a ferroresonant transformer-based uninterruptible power supply (UPS) or any other UPS technology that can deliver no-break power output.

Data-corrupting transients typically do not have much energy, but, due to their high frequency, can travel into microprocessors or electromagnetically couple into nearby data cables. Many data-

BEHIND THE OUTLET

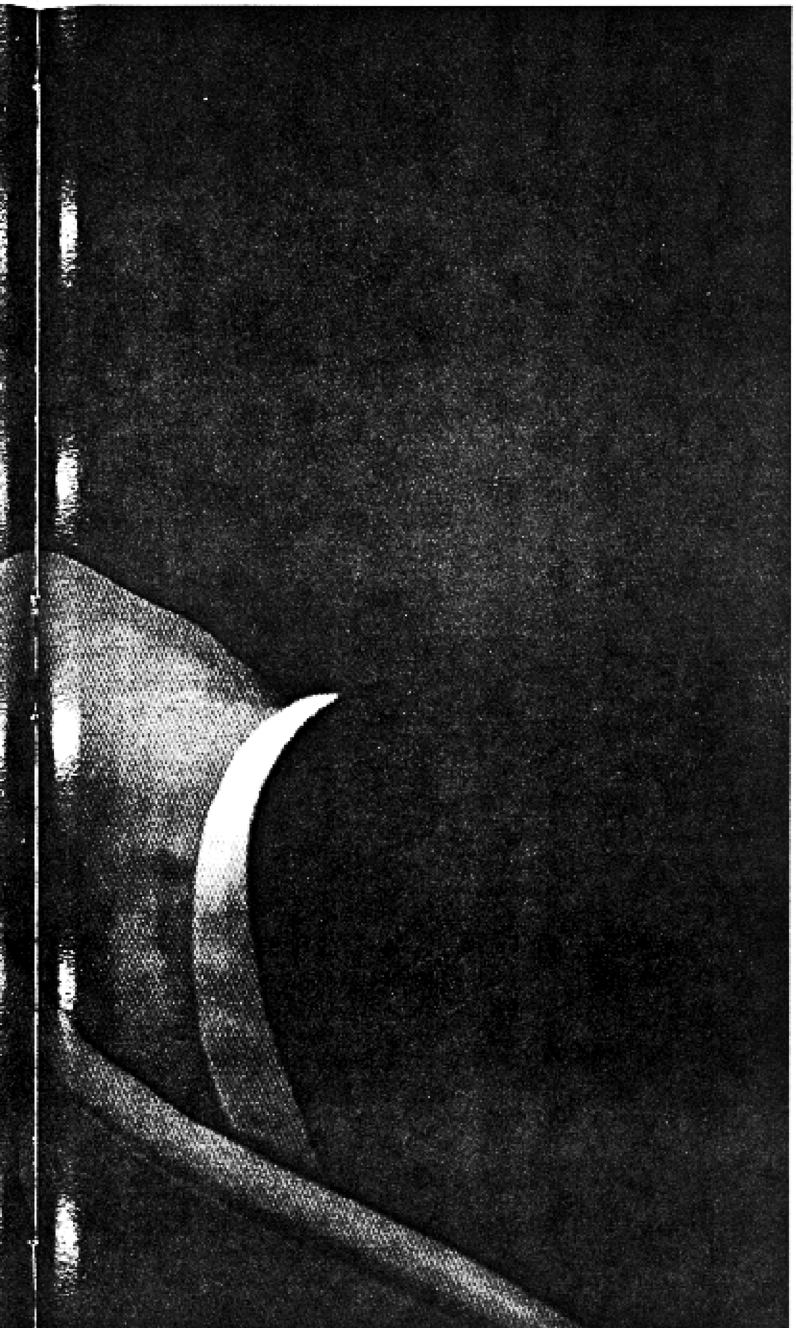

PHOTOGRAPH: RUSSELL INGRAM

corrupting transients are created when equipment inside the building switches on and off. Since transients typically last less than one ten-thousandth of a second, a ferroresonant or other UPS device with no-break output can keep these split-second events from corrupting data.

SWELLS/OVERVOLTAGES. Technically, swells and overvoltages are events of more than 127 volts lasting longer than 2.048 microseconds. They can cause component overheating or destruction, depending on duration and the percentage above nominal power values. Swells also may trigger protective devices, resulting in memory and data losses during equipment reset.

Overvoltages usually happen when large loads like air conditioners and motors are shut off, the utility voltage tap at the building entrance is set too high, or an auxiliary generator supplies high voltage during start-up. High-voltage disturbances recorded by the NPL occurred at a rate of 164 hits per year at the average site.

Recommended protection devices include power-line conditioners or UPSes with overvoltage compensation. A line conditioner with a ferroresonant transformer can regulate high voltages down to normal output levels, regardless of how high the voltage becomes. On the other hand, most UPSes with overvoltage compensation detect an overvoltage and then transfer to battery power to protect attached loads. Some UPSes, however, use ferroresonant transformers, allowing them to avoid battery usage during high-voltage conditions, conserving power reserves for true outages.

Systems may absorb 20 or 30 power deviations without failing, but repetitive, accumulated events will ultimately destroy systems and stored data.

SAGS/UNDERVOLTAGES. When less than 104 volts RMS of power is delivered over a period of more than 2.048 microseconds, a sag, or brownout, occurs. These events can cause computer resets, memory loss, data loss, electronic component overheating and process-control shutdowns. These happen because many microprocessor-based systems try to continue operating, even when voltages fall below levels where logic chips start to behave erratically.

Power sags often occur when large loads, such as elevators, air conditioners or industrial motors, turn on. They also happen when short circuits occur on utility power lines. NPL recorded 90 undervoltage events per year at the average study site.

Solutions to undervoltages and brownouts include power-line conditioners and UPSes. Stand-by power systems (SPSes) with short transfer times also may be acceptable for non-critical applications. Ferroresonant-based power-line conditioners can attain normal output voltages even when input voltages drop down to 70 percent of normal. Any SPS or UPS systems purchased should include brownout protection (a ferroresonant transformer) that boosts the output voltages without switching to the unit's battery.

HARMONIC DISTORTIONS. Normal AC current is an electromagnetic waveform with an optimum frequency. This waveform can be distorted by the normal operations of computers, peripherals and adjustable-speed disk drives, and it can travel throughout a building's power network. Repeated distortion events can overload a building's electrical-service transformer and create a fire hazard. About one out of every ten locations monitored by NPL had problems with power harmonics.

ADDITIONAL POWER-EVENT STATISTICS

For those who want more data on the major types of power events recorded by NPL, please read on:

Low-voltage sags ranged from 10 to 103.8 volts "root mean square" (RMS), and lasted from 0.010 seconds to 1.75 hours, with a mean value of 2.1 seconds. The worst locations experienced 1,659.7 such events per month, against a study-wide average of 27.9 events per month. More potentially damaging sags below 96 volts RMS occurred during the four months from May through August, which accounted for over half of all of the dangerous undervoltages recorded.

High-voltage swells at monitored sites ranged from 127.2 to 149.8 volts RMS. They lasted from 0.010 seconds to 169.9 hours, with a mean value of 44.2 minutes. The worst locations experienced 1,450.0 such events per month, against an overall average of 13.9 events per month.

Transient peaks ranged from 100 to 1,764 volts. They lasted from less than 1 microsecond to more than 2,048 microseconds, with a mean duration of approximately 63.4 microseconds. The worst locations experienced 1,165.8 transients per month, against a study-wide average of 63.5 events per month.

Power interruptions totaled from zero to 145.1 hours per year, with a mean value of 6.17 hours per site, per year. The worst sites experienced 10.2 interruptions per month, against an average of 1.3 interruptions per month.

—James J. Stanislawski

Ferroresonant-based transformers make ideal correction devices. However, some double-conversion UPSes (which do AC-DC-AC conversions) can actually intensify harmonic distortions. Be sure, therefore, that such systems feed input AC through a harmonic-correction circuit. Other protective systems claim to correct harmonics by not using the power network's neutral wire, but this will not eliminate harmonic distortions seen by other loads on the same line, and may still overload electrical-service transformers if enough computer loads are active.

HIGH-FREQUENCY NOISE. Noise problems originate at a variety of electrical sources, such as fluorescent-lighting ballasts, air ionizers, and any other equipment that emits radio-frequency or electromagnetic interference. The magnetic fields from this interference may induce wave distortions in the building's power lines, with the potential for equipment resets and data errors.

Isolation transformers can help prevent this, especially transformers with an output neutral/ground bond plus electronic filtering. Also, try to keep sensitive loads off of the lines where noise may be present, or to run power lines away from potential noise sources.

LIGHTNING. Lightning strikes were among the rarest events recorded in the NPL study, but they were usually the most damaging. Lightning also accounted for many of the short-term power outages seen in the study. These outages were usually induced when the building's utility-protection devices opened to keep a lightning surge out.

The best power-protection devices for diverting lightning surges to ground comply with two standards. Those meeting the IEEE/ANSI C84.1 specification can survive a 6,000-volt simulated lightning hit. To meet another specification—UL1449—a device must pass a grueling series of twenty-six 6,000-volt hits. This measures both protection ability and product safety. Almost all available protection devices will survive the C84.1 test, but fewer also meet UL1449 requirements.

While meeting the UL1449 rating will pretty much guarantee protection against lightning surges, it will do nothing against other power problems. That requires other protection devices, such as UPSes. Many of these include their own surge-protection board, and if it meets UL1449, it will protect the system from lightning surges.

INTERRUPTIONS/OUTAGES. An interruption occurs when 0 volts RMS are delivered for at least 4 milliseconds. Though these outages are the least frequent events recorded in the NPL study, they are likely to cause the majority of memory loss, data loss and costly system failures.

The average NPL study site experienced more than 16 power failures per year. The study also found that 95 percent of recorded interruptions lasted 5 minutes or less, and the majority of these lasted 1 second or less. The real problem for businesses is the remaining 3 percent of outages that continue for at least 30 minutes (some lasting more than 70 hours).

Therefore, a battery back-up with at least 5 minutes of runtime should cover 97 percent of all power interruptions that will be seen. However, especially critical devices, processes or networked systems may not tolerate any transfer time from line current to battery back-up, and will require the protection of a no-break UPS. This will minimize short-term outage effects like system lock-ups, unexplained resets or hard-disk crashes. To protect against outages lasting longer than 30 minutes, an AC generator or an uninterruptible battery system (a DC generator feeding a UPS inverter) are practical solutions.

Power interruptions actually accounted for less than six percent of the disturbances recorded in the NPL study. This fact emphasizes the importance of buying equipment and protection devices with wide operating-voltage tolerances, and not depending solely on stand-by power systems that only protect against outages.

No statistically significant difference was found for the average power downtimes in the study's ten geographical areas, but southern regions experienced a larger actual number of outages,

THE LETHAL 90-VOLT FLOOR

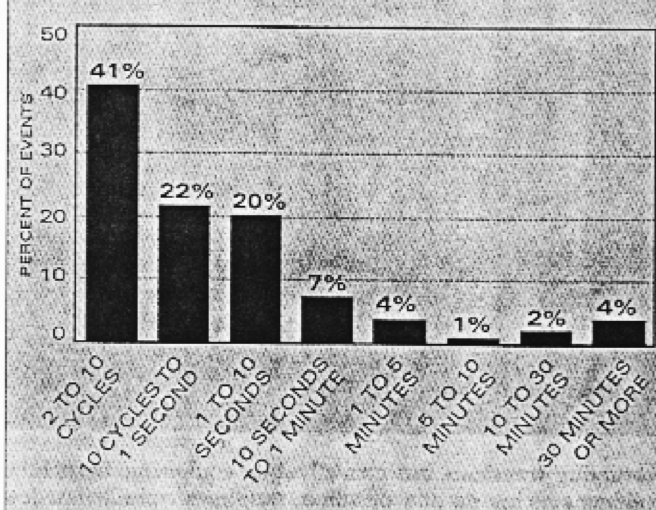

Most electronic equipment will have problems coping with power sags below 90 volts. The average monitored site experienced 38 of these dangerous events per year, charted here by duration in cycle length.

Source: NPL 1990-1994 Study

probably as a result of weather. Rural locations also experienced four times as many power interruptions as urban locations. This is probably due to the larger percentage of overhead transmission lines in rural areas and their susceptibility to weather activity.

MANAGERIAL CONSIDERATIONS. To accurately interpret all the above statistics, managers should first realize that there is no fixed baseline voltage throughout North America. The typical nominal voltage seen in the NPL study actually varied from 112 to 128 volts between locations. Therefore, a business will experience more system-endangering low-voltage events at sites where 112 volts is the nominal level, and more high-voltage events where the nominal voltage is already above average. NPL also found that sites with more-controlled power levels (between 118 and 122 volts) experienced the fewest events outside safe thresholds.

Another finding of importance to facility managers is that sites with unusually high nominal voltages probably had their service-line transformer taps set too high. Unfortunately, site-wiring and line-voltage inspections usually do not precede the installation of sensitive computer and network equipment, so dangerous tap settings may be common in user sites.

Power events also seemed to vary based on the types of load equipment at study locations. NPL estimates that at least half of all recorded events were caused by load equipment in the same building. This means that steps taken to clean up the power coming into the building are only half-way measures. Equipment also needs to be protected from power events originating in other floors, departments or even within the local-area network.

The "Event Voltage" table summarizes 32,225 aggregate days of monitoring at 112 locations, and lists the number of events recorded throughout North America by voltage category. Interruptions (zero-voltage events) are shown separately from the sags and overvoltages, in the table's bottom line.

This table can help system managers to define needed power-protection levels and determine the tolerance of specific electronic equipment. For example, a commercial switch-mode power supply rated to operate between 90 and 130 volts RMS would have protected against 91 percent of the real events shown in the table. Similarly, other immunity levels can be defined for various equipment and types of protection devices. By doing so, any location should be able to significantly limit its exposures and improve system uptime.

The NPL also recorded deviation from nominal AC frequency. The only sites in the U.S. that showed frequency deviations not associated with other power-quality events were those that used stand-by power generators. At these locations, we noticed that stand-by generator tests were sometimes performed in the middle of weekday afternoons, without regard for the possible effects on sensitive electronic equipment in the building. This situation shows that there is a real need for power-quality awareness education.

This article has profiled several types of "electronic rust" events. Electronic systems may absorb 20 or 30 power deviations without failing, but repetitive, accumulated events will ultimately destroy systems and stored data. When building one's defenses, no holes or back doors should be left open to power anomalies. Modem telephone lines and LAN cables are especially often overlooked, but they too can carry damaging noise or surges between your systems.

The NPL data on the types of power problems occurring today will help system managers build their defensive strategies. For some applications, a simple surge suppressor will suffice, whereas eight hours of back-up power may be required for other critical applications. So, learn about the different types of power problems and the variety of protection modes available. Then, identify the specific costs of downtime, lost data, data errors and hardware damage. Once this dollar figure is calculated, it is a simple matter to determine how much "power-protection insurance" each system requires. ∎

James J. Stanislawski is product evaluation manager for National Power Laboratory, an independent operating division of BEST Power Technology Inc., Necedah, Wisconsin.

EVENT VOLTAGE

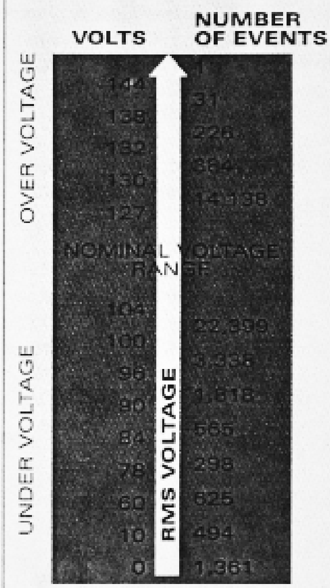

Most surges and sags fall close to the nominal voltage range. Only in one of every five incidents did the voltage stray above 130 or below 100.

Source: NPL 1990-1994 Study

HOW NPL DID IT

The National Power Laboratory in Needham, Wisconsin, recorded events in 130 randomly selected sites in ten regions of the United States and Canada. Monitors collected single-phase, line-to-neutral data at the standard wall receptacle, the normal power interface with computers and other electronic equipment.

Study locations represented users from heavy industry, light industry, office and retail stores, residential and mixed-use sites. The sites ranged from rural to industrial, building types from residential to multi-story, building ages from new to more than 51 years old, and local population densities from less than 2,500 to more than 500,000.

Some past power studies began with preliminary site inspections to verify that all wiring was in compliance with electrical codes, all connections were tight and proper grounds were present. But this type of inspection rarely precedes normal installation of products in the real world, so NPL sites were not "scrubbed" in this way before monitoring began. As a safety precaution, however, a plug-in circuit tester was sent to each site to verify proper conductor relationships at the monitored receptacle.

A dedicated phone line also was installed at each site to provide a communication link for monitor data. Each location was polled for this data at least twice a week. Additionally, all outages lasting more than four hours were verified with a site contact, to insure that they were real.

NPL researchers set monitors to detect power events based on voltage-level and duration thresholds. In general, the monitors were programmed to record any events for which the mathematical RMS voltage over one AC power cycle deviated outside the 104-127 volt range. This range conforms to both the ANSI C84.1-1989 limits and CBEMA recommendations for safe computer-input voltage. The monitors also were programmed to record any instantaneous voltage deviations of more than 100 volts positively or negatively. The low, 100-volt, threshold was chosen to catch transients that could affect Pentium chips and other high-speed, low-voltage processors.

—James J. Stanislawski

PC Power, Part 1:

POWER PROTECTION

Just what do those power protection devices do, and how well do they do it?

Mark Waller

Mention power protection, and the first thing many people think of is lightning. But, as someone who spends every day solving computer-related electrical-power problems, I think of money—protecting your investment in computing without wasting your money in the process.

To be sure, you must protect your computers from lightning. But you don't want to spend hundreds of dollars on a product only to find that it won't solve your problem. Neither should you deceive yourself into thinking that by spending just a few dollars on a surge suppressor, you have absolutely protected your computer from being damaged by a thunderstorm.

The Problem

Computer equipment is designed to operate with a steady stream of uninterrupted sine waves of 120 volts root mean square (RMS). The nature of utility power is such that, as often as twice a day, you may experience some electrical disturbance that falls outside your computer's acceptable limits. In major data centers across the country, study after study has shown that surges, sags, brownouts, blackouts, and damaging impulses happen with dismaying frequency.

Over the last 10 years, the quality of power has steadily declined. Microcomputer users are especially vulnerable to this degradation. While mainframe computers have the advantage of employing a dedicated power source, microcomputers live off power straight from the local power company. However, there is one alleviating factor in this situation.

Since you plug your computer into a nearby outlet, your machine is normally located a good distance from the building service entrance (i.e., the meter, or the place where power enters your building). Thus, in order to reach your equipment, potentially damaging impulses generated outside your location must travel through the impedance of lots of copper wire. This barrier serves to dampen out many of those disturbances, but you can derive only small comfort from this fact.

The real problem occurs along the electrical path from where the power enters your building to your machine. Between these two points are all kinds of devices, such as elevators, air conditioners, coffee makers, and so on. The ignition of an oil furnace, for instance, produces an electrical spark that can generate an impulse that might be more than 1200 V. The starting transient of an air conditioner is strong enough to interfere with any electronic equipment that may be connected to the same power-source transformer. Copiers are notorious as a source of noise that creates soft errors in computers that share circuits with them.

Any equipment that arcs, cycles on and off, or draws excessive bursts of current is a potential hazard to your computer. There are far more pervasive culprits residing inside your building than any potential lightning strike, and they should be the prime focus of your protective strategy. Lightning-caused surges are rare events. When protective devices such as gas tubes (lightning arresters) are shorted across a power line, lightning is diverted to ground. When this happens, you and other users down the line will experience a momentary power sag. This is why you will see lights flicker during storms.

It's more important to protect your computer from the more common electrical malfunctions caused by equipment in your building than to protect it from infrequent lightning surges.

Cause and Effect

The pressure to put computers into smaller and smaller packages caused a quiet revolution in power-supply design. Until the 1980s, computers used what is called a linear power supply (see figure 1). Its most prominent feature was a 60-Hz power transformer connected across the input (between line and neutral). After the line voltage was transformed from 120 V to 5 V, or whatever level was necessary to satisfy the DC logic, the power was rectified and filtered. (A rectifier is a device that converts AC current into DC current.)

Those 60-Hz transformers made linear power supplies big and heavy. Out of the need for smaller, lighter power supplies,

continued

the switching power supply was born (see figure 2). This design change eliminated the power transformer. With the new circuit, the incoming power is applied directly across the bridge rectifier. The resulting ripple DC is then pulsed at between 20 kHz and 100 kHz, depending on the specific supply design.

The action of chopping up the rectifier's output into high-frequency segments allows designers to use a high-frequency transformer, which is smaller and thereby reduces the size of the power supply.

The use of switching power supplies also dramatically affected computers' susceptibility to noise. A linear supply draws current in step with the voltage sine wave. In other words, as the line voltage rises and falls, the power supply's current demand rises and falls along with it. Linear power supplies are voltage-sensitive, however. If the supply voltage varies more than a few percent plus or minus, problems will develop.

On the other hand, a switching power supply (sometimes called a switch-mode power supply) is not voltage-sensitive. Such power supplies draw current in huge gulps once every half cycle. For this reason, the power source's internal impedance can be a problem because if the impedance is too high, the power source cannot deliver power easily. But while you must be concerned about current, you do not have to concern yourself with voltage regulation as you do with the linear power supply. Switching power supplies regulate the level of voltage by varying the amount of current that is drawn. This action is basically independent of the voltage of the power source.

Because they contain switching power supplies, microcomputers can operate over a wide voltage range. This range can be from as low as 80 V to as high as 140 V.

There are devices on the market, such as ferroresonant transformers, that regulate voltage to microcomputers. However, since your computer's power supply does not need voltage reg-

continued

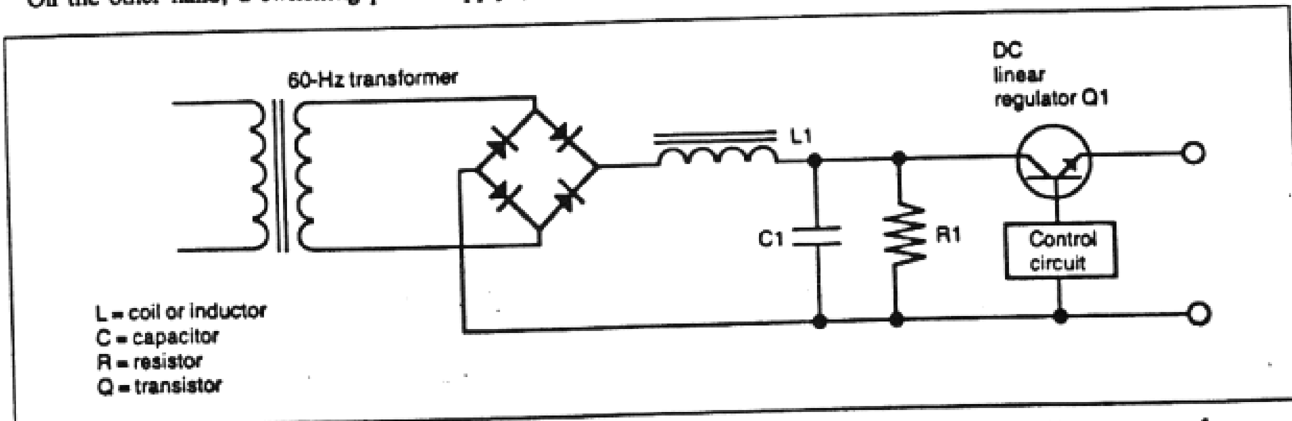

Figure 1: *Linear power supplies, used in small computers up until a few years ago, featured a large 60-Hz power transformer connected across the input. Such power supplies were sensitive to variations in voltage and made power supplies bulky.*

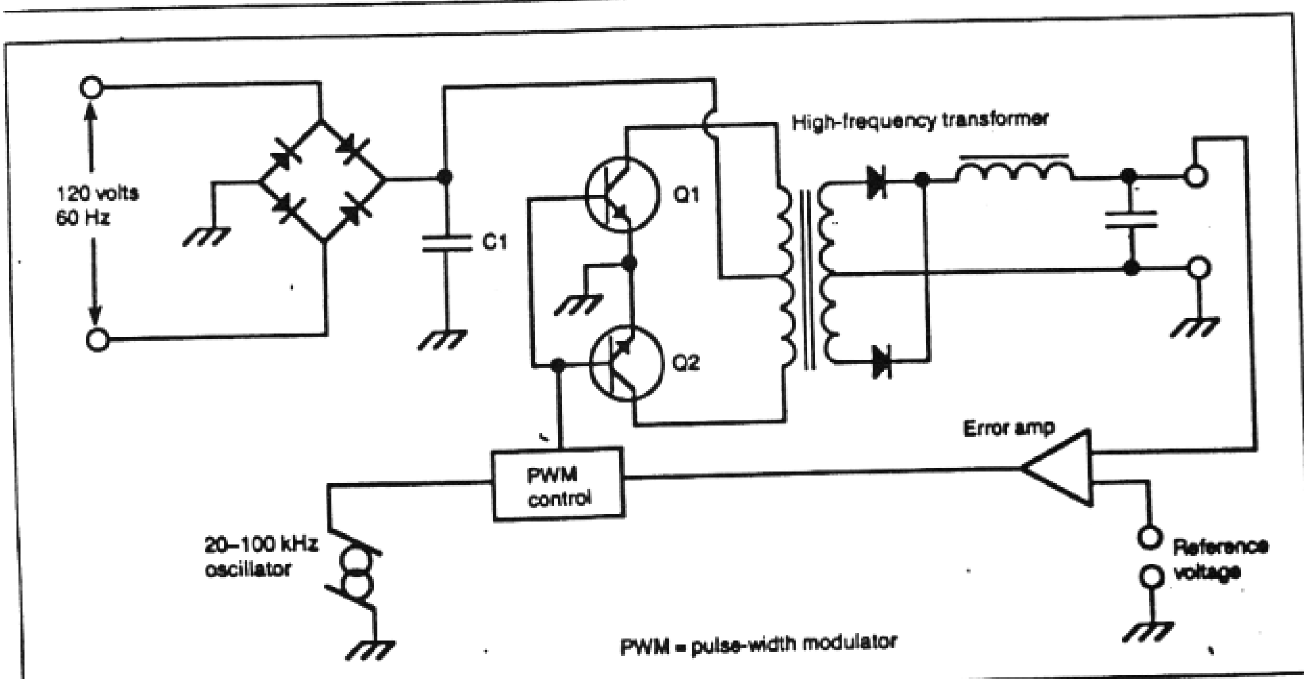

Figure 2: *The circuit of a switching power supply. The use of small, high-frequency components allows such power supplies to be smaller but makes the computer vulnerable to common-mode noise.*

ulation, these devices are unnecessary. In addition, such devices limit the amount of instantaneous current that can be delivered to your machine's power supply—an undesirable attribute, for the reasons explained earlier.

Noise About Noise

If you look behind the faceplate of the nearest wall plug, you will see either two or three wires. The black wire is called the phase wire, sometimes termed the "line" or "hot wire." The white wire is called the neutral wire. If you see a third wire, it will be the ground wire and will be either green or bare copper. If you do not see a third wire, the installing electrician may have used the metal conduit as the ground path.

Where your service enters your building, you will find that the neutral and ground are bonded. If you measure the voltage between neutral and ground at the outlet, it will usually be zero. If you measure from line to ground, or from line to neutral, it will read 120 V. These three wires not only provide power to your computer, but are the path through which electrical noise travels.

Let's define electrical noise as any signal, other than the desired signal, that appears in a circuit. Noise, then, can be either minor or major. Noise can include large transient events or damaging impulses, or it can be continually oscillating signals from spinning motors and other kinds of interference. There are two kinds of noise: *normal mode* and *common mode* (see figure 3).

Normal-mode (or transverse-mode) noise appears as a voltage between line and neutral. The word *normal* is used because that's normally where utility power is transmitted, between line and neutral.

Common-mode noise can be measured from line to ground or neutral to ground. This type of noise appears on both the line and neutral with respect to ground; in other words, it is common to both lines.

Basically, your computer's power supply is vulnerable only to high-energy impulses that appear in the normal mode (normal-mode noise). And generally speaking, a computer's chips and logic are vulnerable only to common-mode noise.

Power-supply components are designed to take line voltage (normal mode) with peaks of up to 170 V and convert it to DC. Because power-supply components are so rugged, they have a high degree of immunity to normal-mode noise. An oncoming impulse would have to be several hundred volts before it would damage your computer's power supply.

The old linear power supply, with its big power transformer, was immune to common-mode noise. Noise appearing along the line and neutral would cancel in the primary winding, because they are 180 degrees out of phase. If the cancellation process was imperfect, the magnetic transformation would convert it to normal-mode noise. Not so with switching power supplies.

Switching power supplies have no up-front transformer. And, because their components are tightly packed, they offer many capacitive paths at various frequencies. Stray capacitive coupling inside your machine and ground loops between other devices can let common-mode noise slip into, around, and through the power supply and reach the computer's chips and logic. Also, your logic chip's ground reference is usually tied directly to power ground—a sure recipe for disaster. What this means is that at various frequencies, common-mode noise may appear across the logic circuits themselves.

Because the distance between connections on the chip is only a few microns, ICs can tolerate only a fraction of the voltage that the rectifiers inside the power supply can tolerate. Noise from a few volts to a few dozen volts will interfere with your processing. Common-mode noise exceeding a few dozen volts could destroy your computer's chips.

Ground Yourself

Ground, as it relates to computers, is probably the single most misunderstood electrical concept. As far as your computer is concerned, ground is not earth. Grounding something has nothing to do with driving a copper rod into your flower bed. The earth is not an electrical septic tank into which we flush unwanted noise to make it disappear forever.

Electricity travels in circuits, and current flowing to a point will flow away from that point. If current is directed to a ground wire, it will reemerge somewhere else along any electrical path that might be part of the ground circuit. This circuit may take different paths at different frequencies.

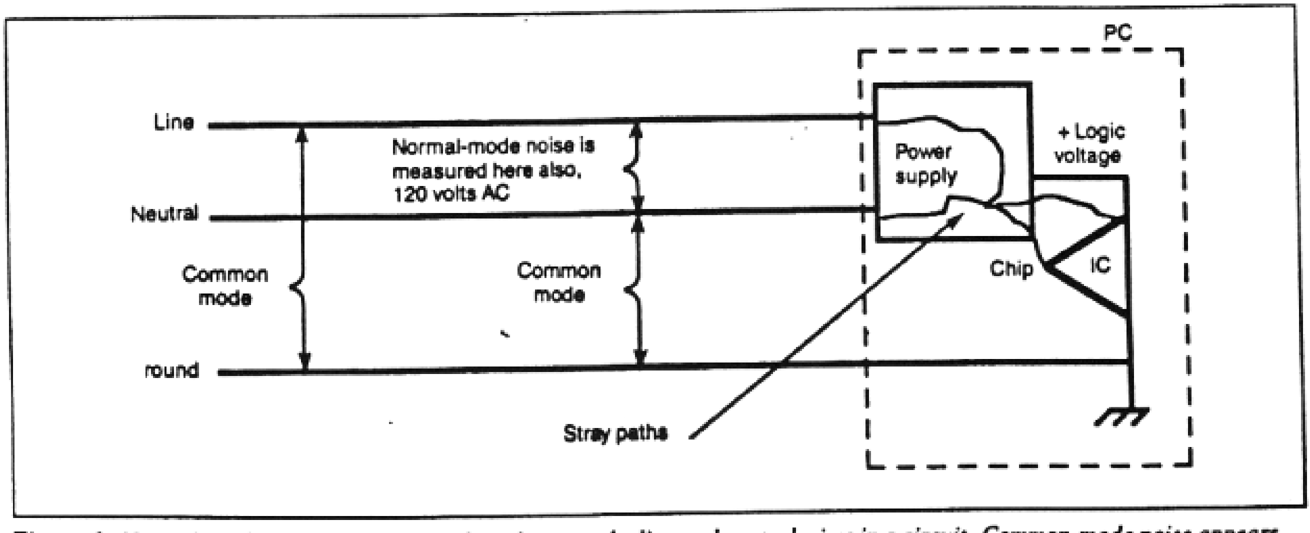

Figure 3: *Normal-mode noise appears as voltage between the line and neutral wires in a circuit. Common-mode noise appears between the neutral and ground wires. If common-mode noise can find a stray path (and it will, especially through a switching power supply), it will appear across the chip from one of its pins and the logic ground pin. Normal-mode noise appears across the power supply just like utility power.*

If a power glitch occurs in your computer at normal power frequencies, electricity directed to the ground wire should travel back to the electrical panel to trip a breaker. At higher frequencies, however, a noise signal may find stray paths through boards, cables, or between cabinets to be a far lower impedance route back to its source than the power ground wire. This is called a ground loop (see figure 4). Ground loops can be a source of processing errors as well as actual hardware damage.

Local area networks are extremely susceptible to ground loops. In such an environment, current will flow because of the electrical potential difference between the ground connections of different workstations. This undesirable current flow may induce dangerous voltage levels in nearby electronic components.

An IC is referenced to ground. It operates by detecting a logic level of so many volts with respect to ground. If the ground reference point changes in relation to the logic level, errors will result. If this voltage difference exceeds the *withstand* rating of the chip, current will bridge the substrate of your chips and destroy them.

Suppressing Those Surges

Before looking at the actual circuit elements involved in the common surge suppressor, let's look at what it is supposed to suppress. Typically, you think of a surge as a spike or an impulse. Figure 5 shows what an impulse might look like. It initially rises to a peak and then oscillates in a diminishing fashion until it dissipates.

There are two vital features to an impulse. The first is its kinetic energy (joules or watt seconds) determined by its peak voltage, current, line impedance, and time span. The second is its rise time, or the time it takes to rise from nominal voltage to its peak voltage.

It is the front slope of the impulse that causes damage to your computer. This rapid rate of change is full of energy at various frequencies. The faster the rise time, the more high-frequency components the spike contains. It is these high frequencies that find those stray paths and cause all the damage. Lightning, arcing, and sparking have extremely fast rise times. At these high frequencies, the physics of electricity and the paths it follows are very different from 60-Hz utility power. Your computer's circuitry was never designed to digest this kind of high-frequency energy.

Scientists have tried to quantify and define what the typical spike might look like. The result of their findings is a standard that has come to be known as the IEEE 587 ring wave (see figure 5). It is a waveform with strict parameters and is a test-measuring criterion for surge-suppression equipment. This is why so much good power equipment states proudly on its package that the product can withstand so many hits of the IEEE 587 test wave.

Recently, UL introduced a testing standard of its own, called UL 1499. In most respects, this waveform is similar to the ring wave. When you are in the market for surge-suppression products, look for these standards to tell you that the product actually performs as advertised.

But will surge suppressors really protect your computer?

Diversion Tactics

Actually, a surge suppressor doesn't suppress unwanted electrical energy; it diverts it. Rather than suppressing, absorbing, arresting, or otherwise making unwanted impulses disappear, these devices actually divert the energy from one path to another.

Transient suppression devices come in four different varieties: metal oxide varistors (MOV), zener diodes, filters, and gas tubes. By far the most popular device is the MOV. The term *varistor* means variable resistor and describes the MOV's basic function. As voltage builds up across this device's terminals, it reaches what's called the breakdown voltage. At this point, the varistor changes from a highly resistive device to a low-resistance device, and large amounts of current can then flow through it.

If you connect a MOV in parallel to your machine, when a spike comes along, the MOV will clip it. In other words, that portion of the impulse that rises above the MOV's breakdown voltage is clipped off and diverted through the MOV. This clipping level is usually around 140 V RMS. The peak let-through

continued

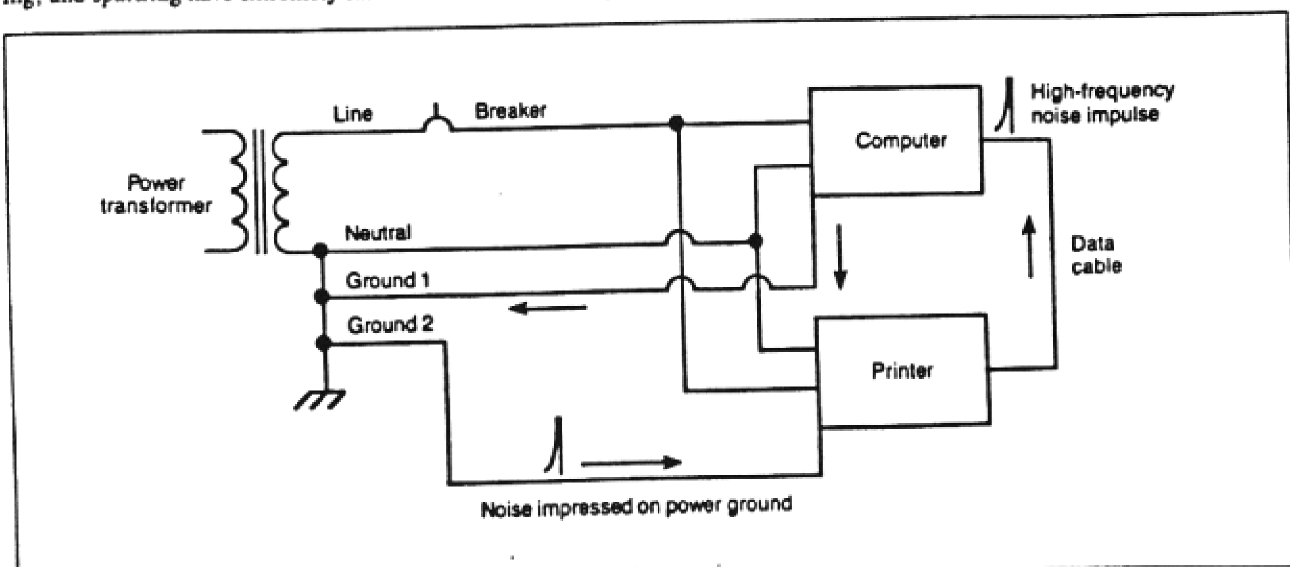

Figure 4: *Noise current will take the path of least resistance, a situation that may interfere with the transmission of data between devices or even cause damage.*

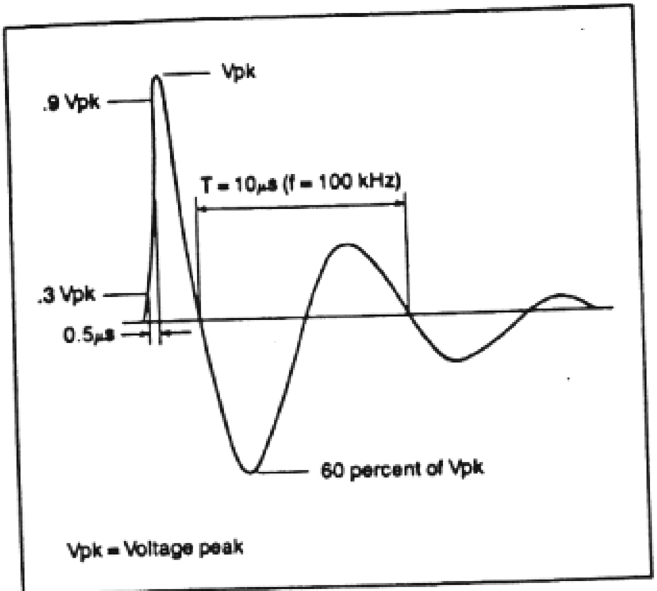

Figure 5: *The IEEE 587 ring wave is a testing standard for surge-suppression products. Engineers have found this wave shape to be typical of what might appear on 120-V circuits leading to your computer.*

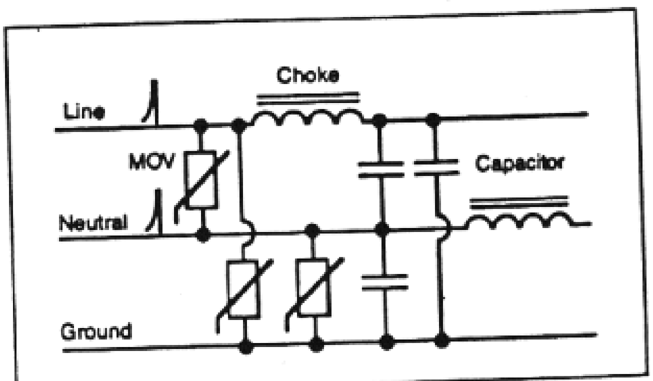

Figure 6: *Typical circuit of a surge suppressor. Most simple surge strips have only a MOV (metal oxide varistor) and capacitors from line to neutral.*

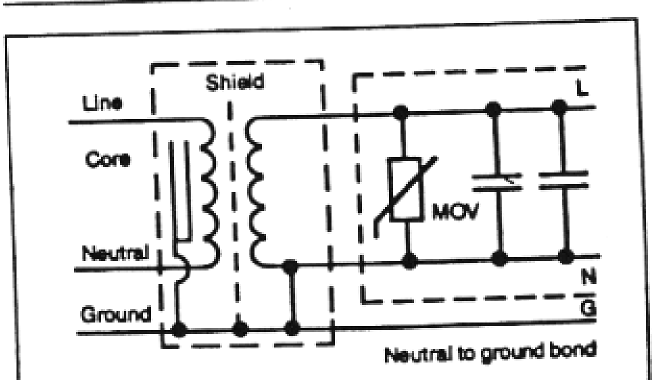

Figure 7: *A power-line conditioner contains an isolation transformer with secondary surge suppression and the neutral and ground bonded.*

voltage is likely to be as high as 340 V in some cases. Most often, you may think of a spike as appearing at the peak of the sine wave. But if the spike appears in the valley of the sine wave, the level of the voltage before clipping will be high. This is one of the weaknesses of this type of transient suppression device.

Zener diodes, sometimes called avalanche diodes, act similarly to MOVs. They do, however, have different performance characteristics. Zeners have a faster response time and come in sizes with a lower breakdown voltage than MOVs. MOVs, however, can usually handle more current than zeners. In order to take advantage of these complementary qualities, manufacturers often place both devices in surge suppressors.

Filters, in the form of capacitors and chokes (coils), are used in surge products to block the flow of noise current at the design frequency and to divert it through a lower-impedance path. Most surge suppressors have one or more capacitors. The better ones have chokes in series on the line and neutral wires.

Another device common to some suppression products is the gas tube. When voltage builds up across a gas tube's terminals, the gas inside the tube ionizes and becomes a conductive path. Through arcing, the path ionizes, and the energy is bypassed.

The arcing action of a gas tube, though, creates very undesirable high-frequency characteristics that make it inappropriate for placement near your computer. In addition, a gas tube can take a seemingly unimportant impulse and turn it into a damaging impulse. Yet, the market has seen the proliferation of tiny gas tubes inside surge suppressors. Evidently, designers think that including a gas tube in a surge suppressor will give you the illusion that it can handle enormous amounts of energy.

The proper use for a gas tube is in a lightning arrester placed near your building's service transformer. Enough wire exists between this point and your machine to block the passage of the high-frequency effects of gas-tube firing.

In figure 6, which shows a typical surge-suppressor circuit, notice the MOV that is placed between line and neutral. As this MOV conducts a high-energy impulse from the normal mode, current is dumped onto the neutral. This current flow creates a voltage drop between neutral and ground. By this process, the surge device has used normal-mode noise to generate common-mode noise. Photo 1 shows that the impulse created by this current flow is nearly as large as the one from line to neutral.

Notice that to protect your computer from common-mode noise, figure 6 also shows MOVs connected from line to ground and from neutral to ground. This is a good feature. But remember that common-mode noise sensitivity is significantly higher than that for normal mode. You must be concerned with the logic and any voltage that might appear across it. A MOV will allow up to several hundred volts to pass through before it activates.

Suppressor Circuit Caveats

In the surge-suppressor circuit (see figure 6) you see filtering elements made up of chokes and capacitors. This is a fairly well-engineered circuit. Someone has taken the time to worry about both normal- and common-mode noise and has included filtering as well. Unfortunately, simple surge strips that go for about $10 to $20 usually have only one MOV between line and neutral. Obviously, you should be concerned about what's inside the surge suppressor, though it is difficult (if not impossible) to tear open a product before you buy it.

There is still another problem. Not only does the common surge protector convert one kind of noise into the kind your computer finds least tolerable, but when parts of your device

continued

fail, the device won't give you any indication that you no longer have surge protection. Because they are connected in parallel to your computer, when MOVs or zener diodes fail, your machine will still run and you won't know that the surge device has passed away.

Perhaps to make you feel better, some manufacturers build into surge protectors a status indicator—usually a little green light. A green light tells you everything is OK, right? Wrong. Most status indicators just tell you that power is flowing. Thus, you may think your surge strip is protecting you, but you don't know for sure.

continued

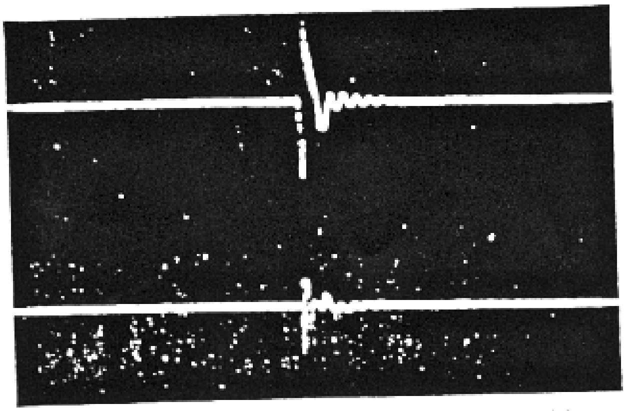

Photo 1: *As the impulse in the normal mode (top trace) is conducted by the MOV from line to neutral, another impulse (bottom trace) appears between neutral and ground, the common mode.*

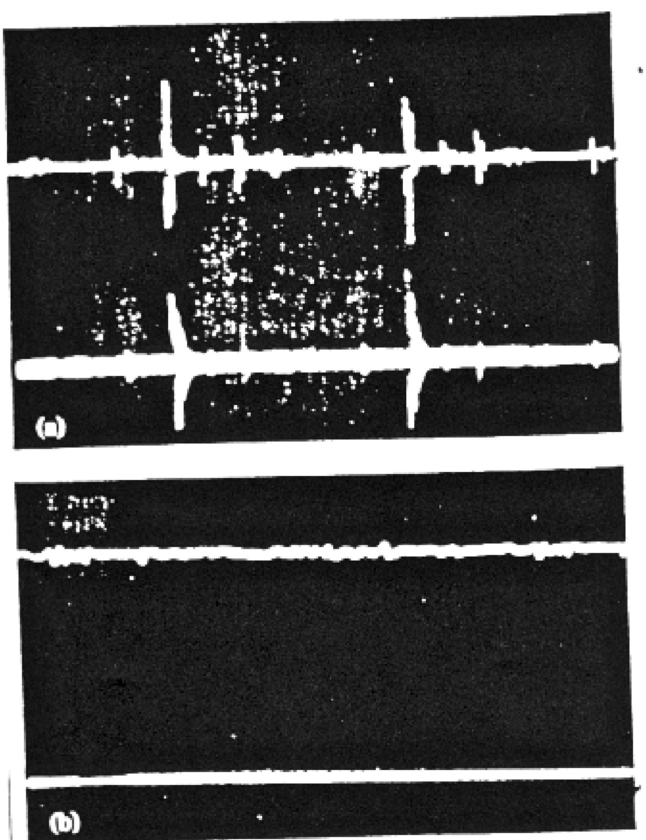

Photo 2: *These photos show noise in an electrical circuit,* (a) *before and* (b) *after insertion of a power-line conditioner. In both photos, the top trace is normal-mode noise and the bottom trace is common-mode noise.*

So, is a surge suppressor the answer to protecting your equipment? Not really. There is an alternative that will protect your computer investment much better than a surge suppressor.

A Better Solution

If you want to protect your investment in computing without wasting money on products of dubious utility, or if you are trying to solve power problems you already have, I recommend a power-line conditioner with a built-in isolation transformer at its heart (see figure 7). Properly designed, the transformer, along with a couple of capacitors and a MOV across the secondary, will give you security far superior to that of a surge protector. Photos 2a and 2b (before and after insertion of a power-line conditioner) show how effective this design can be in protecting your computer from undesirable voltage impulses.

The isolation transformer acts as an inductive cushion, stripping away high-frequency components of normal-mode noise. Any remaining normal-mode noise will be shunted by the filter capacitors connected across the transformer's secondary, or by the MOV if it contains high energy.

Perhaps a power-line conditioner's most important feature is the neutral-to-ground bonding on the secondary side of the transformer. This is a requirement of the national electrical code that has some very happy consequences for all microcomputer users. This bonding is a short circuit for common-mode noise, and, since there is no impedance across a short circuit to allow a voltage to develop, common-mode voltages do not occur (Ohm's law: current × impedance = voltage). With this type of device, no voltage will appear across your logic circuits.

Suppress or Condition?

When all's said and done, then, what kind of device will really power-protect your computer? If you opt for a surge suppressor, a device that is relatively inexpensive and easily available, what features should you make sure it has? You want filtering as well as surge suppression. You have to have both normal- and common-mode protection. And you should have some way of determining the state of the device's internal components. In addition, be sure that it has been tested to UL 1499 or IEEE 587 standards. To obtain this type of surge suppressor, you will probably have to pay more than $100. But even if you do choose this route, you have hardly obtained the ultimate in power protection for your computer.

If you opt for the alternative, a power-line conditioner, you may need to ask the advice of a power professional to help you make the best choice, or you can purchase your device from an industrial or commercial dealer. This more effective product costs around $250, much more than a simple surge strip.

Computer power protection is not as easy or inexpensive as you might think. Protective devices are like insurance—a trade-off between cost and risk. In most cases, a quality choice, while it may not be the least expensive, is the best choice. ■

Editor's note: *Next month, in Part 2 of this series, Mr. Waller will discuss backup power devices.*

Mark Waller is a computer facilities consultant and the author of Computer Electrical Power Requirements *and* Mastering PC Electrical Power, *both published by Howard W. Sams. He can be reached on BIX as "editors."*

PC Power, Part 2:

BACKUP POWER

How do you provide reliable backup power for your computer without creating new electrical problems—and how much will that cost you?

Mark Waller

When the utility power fails, your computer won't work. You may think the solution to this problem is an uninterruptible power supply. A UPS device supplies continuous power to the computer whether the utility power is flowing or not. However, in the world of personal computers, backup power systems are generally standby power systems, or SPSes. These devices switch on when utility power fails. This distinction of switching or not switching is the basic functional difference between a UPS and an SPS.

The main task of these products is to keep your computer running when utility power fails, as opposed to surge suppressors or power conditioners whose main task is to protect your computer. Backup power and power protection are areas of vital importance—especially if your investment in computing includes a local area network (LAN) or desktop publishing system. It is vital that you know the difference between backup power and power protection. Since SPSes are not all alike, you must be vigilant; otherwise, not only will you waste your money, you may end up with additional power problems.

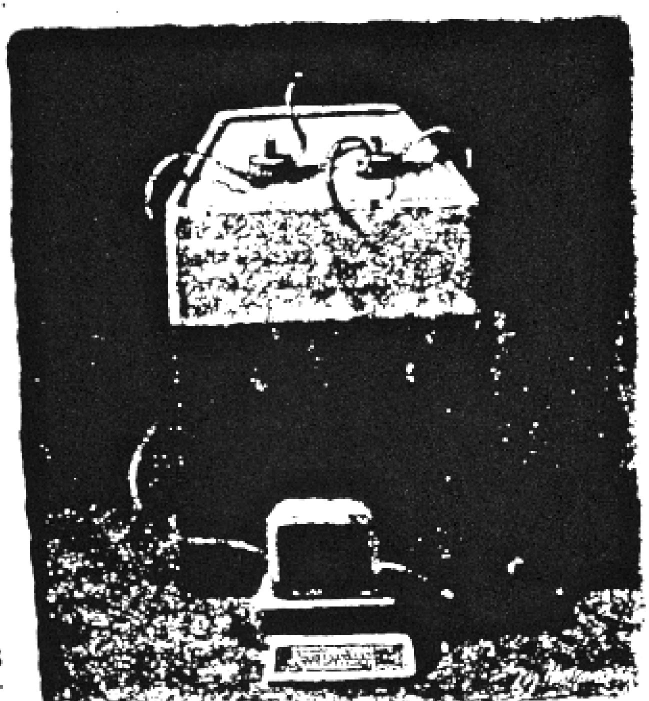

Standing by . . .

The building blocks of an SPS include the battery, the battery charger, the transfer switch, and the inverter (see figure 1). Most of the time, raw utility power is fed into the computer. All the while, though, the battery charger keeps the battery ready in case the power should fail. If power does go out, the transfer switch senses the outage and turns on the inverter, which provides AC power by drawing energy stored in the battery. After power is restored, the switch turns off the inverter and transfers the computer back to utility power.

That process sounds simple. Theoretically, it is. But different brands of SPSes handle this situation with varying degrees of efficiency. There are many functions that an SPS must provide during this simple-sounding process. It must, for instance, recognize at what point the utility power has failed and whether or not the voltage must drop out completely before it makes the transfer. It must also decide how fast it will make the transfer once the process is initiated. The way an SPS deals with these tasks determines the effectiveness of the product.

A better SPS will switch over to battery any time the utility voltage drops below a certain preprogrammed level; say, 103 volts. Low-budget units often have less-expensive sensing circuits that may not transfer to battery until the power drops significantly. Since every manufacturer

continued

rates its SPS in terms of transfer time, how they calculate this time is critical.

Your computer's power supply has some capability to store electrical energy so it can ride-through extremely short-term outages. The term *ride-through* refers to the power supply's ability to deliver stored energy to the computer even when its power supply has lost incoming power. As a general rule, your computer can tolerate outages from 20 to 30 milliseconds before it goes down. If your SPS can transfer to battery power within that time frame, you should be able to stay operational. Generally speaking, the less-expensively constructed your power supply, the less time it will be able to sustain the computer through an outage.

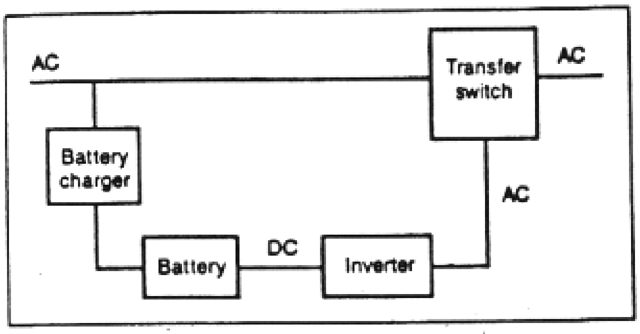

Figure 1: *The building blocks of an SPS are the battery and the inverter. The battery powers the inverter, which converts DC into AC. Notice that these blocks only operate when the transfer switch senses that utility power has failed.*

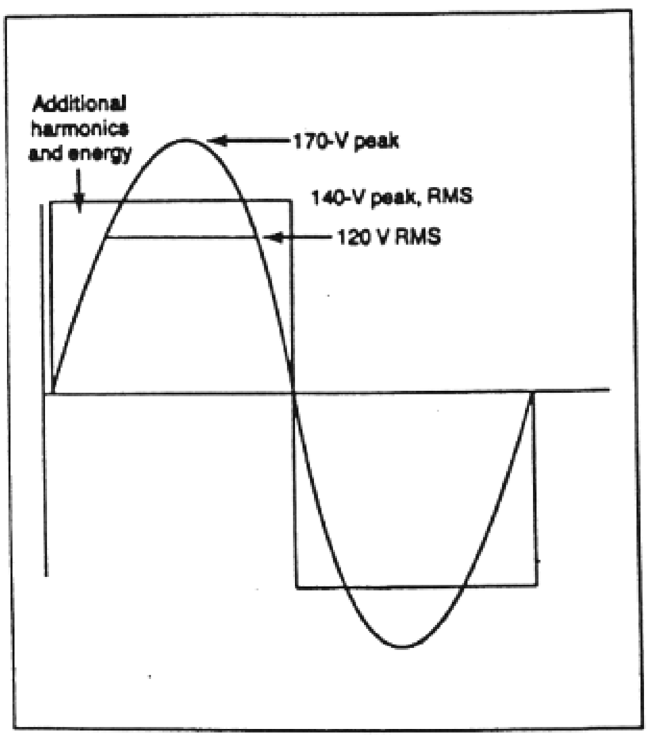

Figure 2: *A comparison of the sine wave generated by utility power and the square wave produced by many SPSes. At any given point along the graph, a square wave contains more energy than a sine wave. Its peak voltage is equal to its RMS voltage.*

Most manufacturers claim their devices have from 4- to 10-ms transfer times, which is within the ride-through time of the computer's power supply, and you should have no problems keeping your system up when the power goes down. If, however, the voltage drops for several milliseconds before the unit senses the outage and this time increment is not taken into account, the stated transfer time may be well under the actual transfer time. Reputable companies include the time it takes to sense an outage in their total transfer-time calculation.

If 4 to 10 ms is good, wouldn't 2 ms be better? Some companies boast that their devices have transfer times of 2 ms or less. Reaction times as fast as these can cause the inverter to kick in and out, constantly draining the battery every time a short-term drop in voltage occurs. The point here is that you don't need your SPS continually reacting to events that are not outages but simply fast fluctuations on the line.

Thus, transfer time is a trade-off between the computer's ability to ride through small power glitches and the need to provide quick backup power in case of a real outage. The fact that computers will tolerate such a long time without power leads most manufacturers to advertise their products as UPSes instead of SPSes—a questionable practice, and you should know the difference before you buy.

Riding the Wave

Utility power is a 60-Hz sine wave. You might just assume that your shiny new SPS will generate just such a sine wave. But, if you look at the fine print in the specifications, you may very well see terms such as *square wave*, *rectangular wave*, or *modified square wave*. Sometimes the literature will even show a picture of a great-looking square wave. Of course, you won't know if, when the picture was taken, the SPS was under load, or if the load was a switching power supply such as that inside your computer, or simply a plain old light bulb.

Figure 2 shows a comparison between a sine wave and a square wave. A square wave is a compromise between cost and quality. Since an inverter incorporates digital switches that turn on and off, it is less expensive to design an inverter to produce a square wave. But since a square wave's peak voltage is equal to its root-mean-square voltage, designers must compromise on a level somewhere between the normal 170-V peak of the utility sine wave and its RMS voltage of 120 V—usually around 140 V.

A 140-V square wave provides too little voltage with too much energy for the standard computer loads. A square wave at any given point along its curve contains more energy than a sine wave. Some engineers claim that your computer's power supply needs a minimum of 148 V RMS. The additional energy in a square wave will cause power-supply overheating and stress.

One school of thought preaches that a well-designed square wave is the best waveshape for switching power supplies since the switching power supply draws current in a nonlinear fashion. Engineers may argue, but you must be able to determine if the product is well designed, or it may produce the kind of waveshape shown in photo 1. This is the "modified" square wave generated by dozens of SPSes on the market. This waveshape will change as the load increases in order to keep the RMS voltage at the proper level, and it may even look more like a square or a rectangle.

On the other hand, when manufacturers go to the trouble of producing a sine wave, the output will likely be electrically cleaner than the average square wave.

A potential problem with an SPS inverter's waveform output is its high-frequency noise content. This noise can damage your computer components or interfere with your processing. Many inverters use pulse-width modulation. PWM is a means of producing a desired waveform using high-frequency switching. The SPS filters out high-frequency components of the resultant signal, and the effect is a sine wave of low harmonic content with relatively little noise output. (Harmonic content refers to multiples of the fundamental frequency of the intended waveform. These multiples may cause distortion of the waveform, or high-frequency harmonics may appear as noise.)

In spite of this filtering, some of the switching noise will leak through and appear on the output. Square and rectangular waveshapes have a greater tendency to produce noise, since the waveshapes contain harmonics of the fundamental 60-Hz power signal and are created by large switching pulses.

Photo 2 shows the high-frequency content of the modified square wave in photo 1. You may think this looks harmless. Consider the fact, however, that these noise impulses occur about three times every cycle and are about 350 V in magnitude. If you buy this model, you may save around $50 over a better-engineered model, but you may put your computer in danger every time your utility power fails.

Synchronicity

After the SPS has been operating off the battery and utility power returns, two things must happen. First, the sine-wave output of the SPS must synchronize with the incoming utility power. This process is sometimes called *phase matching*. Then the unit must switch from battery power to utility power.

These operations sound fairly simple; the term describing them is *retransfer*. Most inexpensive units do not synchronize, however, and synchronization is important. Your computer's power supply is designed to expect the peak voltage of the sine wave to occur at regular intervals. If the peak of the sine wave is missing for very long, your computer may crash. When the SPS is on battery power, it generates its own sine wave or square wave according to its own internal clock. When utility power is restored, the sine wave's phase may not match that of the SPS. In order to prevent a mismatch, and perhaps a system crash, the SPS must slip sideways: *slew* its sine wave to match the phase of the utility wave before retransferring.

The SPS must also decide at what voltage level it should initiate retransfer. Normally, this level is a different, higher voltage than the transfer voltage. For instance, if the SPS transfers when the voltage falls below 103 V, it may wait until the voltage rises to 108 V to retransfer. This process avoids "dancing" on and off the inverter if the voltage hovers at a level near the transfer point.

Many SPSes provide selectable transfer points for site-specific considerations. If your location has periodic brownouts, you may want to set your SPS to a low transfer point, such as 90 V, to avoid constantly draining its batteries.

Battery Basics

The concept behind providing backup power for your computer is to give you enough time to save your work in volatile memory and bring your system down safely. SPSes are not designed to give more than a few minutes of backup power. If you need much more backup time than that, you may have to modify an existing unit to accommodate a larger number of batteries. There are some models that allow for this option.

The amount of battery time you need is a function of the kind of processing you do. You may, for instance, have an accounting package that needs several minutes to complete a task. Many manufacturers gamble that you are only going to load your SPS to about 60 percent of capacity, and in order to achieve smaller, sleeker packages, they may undersize the battery. If your actual load is 300 watts, you may want to give yourself additional capacity and buy a 500-W SPS. Compare one manufacturer's cabinet to another's. All things being equal, the size of the cabinet itself will tell you which unit will give you every minute you need.

Another battery-related feature you should consider is called automatic shutdown. If your SPS's battery completely drains, it may fail prematurely and not last nearly as long. And, short of total drainage, there is a point, called the end voltage, beyond which additional discharge will cause damage to the battery's cells.

SPSes that come with a built-in automatic shutdown function will probably cost a bit more, but this type of device may prove to be a bargain if your batteries last twice as long as those in a less expensive model.

continued

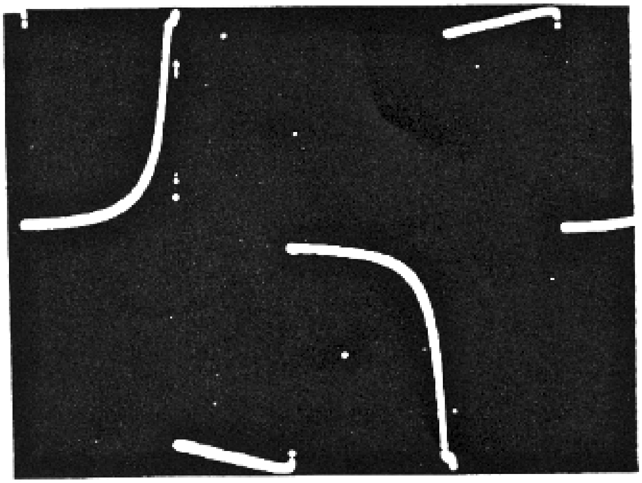

Photo 1: *The modified square wave produced by many inexpensive SPSes. This waveshape changes to correspond with the increasing demands of the load in order to maintain the designed RMS voltage.*

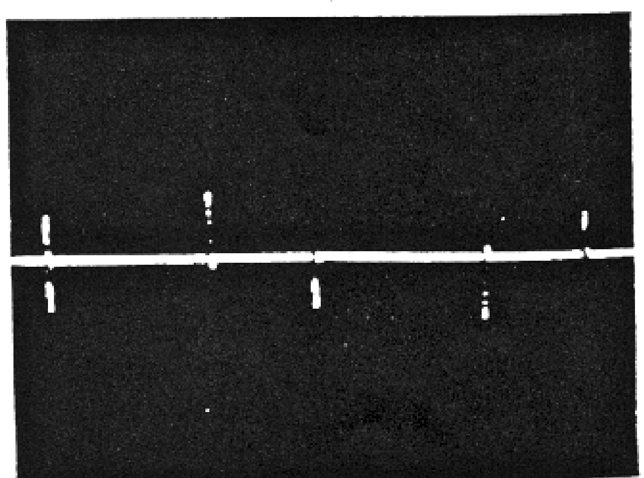

Photo 2: *These innocuous-looking spikes are the output of the typical inexpensive SPS. They are actually 350 V in magnitude and can be the source of stress on your computer's power supply; they may even cause processing errors or damage to your computer's ICs.*

Spike Those Rumors

Many manufacturers advertise that their SPSes are also power conditioners, which might make you think that you are protected from noise and spikes as well as blackouts. But SPSes are not cure-alls. This "conditioning" usually means placing one or more metal-oxide varistors deep inside the SPS (see "PC Power, Part 1: Power Protection," October BYTE). This is an inexpensive way for manufacturers to claim they have provided you with a surge suppressor. But a surge suppressor is an inferior form of power conditioning. To expect this conditioning to be much more than window dressing is wishful thinking.

The Ideal Product

You might think that the ideal product would be an on-line UPS. This concept is partially correct. The difference between a true UPS and an SPS is that in the UPS, the inverter powers the computer at all times. The incoming utility power is converted to DC by a rectifier/charger that transfers power to the inverter over a DC bus. The batteries are connected to this DC bus and, if the utility supply should fail, can provide instantaneous power to the computer—there is no switching time to worry about.

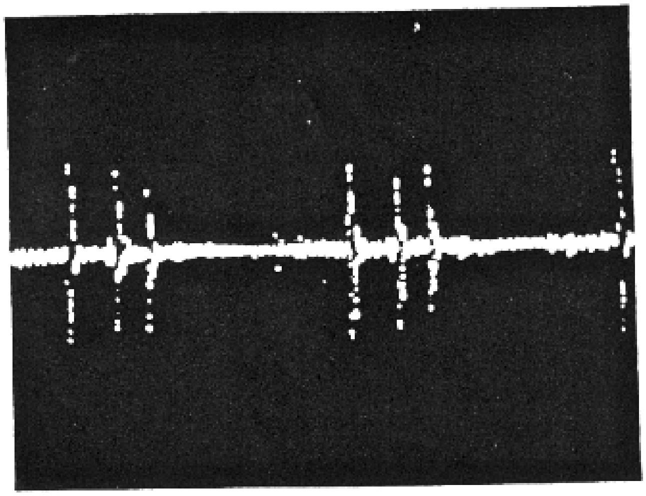

Photo 3: *This photo shows common-mode noise present on the line before insertion of an on-line UPS. Common-mode noise can be the most damaging form of noise.*

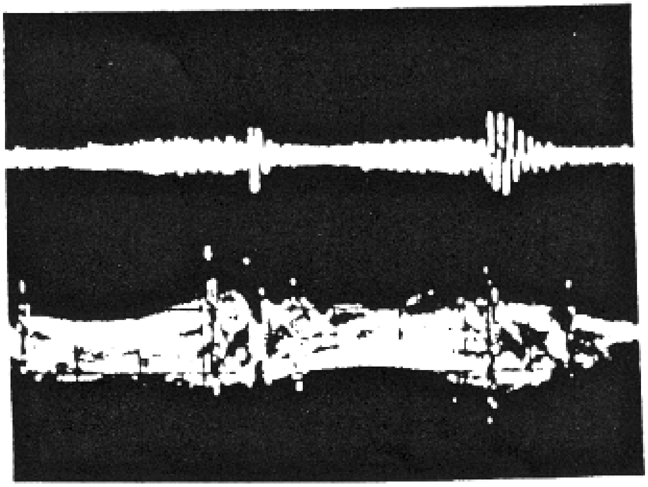

Photo 4: *Normal-mode noise (top trace) and common-mode noise (bottom trace) after the insertion of an on-line UPS. The top trace contains a small amount of noise from the UPS's inverter. The bottom trace shows inverter noise superimposed over the common-mode noise present in photo 3. As you can see, the UPS has contributed to, rather than eliminated, common-mode noise.*

In a UPS, the inverter and rectifier/charger are on-line all the time. Thus, they are larger and more substantial than those found in an SPS. This fact makes the cost of a comparable UPS three to five times that of an SPS. Specialized peripherals such as some large external hard disk drives might not be able to tolerate the switching times of an SPS.

Because of the AC to DC to AC conversion, the on-line UPS provides an excellent barrier to normal-mode noise (again, see last month's article). But the on-line design does little to suppress common-mode noise.

Photo 3 shows the common-mode noise present on the power line. Photo 4 shows the common-mode and normal-mode noise present after the insertion of an on-line UPS. Notice that the PWM inverter inside the UPS produces only a few volts of normal-mode noise (top trace). But the inverter generates about 30 V of common-mode noise, which you can see superimposed on the noise already present in photo 3.

In my previous article, I concluded that a power-line conditioner is the best product to protect your computer. The heart of this device is an isolation transformer with the neutral to ground bonded on the secondary. On the other hand, if you are concerned with blackouts, a properly designed SPS will carry you through outages. But, since it is not on-line all the time, there is no ongoing power protection.

The ideal product would seem to be a combination of power conditioner and SPS. In fact, one company recently announced a product that combines the two in one cabinet. But at $1500 for approximately 500 W, it is out of reach for most of us—although this price is less than that of many on-line UPSes of the same size.

A marriage between a power-line conditioner and an SPS is the best solution. The most effective way to connect the products is to plug the SPS into the wall, and the power conditioner into the SPS. This way, you are conditioning your power even when it is provided by the battery and inverter. The conditioner takes care of any SPS-generated noise.

Protect your LAN

The industry is just beginning to see the results of the computing revolution that networking has brought about. The economic value of data being handled via LANs is greater than ever before because there are multiple users who depend on the data. Networked data is often constantly updated by a file server and a large hard disk drive. They must be protected at all costs. In addition to mere protection, you must also be concerned with outages because they are the source of head crashes, downtime, and lost data.

Therefore, it is essential that you obtain an SPS to back up the file server and associated hard disk drive. Backup power is a necessary element in the success of a LAN because the network will not work for you if you don't learn to rely on and trust its performance.

Not only should you use an SPS for backup power, you should also use its external alarms, which will trigger LAN software to alert other users that they are on battery power and have only minutes to close their files and shut down. Remote users, whose power may be coming from a source that hasn't

continued

cially need this feature.

sers should also have backup power, or the file server be able to use battery time to bring users back to their gracefully.

Of course, you should use a power-line conditioner to protect your hard disk drive and its valuable data from destruction. Essentially, a LAN is just one big ground loop with several power sources and cables running all over. You should make every effort to provide a noise-free power environment so you will avoid degradation of your data.

> *With an expensive computer, you shouldn't scrimp on power products.*

Publish or Perish

A desktop publishing system is usually a state-of-the-art computer with a high-speed processor, a large hard disk drive, a tape backup, and a laser printer. With an investment of this magnitude, you shouldn't scrimp on protection; what comes down the power line can seriously damage your system.

The typical laser printer draws as much as 1000 W, and this can be a problem. Thousand-watt power conditioners are expensive. Your only less-expensive option is a surge suppressor. If it is properly designed, a surge surpressor should protect the printer well enough to prevent any major damage. A laser printer, after all, is a lot more like a copier than a computer on the inside.

In a desktop publishing environment, backup power is essential for the computer, but not necessarily for the printer. If a printing function is disrupted by an outage, it may be a nuisance, but the job can be done again. The computer, on the other hand, has large and rather lengthy processing chores. It makes sense to provide it with an SPS to ensure that jobs finish and files are closed before power shuts down.

Buy Safe, Not Sorry

Many people spend thousands of dollars on hardware and software, only to dash out and buy the least-expensive protective and backup power products they can find. You will waste your money on this strategy and leave your system vulnerable, and you may even introduce undesirable noise into the electrical environment.

To obtain devices that provide you with a safe personal computing system, you should invest the same amount of thought and evaluation as when you selected your computer to begin with. The process will vary depending on your needs and your particular system configuration.

Armed with the facts, you can make informed buying decisions and then relax and let your investment in computing work for you. ■

Mark Waller is a computer facilities consultant and the author of Computer Electrical Power Requirements *and* Mastering PC Electrical Power, *both published by Howard W. Sams. He can be reached on BIX c/o "editors."*

Secrets in Your Socket

Mr. Machrone? I was reading your columns on surge suppression and power management recently, and I represent a company that has some new and exciting ideas about power quality, and I was wondering if you would like to take a look at the product.

I love conversations that begin that way.

The company turned out to be Perma Power Electronics, hardly a newcomer on the surge-suppression scene. The product was the Socket Sentry Site-Wiring Quality Monitor and Surge Suppressor. What they were telling me seemed impossible. It looks like an ordinary power strip, but it has an active monitor that continuously checks the quality of your building's electrical ground.

I've written before about the importance of good grounding in the proper operation of computer equipment, especially networks. Bad grounding introduces a safety hazard. The "hot" side of the circuit has one connection at the panel. The "neutral" side has another connection. The ground is connected to the same place as the neutral line. For safety and freedom from noise, the neutral and ground lines should have no resistance. If there's any resistance between the two, a voltage differential can build up across the resistance and become a source of noise, and can even build up to hazardous levels, where you can get a shock by touching two pieces of grounded equipment that happen to be plugged into different outlets. Ground currents that can't find their way back to the electrical panel because of high resistance may find an easier path through the shield braid or ground wire on your network connections, through your modem, or through you. These errant AC voltages introduce errors onto the network, and, at worst, can destroy network equipment.

You might think that you could check ground resistance simply by sticking an ohmmeter into the neutral and ground sockets. Don't do it! You know that the round one is ground, but do you know which of the flat sockets is hot and which one is neutral? Furthermore, do you have any assurance that the socket wasn't wired backwards? I find sockets with reverse polarity entirely too often for my comfort. At best, your ohmmeter would blow a fuse; at worst, it would be charred plastic. (I have to admit that I have just such a charbroiled ohmmeter as not-so-living proof.)

There are safe ways to thoroughly check a socket for correct wiring, grounding, and other electrical errors. Ecos Electronics Corp. makes a range of power analysis devices, from the $200 Accu-Test II to the PAK-1B Comprehensive Power Analysis Kit, which, at $1,700, will tell you more about your electrical service than you ever wanted to know. SureTest has its models ST-1 ($100) and ST-1P ($180). The ST-1P (a favorite of fellow columnist John C. Dvorak) tests the resistance of the ground circuit in addition to checking for the usual wiring errors. The drawback to all these devices is that they're designed to verify the proper wiring of new construction or to conduct occasional site audits.

Postinstallation checks and audits would be fine if industrial electrical distribution systems were unchanging. Unfortunately, they're updated, revised, repaired, and expanded all the time. Mistakes creep in. Giant air-conditioning compressors wind up drawing from the same line you thought was dedicated to your server closet. Worst of all, the ground connection you depend on heavily for safety and electrically quiet operation is fragile and easily compromised.

Which is smarter, your computer or the power strip it's plugged into?

In truth, the grounding system in your home is probably better than the one at work. Modern home systems use three-conductor wiring to carry the hot, neutral, and safety ground lines. Modern commercial installations use three-wire, but the vast majority depend on thin-wall conduit to provide the safety ground back to the electrical panel. In home installations, metal electrical boxes are increasingly replaced by plastic, and the safety grounds are attached to the green screw on the outlet, not to the electrical box. Not so in commercial installations, where if there's a green wire at all, it's probably attached to the box, not the outlet. With the three-wire installation, you're depending on the screws that hold the outlet to the box to make the ground. With the conduit, you're depending on slip joints, collars, and screws as well as clean, corrosion-free metal-to-metal contacts for that safety ground. The number of things that can go wrong rises exponentially.

MORE THAN A FEELING

This is not just hypothesis. I loosened the screws holding the outlet to the junction box in the hall just outside my office at ZD Labs which is less than a year old. I plugged in the Socket Sentry, then wiggled the outlet. In some positions, the ground circuit rose above the 1.5- to 2-ohm threshold the Socket Sentry detects. It sounded off with a shrill alarm, causing fellow employees to come out of their offices and cubicles to find me cross-legged on the floor, poking a screwdriver into the near vicinity of the outlet. The things we do in the name of science.

Perma Power's little wonder has a big advantage over all the professional gear I mentioned earlier: You can leave it plugged in. Using the "rust never sleeps" dictum, your power distribution system is getting slowly, progressively worse, and some day it'll fall below safety standards. You may not do another site check for a year—most likely never. The day it goes too far, a Socket Sentry will let out a yell. It has a few other nice features, too: detecting reverse polarity, neutral-to-ground voltage faults, and failure of its own surge suppressors.

This latter trick is one of the best, since the "suppression OK" lights on the vast majority of surge suppressors are all but worthless. You can remove all of the MOVs from them, and the light will continue to glow merrily. The Socket Sentry actually shuts itself off if the MOV suppressors approach the end of their useful life, and will not turn itself back on. If an MOV fails to open, the test circuitry detects it and shuts down the device. Perma Power is confident enough that it offers a lifetime warranty. The suggested retail price is $169, but you should be able to find the Socket Sentry for $9[illegible] with savvy shopping. A flat-out bargain. □