RUDI VAN DRUNEN

the basics of power

Rudi van Drunen is a senior UNIX systems consultant with Competa IT B.V. in the Netherlands. He also has his own consulting company, Xlexit Technology, doing low-level hardware-oriented jobs.

rudi-usenix@xlexit.com

WE OFTEN TALK ABOUT SYSTEMS FROM a "in front of the (working) screen" or a "software" perspective. Behind all this there is a complex hardware architecture that makes things work. This is your machine: the machine room, the network, and all. Everything has to do with electronics and electrical signals. In this article I will discuss the background of some of the electronics, introducing the basics of power and how to work with it, so that you will be able to understand the issues and calculations that are the basis of delivering the electrical power that makes your system work.

There are some basic things that drive the electrons through your machine. I will be explaining Ohm's law, the power law, and some aspects that will show you how to lay out your power grid.

Power Law

Any piece of equipment connected to a power source will cause a current to flow. The current will then have the device perform its actions (and produce heat). To calculate the current that will be flowing through the machine (or light bulb) we divide the power rating (in watts) by the voltage (in volts) to which the system is connected. An example here is if you take a 100-watt light bulb and connect this light bulb to the wall power voltage of 115 volts, the resulting current will be 100/115 = 0.87 amperes.

This equation can be written as follows:

$$I = P/U$$

or, after performing some algebra,

$$U = P/I \text{ or } P = UI$$

where

P is the power (in units of watts [W])

U is the voltage (in units of volts [V])

I is the current (in units of amperes [A])

Note the use of U for the voltage here; this is commonly used to distinguish between the voltage at a certain point and the unit of voltage (volts, V). In the literature the symbol V is also used for both the voltage and the unit volts, which can be confusing.

Ohm's Law

To have current flow, the device you are connecting to a voltage source has to pose some resistance to the electrons that want to flow from one terminal to the other terminal. The more resistance the device has, the less current will be flowing through the device.

We calculate resistance using Ohm's law, which can be written as

$$R = U/I$$

or

$$U = IR$$

or

$$I = U/R$$

where

R is the resistance (in units of ohms [Greek capital omega, **Ω**])

U is the voltage (in units of volts [V])

I is the current (in units of amperes [A])

We can calculate the resistance of the light bulb just discussed by dividing the voltage (115 V) by the current flowing (0.87 A). This results in a resistance of 132 **Ω**.

Note that for a light bulb the calculated resistance is the resistance in the "on" or hot state. The "off" resistance can be very different.

Combining Ohm's law and the power law, we can calculate the power that a resistor as a load to a voltage source will convert into heat (or a motor will convert into both heat and mechanical power):

$$P = U(U/R) = U^2/R$$

which can be rewritten using Ohm's law to

$$P = I(RI) = I^2R$$

This might also be applicable to a power cable. We all know that if you use a too thin (too low a rating) power cable for heavy equipment, the cable will get hot and eventually catch fire on the points where the resistance is the highest (most likely at the points where the plug connects to the wire). If we take a standard power cable (12 gauge), it will have a resistance of 2.0 **Ω** per 1000 ft. If you connect a 800-W piece of equipment to it (a strand of 500 ft of extension wire, totaling 1000 ft of conductor), this cable will dissipate more than 300 W of heat. (The resistance of wire can be found at http://www.powerstream.com/Wire_Size.htm and seen in Table 1.)

Rating	Type	ohms/1000 ft	ft/ohm
12 AWG	stranded	1.65	606
14 AWG	stranded	2.62	381
16 AWG	stranded	4.17	239
12 AWG	Solid	1.59	629
14 AWG	Solid	2.52	396
16 AWG	Solid	4.02	249

TABLE 1: RESISTANCE RATINGS FOR COMMON GAUGES OF WIRE

We first calculate the resistance of the device: 800 W/115 V = 6.9 Ω. We then add the power cable resistance to it, for a total of 8.9 Ω. Now we calculate the current flowing through the complete system: 115 V/8.9 Ω = 12.9 A. So both through the device and the cord, we see a current of 12.9 A. With these values we can both calculate the power dissipated in the power cord and the voltage that drops over the power cord. A schematic drawing is shown in Figure 1. The power equals $(12.9)^2 \times 2.0 = 333$ W and the voltage drop is $12.9 \times 2.0 = 25.8$ V. Now as we connect to the mains of 115 V and we have a total voltage drop of 25.8 V (12.9 V per conductor) over the power cable, so only 115 – 25.8 = 89.2 V remains for the device.

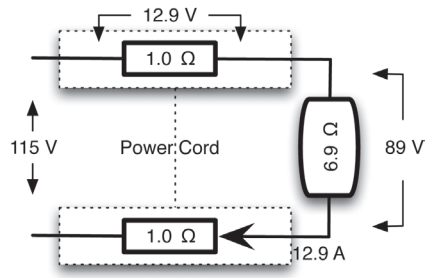

FIGURE 1: THE VOLTAGE DROP WHEN USING 500 FT OF 12 GAUGE EXTENSION CORD WITH 115 V WILL BE 25.8 V, LEAVING ONLY 89 V FOR THE DEVICE.

This calculation shows that it is important to have the correct rating of the power cable, to ensure that as much power as possible flows to the device where it does the work and to minimize the losses in the power cabling (which is just wasted heat).

Signals

Previously we assumed that the voltage applied to the resistor (or light bulb) was constant. However, the mains voltage in the United States varies between 162.15 V and –162.15 V. Figure 2 shows the sine wave for this alternating current (AC) voltage. A number of parameters can be derived. We have the top (maximum) value of the voltage (162.15 V), the average (over a integer number of periods) voltage (0 V), and the root mean square (RMS) voltage (which is the rated voltage of the mains, i.e., the effective voltage). For a sine wave the maximum voltage is $\sqrt{2}$ times the RMS voltage. If the curve is not a sine wave the factor is different. The ratio between the peak (max) value of a signal and the RMS value of this signal is called the crest factor. A direct current (DC) voltage has a crest factor of 1, and a pure sine wave 1.41 ($\sqrt{2}$).

To do power calculations as described previously, we use the RMS value.

We can determine the period and frequency. The frequency is defined as 1/period. With a 16-ms period, the frequency equals 60 hertz [Hz].

A signal can have both AC and DC components. The mains voltage has no DC component, but, as we see later, we can have an AC voltage biased by a DC voltage, as shown in Figure 3. The average of this signal (over an integer number of periods) is not 0, of course, and the RMS value equals the original RMS value (without bias) plus the bias voltage.

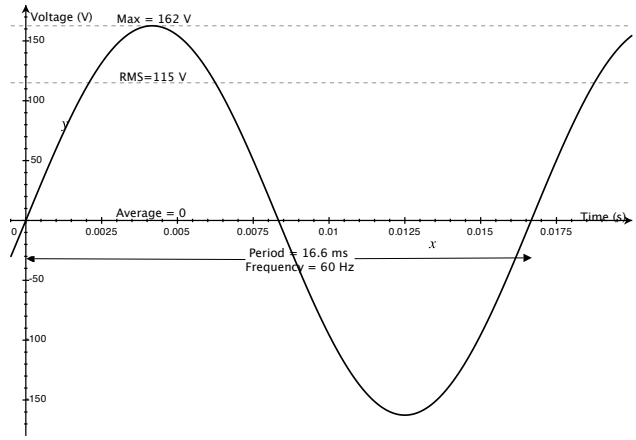

FIGURE 2: AC VOLTAGE SHOWING MAXIMUM AND MINIMUM VALUES ALONG WITH THE RMS VALUE, PERIOD, AND FREQUENCY

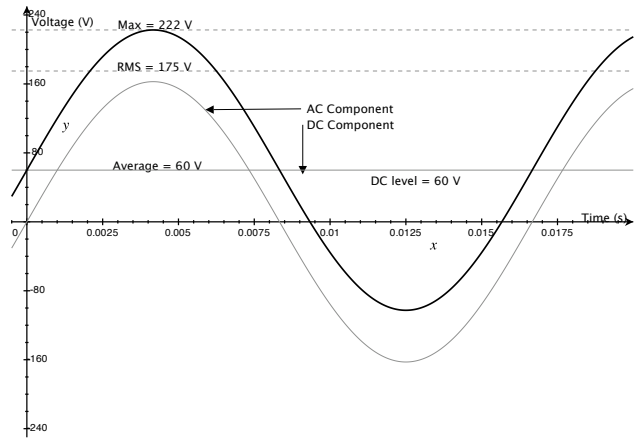

FIGURE 3: AN AC VOLTAGE WITH A DC COMPONENT

Components

In electronics we distinguish passive components (resistors, capacitors, and inductors) from active components (diodes, transistors, or more complex semiconductors). The active components can take actions, such as preventing current from flowing in one direction, or switching current on or off, whereas passive components cannot. In this article we introduce the three basic passive components. Figure 4 shows the corresponding schematic symbols.

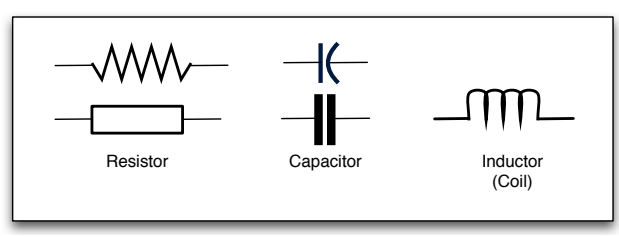

FIGURE 4: THE THREE BASIC PASSIVE COMPONENTS

RESISTOR

Earlier, I silently introduced the resistor as a component. A resistor is often a small component found in almost all electronic circuits, which reduces the

flow of electrons. You can compare it to a obstruction in a garden hose. The water (electrical current) cannot flow freely, and a pressure difference (voltage) builds up over this obstruction. Also the power cord acts as an obstruction. The unit of resistance is the ohm (denoted by the Greek letter omega, Ω).

CAPACITOR

A capacitor can hold electrical charge. Capacitors are generally very small when used in a circuit or take on (very) large forms in power systems as described later. If a current flows to a capacitor, it fills it with electrical charge over time and a voltage builds up over it. Compare this with a bucket that is being filled by the garden hose. The unit of capacitance is the farad (denoted by F).

INDUCTOR

The inductor is a basic coil. It acts about the same as a capacitor, but instead of an electrical field it builds a magnetic field inside the core, as a result of an applied voltage. An applied voltage causes a current to flow and this is used to build up the magnetic field; after the field is in place, current will flow through the inductor. Inductors are often used together with capacitors to filter signals or generate oscillating signals. We use the unit henry (H) for the inductance.

Complex Power

All of what I have discussed so far in power calculations is true if we have a perfect resistor as a load (e.g., a heating element or a lamp), but the calculations get more complicated when we have more complex loads. A piece of electronics has not only resistors but also capacitors and inductors. We can describe the behavior of a capacitor and an inductor by the following expressions: For an inductor,

$$U = -L(dI/dt)$$

and for a capacitor,

$$U = C \int I dt$$

where

L is the value of the inductor (in henries)

C is the value of the capacitor (in farads)

So, if we connect a circuit containing reactive components (capacitors and inductors) to an AC voltage source, some time will be needed to build up energy in capacitors, for example, so the result will be a time difference (phase difference) between the applied voltage waveform and the resulting current waveform. The stored energy [in electric fields (capacitor) or magnetic fields (inductor)] cannot result in "work" and flows back into the applied source. The factor between the current able to do actual work and the current that is used to build up charge (blind current) and not to perform work is called the power factor (*Pf*).

Real (pure) resistive loads (heaters, lamps, etc.) have a power factor of 1, and the voltage and current are in phase. But if we observe motors, for example, [containing coils (inductors)] we see that the power factor is below 1. In general:

$Pf = P/S$

where

P is the real power (in watts)
S is the apparent power (in volt-amperes)

We define the power factor as the cosine of the angle (phi) between the voltage and current [Pf = cos(phi)]. I need to explain the angle here. I define a full wave as 360 degrees, or 2π radians; hence we can relate a time period to an angle. Figure 5 shows the waveforms of the voltage, current, and power and their relationship over time: if phi = [1.8 ms (phi)]/[16.6 ms (full wave)] × 360 degrees = 39 degrees, the power factor Pf = cos(39 degrees) = 0.77.

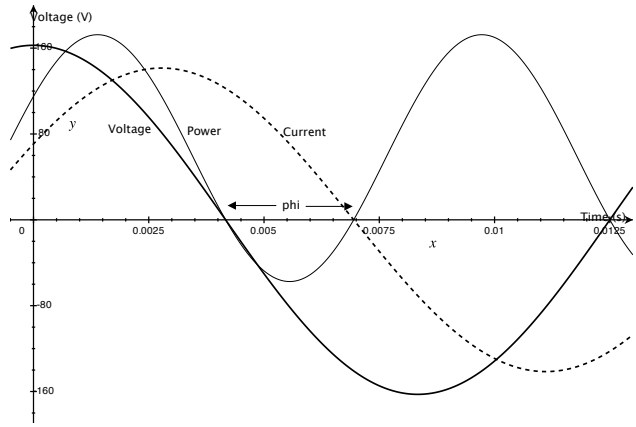

FIGURE 5: RELATIONSHIP AMONG VOLTAGE, CURRENT, AND POWER OVER TIME

If you have a load of 10 kW and the power factor is 1.0, your power supply has to supply 10 kW of power to the system. However, if the power factor is 0.6, the supply system in place needs to be set up for transport of 10/0.6 = 16.6 kW. The power company wants you to have your power factor as close to 1.0 as possible, so (as they bill you only for the power that actually does do work) they can design their systems and bill you for the power properly. Mostly in your contract with a power company it states that you (your datacenter) should have a power factor of 0.8 or better. The UPS that you probably have installed in your datacenter is specified in volt-amperes (VA), which is the total (apparent) power you need.

To get your power factor in shape (which generally means $Pf > 0.8$) you might have devices that contain capacitors and/or inductors installed to cancel out (part of) the effects of the loads you have. Mostly they consist of some measuring circuitry to measure the phase angle between the voltage and resulting current and a system to add capacity or induction to the system.

The power factor of most modern computer equipment (with switched mode power supplies) is close to 1, but other equipment used in the datacenter (for, say, cooling or ventilation) involving large motors can change the power factor dramatically.

Multiphase Systems

Generally, the power company delivers your power in three phases. The wire that enters your building (not residential) contains three or four conductors, each (of the three) carrying a single phase of 120 V power. A common return line can be provided. The 120 V is measured against this common return line. The sinusoidal voltage on each of the three phases is shifted 120 degrees.

The voltage between the phases equals $\sqrt{3} \times 120 = 208$ volts.

Using three phases is a more efficient way of transporting and using electrical power. A motor can more easily be driven by a magnetic field that rotates by itself. This can be accomplished by using three or more coils in the static part of the motor that are connected to the different phases. As the different coils are driven 120 degrees apart electrically, the magnetic field will rotate accordingly and drive the motor. Large motors (larger than 5 kW) are only available in three-phase versions. Using three-phase systems for large loads can be up to 75% more efficient than single-phase systems.

You can use each individual phase to act as a single-phase supply and connect part of your racks to that by using a step-down transformer (converting 208 to 120 V, as often the "common return" line is not available). The power company will request that you distribute your load evenly over the different phases. Some three-phase UPS devices will do that for you.

In a residential environment a number of buildings are connected to one phase, using a converter transformer often found on the pole, yielding an evenly distributed load.

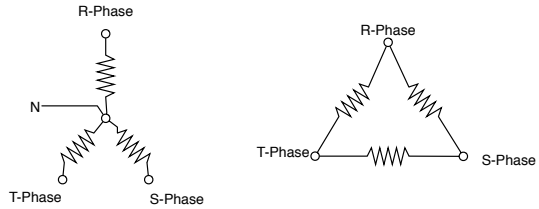

FIGURE 6: DIFFERENT WAYS OF USING THREE-PHASE POWER IN MOTORS

There are two different wiring schemes for three-phase systems, called star and delta (triangle) configurations (see Fig. 6). These configurations are the layout of how the coils in a motor are connected to the phases of a three-phase system. The star configuration is used to start the motor and deliver high torque; after that some smart electronics switch over to delta mode when running. As the different phases (coils in the static part of the motor) are driven with a time lag (see Fig. 7) the magnetic field inside the motor rotates, causing the rotor to rotate. You will see this in large UPS systems, flywheels, and air-conditioning systems comprising large motors. Most "normal" computing systems will connect to one phase. However, large machines and disk cabinets will eventually use three-phase input and often use the three phases to feed three power supplies.

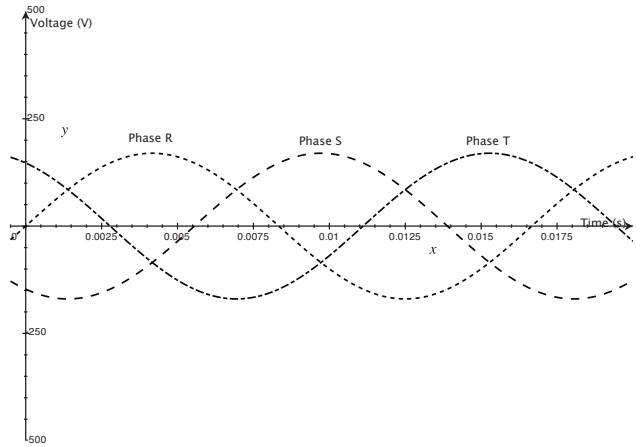

FIGURE 7: THREE-PHASE POWER VOLTAGE GRAPH

Grounding

Each system (computer, rack, etc.) should have an adequate connection to the ground wire. Exceptions to this rule are some double-insulated consumer devices, which we will discuss later. The ground wire (often a green/yellow striped wire) is connected to an electrode that is physically driven into the earth. The ground connection prevents the physically exposed conducting parts of a device from getting exposed to dangerous potentials in case of a failure of the insulation. A special kind of circuit breaker will signal such a fault (by comparing the current in the power line with that in the return line and noting any difference) and switch off the power. This is commonly called a Ground Fault Interrupter (GFI).

Also, the connection to ground prevents the buildup of static electricity (e.g., from air friction from the fans inside a cabinet). I will discuss this in a later article in which we evaluate static electricity and the issues concerned with it.

A common issue is the use of so-called transzorbs (overvoltage protecting devices) in power supplies. These devices make a short circuit from the mains line to the grounding pin when a power surge occurs, thereby protecting the device from getting fried. These devices can go bad (which often happens when they have already done their job once) and leak just a little of the mains voltage to the ground and the metal chassis. They then will act as a capacitor and have a noninfinite resistance to AC voltage.

If this is the case (and you might have an improper connection to earth) your metal case will carry (part of) the mains voltage. This is particularly dangerous if you touch both the case and, for example, a central heating pipe (which is supposed to be grounded), or if you connect this system to another system that is properly grounded.

So, in short, make sure all equipment is securely grounded, including the racks. Be very careful to have the ground connection itself done properly.

Ground also refers to a reference level against which all voltages are measured. In a computer system we have the ground level, which represents the 0-V rail on the power supply, and the mains ground, which is connected to the chassis. They may or may not be tied together. (If not, we refer to this as a system with a "floating ground.") Sometimes there is an AC coupling (capacitor) between the circuit ground and mains ground to get rid of high-frequency noise on the power line.

Adequately insulated equipment, such as some consumer equipment, is allowed by the regulatory body not to have ground connections. Often this equipment does not have a metal case or metal knobs.

GROUND LOOPS

Ground loops occur when the ground (0-V rail) potential of one device is not the same as the ground potential of the connecting device. Because of the resistance in one of the ground leads and the current flowing through it, the 0-V rail in device A will not have the same potential as the 0-V rail of device B. Figure 8 (next page) shows the voltages in time at two points in the circuit. Device A "sees" a different voltage compared to device B.

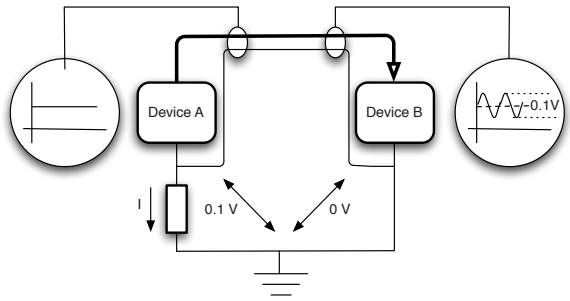

FIGURE 8: PRINCIPLE OF A GROUND LOOP

Figure 9 shows the circuit diagram, where the output voltage is not the same as U2 (as you would expect) but is the addition of U2 and the voltage generated through the ground current Uground:

$$Uout = U2 - Uground = U2 - U1[Rground/(Rground + R1)]$$

$$= U2 - (I1 \times Rground)$$

Often, as the bias voltage is a result of the mains (ground) connection current, the bias voltage will be a sinusoidal waveform with the mains frequency. This is the "hum" you can encounter when connecting two mains-powered devices to each other (such as a computer and an audio amplifier). In a video system we can see this as bars scrolling vertically over the screen.

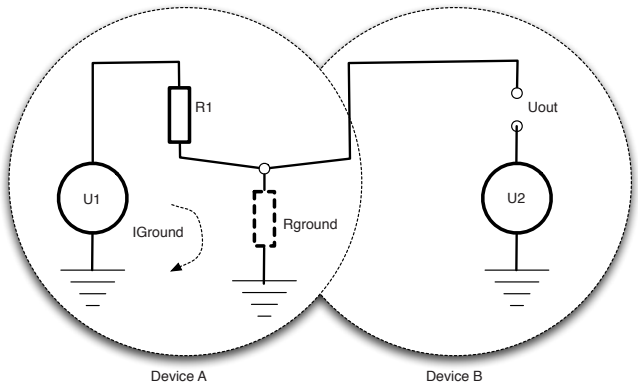

FIGURE 9: SCHEMATIC DIAGRAM OF A GROUND LOOP

To prevent ground loops, there should be one point at which a complete system is connected to ground, so we end up with a star topology. In practice that means that you should plug all parts of a system into the same outlet (keeping in mind that you do not want to overload the outlet or power cord).

Another source of hum, to be explained in a later article, is that from the bias signal as a result of the magnetic and electric field around a conductor that carries a current. When you run a power cable along an audio cable, the electromagnetic radiation coming from the current flowing through the power cable can induce a voltage in the low-signal (audio) cable. Therefore audio and other small signal-carrying cables are often shielded [i.e., have a metal shield that is connected to (signal) ground surrounding the conductors].

Conclusion

In this article I have described the way to calculate power requirements, how to lay out power cabling, and the various other issues that become important when scaling up from a couple of boxes to a datacenter. Using this

knowledge, you can understand why it is important not only to build a nice software design or a systems design but also to take into account the way you connect these together and have the electrons flow through them. This is the first of a series of articles in which I will give some background on systems hardware, with the purpose of bringing you some more insight into what is happening behind the faceplate of your system, with the ultimate goal of helping you troubleshoot some of the hardware problems that you may encounter.