HARDWARE HACKER

Using Double-Layer Capacitors; Computer Power Interface, Underground Radio Technology

By Don Lancaster

How Big is a one-Farad capacitor? Would you believe less than a cubic inch? Conventional wisdom and poorly done cartoons notwithstanding, a new capacitor technology is now available that crams incredible capacitance values into very small volumes. As Fig. 1 shows, these are called double-layer capacitors. They are both reasonably priced and available from distributor stock. For instance, a 1.0-Farad, 5-volt *NEC # FZ0H105Z* is available from *Mouser Electronics* for $4.59 each (part 551-1.0Z05).

Being a brand new and "hacker friendly" part, there's lots of exciting new possibilities here. But there are also some definite use limits and restrictions.

One major use for double-layer capacitors is as power backup for CMOS memory chips. The capacitor can fake memory nonvolatility during a short power outage. Unlike most batteries, these devices can be cycled indefinitely, have at least a 10-year life, can easily be soldered in place, are not polarized, usually cannot leak or explode, and readily indicate their charge.

Back to the basics. A traditional capacitor consists of two conductors separated by an insulator. When a current is routed through a capacitor, a charge will build up and will be stored. The amount of the charge depends on the current and how long the current is applied.

The charge also depends on the surface area of the conductors, the spacing between conductors, and a property of the insulator material known as the dielectric constant.

Capacitors have many electronic uses, ranging from storing dc in power supplies; coupling/bypassing ac signals; resonating coils for frequency selection; and for integrating precision waveforms.

The storage size of a capacitor is measured in Farads. A 1-Farad charge will supply a current of 1 ampere for 1 second. Most capacitors have traditionally been much smaller than this, and are measured in microfarads (μF) or picofarads (pF) instead of Farads.
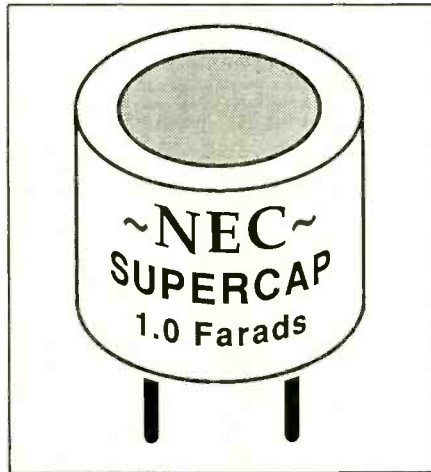

Fig. 1. A 1.0-Farad capacitor!

For instance, if you hold two fingernails one fingernail apart, you will have created the equivalent of a 1-pF capacitor. Among other difficulties, there are not nearly enough people in the world to build a 1-F capacitor by this method.

Some ways to get more capacitance are to increase the plate area, to decrease the spacing between the plates, and to raise the dielectric constant. All of these techniques have been used in combination to create the various capacitor families available today.

For example, ceramic capacitors use special insulators with a dielectric constant of 8000 or more. Electrolytic capacitors build up an extremely thin insulating oxide film. They also increase the conductor surface area by etching. Sometimes a double etch is used where big pits are first added to the original surface, and when this is done, small pits are added to the big pits.

But the double-layer capacitor goes far beyond what traditional capacitors are able to do. It uses a fundamentally different principle, first discovered by Helmholtz over a century ago. It turns out that if you have a boundary between a solid and a liquid, a double layer of charge builds up at the surface between the two. To build a useful double-layer capacitor, you use a solid plate, a liquid electrolyte and a second solid plate. To prevent the two solid plates from touching, you place

a porous, insulating separator between the two. The separator allows ions in the electrolyte to move back and forth as charge is accumulated, but prevents any direct short-circuit between plates.

The solid plates used for today's double-layer capacitors are made from activated charcoal which has an incredible surface area, since it is made up of countless tiny particles. A gram of activiated charcoal may have a surface area equal to one-quarter of a football field! And therein lies the secret to extreme capacitance values in a very small package.

In exchange for the extreme capacity in a very small space, there are some definite tradeoffs. The breakdown voltage of the electrolyte used in today's double-layer capacitors is only 1.2 volts. To get a useful +5-volt dc working voltage, six capacitors are placed in series inside the package.

Instead of directly building a 1-F capacitor, there are really 36 farads of capacitance in the package, since six series connected capacitors each yield a 1-F series equivalent capacitance.

So, you *must* be careful to never exceed the recommended working voltage of a double-layer capacitor.

In addition, these are not precision components, since activated charcoal itself is not a very precise material. Tolerances are in the −20% to +80% range. Also, charge is inherently an "undense" way of storing energy, compared to chemical conversions. Large single-layer capacitors are simply not in the same energy storage league as batteries, and I don't thing you will ever see the day when you start a lawnmower with a double-layer capacitor.

So, while double-layer capacitors are incredibly better than regular capacitors for storing bunches of charge in a small volume, they are not nearly as good as a battery when mucho dense energy storage is required.

Finally, charcoal is not a very good conductor, and conductive elastomers are also used in the construction that are only moderately conductive. Thus, a double-layer capacitor is *not* a low-impedance device and can *not* be rapidly charged or discharged.

One interesting side effect here is that the inherently limited charge and discharge rates of a double-layer capacitor give automatic short-circuit protection. Some batteries, especially Ni-Cds, instead tend to get downright violent when you short them.

Several manufacturing options are available that trade off maximum possible charge and discharge currents against cost and size. The smallest and cheapest versions are intended for long-time, low-current memory backup uses.

Let me know what interesting uses you can come up with for this exciting, yet imperfect, new hacker component.

OK. How did you do it?

If there's one thing I can't resist, it's a good hack, especially if it lies well beyond the far side of "disgustingly elegant." Check the capacitor lettering in Fig. 1. Note that we not only have isometric lettering, but the lettering is also correctly wrapped around a cylinder! This was quickly and easily "drawn" with the stock and standard Apple Writer word processor, as were *all* parts of *all* figures in this column. Oh, yes, I did add a short little custom text routing of my own that I'll be happy to share with you. The same routine also lets you set text in a helix or an inwardly closing spirial.

Write or call for you free copy.

How can I Control AC Power With a Personal Computer?

There are two problems involved in controlling high-power ac loads with a personal computer or microprocessor parallel port. The first is that the weak port signals must somehow be *amplified* to make them "strong enough" to control a high power load. The second hassle is that there is a very deadly shock hazard involved with ac power control, which means that some sort of *safety isolation* is an absolute must.

While you could use a relay driver IC and a relay, the usual way to handle 110-volt ac high-power interface today is to connect a phototriac optoisolator to a triac power control device. Figure 2 shows details.
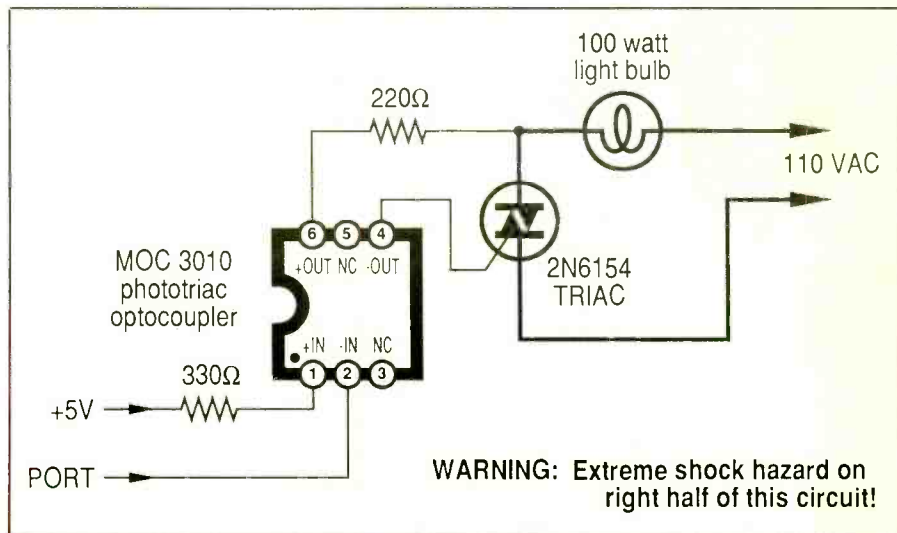

WARNING: Extreme shock hazard on right half of this circuit!

Fig. 2. A 110-volt ac power control interface for a personal computer.

The triac optoisolator consists of a light-emitting diode that drives a photosensitive ac power control device called a phototriac. When the LED lights, the phototriac turns and activates the main high-power triac to switch the load.

Because only a light beam connects the input to the output, you get total safety isolation. The "amplification" you need comes from both the internal phototriac and the main power triac.

There are some important details you have to be aware of when you use this circuit. First and foremost, note that there is a deadly shock hazard anywhere to the right of the triac optoisolator; so use extreme caution when working on *any* circuit of this type.

The power triac will switch on or off most any ac load, such as lamps, motors, heaters, or whatever. A minimum load must always be provided for proper operation, say 7 watts or so. The maximum load depends on the triac and the size of the heat sink you use.

One classic handbook for triac power control is the *SCR Manual* from *General Electric*. My issue is sort of dated, so it may have a new name and a new price by now. Cost should be around $5.

You also *must* use a phototriac optoisolator, rated at least 200 volts for 110-volt ac operation, or 400 volts for

220-volt ac operation. Note that the other types of optocouplers will instantly self-destruct if you try to use them for ac power control. The *Motorola* MOC 3010 is often a good choice and costs slightly over a dollar.

I have also shown the inputs to this circuit connected "backwards" from some other circuits you may have seen. There is a very important reason for this. Most parallel ports on microprocessor systems are much better at sinking current to ground than they are at sourcing current from a positive voltage.

With this circuit, the output port either sinks current to ground, activiating the phototriac optoisolator and powering the load, or doing nothing and not delivering power to the load. Connected "backwards," the optoisolator is much better matched to the parallel computer port and will work much more reliably.

One little gotcha with this "backwards" connection: Your port will now be active low, so a zero at the port will *activate* the load. Your software must, of course, recognize this fact, and write zeros to the port locations when you want to power the load.

Note that a machine-language command of EOR #$FF will automatically convert all of the ones in a word to zeros and vice versa. From BASIC, an XX =

(255 – XX) will also convert all ones in a word to zeros or vice versa. Details vary with the application, of course.

Typical cost of one channel of power control should be under $3, unless you are using a very large triac.

More info on microcomputer power interface appears in my *Micro Cookbook*, Volume II (SAMS #21829).

Is There an Easier Way?

Well, if you want to pay a little more, you can cut the parts count to one. Figure 3 shows how to use a Crydom D2W202F solid-state relay by *International Rectifier* and do all of your power control with a single part. *Newark Electronics* sells them for $7.95 in singles as stock #81F4904.

The circuit is still the same as before, since the resistors, triac and optocoupler are all inside the package. Also as before, the + input goes to the microcomputer's +5-volt supply line, and the – pin goes to the port output. A zero turns on the load, a one turns it off.

The device is good for 2 amperes and might be pushed a little further with some sort of clip-on heat sink. Around 220 watts of light bulbs, or possibly a quarter-horsepower motor can be controlled. Note that a severe shock hazard exists on the power side of this circuit.

What About Proportional Control?

If you have a proportional control interface, you not only can turn on/off motors, heaters, and lamps, but you can control their brightness, temperature, or speed as well.

It turns out that the Fig. 2 and 3 circuits can also be used for proportional control, but the software gets more complicated. And you have to be aware of exactly what you are up to.

First, the load you use *must* respond to voltage in a proportional manner. Incandescent lamps will, but fluorescent lamps will not. Most heaters or soldering irons will respond proportionally and present no problems.

Universal type motors (those with brushes) will alter speed with changing
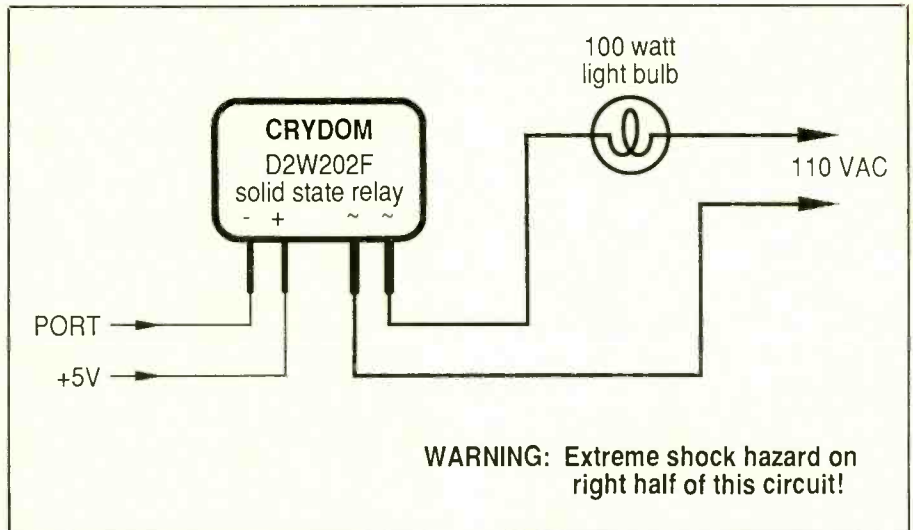
Fig. 3. A simpler but more costly power interface option.

voltage, but the more common ac induction motor can destroy itself if you try voltage control. Even when you control the speed of a universal motor by varying its voltage, you may get very limited torque. There are fancier techniques involving feedback that can be used to get around this problem.

So, your initial experiments at proportional power control are best done with incandescent lamps, heaters, etc.

The trick to proportional power control is to deliver a turn-on pulse that is precisely delayed from each ac-line half cycle. Figure 4 shows details.

If turn-on occurs very early in each ac half-cycle, nearly all the power will appear across the load. If turn-on takes place in mid-cycle, around half of the power gets to the load. Shoud the turn-on occur very late in the cycle, only a small amount of power gets to the load. Finally, with no turn-on pulse at all, the load gets no power and stays off.

Your computer must somehow sense the beginning of each half line cycle. This can be done with a line-derived interrupt signal or by routing something similar to an input port. The software inside the computer or microcontroller then will
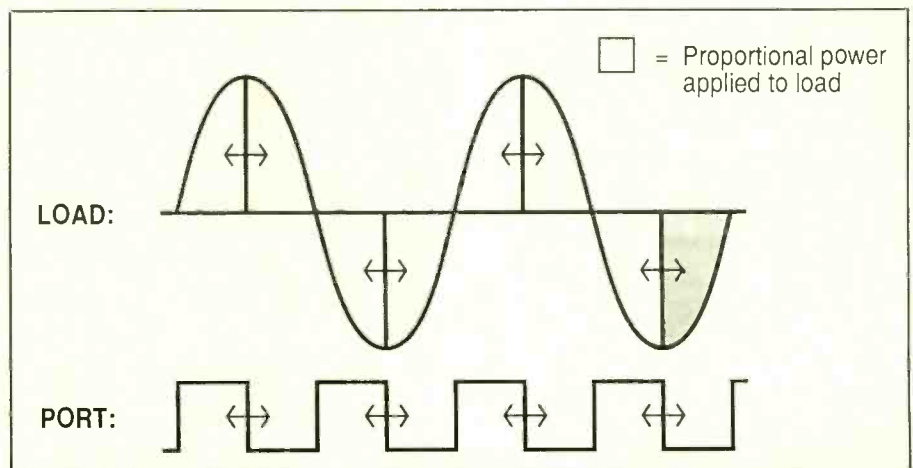

Fig. 4. Proportional power control requires precise timing.

It's called *Speleonics* and is published quarterly (four times a year).

Included in each issue are a dozen how-to articles and a subscriber list of others interested in similar underground tools, concepts, and techniques. Price is — better sit down for this one — a dollar per issue or $4 per year!

Should you want information on caving and speleology in general, you can instead contact the *National Speleological Society* and they can put you in touch with individual local cavers and an NSS grotto or two in your area. **ME**
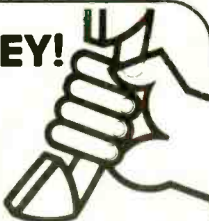
automatically convert a speed, brightness, or temperature value to the proper turn-on delay.

The timing is such that machine-language code *must* be used, particularly if you are controlling more than a few loads with the same computer. The precision and linearity of your outputs can be greatly improved by using some sort of feedback. This, of course, ups your system complexity.

You should be able to use a computer the size of an Apple IIe to control something like 32 theater lights to 32 brightness levels. This would make a very interesting programming project.

A totally different way to handle proportional control is with those *BSR* power modules sold by *Radio Shack* and others. These are fairly easy to interface to most computers, besides offering remote-control options.

Is There any more Information on Underground Radio Technology?

I've found another most interesting publication on underground technology. This beauty covers such things as underground radio communications and position measurement, ultrasonic and laser rangefinding, self-climbing ladders and even automatic spare-lamp switchers. All hands-on and real-world with full schematics, parts lists, the whole bit.