

THE COMPUTER JOURNAL[®]

For Those Who Interface, Build, and Apply Micros

Issue Number 4

Vol. 1, Mo. 4

\$2.50 US

Optoelectronics page 2

Multi-user page 10

Making the CP/M User Function
More Useful page 11

Part Three:
Build a Hardware Print Spooler page 12

Beginner's Column

Part Three: Power Supply

Anyone for a Little "KISS" Electronics? page 15

Editor's Page

Is The Home Computer Fad Over?

"We'll put a computer in every home" was the motto of the mass merchandisers as they visualized the money to be made in microcomputers. They did succeed in placing a lot of low-cost computers in homes, but many of these computers are now gathering dust in the closet and the home computer market is stumbling. What went wrong? Is there really a place for a computer in every home?

It is common to make marketing errors in a rapidly exploding area such as we have seen in the recent microcomputer field. Typically, the field is started by technically involved individuals with a strong entrepreneurial spirit. These people are knowledgeable of the field, and put out a product well suited for other people with similar interests. These little businesses are often successful because they understand the product and the people who will be using the it. As the market grows, it attracts the attention of venture capitalists who make spectacular profits thru investment in some of these young companies. Of course, some also lose their investment, but this is little noticed in a rapidly expanding market.

Eventually the established large corporations decide that they should jump into the market and skim off the fantastic profits which those upstarts with no business experience have been rolling up. It is obvious that if those businesses started in a garage with no business experience and little capital can make such large profits, then the large corporations with years of experience, lots of money, and dozens of vice-presidents, can surely do even better. Or is it?

The original microcomputer related businesses supplied products for a reasonably well defined market. They were used by hobbyists, small businesses, technically advanced people, and for games. These people either had some use in mind for the computer, or they purchased it for the challenge and the opportunity to learn about something new. Early products were designed by technicians for technically orientated people, and sold well even though the products were clumsy and hard to use. The early buyers were willing to put up with poor documentation and hard to use programs because they really wanted to use the computer, and they had some kind of use for it.

When big business discovered the micro and decided that there should be a computer in every home, some people wondered what everybody was expected to do with *all* those computers. The answer was that they could play games, balance their checkbook, and write letters. The computer industry failed to realize that playing games only lasts so long, and that vast numbers of the population rarely write letters and don't even have a checking account. Industry should have taken a look at what people are doing without a computer, and presented a product that would help people do these things better. Someone whose only interest is spending their evenings in front of the boob tube with a six pack is not very likely to make use of a computer, while someone who is active in other interests will probably make use of a well designed computer/software package designed around their interests.

The computer industry has been so busy developing home computer products that they have not had the foresight to realize that before computers can be really useful (or essential) in the home, they must first create the the need as well as the software to fill the need. Up to now, most computer applications have been designed either for people interested in computers or for businesses. Now is the time to develop uses for the average home.

Editor/Publisher..... Art Carlson
Art Director.....Joan Thompson
Technical Editor..... Lance Rose
Technical Editor..... Phil Wells
Production Assistant..... Judie Overbeek

The Computer Journal^{P1} is published 12 times a year. Annual subscription is 224 in the U.S., \$30 in Canada, and \$39 in other countries.

Entire contents copyright © 1983 by The Computer Journal

Postmaster: Send address changes to: The Computer Journal P.O. Box 1697, Kalispell MT 59903-1697.

Address all editorial advertising and subscription inquiries to: The Computer Journal P.O. Box 1697, Kalispell MT59903-1697.

OPTOELECTRONICS

by Roger Johnson

Light is playing a role of ever increasing importance in the field of electronics. In order to learn how to use optoelectronic devices, it is essential to understand some of the characteristics of light. Once these fundamentals are known, they will aid in the understanding and design of optoelectronic circuitry. At the end of this article you will find several good references available at no or low cost. I believe these are the finest aids available to the designer of good, practical opto-electronic circuits.

Light is electromagnetic (EM) energy that humans are capable of detecting. Other portions are invisible to us, but are capable of being detected by antennas, solid state crystals, or by thermal means. All EM energy consists of packets of energy called photons. Photons are best thought of as particles of energy with wave-like properties. To go into deeper discussion of this wave-particle duality would be to go beyond the scope of this article. The interested reader will find plenty of history and an account of the discovery of the photon in any good college physics text.

Since EM energy travels at the speed of light, there is the familiar equation which relates to all wave propagation. It is:

$$c = \text{speed of light in meters/sec}$$

$$c = f \times X \quad f = \text{frequency in Hz}$$

$$X = \text{wavelength in meters}$$

In most application literature you will find wavelength given in units of nanometers, microns, or Angstroms. The relationship between these units is:

- 1 nanometer is 10⁻⁹ meters, abbreviated nm
- 1 micron is 10⁻⁶ meters, abbreviated μ
- 1 Angstrom is 10⁻¹⁰ meters, abbreviated A


Humans see in the wavelength region ranging from 350 nm (violet) to 700 nm (red). There are other units of interest when describing light or EM waves. One is intensity. Because EM energy radiates spherically from a source, it is important to distinguish between the total power of the source, and how much is radiated in a given direction in space and intercepted by a detector. Names and definitions have been established for use in the field of optics, and are encountered in product literature. We will first direct our attention toward understanding these definitions. Units in the radiometric system measure how a perfect instrument detects EM energy, and photometric units measure how humans detect it. See table 1.

Let us now pictorially describe how each of the radiometric units are established. In Fig 1 we show a 1W

RADIOMETRIC UNITS		
Name	Description	Units
Radiant power	rate of transfer of radiant energy	watts (W)
Irradiance	radiant power per unit area incident upon or leaving a surface	W/cm ²
Radiant intensity	radiant power from a point source per unit solid angle. Another name for solid angle is steradians (sr)	W/sr
Radiance	radiant power incident upon or leaving a surface per unit solid angle per unit area	W/sr/cm ²
PHOTOMETRIC UNITS		
Luminous flux	rate of transfer of luminous energy	lumens (lm)
Illuminance	luminous flux per unit area incident upon or leaving a surface	lm/cm ²
Luminous intensity	luminous flux from a point source per unit solid angle	lm/sr- candela (cd)
Luminance	luminous flux incident upon or leaving a surface per unit solid angle per unit area	lm/sr/cm ²

Table 1

point source radiating EM energy into space. For the time being, let's assume that it is operating at one particular wavelength. You will see later how complicated it gets if this assumption is not made. In Figure 2 we show a detector at some distance R away from the source. The irradiance at this point is the total power divided by the area of a conceptual sphere of this radius. That gives power per unit area. Note how irradiance falls off as 1 over the distance squared. In Figure 3 the idea of "solid angle" is presented. Here a cone with its apex at the source is drawn. Solid angle is defined as the cone generated by a line passing through the apex and a point that is moved along the surface of the sphere that is generated. The solid angle, which is measured in units called steradians, is equal to the area intercepted by this cone in this imaginary unit (R = 1) sphere. Thus the total solid angle around a point in space is 4pi steradians. For our 1W source, the radiant intensity is 1W/4pi steradians. Note




that this definition does not depend on distance.

It turns out that it is useful in opto-electronics to be able to calculate the solid angle that a receiver or detector "subtends." There is a simple equation derived from calculus and the definition of solid angle that gives solid angle in terms of the more common linear angle. Figure 4 shows how it is used. The solid angle that an object subtends is:

$$\Psi = \pi \theta^2$$

θ = the angle in radians
 Ψ = the solid angle in steradians

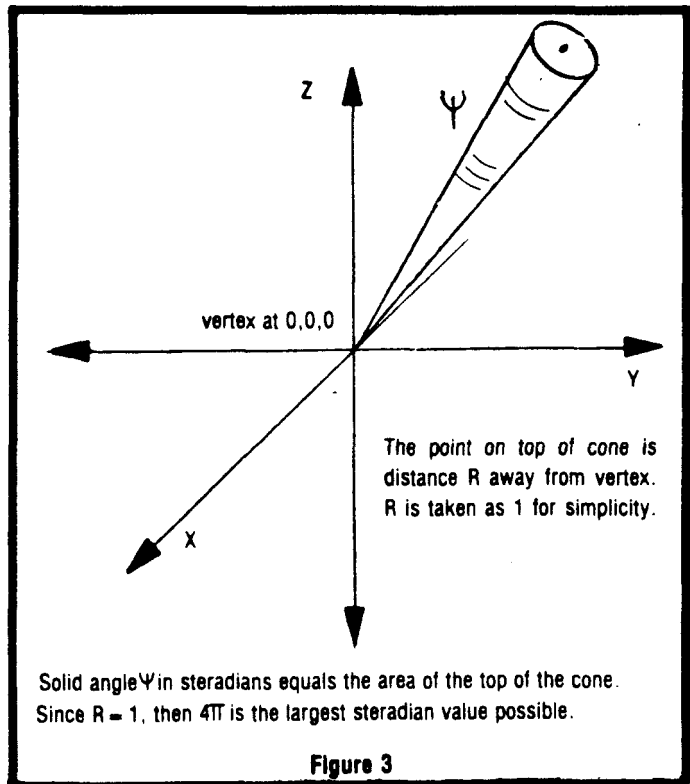
This is an approximation that is good if we are at least ten detector dimensions away from the source. This is almost always the case in communications via optics anyway. This formula is accurate to within 1% for cone half-angles of less than 20 degrees and to within 1% for cone half-angles of less than 6 degrees. Be sure to note how the half-angle is described and used in Figure 4. This is a very commonly



used term in opto-electronics and is widely used in fiber optics.

The photometric units seem almost identical and you might well wonder why. As I stated before, photometric quantities are not strictly physical, but involve how humans detect light. For example, humans "see" some wavelengths better than others. For two sources having the same power output but operating at different wavelengths, one will seem brighter than the other. Figure 5 shows the internationally accepted eye response curve for humans. Note how poorly red is seen compared to green. Now you can see why it is so complicated to convert back and forth between radiometric and photometric units; if the source operates over a large range of wavelengths, then you must "weight" every wavelength with its appropriate response. With the aid of a computer, you might be able to calculate the response every 0.1nm. The sum would be the total response, but it would be a tedious task.

An arbitrary standard has been established for photometric units. These are the units routinely used in



opto product literature. 60 candela (cd) is defined as the luminous intensity of one square centimeter of platinum at solidification (2315 degrees Centigrade). By definition, this is equal to:

$$1 \text{ candela (cd)} = 1 \text{ lumen per steradian}$$

also,

$$1 \text{ lux} = 0.0001 \text{ lumens per square cm}$$

We will now leave this seemingly bewildering world of units. Optical engineers spend a lot of time designing sources and detectors with specified optical characteristics, and even they sometimes get units and definitions confused. But there has been a purpose in exposing you to these

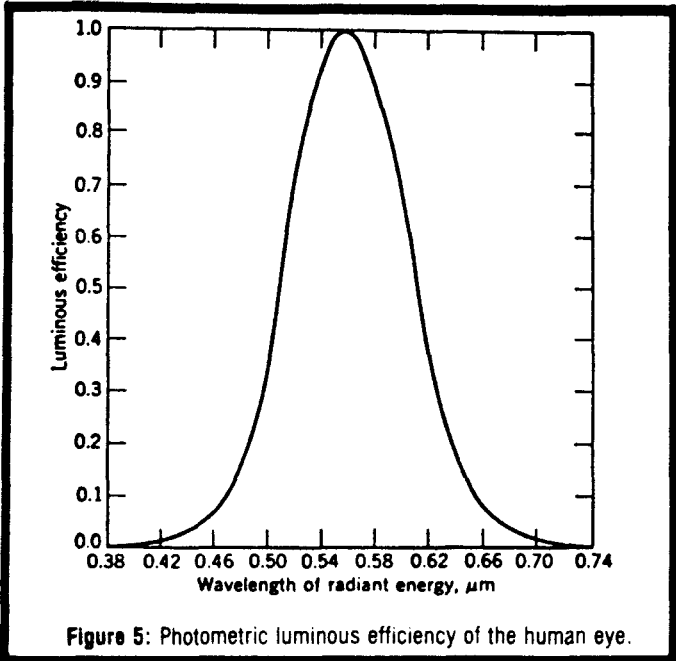
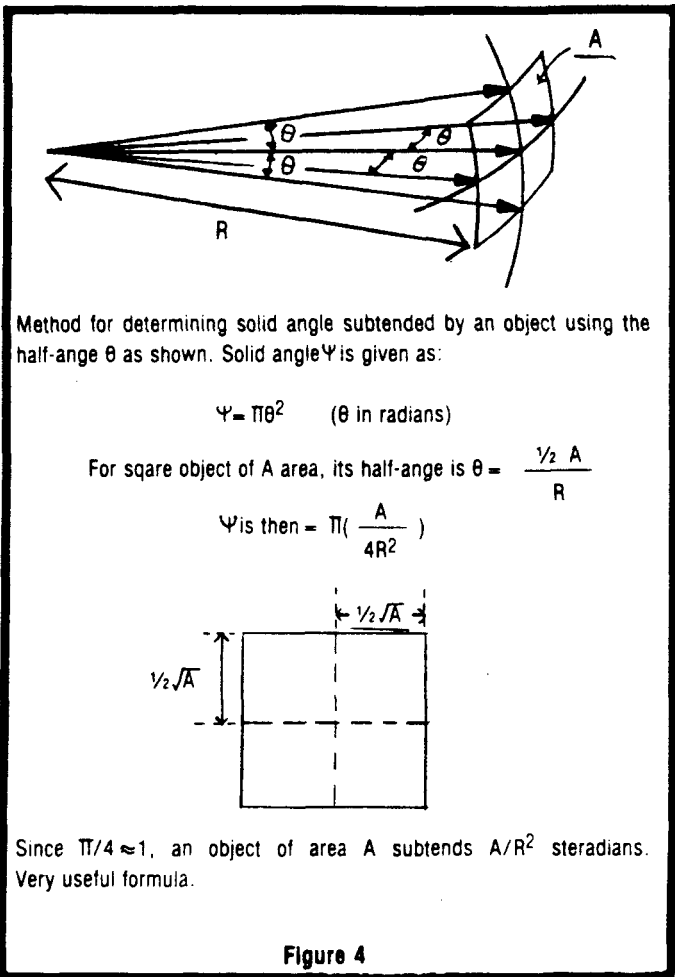
terms. Manufacturers have not yet agreed on a common set of terms, so you must be somewhat adept at converting back and forth among them. And Beware! Some firms give an emitter's luminous intensity figures in the axial direction, that is, straight out from the mechanical axis of the case. It may have little energy beyond a 20 degree half-angle, so be sure to look at angular displacement vs intensity curves for the source or detector of your choice.

*s sort of an aside, optics is much more a mechanical world than electronics in many ways. For example, in choosing a resistor, all you have to be concerned with is its resistance and power value. But with a detector or lens or emitter, you also have to be concerned with what direction it is pointing. Is its surface clean? What is the quality of its surface? Is it mounted securely? Is excess light hitting it? The list goes on. This is not meant to scare you, but to give you an appreciation for the kind of thought process that goes on in opto-electronics. Most of the time it is far less demanding than this, but you should be aware of these characteristics - it pays off handsomely when designing and troubleshooting.

Finally, Table 2 gives some cd values for some sources and backgrounds. Note the position of the ubiquitous LED.

Uses of Opto-electronics

Just where does electronics use light? There are 3 main areas: displays, measurement, and signal transmission.



•Displays

The LED (light emitting diode) is used by itself as an indicator, or several can be grouped together to form alphanumeric and bargraph displays. Their advantages are: long lifetime, resistance to shock, and little generated heat. LEDs can also generate infrared (IR) light. Liquid crystal (LC) displays use the action of twisting long chains of molecules to stop light or to admit it through a front polarizer. They use extremely low amounts of power, unlike current-hungry LEDs. Typical values for LC displays are 1 microamp per square centimeter of display area. They are coming down in price and a plastic LC display has recently been introduced. It will replace the glass types in most applications at a much lower cost. Also, more colors are available than with LEDs.

Of course there are the popular video displays which use electrons striking phosphor to generate light. A display using a neon gas-filled tube with wire stretched in the horizontal and vertical directions is called a plasma display. A potential exists on both wires which is below the ignition voltage of neon. When a small voltage which rides on top of this bias voltage is impressed on both the horizontal and vertical wires, the neon will ionize where they intersect and stay on even after removal of the two WRITE voltages. Because of this type of non-raster display, memory requirements are smaller and there is no flickering of the display. Vacuum florescent and incandescent displays are other examples.

Lasers are growing more and more popular for both display generation and for document scanning and reproduction. Lasers which are scanned by fast moving mirrors or by acoustical beams in crystals produce displays for light shows and in scene generation for aircraft flight simulators. Laser printers using the Xerographic process can scan a document, turn it into binary numbers and reproduce it elsewhere. Even more impressive is the ability to "create" a document without an original. An electronic

database does type font generation and graphics. It then scans the toner drum inside the machine with a scanned laser beam, creating thousands of copies. As an interesting practical example, a restaurant could generate an entirely different yet current menu every day of its operation. This type of printer exists now.

•Measurement

Using a source-detector pair, time, speed, or velocity can be measured by noting how long a vane interrupts the light between the two. Temperature can be measured by noting how the spectral content of an LED changes with temperature. The position of a light beam on a quad array of photodetectors can be used for alignment and positioning purposes. This is how the tracking and focus opto-electronics works on the popular laser video disks. Lasers have had their properties exploited to measure down to a half wavelength. Such an instrument is called an interferometer. Although it has been around for a hundred years, it wasn't until the advent of the laser and modern electronics that such measurements could be practically done.

LUMINOUS INTENSITY VALUES	
Source ¹	Value in candela* (cd) med - 0.001cd
Hydrogen bomb, 100 microseconds after detonation, 30 meter fireball diameter	2x10 ¹²
Surface of sun	2x10 ⁸
Flashlamp	2x10 ⁸
Sun as viewed from Earth	6x10 ⁶
Surface of 60W bulb	6x10 ⁵
Clear sky	1x10 ⁴
Surface of moon on clear night	1x10 ³
Overcast sky	1x10 ²
Twilight	5
Clear moonlit night	30 mcd
Average LED	1 mcd

Table 2

•Signal Transmission

The use of high output IR LEDs and sensitive photodiodes has resulted in high speed data communications. The transmission medium can be air, water, plastic, or glass. The later two materials encompass fiber optics. We shall return to this field later.

Why use light? Light does not interfere with anything it

is trying to measure. It is non-contacting in principle, so we don't have to physically touch the thing we are trying to measure. It is very fast. It is safe, easy to work with and its components are relatively inexpensive. It is non-RF (radio frequency) and will never be regulated by the FCC. In the following sections we will start with some basics and build towards more sophisticated designs.

Parts

It helps to have a quick supply of opto-electronic parts. Where possible I have tried to stick to readily available components. The following list of parts houses offer extremely good and fast service. I have dealt with them for years and they are above reproach when it comes to service and honesty. All except Radio Shack are single location centers. Most have a toll-free number for ordering and all take VISA and MasterCard. They have free catalogs for the asking which are very comprehensive. Most have a \$10 minimum, but your projects will probably cost that anyway.

NAME	LOCATION	PHONE
Radio Shack	everywhere	Yellow Pages
Digi-Key Corp.	Highway 32 South P.O. Box 677 Thief River Falls, MN 56701	1-800-346-5144
Jameco Electronics	1355 Shoreway Rd. Belmont, CA 94002	1-415-592-8097
Priority One	9161 Deering Ave. Chatsworth, CA 91311	1-800-423-5922
Mouser Electronics	11433 Woodside Santee, CA 92071	1-619-449-2222 Sat. phone orders

Books

There are lots of books on electronics out there, all having their own particular bent. In fact, you can purchase most of them from the companies listed above. But there is one book that stands far above them all. I regard it so highly that it is difficult to put into words. With this one book, I maintain that you can understand 100% of anything electronic and with some practice you can learn how to design virtually 90% of all electronic and opto-electronic circuitry. It is an extremely practical book. The authors' philosophy is that electronics, as currently practiced, is basically simple, with some elementary laws, rules of thumb, and a large bag of tricks. Their treatment is largely non-mathematical (and what there is is just algebra) and they encourage brainstorming. I like the authors because one is a professor and the other has his own electronics company on the side. Together, these guys are very aware of what works. They know the difference between theory and application. There are chapters on construction practices, low-noise and shielding techniques, and microprocessor interfacing to the real world. But the most interesting thing about the book is

the end of each chapter. Here, on gray colored pages, are a bunch of circuits that perform useful tasks based on the concepts and principles taught in that chapter. They come under the heading of "Circuit Ideas". After that, there are circuits under the heading of "Bad Circuits" but they don't tell you what is wrong with them. Based on the preceding chapter you should be able to figure out what it is. I guarantee that if you can solve all the "bad circuits", you will have mastered most of electronics. Some of the bad circuits are very funny and some are subtle.

The book has been out for three years and it is my hope that they will revise it so that it doesn't get out of date. The authors will have the electronics book market wrapped up for years if they do this. Coupled with their efforts to teach you how to design and why certain components are chosen, you couldn't have a finer book in your technical library.

The Art of Electronics

by

Paul Horowitz and Winfield Hill

published by Cambridge University Press
1980

Library of Congress Number TK7815.H67

Dewey Decimal Number 621.381

ISBN Number 0 521 23151

Try your local library or a college or university bookstore to purchase this gem.

Display Circuits

LEDs are PN junction devices with an energy band gap determined by the crystal material and structure. Some common materials are gallium arsenide (GaAs), gallium phosphide (GaP), and gallium arsenide phosphide (GaAsP). LEDs are available today in IR, red, orange, yellow, and green. Blue is different because the bandgap is large and the number of blue photons generated is small. Current injected into the PN junction eventually causes electrons to drop to their lower energy states and emit a photon in the process. So, the current must be limited to a typical value of 20 ma. At this level, most LEDs will exhibit a forward voltage of 1.6 across itself. Let's calculate the current limiting resistor needed to drive this LED from 5v. See Figure 6.

In order to drive an LED from TTL logic, you must know that TTL can only sink current. This means it can't output both a high voltage and high current. To sink current means to do it with the output at or almost at ground potential. Normal old-fashioned TTL can sink 16 ma per gate and low-powered Schottky (LS) can only sink 8 ma. So this particular LED can't be driven directly from a TTL or LSTTL gate. (Some newer LEDs put out lots of light at a few ma so the circuit in Figure 7a will work). We will use the most popular low power signal transistor around to drive the LED. All we have to know is that the collector current of a bipolar transistor is equal to the beta of the transistor multiplied by the injected base current. For those of you to whom this is familiar, be patient.

We assume a worst case beta. This means the lowest beta

the manufacturer claims on the data sheet. For the 2N2222 it is 100. That means for every 0.1 ma of base current we get 10 ma of collector current, which is what our LED needs. So, in order to get 20 ma of collector current, we need only

$V_R \text{ across } R = 5 - 1.6 = 3.4V$
 $R = \frac{V_R}{I} = \frac{3.4V}{0.02A} = 170 \text{ ohms}$

What power is dissipated in R? ¹
 $P = I^2 R = (0.02)^2 (170) = 68 \text{ milliwatts}$
 So we can use a 1/4 watt resistor.

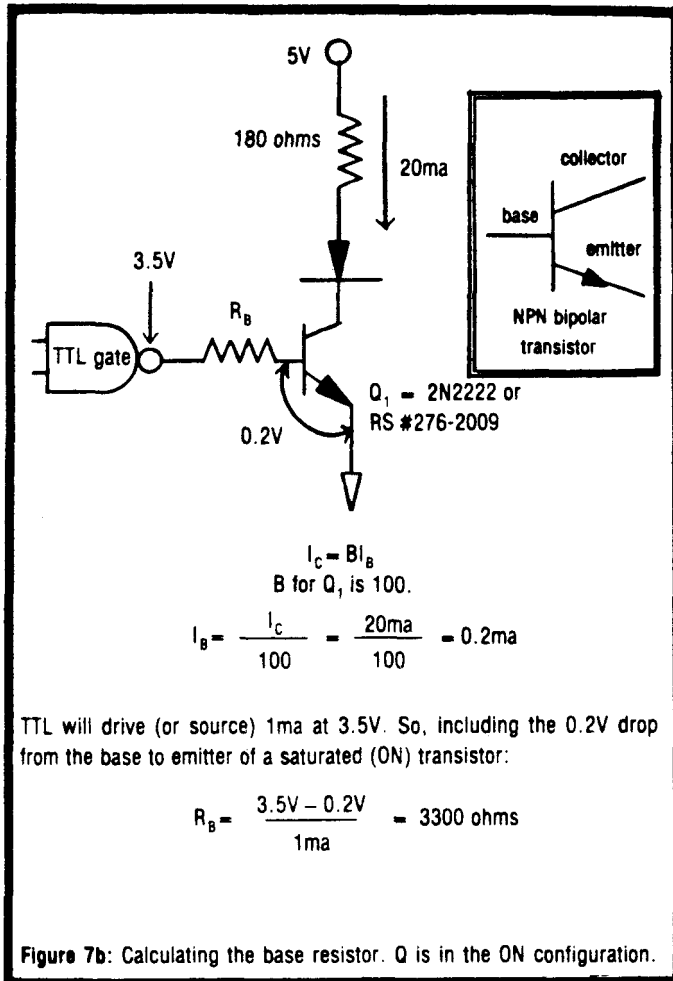
The most popular common value closest to this R value is: 180 ohms.

Figure 6: Calculate the current limiting resistor needed to drive this LED from 5V (LED forward voltage = 1.6)

Any TTL or LSTTL gate

Figure 7a: A load on the output of the gate will sink the 5ma (5V/1 k XL) current in the circuit. This is within the sink current capability (called I_{OL} in databooks) of the gates shown. This circuit only works for LED's that produce light at 5ma of current.

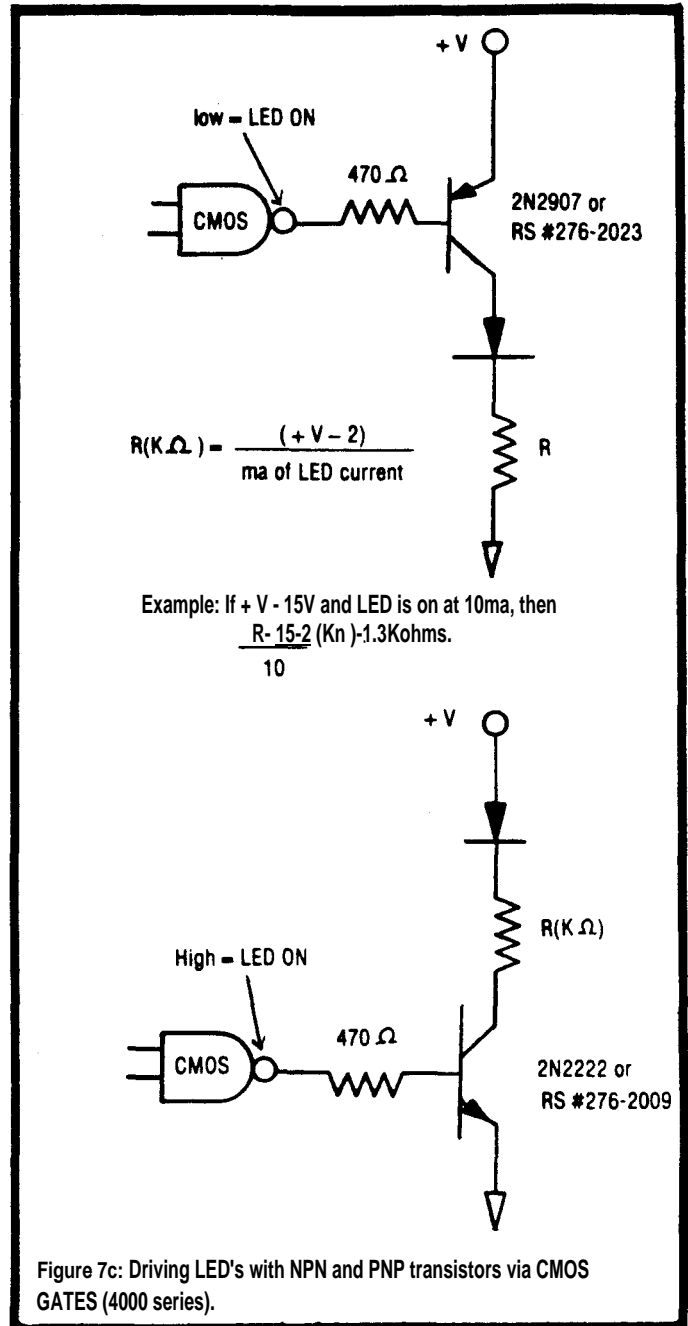
0.2 ma of base current. Normal TTL can actually source a few ma with its output at 3.5v and even LSTTL can source a few tenths of an ma. If we pick the base resistor to be 3.3K ohms calculated as shown in Figure 7b, then the gate can well source the 0.5 ma of current when brought high. We only needed 0.2 ma, we've supplied 0.5 ma, so that transistor is really turned on hard. That's all there is to it.



CMOS can drive LED's directly with the same calculations. B series CMOS can drive LEDs directly with no current limiting resistor. All other families such as the 4000 series need a transistor and 1K base resistor to drive loads of tens of milliamps. See Figure 7c.

The way to drive 7 segment displays is to use the popular decoder driver 7446. Its outputs are low when the particular segment is to be on. It has an open collector output transistor and you must complete the output circuit by adding the current limiting resistor. It drives a common anode type display. Some chips have all the drivers and resistors built right into the display. All you have to do is present the BCD (binary coded decimal) data on the input lines. An example of this type of chip is the MAN2 from Monsanto or the TIL305 from Texas Instruments.

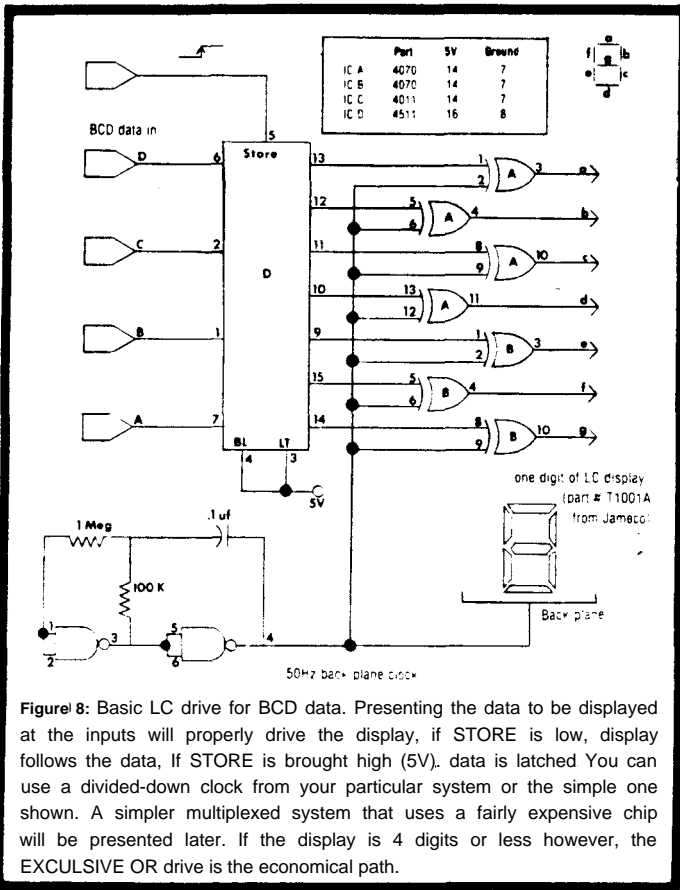
Figure 8 shows the proper and easy way to drive LC displays. There are more expensive chips to do the complicated segment multiplexing, but this is the easiest way and it shows how to drive the display using



EXCLUSIVE OR GATES. It also lets you drive a particular part of an LC display like a decimal point or other character like a bell or colon.

In Figure 9 is a very interesting bar graph chip. It lets you input almost any analog voltage and it automatically breaks it up into 10 equal segments. These type of displays are very popular because they digitally display a quantity in a linear bar. Things like speed or accelerations are difficult to interpret on digit type displays, but bar graphs displays let the observer see trends and accelerations (The bar graph display itself is still digital in its outputs). The chip is now even offered on a small carrier with all 10 LEDs on it. All you have to do is hook up the supply and a few range resistors and you're done. The chip is offered in a logarithmic version and in VU units for audio circuit applications.

Figure 10 hooks two of these bargraph circuits together



to form a bipolar input bargraph display with a center or zero indicator. Input diodes protect each circuit from the negative voltages created when the input crosses zero volts. The center LED lights when the voltage is close to zero. Such a circuit finds use as an indicator of a bipolar signal in a measuring system. You can think of this display as a center zero type analog meter.

In the next article we shall show how to detect light. Timing circuits, photodiode and phototransistor techniques will be discussed. How to measure the speed of a shaft will be presented. Also, an interesting way to measure angles will be disclosed. This "digital protractor" finds use in robotics, at the drafting table, and in the shop.

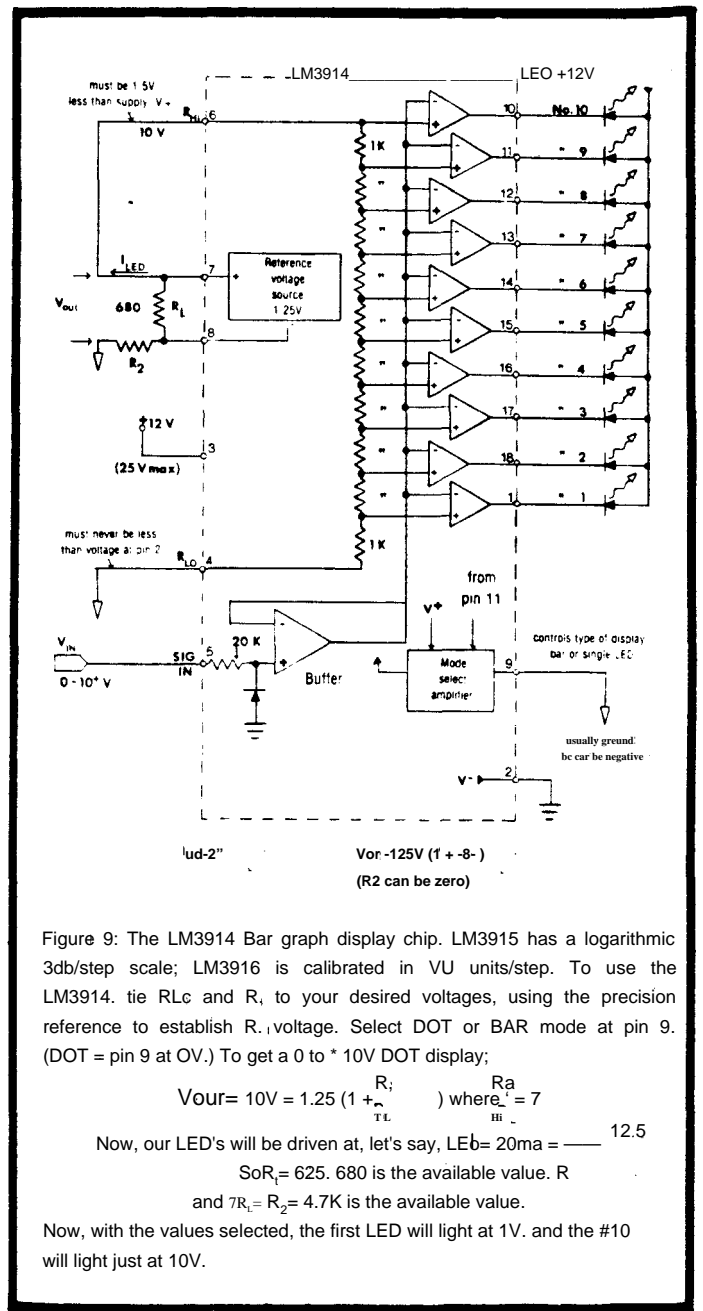
References

The TTL Cookbook
by Don Lancaster
Howard Sams Pub. Co.

This is available through most of the catalogs listed in the article, or from the Computer Journal Bookshelf in this issue.

The CMOS Cookbook
by Don Lancaster
Howard Sams Pub. Co.

Same as above



Optoelectronics Designer's Catalog
Hewlett-Packard
Opto-Electronics Division
640 Page Mill Road
Palo Alto, CA 94304
1-213-970-7500

Absolutely the best text of its kind. Its application section really takes you through the steps required to work with LED's, opto-isolators, displays, fiber-optics and lenses. Get it!




figure 10: This circuit lights 1 LED per volt over a ± 10 volt range. The "ZERO" LED lights whenever the input voltage is between -1V and +1V. Other ranges can be produced, just follow the equations used in Figure 9. U1 inverts the input signal. Only positive voltage at pin 5 of the 3914 results in lighting of the LED string. Negative voltages are clamped by the input diode in front of the buffer (see figure 9). U2a and U2b are open collector comparators wired in the "OR" configuration. They force the "ZERO" LED to light over the -1V to +1V input signal range.

Multi-user

A Column by E.G. Brooner

Different kinds of multi-user computer systems are becoming common enough to cause some worry and confusion among those of us who use personal or desk-top computers. Are we missing something we should be in on?

What do all those buzz-words mean: Is the Source a network? What is the difference between multi-processing and timesharing? Are all of these things just gimmicks? Are they of any value to the personal computer user? That's what we are going to try to clear up in this column.

It doesn't really matter whether you use your computer for business, science or control applications —you are a candidate for some kind of multi-user system if: (a) your operation has grown a bit bigger than can be handled conveniently by a single microcomputer, or (b) you (even occasionally) want to have two or more computer projects going on at the same time. Schools, medium-to-large labs and/or offices and very serious hobbyists are all potential users.

Timesharing Systems

Let's begin with timesharing. If you connect two or more terminals (or micros emulating terminals) to a common computer, and share its CPU (Central Processing Unit) and memory, that is a timesharing system. Big computers have been doing this for years and in fact information services like the Source are actually gigantic timesharing systems.

It has become possible, in recent years, to apply the same technique to small computers. North Star and CPM, to name two products, have made it possible for us to use up to eight or so terminals on a single microcomputer. Most of the minicomputers, such as the IBM System 34, permit the same sort of thing. If you have used such a system, you are probably aware that the 'throughput,' or the operating speed of each individual user, suffers quite a bit.

There are three major reasons for this degradation. First of all, every user shares the same CPU, which is actually doing the work for everyone, in 'chunks.' Secondly, the same memory system also has to be shared. The third reason is that the terminals are usually connected to the computer by some relatively slow communication path.

There have been schemes for extended addressing and 'virtual memory' (shuffling your stuff back and forth from a disk drive), but in each case there is a practical limit to the number of users that can be accommodated at an acceptable speed. The limit is set by the speed of the single shared CPU and the practical amount of available memory, no matter how it is described. When we deal with microcomputers, two or three users are usually practical, and perhaps a dozen (with a lot of degradation of

performance are possible) in some circumstances. You might enter some characters, for example, and experience a noticeable delay between the time you do so and the time they appear on the screen of your terminal. The more expensive minicomputers do not do much better.

There are really just two choices when it comes to timesharing: you can have a large, expensive system or a small one which won't be able to grow with your business or application

Networks

We turn our attention next to nets, networks, or LANs (Local Area Networks.) The manufacturers talk confidently about 64 users, or 1,000 users, or even 64K users all interconnected and 'sharing' various amounts of 'power.' How much truth is there to these claims, and how do they even propose to go about it? The simplest answer is that each user has his or her own complete microcomputer. They are capable of operating in a completely independent manner and, as a matter of fact, they usually do. But they are connected by way of some kind of communication system, so they can share printers or large disk systems and databases, upload and download one another's programs and data, and send messages to one another. They can operate at speeds of up to 1,000 times that of a timesharing system. Some transfer data at a rate of 10 megabytes per second. The communications system they use between devices is nothing at all like the ones we normally use with terminals and other computer peripherals.

The essence of a true network is that each user has at his disposal, all of the time, a complete computer; but in addition, he has access to the facilities of everyone else with whom he is 'networked.' With a network you can start small, with two or three microcomputers, as you might with a small timesharing system. The major difference is that as your system grows the performance doesn't diminish, as it well can with timesharing. Each added 'workstation' can be a complete computer that can perform entirely by itself. It can, though, call on the others to which it is connected to obtain additional data or other services. Each added device adds to, rather than detracts from, the 'power' available to the existing users.

For example, a dozen or more desktop computers might share one or two printers or a large hard disk storage system. Software can be shared and messages can be sent between offices, just as we do now with inter-office memos. The transfers between users happen very rapidly, and during the remaining time each user's resources are entirely his own.

The initial cost of 'going network' is slightly higher than

continued on page 22

MAKING THE CP/M USER FUNCTION MORE USEFUL

by M. Mosher

CP/M 2.x from Digital Research contains several nice features not found in the original CP/M 1.4. Among these is the ability to divide the disk into different "users." I put this in quotes because CP/M is not a multiuser operating system so the term user is perhaps a bit of a misnomer. The partitioning of the disk into different areas is not done on the data areas but is confined to the directory. Encoded in each directory entry is a number from 0-15 which indicates what user that file belongs to. Files can exist under different users with the same names and be totally different files. In displaying a directory, only those files belonging to the current user are displayed. This is analogous to the display of only those files on the currently logged-in drive when you give the DIR command.

This makes for a very handy way of partitioning up a disk so as to avoid a large directory where it is hard to find what you're looking for and which takes up (with hard disks, anyway) more than can be displayed on one screen of a video terminal. It also allows for a logical division of program types or of different languages and their respective assemblers, compilers and interpreters. For instance, one possible logical division would be to put all system files and utilities under User #0, BASIC source files and the BASIC interpreter under User #1, word processors and text files under User #2, and so on. With 16 different users, there ought to be ample categories for anyone.

The only fly in the ointment here is that there is little or no provision for interaction between user areas. While this may be necessary and desirable for true multiuser systems such as MP/M, to avoid one user accidentally or deliberately getting into another user's files, it seems totally unnecessary in a single-user operating system such as CP/M. While it's* possible to copy programs from one user to another with PIP, you can't execute a program in User #1 when you are logged in to, say, User #7. Even getting PIP to the user to start with requires a special operation using DDT and SAVE, since PIP can only copy files from other users, not to them. By the way, PIP Version 1.5 has bugs which cause unpleasant results when doing inter-user copying. Version 1.8 fixes these bugs. You'll have to look at the first block or two of PIP with DDT to see which version you have. It isn't displayed in the sign-on message but can be found embedded in the program.

This limitation causes a lot of program repetition since it's nice to have some of the utilities such as STAT and PIP available under all users. It has always seemed silly to me to have to have as many as 16 copies of STAT on the same disk when it would be easier to be able to access the same copy from all users. Fortunately this can be done with a very

simple patch to CP/M.

When you try to execute a COM file, the CCP tries to open the file of that name under the current user. If it isn't found, it gives up. The file might have been in the directory under another user number but, no matter. That doesn't count. Listing 1 shows a short patch to alleviate this situation. What the patch does, in short, is to intercept the open file request and try first to find it under the current user. If not found there, instead of giving up, it searches all the other user numbers for the file of that name starting with User 0 and going on up to User 15. Alternatively, you can limit it to search only User 0 by changing just one line of the program. Only if the program doesn't exist under any user does the patch allow a failure to be reported to the CCP.

The patch can be located anywhere there is room for it in the CP/M system. There is a 127-byte console buffer located near the beginning of the CCP. Since few of us type in command lines that long, the listing as shown borrows some of this space for the patches shown. The length of the buffer is reduced accordingly. Alternatively, if you simply have to have that long console buffer, you can locate the patch in your BIOS or user customization area where I/O drivers are found. To do this, the ORG must be changed accordingly.

To install the patch, you must know the location of your BIOS. By BIOS I don't mean the custom user I/O area since the latter is usually higher in memory. To find the BIOS, load DDT and look at location 2. This is the high byte of the starting block of the BIOS. A statement that equates the variable BIOS to this value is needed at the beginning of the PATCH before it is assembled.

To get the patch into your CP/M, first assemble it, creating a HEX file. Use SYSGEN or MOVCPM to get a system memory image. Most systems tell you where this is but if you want to be sure of installing this correctly, you had better have some facility with DDT and know a little bit about the inner workings of CP/M already. After loading the system into memory, calculate the offset. Offset = (BIOS location in system memory image (Actual BIOS location when CP/M is running)). Load the HEX file containing the patch using the DDT command R(offset). Then go ahead and save the altered memory image with SYSGEN or SAVE.

If this seems a bit sketchy, I'm assuming you're already something of a software hacker and can pick up on the smaller details of the operation. If it's totally beyond you but you still want to put the patch in, try to find a friend who knows more about software to help you.

Once installed, I think you'll be surprised at how much easier it is to use the USER function. Before I figured out

BUILD A HARDWARE PRINT SPOOLER

Part Three: Enhancements

by Lance Rose, Technical Editor

In this final part of the series on hardware print spoolers, I would like to talk a little about some possible modifications to the basic design to make it possible to adapt the spooler to a wider variety of systems. These enhancements include the use of different handshaking protocols and different interfaces. As you may recall from Parts 1 and 2 of the series, our spooler was designed to use RS-232 communication and DTR handshaking on pin 20 of the interface connectors. I will discuss the possible modifications in approximate order of increasing complexity.

RTS Handshaking

Request-To-Send handshaking is very similar to Data-Terminal-Ready handshaking and would require only some simple wiring changes to implement. The difference here is that the handshaking occurs on pin 4 of the interface connector instead of pin 20. To change to RTS handshaking to the computer, all you need to do is reverse the connections on the spooler to pins 4 and 20 of J1. To use RTS handshaking to the printer, simply move the wire in the spooler from pin 20 to pin 4. There are no software changes needed.

X-on/X-off Handshaking

This is a method of software handshaking whereby the sending device keeps transmitting characters to the receiving device until the receiving device sends an ASCII X-off (Control S) to the sending device. The sending device then ceases transmission until it receives an ASCII X-on (Control Q) from the receiving device. It then resumes transmission. There are some advantages and disadvantages to this protocol. On the positive side, there are no additional wires needed in the interface other than for the serial data signal and ground. This eliminates the condition of non-standard pin usage for hardware handshaking functions, since some manufacturers use reverse channel hardware handshaking on pin 11. The serial data is always transmitted on pin 2 (or pin 3, depending on whether the device is configured as a modem or a terminal).

On the negative side, any software handshaking method requires that the spooler have transmit and receive hardware for both the computer and the printer. In the case of our design here, this would necessitate adding another UART for the additional capability. A second UART would also require a few more logic gates to allow selectively enabling one UART or the other during an I/O read or write. The software for the spooler ROM would be slightly different, too. Instead of watching the printer for a DTR low, it would have to test for a character received from the printer, then read the character and see if it was an X-off. If

it was an X-off the spooler would have to stop sending. Conversely if the spooler buffer filled up, instead of setting DTR low to the computer the spooler would transmit an X-off character to the computer. Sending X-on characters would also replace the process of setting DTR high again.

If you have a large enough chassis box and don't mind adding a couple more chips, this implementation should be well within the capabilities of the serious hacker.

ETX/ACK Handshaking

This handshaking convention is similar to the X-on/X-off protocol in that it would require bidirectional data flow to both the computer and the printer. The same hardware additions mentioned above would also hold true here. The software modifications would be a bit more complicated. In this protocol, the computer periodically inserts ETX characters (Control C) in the text stream. When the printer encounters one of these, it transmits an ACK (Control F) back to the computer (spooler in this case). The question here is, how often do we insert these ETX characters? This depends on the size of the printer's internal buffer. Since different printers have different buffer sizes, the software for the spooler has to be tailored to the printer for proper ETX/ACK handshaking. Changing printers might necessitate reprogramming the EPROM in the spooler.

An additional complication is that many printers use a sequence of characters starting with an ESC (IB hex) to initiate special functions such as change to graphics mode, change line spacing and so forth. Unless the driver software in the spooler is smart enough to recognize all these special combinations and not insert an unwanted ETX in the middle of the string, some very strange results can occur.

Luckily, ETX/ACK handshaking is becoming less common these days, being largely superseded by X-on/X-off which puts more of the burden of intelligence on the printer and less on the spooler. If your printer uses only ETX/ACK handshaking and you are determined to build a hardware spooler, you have your work cut out for you.

Open Loop (No Handshaking)

This is obviously the least desirable method to use. However there are some printers (for example some Selectric conversions) that have no handshaking ability at all. To utilize the spooler with one of these will require no additional hardware but a great deal of software. The focal point of the operating program here would be to include some software timing loops to set a transmission rate of characters to the printer. Doing this while still handling all the other functions of the spooler would require some finesse with software. Once again, it's only for people who

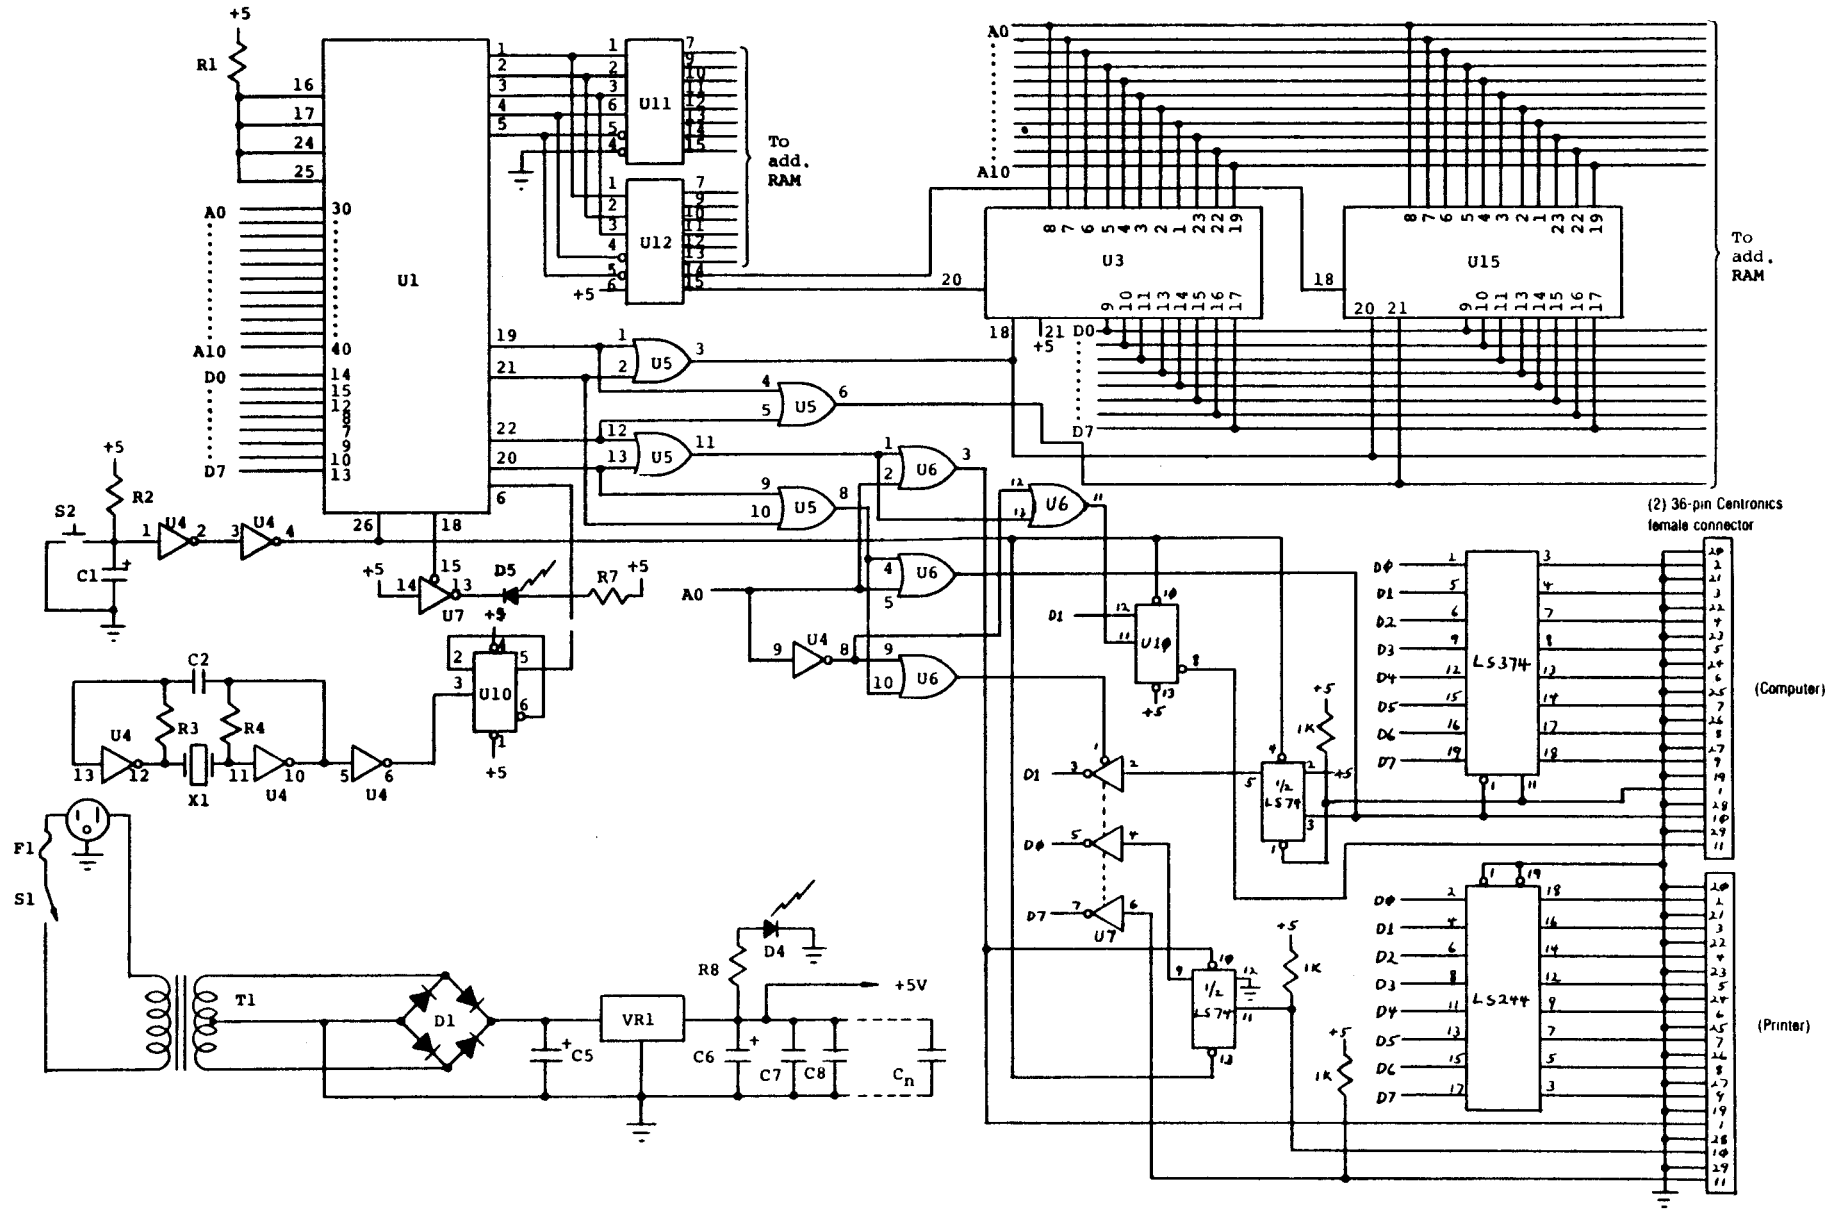


Figure I

have no other choice.

Centronics Interface

The principle of hardware print spooling is, of course, independent of the actual interface convention chosen. The original design was built around the RS-232 interface for reasons already given. The Centronics interface is another widely used convention especially among lower-cost computers. Building a version of the spooler for this interface is no harder than for the RS-232 convention. In fact, it is somewhat easier since some of the hardware used for the serial I/O can be eliminated.

Figure 1 shows a modified schematic for a proposed Centronics version of the hardware spooler. Note that some of the ICs needed for the RS-232 version are absent here. The basic clock circuit is still necessary for the Z80 but the two baud rate generators (U13 and U14 in the RS-232 version) are not needed. Also, since all signals are TTL levels in the Centronics interface, the EIA line driver and receiver (U8 and U9) are not needed. The ± 12 volt power supplies are not needed either since they are used exclusively for the line driver IC.

The select circuitry is unchanged since we still need to read status and data, and write control signals and data here. The half of U10 that previously held the DTR status to the computer is now used to hold the BUSY status which is set whenever the spooler fills up. An additional 74LS74 is needed to store the input waiting status from the computer and the output accepted status from the printer. The first of these is set whenever the computer sends a character to the spooler. At the same time the actual data is latched into the 74LS377 octaHatch. When the computer reads the character in, the input ready flag is reset and an ACK pulse transmitted back to the computer. Tri-state U7 gates the

pertinent status flags onto the bus during a status read for examination by the program.

To send data to the printer, the computer writes to the data port which sets the output busy flag bit. When the printer accepts the data, its ACK pulse resets the busy flag. In addition, the spooler can test the busy status from the printer directly. Most printers delay the ACK pulse during a busy condition so the busy status might not be necessary, but a good operating program should examine both, especially since the hardware is there. The software should be quite similar, especially since the design shown keeps the flag bits and polarities in the same place as the RS-232 version. The primary difference would be that when the spooler was almost full, it should set the busy line to the computer (the same as the RS-232 version) but should also not input the character, thus delaying the ACK pulse to the computer in case the software driver on the computer end relies on ACK only and doesn't test for BUSY.

Construction and checkout of a Centronics version would be no different than for the serial communications unit. If your printer has a Centronics interface you can still add hardware print spooling to your system.

Summary

In this series of articles we have progressed from the design of a hardware print spooler to the construction and finally to the enhancements. There is an enormous variety of systems in the hands of hackers and I'm sure that I haven't covered all possible enhancements or all possible applications. If you adapt the spooler to some system I haven't mentioned or have modified it in some other way, why not write in to the *Computer Journal* and share your experiences with the rest of the readers. We can all benefit from such an exchange. Happy spooling!

"Making the CP/M User Function More Useful," from p.11

the patch, I didn't bother with it at all for the reasons I mentioned above. Now it's a handy way to keep my disk files organized.

```

;
; Put BIOS equate here, e.g. BIOS EOU 5C0BB
;
BEGCPM|EQU BIOS-IM>0H
;
; Patch to search all users for COM files
;
ORG BEGCPM+6606H | :Make consol*? buffer shorter
;
DB 56B
;
ORG BEGCPH+085PH | :Change this if you patch BIOS

```

```

TRYALL:: CALL BEGCPM+0113H | ;Get the current user
          STA RSTUSR+1. ;Save it for later
          CALL BEGCPM+00D0H ;Try open under current user
          REZ ;Return if successful
          MVI E,00H ;Start with user 0
          TRYLP: PUSH D ;Save user being searched
                CALL BEGCPM+0115H | ;Set to this user
                CALL BEGCPM+00D0H ; Try to open the file
                POP D ;Restore user
                RBZ ;Return if found
                I MR E ;Go to next user
                MOV A,E
                CPI 16 ;Use CPI 1 for user 0 only
                JC TRYLP ;Try all users
RSTUSR:: MVI E,0H |
          CALL BEGCPM+0115H | ;Restore current user
          XRA A ;Set flag for file not found
          RET
GETUSR: CALL RSTUSR ;Get back current user
        JMP BEGCPM+011AH ;Continue as before
;
; ORG BEGCPM+06D8H ;Look for it under all users
;
; CALL TRYALL
;
; ORG BEGCPM+0759H ;Restore current user
;
; CALL GETUSR

```

ANYONE FOR A LITTLE "KISS" ELECTRONICS?

Part Three: Power Supply

by Phil Wells, Technical Editor

There are many different kinds of power supplies, each most suitable in particular situations. High-efficiency switching-mode supplies with multiple output voltages, fold-back current limiting, crowbar over-voltage protection, etc., are complex things of beauty now available at low cost. But not at a cost below \$15. And not needed for many small projects. Simplicity, ease of construction, readily available parts and low power outputs are not the strong points of such wonders. You can design your own power supplies with ease by keeping things simple, even if you've never designed anything before. If you collect a "junk box" of parts as good deals come along, you can put together just the right power supply the day you need it.

In this KISS (Keep It Simple, Stupid) Electronics article, we'll cover some ways to design small supplies using components that are easy to obtain at low cost. Our goal is to give the project builder control over the project. Remember, this is for a one-shot project; if you plan to mass-produce 100,000 of these units, you must be much more careful in your design.

The General Concept

You are designing a stand-alone electronic box that does something useful, entertaining, or enlightening. The box uses digital integrated circuits, which need a de power supply of 5V at 1A, or $\pm 12V$ at 0.25A, or 24V at 0.5A, or the like. You can buy a supply, but ready-made ones are bulky or expensive or don't fit into your box. You would like to design your own supply.

Figure 1 shows the building blocks involved in a standard, linear, series-regulated power supply. The purpose of

the supply is to tap into the unlimited power available from the 110 vac power wall outlet and convert that high voltage alternating current into low voltage direct current.

The transformer converts the relatively high voltage into a lower voltage, still alternating current. The rectifier converts alternating current into pulsating direct current. The filter smooths out the current pulses to provide a direct current which still has some degree of ripple voltage. The regulator eliminates the ripple voltage and maintains a constant output voltage when the input (power line) voltage fluctuates and when the current drawn by the load varies.

The regulator can now be thought of as a single building block; not many years ago we had to build this block out of a dozen or so parts even for low-current supplies. The idea behind the popular three-terminal series regulator is illustrated in figure 2.

The regulator acts like a variable resistor in series with the load resistance. You don't need to understand the internal workings of these devices to use them. Simply put, you have a voltage divider; if the input voltage increases, the regulator must increase its "resistance" so that the output voltage at the divider remains constant. If the load resistance drops, the regulator's "resistance" must automatically compensate by dropping proportionally to maintain constant output voltage.

The main thing to understand is that all the load current passes through the regulator (hence the name "series" regulator), and any excess input voltage must be dropped across the regulator. The product of the load current and the excess voltage is power, which means heat dissipated by the regulator. This heat can destroy the regulator.

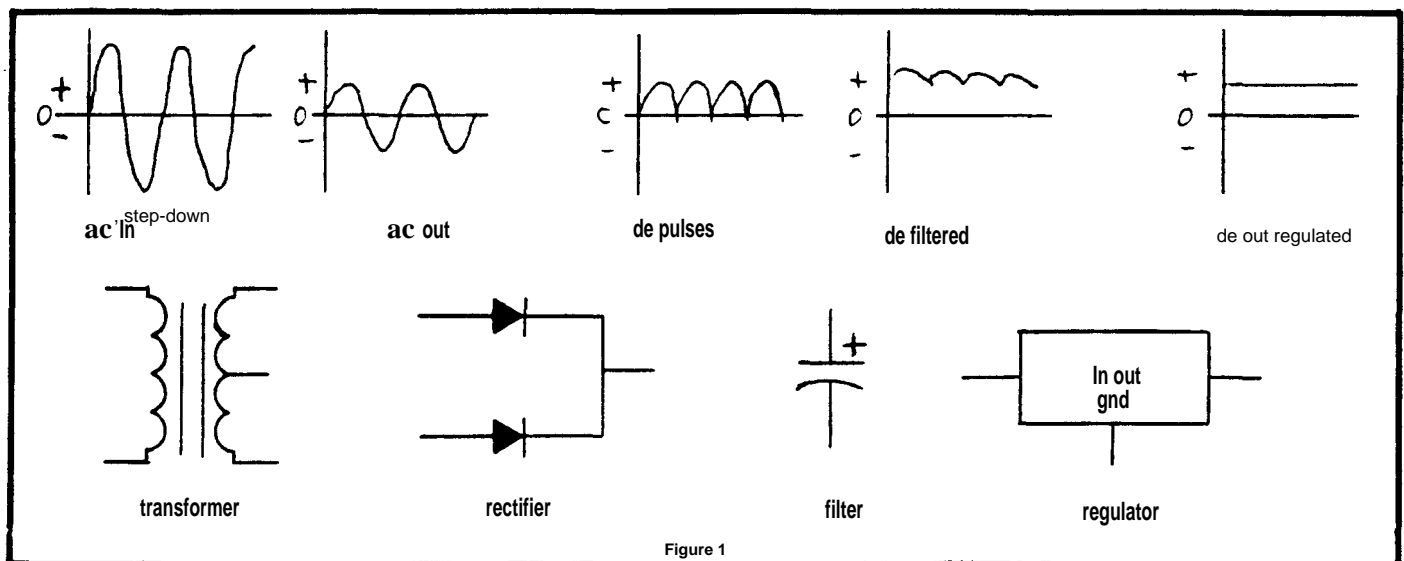



Figure 1



Design Backwards

The practical way to design these supplies is to work backwards, from output to input.

Figure 3 shows the two most common low voltage dc power supply circuits for output currents below three amperes. Figure 3a is a *Full-Wave Center Tapped* or *dual-diode* layout, while figure 3b is a *Full-Wave Bridge* circuit. How do you select the components to build such a supply to meet your own requirements?

There is no real "circuit design" involved; you simply select the components, wire it up according to figure 1a or 1b and see if it meets your needs. To select the components:

1. Specify your required output dc voltage and current.
2. Select a three-terminal IC regulator which meets your worst-case needs.
3. Select a filter capacitor.
4. Specify your desired ac line input voltage range.
5. Select a transformer and rectifier combination.
6. Provide adequate heat sinking for the regulator IC.

Let's take these steps in order:

Define Output Voltage and Current

The kinds of circuits you are supplying power to will determine your voltage requirement. Standard bipolar digital integrated circuits usually require 5.0 Vdc. CMOS devices work over a wide range and don't require precisely regulated supplies. RS-232-C interface circuits require dual (positive and negative) supplies of 12 to 15 volts. The three-terminal regulators will hold the output voltage to within a few per cent of nominal, so you rarely will have to worry about regulation.

What output voltage do you need? Let's assume you need to power a circuit made up of LSTTL parts and other microprocessor-type chips, so you need +5Vdc.

Output current again depends on the circuit you are building. For most experimenters, the selection procedure is to use an existing *bench supply* to power the breadboard version of a project, then measure the current drawn to determine the current rating for the dedicated supply. Aim for a supply current capability 25% or so higher than measured on the prototype to allow for worst case conditions.

What output current do you need? Let's assume you need something more than 0.5 amp but probably less than one amp. If you need more current, you can go as high as five amps with some of the newer three-terminal devices, but I think for anything over three amps you should consider buying a good switching supply because it will be smaller, cooler and more reliable.

Just a caution at this step: the output specifications can be considerably more complex than just nominal voltage and current. If you need a specific temperature stability or load transient response or some other such requirement you should know more about what you're doing than we're covering here.

Select a Regulator IC

Pick a three-terminal regulator that is inexpensive enough for your budget and is easy to get. The TO-220 case style is the easiest to handle physically, but you may find you need the TO-3 case style to handle the power dissipation after you select your power transformer.

Look at the National Semiconductor "Voltage Regulator Handbook" or the equivalent handbook from Motorola or Texas Instruments or your favorite semiconductor candy factory. Or check out the catalogs from such mail order houses as Jameco or Priority One. Or look over the stock at your local Radio Shack. This first selection may not be your final choice, but my experience is that the design process is smoother if you make a preliminary choice here. Make sure the regulator's maximum output current rating is at least 20% greater than you think your circuit will require. We can use an MC7805, LM7805, LM309K, LM340 or similar IC regulator.

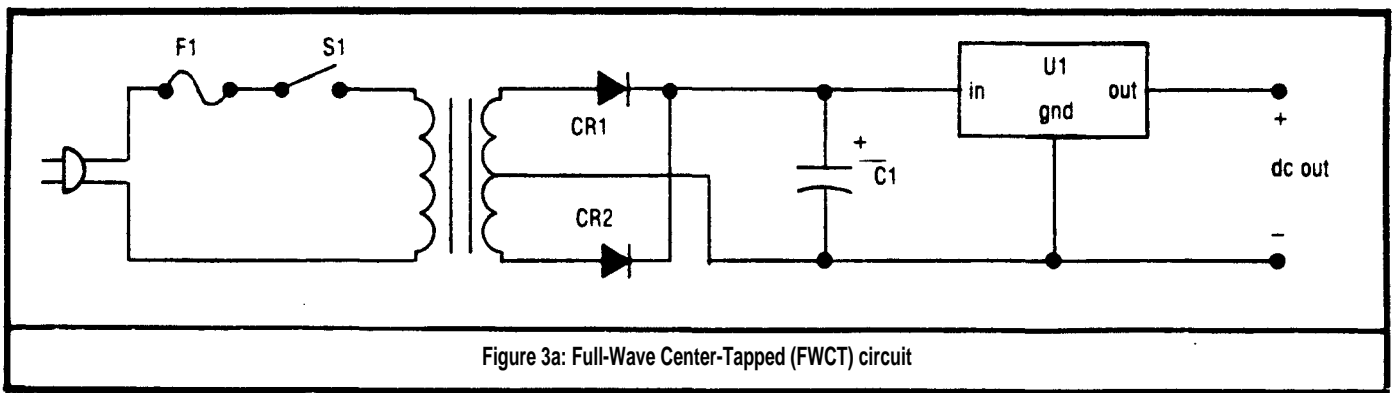


Figure 3a: Full-Wave Center-Tapped (FWCT) circuit

Select a Filter Capacitor

What value of filter capacitor do we need? There are formulas for selecting the capacitor, but for output currents below an amp or two, try to use 4,000 microfarads per amp. For a one-amp supply you might be able to go as low as 2,000 microfarads, but 4,000 would be much better. For a 100-200 milliamp output, 470 microfarads would work and is a common value.

The capacitor's voltage rating should be at least twice as high as the nominal dc voltage you expect to see across it.

We'll select Radio Shack's #272-1022 capacitor; 4700 microfarads at 35 volts.

Select the Transformer and Rectifier Combination

First calculate the needed transformer voltage from the formula in figure 4.

In this formula, V(REG) is the minimum input-to-output differential required by your regulator IC, from the data sheet. For almost all three-terminal regulators this will be at least 2 V. Three volts is conservative.

V(RECT) is the voltage drop across the rectifier diodes. Each diode in series with the load costs about 0.8 to 1.25 volts in transformer rating. The FWCT (Full Wave Center Tap) circuit has only one diode conducting at a time, while the FWB (Full Wave Bridge) circuit has two diode voltage drops. The voltage dropped across a forward biased silicon rectifier diode varies with current and the type of diode; 1.25 volts per diode is conservative.

V(RIPPLE) is the peak-to-peak ripple voltage at the input to the regulator, which depends on the size of the capacitor.

The ripple voltage can be calculated from:

$$V(\text{ripple}) = .006 / C(\text{microfarads})$$

Our 4700 microfarad capacitor yields about 1.3 volts peak-to-peak.

V(NOM) is the nominal or *normal* value of the input ac power line voltage, and V(MIN) is the lowest line voltage at which we require our supply to maintain regulation. The problem here is that as the power line voltage decreases, the transformer secondary output voltage decreases in direct proportion. If your facility has low-line-voltage problems, use a smaller value. If I were designing for production I would use the conservative value of 95 vac for V(MIN), but since I don't often encounter low-line conditions where I live, I used 105 vac. Using the more conservative low-line value results in selecting a higher secondary voltage, which means that under typical 115 vac conditions, the regulator IC will have to dissipate the extra voltage and so will run much hotter.

From Figure 4, if we use the FWCT circuit we need a transformer secondary voltage of 18 vac (9 vac for *each half of the secondary*). For a FWB circuit we need a 10 vac secondary (for the full winding).

The transformer's current rating again depends on which rectifier circuit we use. The FWCT requires 1.2 times the maximum dc output current, or 1.2×1.0 amps - 1.2 amps. The FWB requires 1.8 times the dc output current, or 1.8 amps, for our 5 volt, 1 amp supply.

Notice that these are the minimum voltage ratings you want *at full load current* and lowest line voltage. Let's not

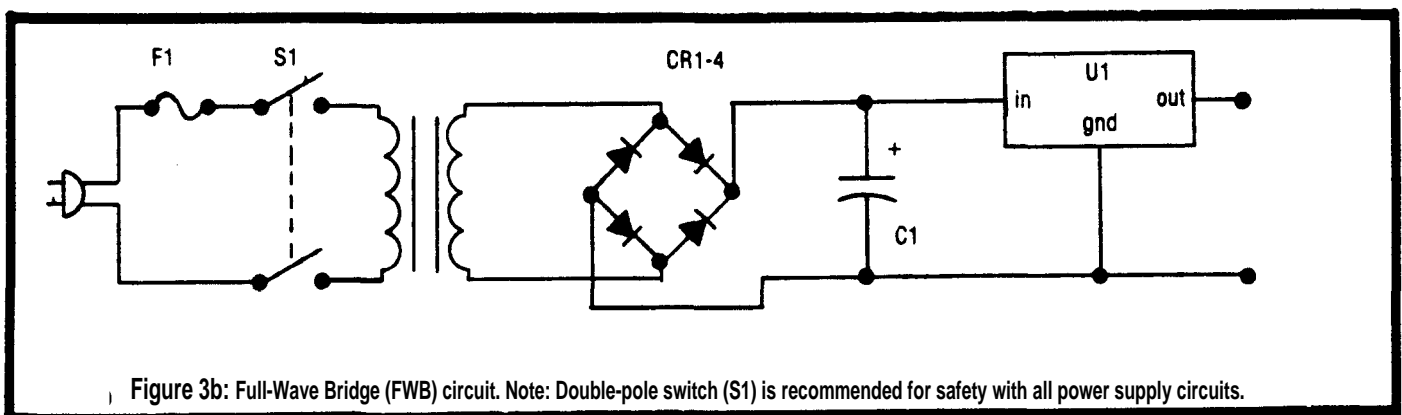


Figure 3b: Full-Wave Bridge (FWB) circuit. Note: Double-pole switch (S1) is recommended for safety with all power supply circuits.

Calculate transformer secondary voltage:

$$V_{ac} = \frac{V_{out} + V_{reg} + V_{rec} + V_{ripple} + V_{nominal}}{0.9 \cdot V_{ow/line} \cdot 2} \cdot 1$$

Where:

- V_{ac} - RMS output from transformer secondary
- V_{out} - dc output from regulator
- V_{reg} - minimum input-to-output voltage differential required across regulator IC (from spec sheet).
- V_{rec} - rectifier forward voltage drop at full load current
- V_{ripple} - peak ripple voltage across filter capacitor
- $V_{nominal}$ - RMS voltage of power line (nominal)
- $V_{ow/line}$ - lowest RMS power line voltage expected

Figure 4

get too fancy here; an 18 to 24 vac ct (center tap) transformer works fine in practice for the dual-diode circuit, and for the bridge circuit a 9 to 12 vac transformer works fine, as long as you don't push too close to the transformer's maximum secondary current rating or run into a really low line voltage.

Why does the secondary current rating have to be higher than the maximum dc output current desired? This is one of the possible "gotchas" in power supply design. First, the transformer secondary must be able to supply the dc load current we want. Second, the secondary must also supply the *peak charging current* to the filter capacitor. Third, in practice, manufacturers of some low-cost transformers "cheat" a little in the amount of iron they put in the transformer's core; when you try to draw the rated current from the secondary, the core *saturates* (cannot support enough magnetic flux) and the secondary voltage drops.

Radio Shack's #273-1515 transformer, rated at 18 vac center tapped, and 2 amp secondary output current, works well in the FWCT circuit and by my measurements just barely makes it in the FWB circuit when the line voltage drops to 105 vac.

The diodes you pick will be rated in terms of current and voltage. This means the maximum *average* forward current and the *absolute maximum* reverse blocking voltage. Since the current pulses into the filter capacitor can be two to three times as high as the *average dc output current* but each diode (in either type of circuit) conducts only on alternating half-cycles, you need a diode current rating at least *twice* the maximum dc output current of the supply. For a one-amp supply, use diodes rated for at least two amps. The reverse blocking voltage rating must be higher than 2.8 times the secondary RMS voltage for the FWCT and 1.4 times for the FWB. Fortunately, voltage ratings are cheap; use a 50-volt diode in a 5-volt supply.

Provide Adequate Heat Sinking

What about a heat sink? This may be the hardest one to figure out, but a frequent approach is to pick the biggest one

that fits in your project's box, run the supply at maximum load and see how hot the IC gets. If you can hold your finger on it, you're probably safe. If the case of the IC will boil water, you should find a bigger heat sink. If it stays really cool, you might be able to go to a smaller heat sink; just remember that heat will probably be the eventual cause of failure of the regulator IC, and the heat will be maximum when the power line is at its highest voltage and load current is maximum.

It is possible to make a theoretical calculation of the heat sink needed, but this requires an understanding of *thermal resistance* which is beyond the scope of this article. The best (and simplest) exposition of the formulas involved in heat selection that I've seen are on page 5-1 of the 1977 National Semiconductor Voltage Regulator Handbook.

Some Refinements and Cautions

First, be aware that just about everything in your supply can generate heat, and will perform worse with increased temperature. You must be conservative in selecting your components and you must provide an air flow path. At one ampere and low voltages a fan isn't needed, but do provide some inlets and outlets for convection air flow.

Don't skimp on the transformer current rating and be wary of transformers with small cores (the laminated steel stack around which the wire is wound). You can't get double the current rating by using a dual-diode center-tapped circuit because the core and the wire are not designed for it. If you run the transformer at or near its rated secondary current, you can expect it to get very warm (expensive ones may be designed to run very hot). You must get this extra heat out of your enclosure because it raises the ambient temperature of all the components *especially the regulator IC*. Heat is also the most likely cause of transformer failure.

Be aware that transformers are not ideal devices. Their output voltage changes as you draw more or less current from your supply. You can expect a 10 to 20 per cent *regulation* from your transformer. That is, if it puts out 10 volts RMS at very small load currents, it could drop to 9 or even 8 vac at full rated load. Almost always, however, manufacturers add extra turns of wire to the secondary (or subtract turns from the primary) to compensate for this. The result is that, for example, Radio Shack's very nice 18 vac, 2 amp transformer actually puts out 20 vac at a few milliamps of load and drops to 18 vac at 2.0 amps (by my measurements).

Also consider your power line input voltage. This is nominally 110 to 117 vac RMS, but in most residential and commercial buildings can vary from below 105 to above 125 vac RMS. This, coupled with the transformer's higher output voltage at low currents, effects the heat generated by the regulator IC and the voltage rating you need for your filter capacitor. If the power line is running high and your load current is low, the capacitor voltage can easily be 20% above your *nominal* calculations.

The power dissipated by your regulator IC is the product of the voltage across it and the current through it. If you

neglect variations in power line voltage for a moment, you'll find the power dissipated in the IC increases as you increase the load current. If you have 11 vdc at the regulator's input when the load is 0.1 amps, the power dissipated is $0.1 \times (11 - 5) = 0.6$ watts. As you increase the load current, the transformer core saturation will reduce the regulator's input voltage to perhaps 9 volts (average, remember). This gives 1.0 amps times 4 volts for 4 watts.

Now consider the power line fluctuations. You should design the supply so that at 105 vac input, there is still the needed input-output differential across the regulator, *even at full load current*. But this means that if the power line goes up to 125 vac, you'll have almost a 20% surplus of dc voltage into the regulator and instead of 1.0 amps times 4.0 volts you'll have 1.0 amps times ((120% of 9 vdc) - 5.0) or $1.0 \times (10.8 - 5) = 1.0$ amps times 5.8 volts for 5.8 watts. You must pick a heat sink large enough to keep the regulator IC below its rated junction temperature even when the room the supply is at its hottest. This is why it is very important to provide air flow and a good heat sink.

To complete our design we have one more consideration. Most three-terminal regulators require one or two additional small added capacitors to prevent oscillation. You'll have to check the data sheet for your particular regulator. If you can't find a data sheet, use a 0.22 microfarad disk ceramic or 1 microfarad solid tantalum capacitor as close as possible to the input lead of the IC, and a 1 microfarad solid tantalum or 25 microfarad aluminum electrolytic at the output lead.

Figure 5 shows our completed 5 vdc, 1 amp power supply in both versions.

We should add a bleeder resistor across our filter capacitor to provide a discharge path for the capacitor's stored charge when the supply is turned off. Ideally, we would like the capacitor to discharge completely in a second or so, but this would require a resistor value so low that an unreasonable amount of current would be constantly drawn through the bleeder and wasted. To calculate the bleeder's

resistance, you need to know that it takes five RC time constants to discharge the capacitor to 1% of full charge. An RC time constant is the product of the capacitance in farads and the resistor in ohms. For a discharge time to 1% of full charge, we need $5RC = 1$ second, which yields, for a 4700 microfarad capacitor, a resistor value of $R = 1 / (5 \times 0.0047) = 43$ ohms. A 43 ohm resistor would draw $1 - V / R = 8 / 43 = 186$ milliamps (roughly) constantly. The 8 volts is an approximate value for the filter's voltage since the actual voltage will vary with line and load changes. In practice, I usually use a 1,000 ohm bleeder which gives a calculated 24 second discharge time. If you need faster discharge, use a lower value.

You must also calculate the power dissipated in the bleeder resistor. Power is voltage squared divided by resistance, so for the 1,000 ohm bleeder we have: $P = (8 \times 8) / 1000 = 0.064$ watts, so a 1/4 watt resistor will stay cool.

I also recommend a 1,000 ohm, 1/4 watt resistor as a permanent load at the supply's output, to discharge any output capacitance in the supply and any filter capacitors which may be located at the device being powered by the supply. I have seen some regulator circuits which oscillate under no-load conditions. The output resistor prevents this condition. It will draw a constant 5 milliamps of current and dissipate 25 milliwatts of power (you do the calculation).

Figure 5 shows the two versions of our 5.0 volt supply. Notice I've added an optional protection diode across the regulator IC. This is usually not needed, but you should be aware that its purpose is to prevent the output of the regulator from becoming more positive than the input in the event that the load includes a large capacitor which discharges more slowly than C1 when the power switch (S1) is turned off. The diode conducts only if the voltage at the regulator's output becomes more positive than the voltage at the regulator's input, protecting the internal transistors in the regulator IC.

To complete our design we have one more consideration. Most three-terminal regulators require one or two ad-

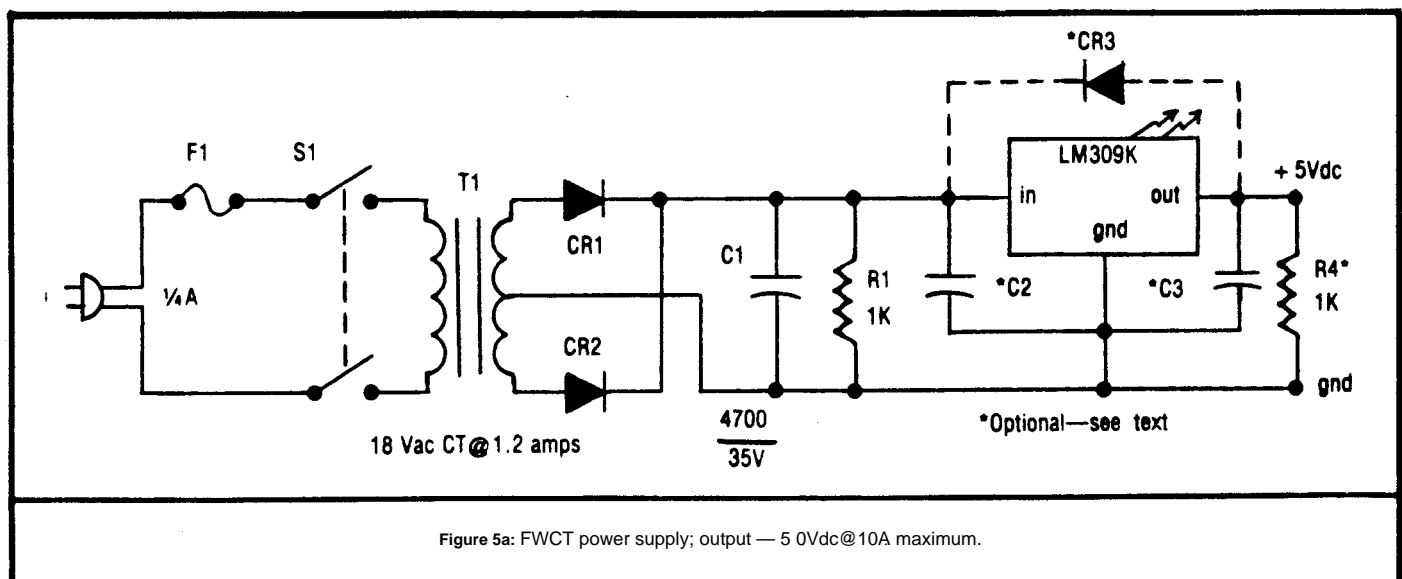


Figure 5a: FWCT power supply; output — 5.0Vdc@10A maximum.

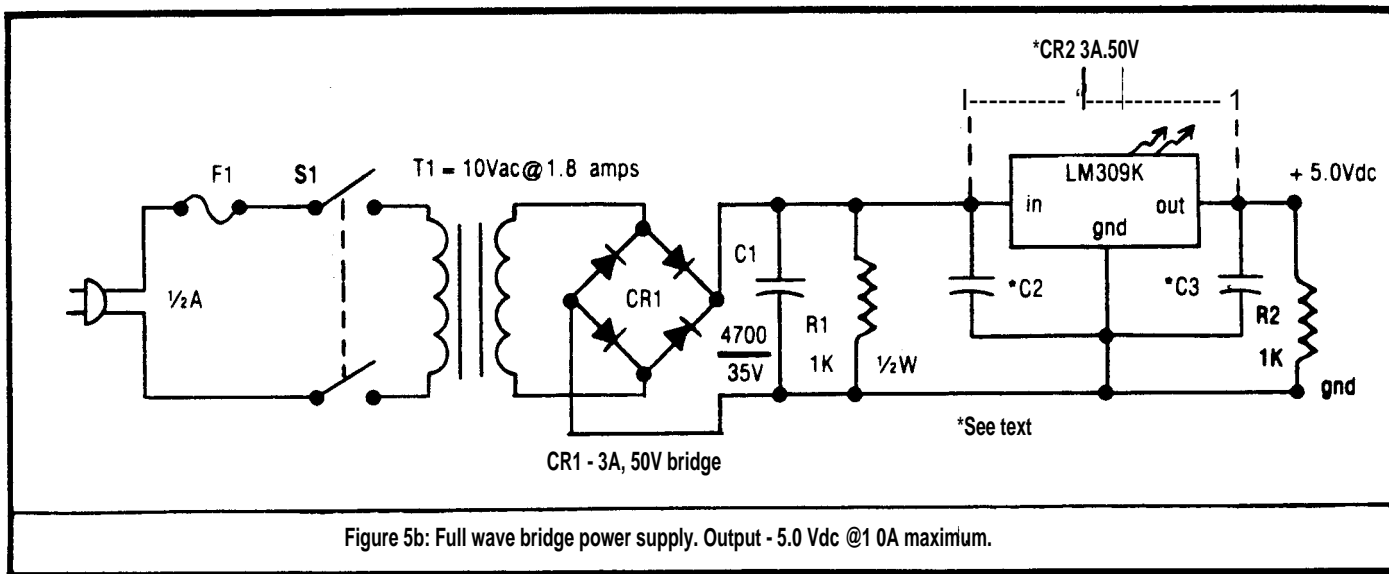


Figure 5b: Full wave bridge power supply. Output - 5.0 Vdc @1 0A maximum.

ditional small added capacitors to prevent oscillation. You'll have to check the data sheet for your particular regulator. If you can't find a data sheet, use a 0.22 microfarad disk ceramic or 1 microfarad solid tantalum capacitor as close as possible to the input lead of the IC, and a 1 microfarad solid tantalum or 25 microfarad aluminum electrolytic at the output lead. The type of capacitor is important, since we are concerned with the capacitor's ability to short-circuit high-frequency oscillations or transients; aluminum electrolytics have much higher impedance at high frequencies than do the solid tantalum types so more capacitance is required if you use an aluminum type.

The switch, SI, is a double-pole type rated at 126 vac, 1 amp. You should always switch both sides of the power line for safety's sake.

The fuse goes on the power line side of the switch, since even the switch can short-circuit. To calculate a fuse value with any accuracy requires graphical techniques; fuses are not the simple devices they appear to be. If you can find a good fuse catalog, take a look at the time versus current graphs. We can make a simple approximation, close enough for our purposes, by calculating the transformer primary current. This is simple, since the product of primary current and voltage equals the product of secondary current and voltage (for an ideal transformer). Solving for primary current yields:

$$I_{\text{primary}} = \frac{I_{\text{secondary}} \cdot V_{\text{secondary}}}{V_{\text{primary}}}$$

$$= \frac{1.8 \text{ amps} \cdot 18 \text{ volts}}{115 \text{ volts}} = .28 \text{ amps (FWB)}$$

$$= \frac{1.2 \text{ amps} \cdot 10 \text{ volts}}{115 \text{ volts}} = .104 \text{ amps (FWCT)}$$


We can use a 1/2 amp fuse for the bridge design and a 1/4 amp fuse for the dual-diode center-tapped design. The fuse must be able to tolerate the inrush current - the initial surge of current to charge the fully depleted filter capacitor when the supply is first turned on. You may find it necessary to use a slow-blow fuse for this reason.

Conclusion

I've tried to present just enough of the concepts and formulas to enable a novice to design simple low-current, low-voltage power supplies of the type most often used in microcomputer circuit projects. You should be able to use the circuits and formulas given here to design a supply meeting other voltage and current requirements. I hope I've presented enough of the concepts to keep you out of trouble. Always, I urge you to draw it out on paper, complete with calculations, then build and test it using dummy load resistors. Measure all the voltages, then calculate all the currents and power dissipations to make sure the real world conforms closely enough to our approximations. If you have an oscilloscope, take the time to check out the waveforms in various parts of the circuit, just for your own edification, but don't connect the scope to the primary side of the transformer or you could get a nasty surprise.

In future KISS articles I'll present dual-voltage circuits, very low-current circuits, variable-output voltage supplies and other embellishments on the basic designs presented this month. I would also like to back-track and cover in more detail some of the devices (rectifiers, capacitors, transformers and regulators) used in these circuits. But before I do so, I would like a better idea of what you need to know. Please write to me in care of *The Computer Journal* with any questions or topic suggestions you may have.

Low Cost, Modular Encoders



from a priori let
you easily build
and control:

- ROBOTS
- XY TABLES
- VEHICLES
- MOTORS
- CONVEYORS
and more.

This modular approach to position and direction sensing of rotary motion lets the OEM or robotic designer fit the encoder to the job, instead of fitting the job to the encoder. PU-2/DISC and IM-1/UNI-DISC form a flexible, low cost, precision dual-track encoder. Both packages output a 50 cycle/rev quadrature TTL square wave signal. OSPB processes these signals and gives up/down pulses at 1, 2 or 4 times encoder resolution. A unique collet design lets DISC/UNI-DISC mount on standard shaft sizes of .187", .250", .375" and .500".

1. DISC-(shaft size) 2.25" encoder used with PU-2 \$29.95
2. UNIDISC-(shaft size) 1.85" encoder used with IM-1 \$24.95
3. PU-2 Retroreflective IR pick-up electronics for DISC. \$29.95
Includes 4-pin connector. 2" by 2" pc board.
4. IM-1 Through-beam IR pick-up electronics for UNIDISC. \$24.95
Includes 4-pin connector. 2" by 2" pc board.
5. OSPB-1X, -2X or -4X use with PU-2, IM 1 or \$29.95
any dual track encoder. Outputs: UP/DOWN and PULSE, or
UP and DOWN (add DC). Includes 4 and 5 pin connectors.

Call or write for brochure. It describes the above parts, plus linear encoders, sprockets and thin, flexible metal belts for drive systems.

Terms: Purchase Order or check. Add \$3.00 shipping and handling Wash, residents add 7.9% sales tax.

a priori

(206) 630-2144

21318125th S.E. Kent, WA 98031

Classified

Rate: \$.50 per word, minimum charge \$7.50. All classified ads must be paid in advance, and will be published in the next available issue. No checking copies or proofs are supplied.

68 Experiments with Heath's Robot "Hero" — \$21.95 plus \$2.00 shipping. Experiment books on 8085/8051. Boyet, 14 E. 8th Street, New York, NY 10003

Wanted: Teletype KSR-35 manuals needed to restore old teletype machine. Also need manuals for paper tape punch and reader. The Computer Journal, PO Box 1697, Kalispell, MT 59903-1697

For Sale: SSM 10/4 board for S-100 bus. Two serial, two parallel ports. \$100. DEC LSI-11 minicomputer. Rack mount. KD11-F processor with KEV11 hardware math chip, DLV11 serial card, DRV11 parallel card. Total of 48K RAM. Paper Tape O.S. \$995. Write Lance Rose, c/o The Computer Journal, Box 1697, Kalispell, MT 59903-1697.

Multi-users, continued from page 10

an equivalent (say, four station) timesharing computer. Each piece of equipment that is added to the network has to have the necessary additional hardware and software that enables it to communicate with everyone else. And each workstation will likely be a complete computer rather than just an intelligent terminal. The only way you can really evaluate the cost is to consider two otherwise similar systems side by side.

There are two other major differences between timesharing and networking, which, at the present state of the art, mutually exclude one another. These are speed and distance. There is no inherent limit on the distance to which a timesharing system can be expanded, using (for example) telephone lines. This is a very slow process. Networks are almost unbelievably faster but to gain all of the network advantages (again, at the present state of the art) the stations must be located within a few thousand feet of one another. If extended much further they slow down to timesharing speeds.

Multi-processor Systems

A multi-processor is functionally somewhere between a timesharing system and a network with regard to both speed and distance, as well as cost and overall performance. Small timesharing computers are very limited, and networks are relatively expensive. The multi-processor gives you some of both worlds.

These three major multi-user concepts have a lot in common application-wise and also exhibit a lot of differences in how the objectives are obtained.

When you think about 'systems' with more than one user, remember the following definitions:

MU/MT (multi-user, multi-task) systems share a common processor and, usually, memory and peripherals.

MU/MP (multi-user, multi-processor) systems provide each user with his own processor, memory, and sometimes peripherals.

LAN (Local Area Networks) are usually completely separate computers, in close proximity and in more-or-less constant communication with one another, who may share some peripheral equipment as well as one another's software and files.

There are advantages and disadvantages inherent to any of the different multi-computing schemes. In future columns we will discuss some of the specific hardware and software that is available to implement them. The next multi-user column will investigate multi-processors in some detail.

Your questions, comments, and experiences with multi-users systems are welcome. Send them to E. G. Brouner in care of The Computer Journal

Help

This space is for you. We encourage you to communicate with other readers through this column by asking for their help with your problems, and by writing in with your solutions to questions like "Where can I...?" or "How can I...?" As a forum for sharing hands-on experience, this section can be an important resource for you. We will try to keep the lead time short for a rapid exchange of information. Let us hear from you!

Dear Computer Journal,

In response to "The CP/M Operating System" by M. Mosher, there are several points which I felt needed clarification. First, the S-100 microcomputer is not specific to CP/M-80 unless it has an 8080 or Z80 CPU; S-100 itself refers to a bus, not a system or CPU.

Second, CP/M-80 is specific to an 8080 CPU and can not operate in any other environment. The Z80 actually has the 8080 instruction set as a subset and thus runs CP/M-80 but CP/M-80 does not use all the instructions available on a Z80 CPU. It actually uses only the 8080 subset in the Z80 as it apparently does with the 8085 CPU.

There is an operating system that runs with CP/M-80 programs but is specific to a Z80 CPU and uses instructions available to the Z80 that are not available to the 8080 as it is running on the superset instructions.

This operating system is called "Turbo Dos" and it comes in two versions, single user and multi user. It can run CP/M programs better and faster than the CP/M-80 operating system but only on Z80 CPU systems. The single user system runs on a single CPU and the multi user system runs on a master and slave or master and multi slave CPU combination. All CPUs are Z80s. The advantages are manifold; more data placed on the disks, faster operating speeds by using the Z80 CPU to its fullest ability, spooling data and other enhancements over CP/M-80.

Third, in transferring CP/M-80 to an Apple II there had better be an 8080/Z80 CPU in the Apple II system to read the instructions as an additional card since the Apple II has a 6800 family CPU which is not compatible with the 8080/Z80 instruction set.

I would like to see some articles on the new 68000 16/32 bit CPUs, Turbo Dos, UNIX, and "C" language as do it yourself or tutorial articles. I felt that Mosher's article was most illuminating and well written. It should be pointed out that there is a wealth of public domain (free) CP/M software programs available from user groups and independent persons.

Sincerely,
Ronald Gillen
Wisconsin

Dear Mr Gillen,

Although it isn't stated specifically in my article, you are correct in noting that the S-100 bus is not limited to the CP/M operating system and is simply a hardware, not a software standard. I would disagree, however, that CP/M

cannot operate in any other environment than the 8080 processor. Although written originally for the 8080, CP/M can still be used in an 8085 or Z80 based hardware system since both of these processors contain the 8080 instruction set as a subset. There may be better operating systems available for Z80 systems but this does not preclude the use of CP/M in those systems. Finally, you are also correct in pointing out that an Apple II (or any other system for that matter) must have an 8080, 8085, or Z80 to run the CP/M operating system.

Sincerely
M. Mosher

Dear Computer Journal,

Along with your information, I am interested in training myself in programming, assembly language, and interfacing microprocessors (esp. Z80) for industrial application. What trainer or self teach method would you suggest?

Rob Babcock
Indiana

Do our readers have any suggestions?????

Micro Use Restricted

We are just beginning to see the effects of Micros on society and our lifestyles, and the bureaucrats are already trying to control the use of Micros.

Many companies are allowing their employees to work at home with a Micro instead of having to travel to the office, and most freelance writers will be using Micros with wordprocessors. It has been predicted that the changes in our lifestyles due to the widespread use of Micros will be as profound as the changes caused by the introduction of the automobile.

The December issue of *73 Magazine* reports a recent case in Chicago where the city declared it illegal for people to use a personal computer in the home for business. And that would include writing articles for pay!

We would appreciate any additional information from our readers on this case or any other restrictions on the use of Micros. We have contacted the Chicago Chamber of Commerce and the Mayor's office for further details and an explanation of the reasoning behind this action, but I wouldn't bet on a speedy or meaningful reply.

Letters From Our Readers

Dear Computer Journal,

As per your invitation, I'm listing my interests relevant (I hope) to your publication. They include the following: stereo remote control (for non-remote units), designing your own digital readout thermostat, microprocessor based burglar alarm systems, voice synthesis and recognition, fibre optics and/or FM intercom for the house, homemade garage door openers with wireless remote, controlling current consumption of house appliances and lights via microcomputer monitoring, robotic arms, legs, hands, vision and hearing, character recognition for the experimenter, and adding computer control to your car.

Although not all topics mentioned would be practical to implement, a discussion of their inherent problems would be satisfying in itself. Thank-you for this invitation and your publication.

Sincerely,
Bryan Greifinger
New York

Dear Mr. Rose,

I read your article "File Transfer Programs for CP/M" in *The Computer Hacker* with interest. Before attempting an Apple/Commodore 64 transfer requiring keying in the programs, I request information on the availability of either

cassettes or floppys. I've enclosed a SASE for your response.

Yours Truly,
Jerry Hulman
Maryland

Dear Jerry,

Regarding your letter on the file transfer programs, I'm not quite sure whether you want to try using them with the Apple DOS or Commodore 64 operating system. The programs only work with the CP/M disk operating system and require an 8080, 8085, or Z80 processor on both machines. They won't work with the 6502 found in both the Apple and the Commodore 64. Also, they can't be used with a cassette system. If you have a floppy disk drive and CP/M running on both your Apple and Commodore, we could probably provide you with a copy of the source programs on an Apple CP/M diskette for the same price listed in the article for the 8-inch floppy (\$15).

Hope this information helps you out. We're glad to see you are reading (and hopefully enjoying) *The Computer Journal*. Thanks for writing.

Regards,
Lance Rose
Technical Editor

New Products

Searchmart Corporation, a South Florida firm specializing in database development and information retrieval systems, will offer a Free Access Software Library that lists, describes and demonstrates tens of thousands of individual applications and systems software packages on-line.

"Printed catalogs and directories of software packages are obsolete in this computer age," says Searchmart's president, Victor Gruneau. "With dozens of new software programs offered daily and hundreds monthly, all print listings are out of date the day they are published."

Searchmart's concept is a simple one—an on-line library of systems and applications software up-dated daily and categorized by manufacturer, publisher or vendor, operating systems compatibility, protocol requirements, program classification, features, price and ordering information.

What makes it revolutionary is the *free* on-line access to the software database. "There are several services with software databases," states Gruneau, "but they charge substantial fees for making searches and they are not

available on-line to software shoppers who want to search the files on home or office terminals at their convenience."

Searchmart's Free Access On-Line Software Library allows anyone with data communications capability to search the software database. The telephone numbers to Searchmart's computers will be publicized in a direct mail and computer magazine advertising campaign to begin in January, 1984.

As Gruneau explains it, software manufacturers and vendors will have the opportunity to describe their products and companies on "pages", each page a 40 character by 20 line CRT screen. "They'll have the opportunity to give the software shopper as much information as they want—even demonstrations—and at a very modest cost per page."

Searchmart plans to be on-line with the Software Library by January 1, 1984. For more information, direct inquiries to Mary K. Hamm, Marketing Services Director, Searchmart Corporation, 636 U.S. Highway 1, Suite 210, North Palm Beach, Florida 33408.