

Designer's Guide to: Temperature control

Nine working circuits and a list of hints emphasize the practical aspects of thermal control systems.

Jim Williams, Dept. of Nutrition and Food Science, Massachusetts Institute of Technology

Once you have selected a temperature sensor (Part 1 of this series, EDN 5/5/77) and developed a measurement scheme (Part 2, EDN 5/20/77), it's theoretically a simple task to create a control loop by well known servo-feedback techniques. Unfortunately, the long time constants inherent in thermal systems present circuit designers with significant challenges. Nowhere is the close relationship between servo systems and oscillators more apparent than in thermal control systems.

A thermal control loop can be simply modeled as a network of resistors and capacitors, where the resistors are equivalent to thermal resistances and the capacitors to thermal capacity. In **Fig. 1**,

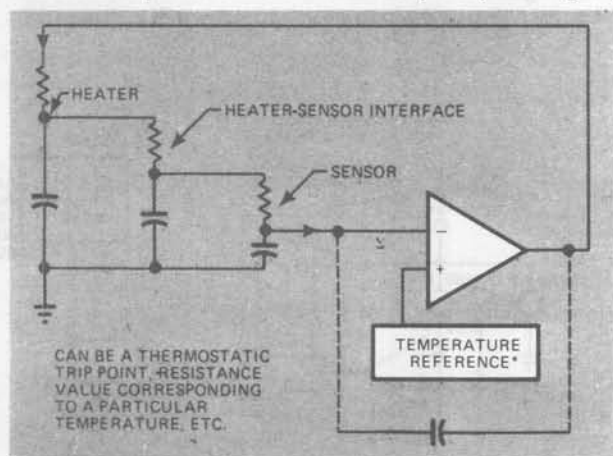


Fig. 1—Electrical model of a thermal-control system shows the RC factors that contribute to the delay in its response to driver inputs.

observe that the heater, heater-sensor interface and sensor all have RC factors that contribute to a lumped delay in a thermal system's ability to respond to its driving device. The designer must minimize these delays, and in so doing he must make a number of tradeoffs. The heater element's physical size and electrical resistivity determine its time constant. And the heater-sensor interface time constant can be minimized by placing the sensor in intimate contact with the heater.

Keep the thermal time constant small

We cannot overstate the importance of a minimized heater-sensor thermal time constant, because this interface is the most commonly mis-designed part of thermal-control systems. When working with a still-air oven, you should make every attempt to mate the sensor very tightly to the heater. With baths the same provision applies, but some relaxation can be tolerated if you use adequate and truly isothermal mechanical stirring of the liquid. You can operate forced air ovens with the sensor somewhat removed from the heater, but only if the air moves fast enough from the heater to the sensor to inhibit significant phase shift.

Generally speaking, in a well designed thermal-control system *the heater and sensor are always at the same temperature.*

You can minimize the sensor RC product by selecting a sensor of small size relative to the capacity of its thermal environment. Clearly, if the wall of your oven is 6-in. thick aluminum, your sensor need not be the smallest available size. Conversely, if you are controlling the temperature of a 1/16-in. thick glass microscope slide, you need a very tiny sensor (i.e., a fast one).

After minimizing the thermal time constants relating to the heater and sensor, you should choose some form of insulation for the system. The type of thermal load will dictate the optimum type and amount of insulation. Remember that the function of insulation is to keep the loss rate down so the temperature-control device can keep up with the losses. For any given system, *the higher the ratio between the insulation time constants and the heater-sensor time constants, the better the performance of the control loop.*

You can compensate thermal loops by adjusting the gain of the control element, limiting its bandwidth to less than the thermal delays, or both. Such practice is consistent with well-known feedback theory. A single bandwidth-limiting element and a single gain-setting adjustment should suffice to stabilize any system. Systems that employ numerous gain and bandwidth ad-

as easily as an electrical heater to achieve minimum overshoot.

The circuit illustrates the latter situation. It functions by comparing the output of the AD590 IC temperature transducer with a stable voltage reference. When the 311 comparator trips at the set-point temperature, the 555 one-shot causes the transistors to turn on the gas solenoid and light the 40 kW burner. When the 555 times out, the burner is extinguished—regardless of the 311's output state. The time constant of the 555 compensates for lags in the system by turning off the heater before the phase-lagged sensor actually hits cutoff value.

As the ambient temperature around the thermal load changes, the amount of time the 555 is on must also change to maintain optimum control. Why? Because the 555 effectively breaks the control loop. As the ambient temperature changes, the dissipation constant of the system varies and the 555's output pulse width must vary correspondingly. The thermistor across the timing resistor serves this function, sensing ambient temperature and providing proper time-constant trimming. It is mounted in an aluminum block whose time constant ensures proper phasing.

Fig. 4's basic approach has been used to control a 5000 gal. vat at 100°C to 0.1°C. It has also provided 0.001°C control for a quartz delay-line oven in a retrofit application where heater and power-supply parameters could not be altered.

Phase-fired systems are efficient

Any thermostatic control loop, even one as sophisticated as Fig. 4, still suffers from its unrealistic correction characteristic. Thermostatic systems tend to control about, rather than at, a desired set point. Since energy is leaving the controlled environment at all times, it follows that even the most carefully optimized thermostatic circuit cannot provide the best possible performance for a given application. Phase-fired control circuits, however, can more closely approximate a linear system.

In phase-fired circuits an ac waveform across the heater is chopped at a servo-controlled phase angle to provide nearly linear response. Since thermal control-loop response crosses over long before 60 Hz, the assumption is that the thermal components of the loop will nearly perfectly integrate the chopped heater drive into a thermal dc level. Typically, phase-fired systems use SCR's as choppers and are capable of good efficiency at high power.

In Fig. 5, the AD580 band-gap voltage reference furnishes power to a bridge circuit, while the platinum sensor provides the sensing function. The AD504 amplifies the bridge output; its 110k feedback resistor allows maximum loop gain without oscillation. The op amp's output biases a 2N2907 whose collector potential controls a 60 Hz synchronized unijunction transistor (UJT) oscillator. The UJT's output drives the SCR gate via an

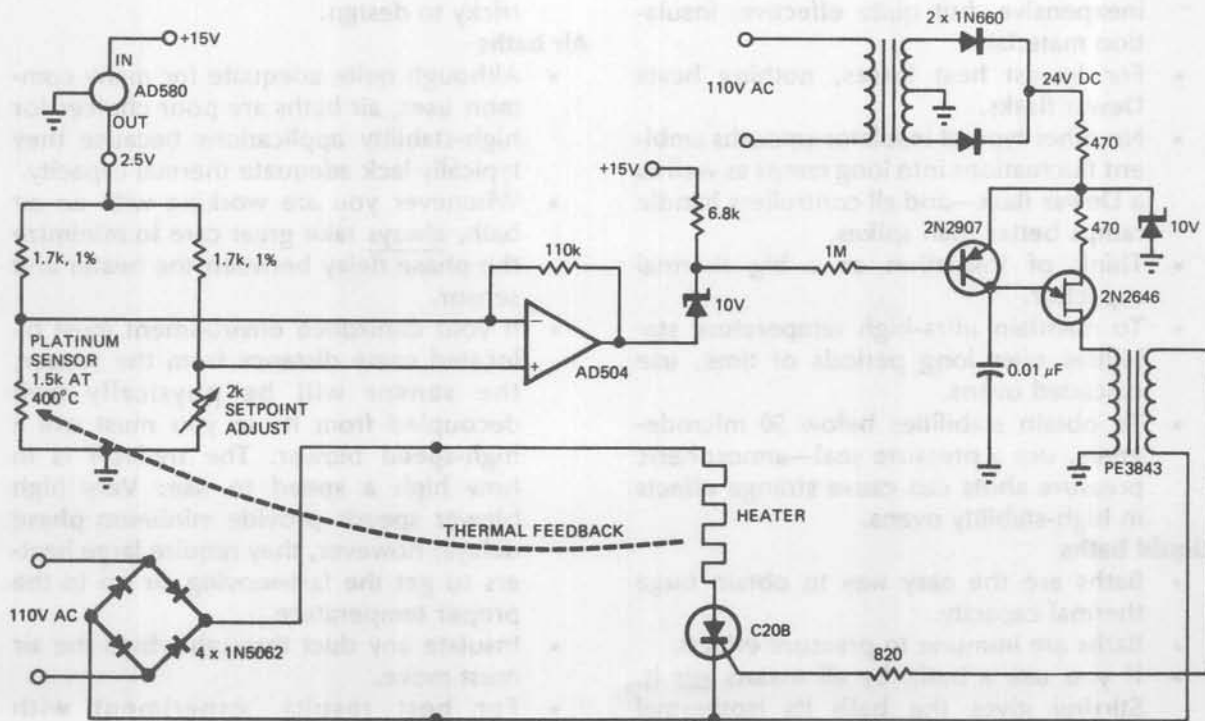


Fig. 5—Servo-controlled phase-angle firing of an SCR can provide linear thermal control to 0.001° stabilities with a properly designed thermal load. Beware of the RFI generated, however.

isolation transformer.

This system provides analog temperature control as the AD504 biases the 2N2907-2N2646 combination to fire at different points on the ac waveform. The 2N2646 always runs at 60 Hz, but the op amp's servo action controls the phase angle at which it fires. The integral of the ac waveform across the heater equals the dc level of the energy leaving the oven.

Phase-fired circuits of this type should be arranged so that the SCR chops at the peak of the waveform, thereby ensuring the highest gain. And the higher the gain, the more closely this circuit can approximate "energy in equals energy out" operation. The gain, in turn, is limited by that eternal problem of thermal control loops—phase delay.

Typically, a circuit of this type can exceed 0.001° stabilities if the thermal load is properly designed.

Thermal-control system design hints

Ovens

- Always pay careful attention to oven geometry.
- Thick walls help reduce heat losses.
- Spheres, though cumbersome, are the best choice for a super stable oven.
- Cylindrical ovens provide good performance, but the heater wires must be wound more densely on the ends to compensate for losses there.
- Styrofoam and fiberglass are relatively inexpensive, but quite effective, insulation materials.
- For lowest heat losses, nothing beats Dewar flasks.
- No other type of insulator smooths ambient fluctuations into long ramps as well as a Dewar flask—and all controllers handle ramps better than spikes.
- Think of insulation as a big thermal capacitor.
- To maintain ultra-high temperature stabilities over long periods of time, use cascaded ovens.
- To obtain stabilities below 50 microdegrees, use a pressure seal—atmospheric pressure shifts can cause strange effects in high-stability ovens.

Liquid baths

- Baths are the easy way to obtain huge thermal capacity.
- Baths are immune to pressure effects.
- If you use a bath, by all means stir it. Stirring gives the bath its isothermal integrating characteristics.

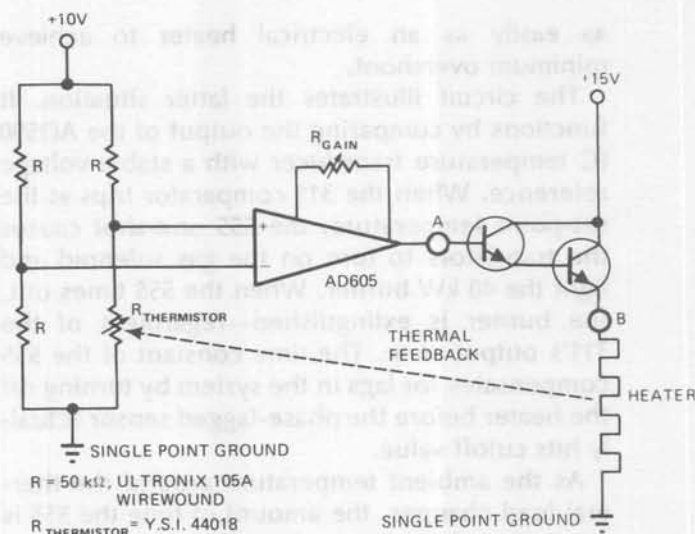


Fig. 6—Instrumentation-amplifier gain resistor sets the loop gain in this linear thermal controller. Circuit can control an oven or bath to stabilities of 75 microdegrees.

- Use caution in the design of a stirred bath. Carefully evaluate propellor type and placement, as well as the type, size, shape and placement of all baffles. In particular, watch out for layering effects.
- Heater placement deserves extra thought and experimentation. Consider evenly distributed heater schemes such as the concentric spiral.
- Multisensor schemes can take advantage of a bath's fast integration property. However, such configurations can be tricky to design.

Air baths

- Although quite adequate for many common uses, air baths are poor choices for high-stability applications because they typically lack adequate thermal capacity.
- Whenever you are working with an air bath, always take great care to minimize the phase delay between the heater and sensor.
- If your controlled environment must be located some distance from the sensor, the sensor will be physically well decoupled from it, so you must use a high-speed blower. The tradeoff is in how high a speed to use: Very high blower speeds provide minimum phase delays; however, they require large heaters to get the fast-moving air up to the proper temperature.
- Insulate any duct through which the air must move.
- For best results, experiment with "shaped" heaters and ducts.

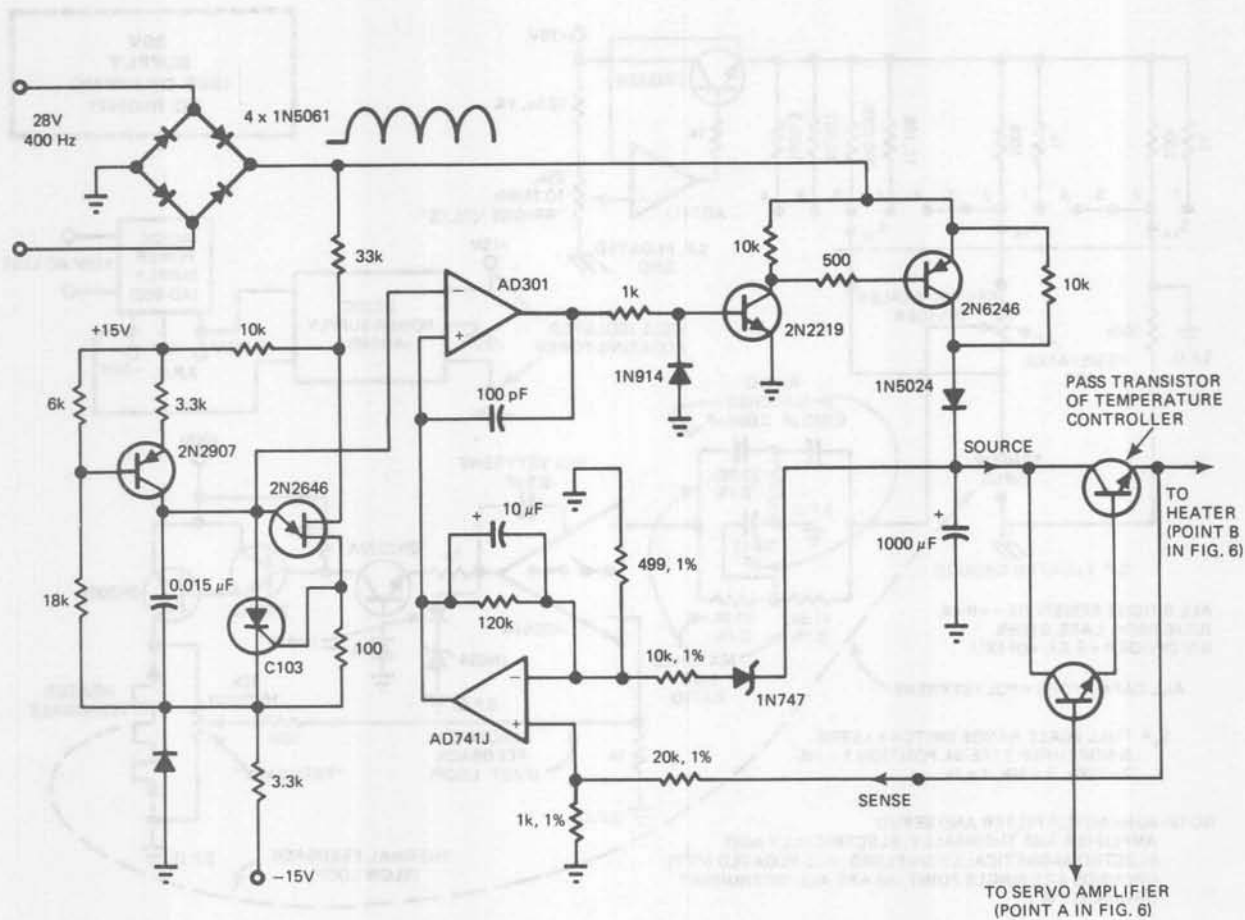


Fig. 7—To reduce the power dissipated in Fig. 6's pass transistor, this circuit maintains a constant 3V across it.

However, note that it generates significant RF noise because the SCR chops in the middle of the waveform. This noise can be especially troublesome if the circuit serves to control an oven containing RFI-sensitive low-level electronics.

For applications demanding the very best

Certain situations require ultra-stable temperature control that can only result from a true linear system. Zener reference and quartz crystal time-drift studies, laser micro-interferometry, crystal growing and (especially) biochemical microcalorimetry have benefitted from the development of ultra-stable thermal control. Temperature controllers with a stability of 3 microdegrees at 50°C have been constructed. That's just 0.06-ppm drift!

To achieve such spectacular performance, you must use components with the ultimate in specifications and pay strict and careful attention to every detail of the thermal system. Fig. 6 shows a circuit that will control an oven or bath to stabilities of about 75 microdegrees at 37.5°C.

In operation, the bridge offset is sensed by an instrumentation amplifier that features high common-mode rejection, low drift and balanced inputs. The amplifier drives a Darlington pair that provides current gain to the heater. Loop gain,

controlled by the gain-setting resistor at the instrumentation amplifier, is adjusted so that the loop is stable and a dc voltage appears across the heater.

One of the penalties of this circuit is that the pass transistor must dissipate any power not delivered to the heater. Since an ample collector supply must be available for turn-on and low-temperature ambient conditions, the device may have to handle a sizeable amount of power within many operating environments.

Fig. 7 shows an accessory circuit that you can add to Fig. 6's controller to eliminate this problem. Developed for an aerospace application that requires high efficiency, the circuit is essentially a servo that maintains a constant 3V across the pass transistor regardless of the controller's demand for power.

The 741 looks differentially across the transistor, with its negative input biased through a 3.3V zener. The op amp's output is compared to a 400 Hz line-synchronized ramp by a 301A amplifier that functions as a pulse-width modulator and drives a 2N2219-2N6246 switch. In turn, the switch delivers phase-controlled power to the 1000 μF integrating capacitor and the temperature-controller pass transistor. The 1N5624 diode

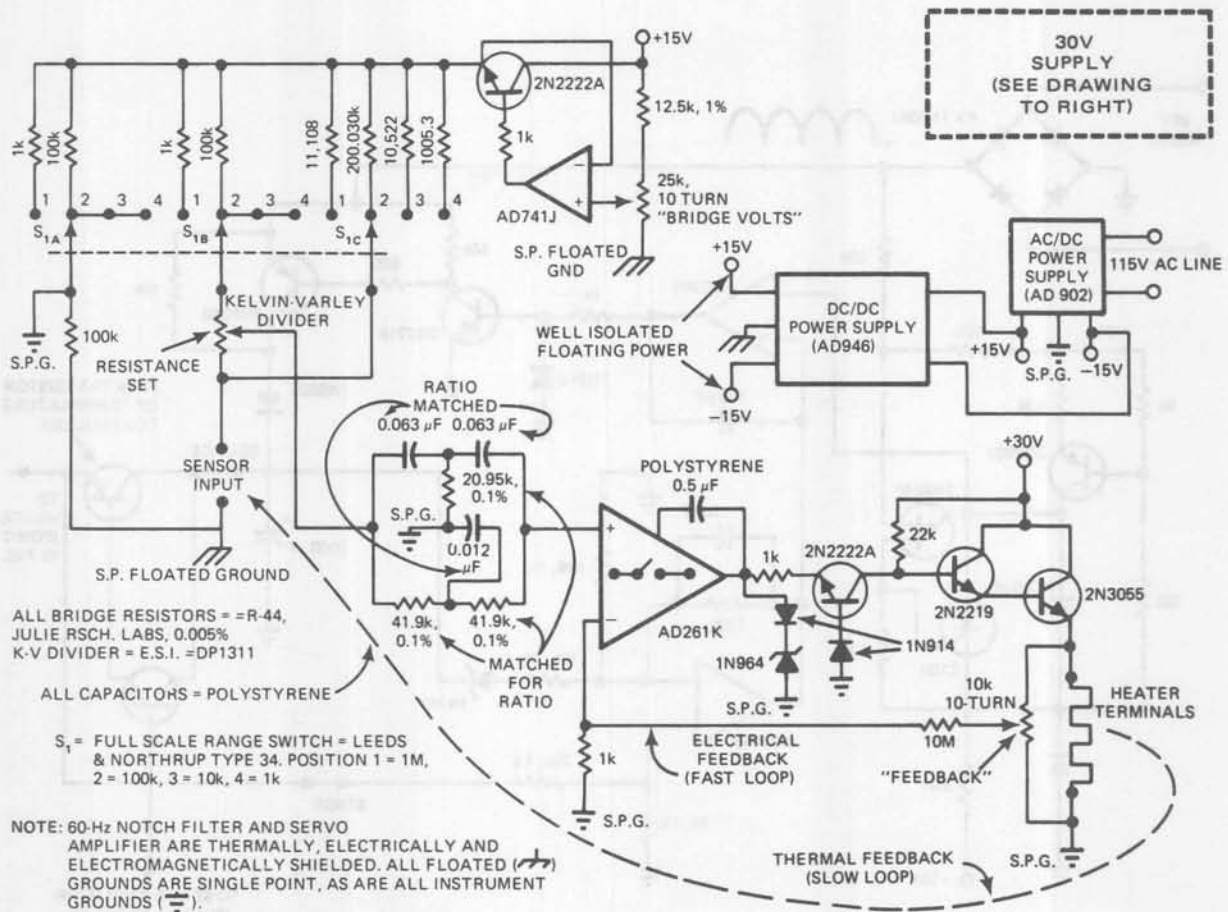


Fig. 8—Kelvin-Varley divider dials out the resistance control point directly in this circuit that accommodates sensor resistances from 10Ω to 1 MΩ. The feedback potentiometer provides adjustable global feedback.


ensures that the 2N6246 will not be reverse biased when the 400 Hz signal falls below the dc value of the 1000 μF capacitor.


This circuit avoids the unpleasant surprises that may arise when a servo loop is run within a servo loop by giving the dissipation-limiting servo a much slower time constant than the thermal-control servo. The 10 μF capacitor across the 301A

serves this function.

Settling in 2.5 sec for a 500 μ° step

Chopper-stabilization techniques and other refinements embodied in the circuit of **Fig. 8** provide the ultimate in temperature control. The circuit features a multiranging bridge that accommodates sensors from 10Ω to 1 MΩ—a Kelvin-





Varley divider dials the sensor-resistance control point directly to five digits.

A floated power supply drives the bridge. This approach allows the single-ended noninverting chopper-stabilized amplifier to take a differential measurement, and also completely eliminates the common-mode voltage error an instrumentation amplifier would contribute. A passive 60 Hz notch filter eliminates the pick-up noise of the floated bridge. The common base-amplifier transistor at the 261K's output provides sufficient amplification to permit the Darlington pair to drive up to 30V across the heater. Note that the feedback potentiometer is across the total of all gain elements in the circuit and supplies adjustable global feedback.

The thermal loop operates separately from the electrical loop, with the latter functioning to set the circuit gain at maximum; i.e., the feedback setting at which the system exhibits a nearly perfect response. Fig. 9 shows an optimum response waveform obtained by step changing the temperature set point by 500 microdegrees.




Fig. 9—Changing the temperature set point by 500 microdegrees causes Fig. 8's entire loop to settle within 2.5 sec, indicating that the gain is set for optimum response.




Fig. 10—High loop gain puts the circuit of Fig. 8 on the verge of oscillation. Ringing resulting from a 10 millidegree step set-point change lasts almost 15 min.!




Fig. 11—Stability of spherically shaped, Dewar enclosed, pressure-sealed oven is measured at 3.3 microdegrees over a 3-hr period, despite the sensor's 3 microdegree (calculated) noise figure. (Only part of the trace is shown.)

The same loop can also generate some extraordinary waveforms. When the loop gain was set so high that the circuit was on the verge of sustaining oscillation, a 10 millidegree step set-point change produced a damped, ringing response of almost 15 minutes duration (Fig. 10)!

Would you believe 3.3 μ° control?

High-quality control loops equipped with thermally optimized ovens and operating under tightly controlled laboratory conditions can provide amazing results. The strip chart of Fig. 11 shows the stability of a spherically shaped, Dewar enclosed, pressure-sealed oven (since $PV=nRT$, ultra-high stability control requires pressure sealing to avoid atmospheric pressure shifts from causing temperature deltas inside a fixed-volume oven). Observe that the servo holds the monitoring sensor in the oven to within just 3.3 millionths of a degree for over three hours. This proves especially significant because the noise figure for the particular sensor used was calculated at about

3 microdegrees.

Fig. 12 shows what happens to the stability baseline when a 20 nW pulse—generated by putting 20 mV across a 1 M Ω resistor for 15 sec—is released into the oven 1 inch from the monitoring sensor. Temperature rise at the sensor is 7 μ° peak; setting time, about 11 min. to within 2 μ° of the initial baseline. Note that this performance is achieved despite the fact that most of the energy released is not picked up by the monitoring sensor. (As a point of reference, a single human cell operates at about the nanowatt level.)

Other controllers provide special functions

Some temperature-control applications require other than simple stability. Fig. 13's circuit furnishes a linear temperature shift from 45°C to 45.2°C over a 10,000-hr period. It drives a controller similar to the one in Fig. 8 and consists of a voltage referenced half-bridge, a 16-bit digital-to-analog converter and an integrator.

Initially, the DAC's output is zero and the oven

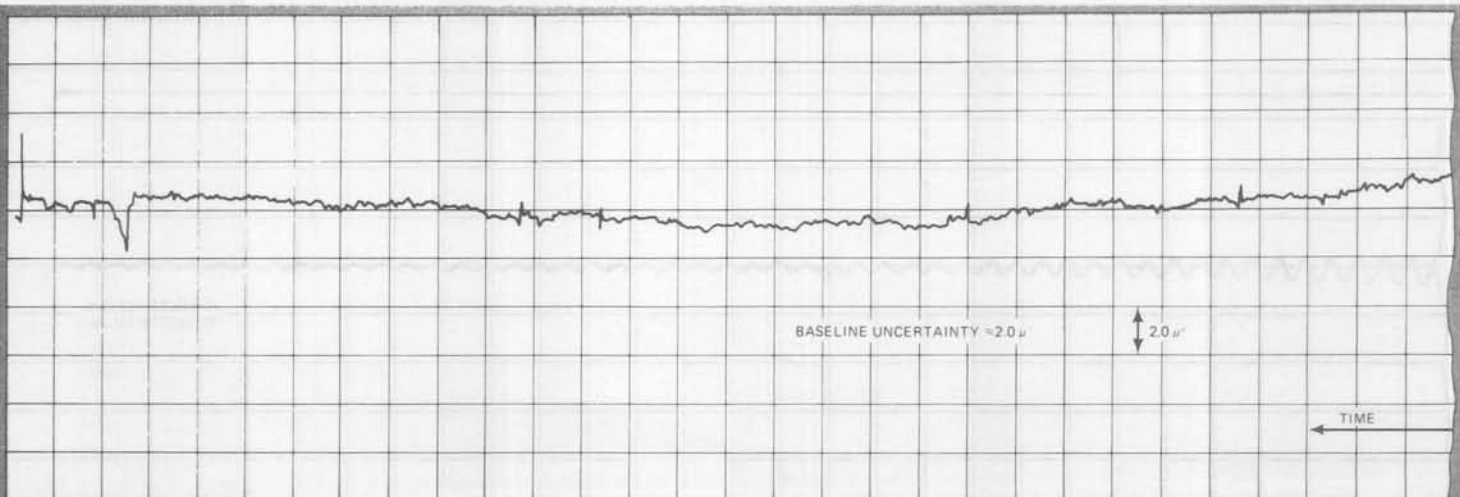


Fig. 12—A 20 nW pulse released 1 in. from the monitoring sensor into the oven described in Fig 11 produced this extraordinary trace. Peak temperature rise is 7 microdegrees.

is running at 45°C and stabilized. Then, every 10 min. the DAC is clocked one step, each time producing a set-point change of 3 microdegrees via the summing amplifier. The amplifier's output is integrated with a time constant slightly greater than 10 min., producing the absolutely smooth ramp that feeds the temperature controller's reference input. Over 10,000 hours time the temperature of the controlled environment slews from 45°C to 45.2°C in a highly linear fashion. Nonlinear functions can also be obtained by clocking the DAC aperiodically.

Other special-type controllers include feed-forward systems and heater-sensed controllers. In a feed-forward system the temperature of the ambient surrounding the controlled environment is sensed, differentiated with the appropriate time constant and summed with the servo sen-

sor's signal. The advantage of this type of system is that it gives the controller the "look ahead" capability it needs to offset large ambient shifts. Feed-forward systems particularly suit situations such as large buildings where the controlled environment lacks good insulation characteristics. In general, they are cost effective when designed properly, which proves no easy task. Results range from excellent to catastrophic.

Because heater-sensed controllers use the heater itself as a sensor, they offer very fast response and are perhaps more reliable and cost effective than separate sensor types. However, they are usually somewhat lacking in precision because most heater materials exhibit only small changes in resistivity with temperature. Most controllers of this type time-share the heater/sensor function. □

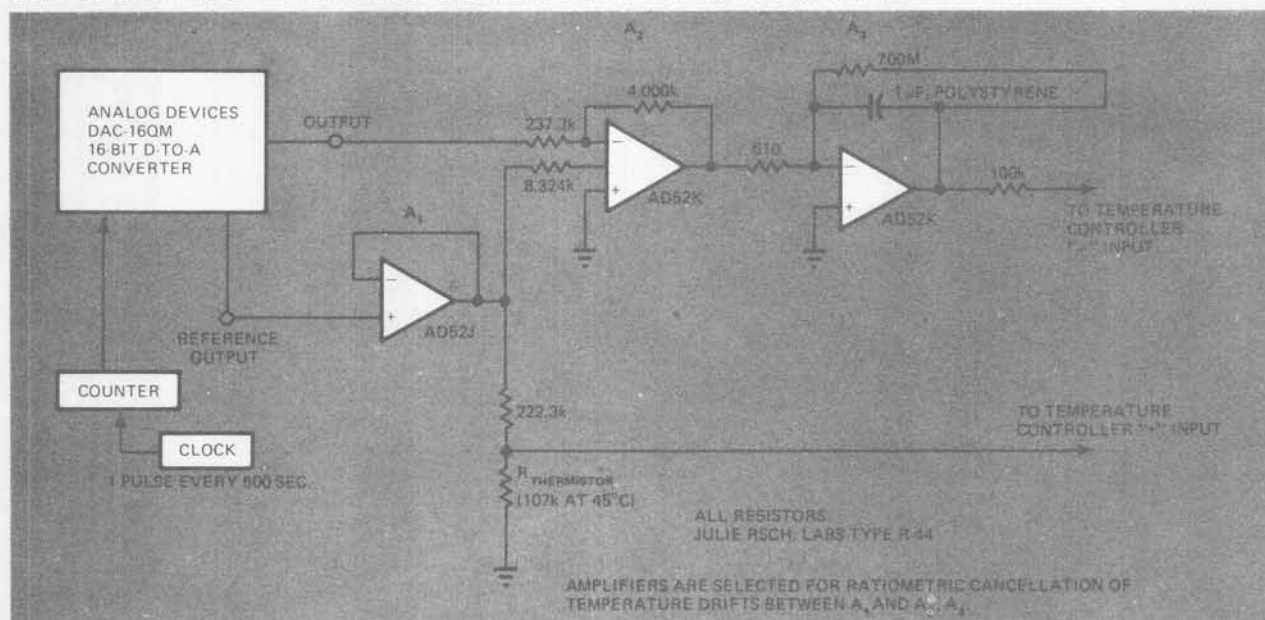



Fig. 13—It takes 10,000 hrs. to linearly ramp the oven temperature from 45.0 to 45.2°C with this circuit. The 16-bit DAC is clocked one step every 10 min. to produce a 3 μ° set-point change.



Bibliography

1. Fulton, S.P., "The Thermal Enzyme Probe," S.B. Thesis, M.I.T., Cambridge, MA, 1975.
2. Harvey, M.E., "Precision Temperature Controlled Water Bath," *Review of Scientific Instruments*, 39-1, 1968.
3. Krasner, M., "Precision Control of Temperature," Dept. of Electrical Engineering, M.I.T., Cambridge, MA, 1974.
4. Olson, J.V., "A High Stability Temperature Controlled Oven," S.B. Thesis, M.I.T., Cambridge, MA, 1974.
5. Williams, J.M., "Temperature Control of Microdegrees," Educational Research Center, M.I.T., Cambridge, MA, 1971.
6. Williams, J.M., "Portable Wide Range Chopper Stabilized Temperature Controller," Dept. of Nutrition and Food Science, M.I.T., Cambridge, MA, 1974.
7. Williams, J.M., "Split a Temperature Degree to 10 Microdegrees," *Electronic Design*, May, 1974.