RF and Optical Bolometer

When RF energy is delivered to a resistive load it dissipates heat. If the load has a relatively poor thermal coupling to its surrounding environment its temperature will rise. By measuring the temperature rise it is possible to determine the average power delivered to the load. There are various problems with the approach though, calibrating it is troublesome as ambient temperatures change and many of the coefficients involved are unknown and would need to be determined experimentally. It was during my design of the calibration system I discovered it is fairly easy to just put the temperature sensor and load into a servo loop and maintain a constant temperature above the ambient. In this way the change in DC power required to maintain the sensor at a constant temperature is directly the amount of power delivered to it by other means.

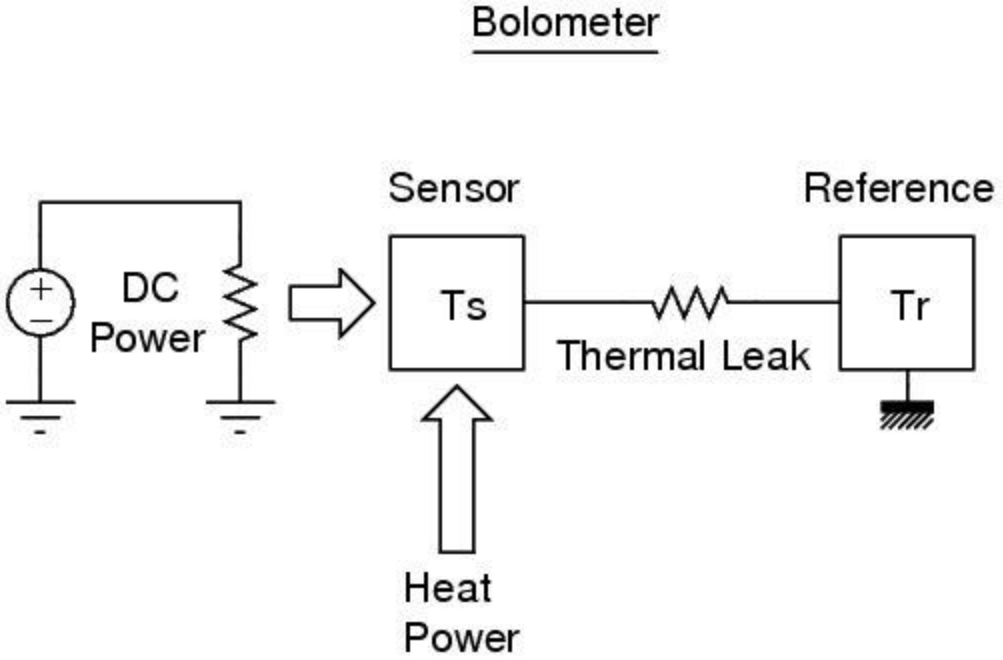

To achieve the thermal control system I started with thermistors but rapidly changed to using common silicon diodes who have a temperature coefficient of about 2 mV/K. An Op-Amp with offset compensation (a CA3130 was in the junkbox) amplifies the drop voltage difference between two diodes and drives a resistor heating element to keep one junction a constant temperature above the other. The amplifier gain is about 30 dB and is not completely open-loop except for the thermal feedback to give it some stability. The circuit in all is very simple and functions adequately for the task.

Thermal Power Balance Bolometer

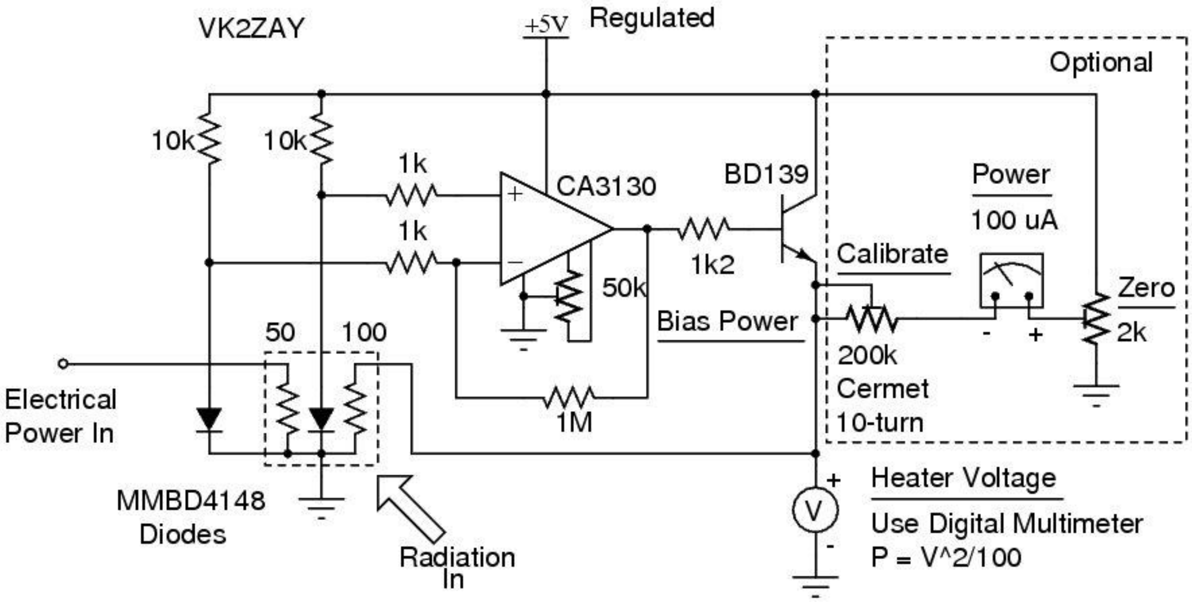

In the first version I used normal leaded components held together with a drop of liquid electrical tape. This arrangement allowed me to test the concept on a solderless breadboard. A centre-nulling meter facilitated "zeroing" the difference by biasing its other terminal to equal the quiescent voltage across the heater. This way I could watch the sensor drift and bias either sensor with my fingers resulting in a swing in the appropriate direction. Such an arrangement **does not** yield direct-reading of power, but can be calibrated at a particular ambient temperature and is handy for trending. For direct reading the absolute power levels in the heater load must be measured and subtracted to yield an accurate figure of power delivered by the external source.

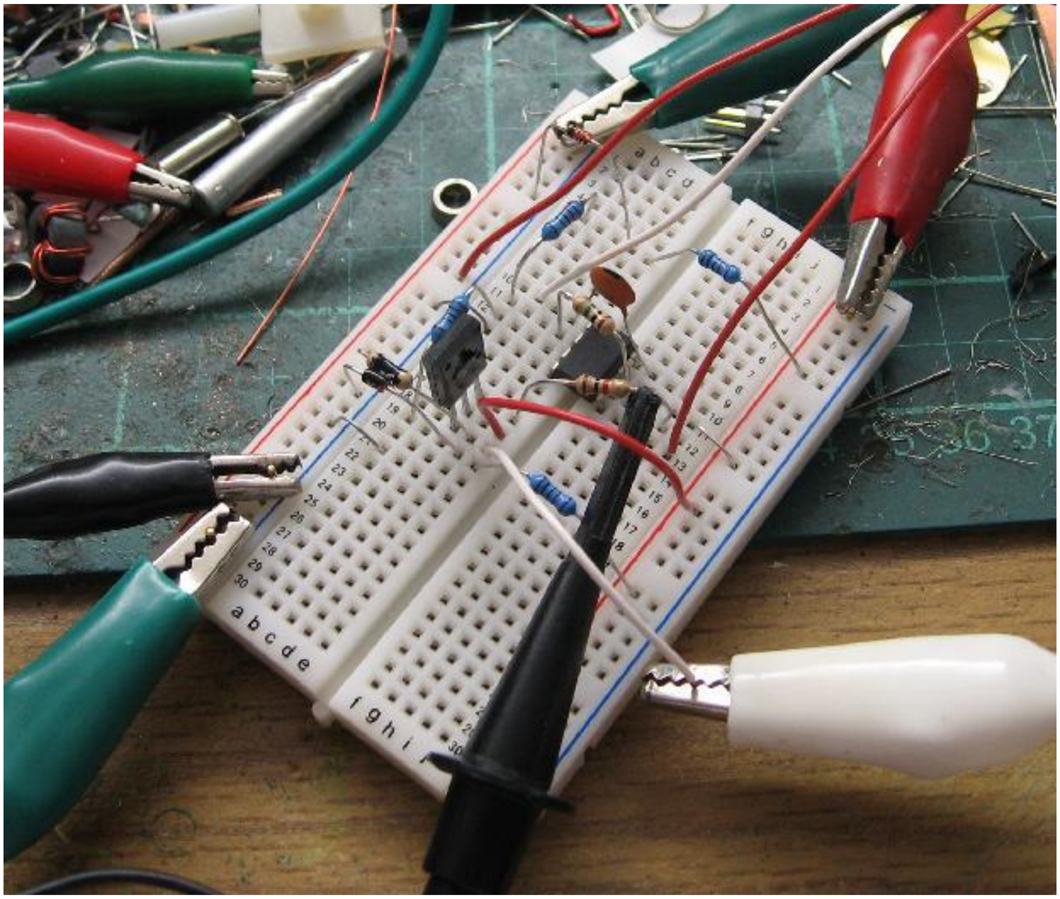

I considered using the same load for RF as the heater. In theory this is quite practical; an RF choke implementing a bias-Tee to deliver the DC heating power while the RF is delivered through a DC blocking capacitor. I wanted to be able to calibrate the device with DC (which is easily measured), so I went for a dual heater design. A separate 50 Ohm load is thermally coupled to the same heater/diode pair used in the control loop. This mandates calibration as the heating effect of the current in each load may not be precisely communicated at the same level to the diode but in practice it was shown to work very well. For RF-only measurement the single heater/sensor system is simple and direct reading (but may require a high heater supply compliance to drive the 50 Ohm load).

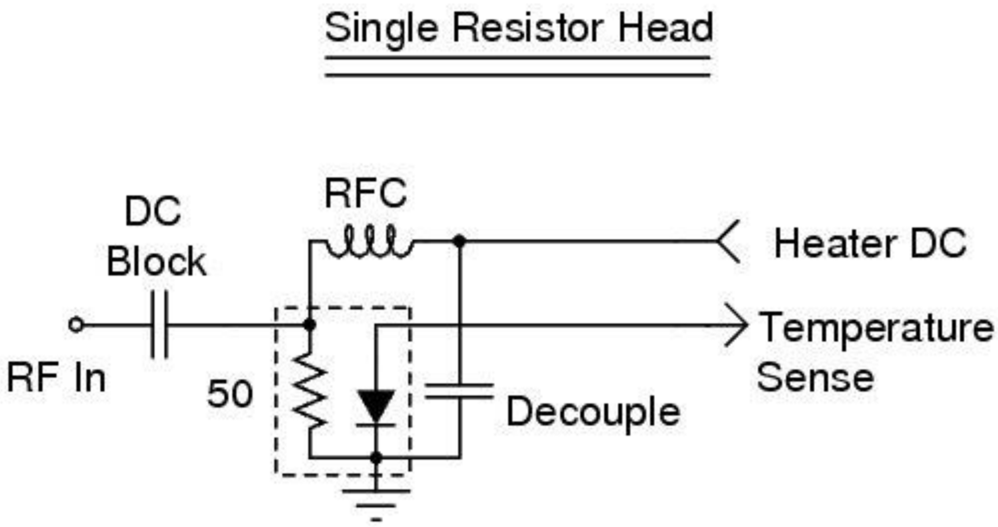

The power delivered to the load/sensor system need not be via the resistor. I discovered external heat sources (like my soldering iron) would produce measurable deflections from quite some distance. This enables direct measurement of electromagnetic radiation that the sensor can absorb. I tried aiming a toy laser pointer at the initial leaded detector arrangement and measured roughly 2 mW of dissipation difference, this seemed consistent with its < 5 mW compliance labelling.

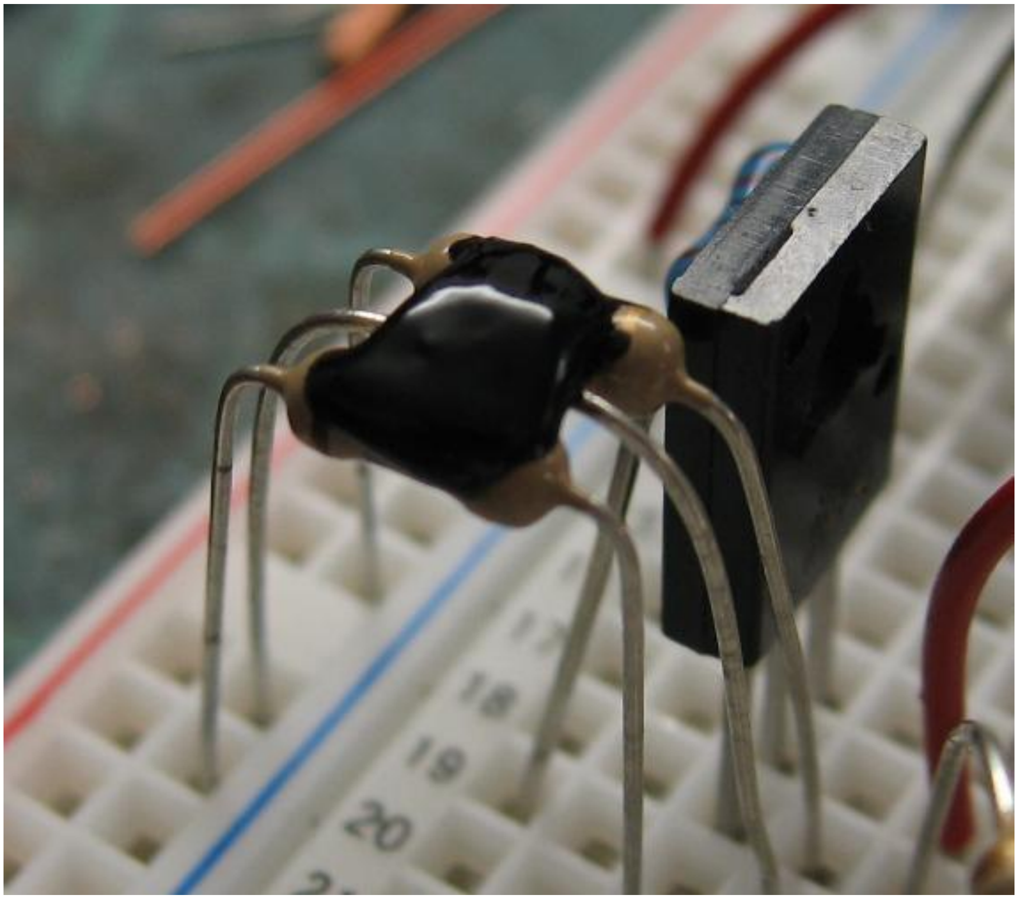

The time constant of the lashed-up detector was fairly long, and to minimise this I decided to build a more physically compact detector. I used SMD components soldered to a small square of brass shim stock. I reasoned the surface area of the plate could be measured and used to calculate optical fluxes. I soldered the SMDs directly to each other after bonding them together with superglue and then soldered their common ground to the plate. The entire assembly was then suspended over a larger brass plate which holds the reference diode. The larger plate can be attached to a physically large heatsink to form a stable ambient reference, while radiation and conduction to the reference mass from the sensor implements the heat-leak required. To facilitate the very low reflectance required for radiation measurements I eventually painted the top side of the brass sensor plate flat-black using Lampblack mixed with a little Red Gum in alcohol to bind it.

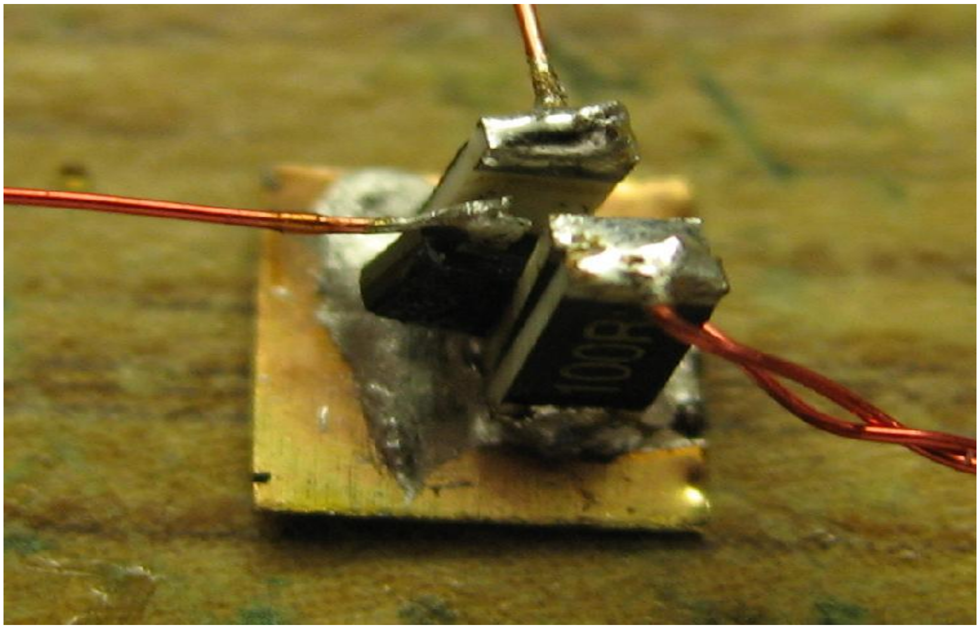

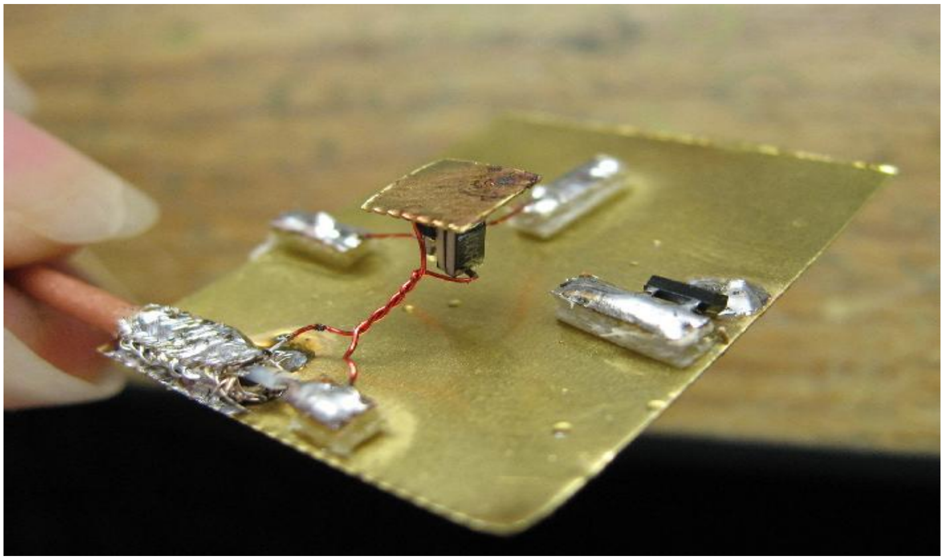

It is important to shield the detector and reference junction from drafts and differences in ambient illumination. The sensor is especially sensitive to drafts and I reason you could calibrate it quite accurately as an Anemometer (kinda neat, no moving parts!), at least at constant barometric pressure. I used a small potting box to cover the detector assembly to exclude drafts. This improved the baseline stability enormously. A small hole in the box allows a radiation beam to enter and hit the sensor plate facilitating its measurement.

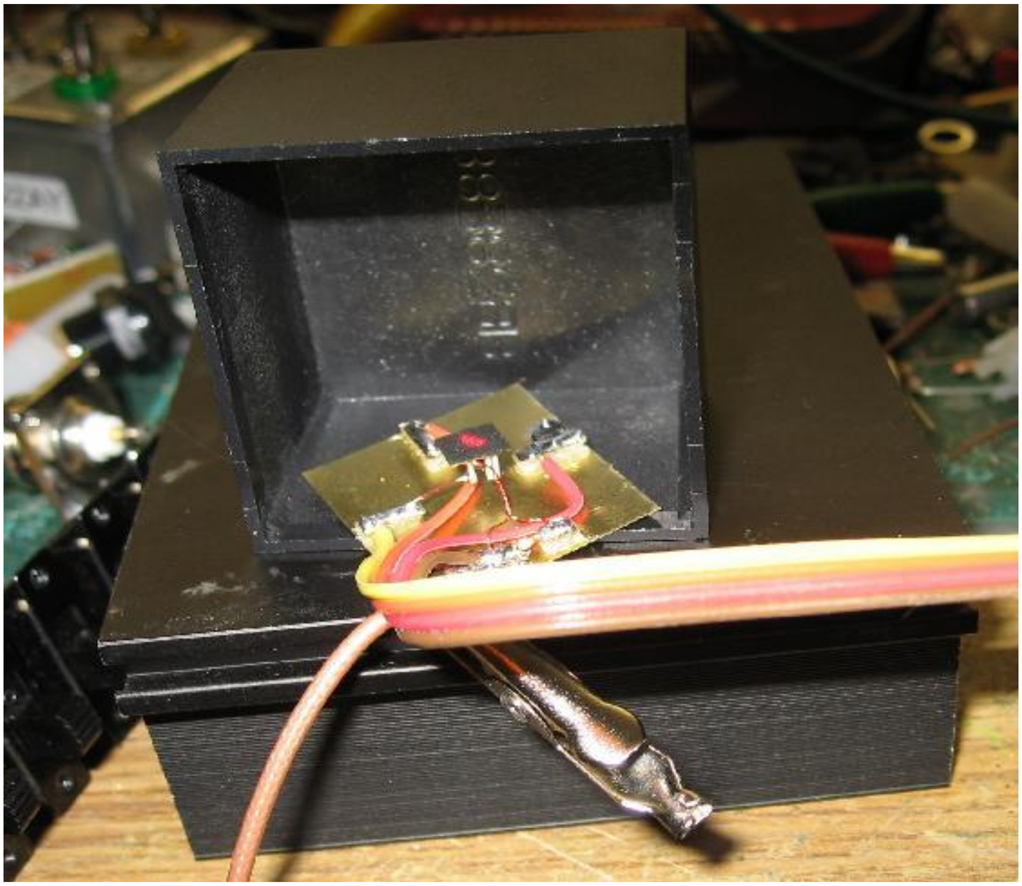

With the improved detector the output of the laser pointer was once again measured. 2.4 mW or 3.8 dBm was measured. The DC input from its battery is 72.8 mW, giving it a rather unspectacular efficiency of 3.3%. This concerned me that the sensor reflectance might still be fairly high, but as the compliance data suggests the output is < 5 mW it seems at least consistent. Diode lasers are usually more efficient than that, but the drive electronics is likely wasting a lot of power. I did not dismantle the laser head to directly measure the power delivered to the diode. I'd need a calibrated optical source to check the sensor at optical frequencies. Lampblack should be fairly flat with respect to frequency, at least compared to other "black" pigments. A green laser pointer measured 3.6 mW (5.6 dBm) and is also labelled as < 5 mW. Its IR local oscillator must be filtered from the output fairly well, I was expecting an unusually large reading from its IR leakage.

At RF the detector performs very well and consistently. I measured my 50 MHz "16 dBm" signal source at 48 mW (16.8 dBm) The previous calibration was by comparison to a DC-calibrated diode detector, so I am amazed by the agreement actually. Similarly I measured attenuation steps of a ~100 mW (20 dBm) signal at 10 MHz the results being quite consistent with my previous attempts to calibrate the poorly constructed attenuator at DC.

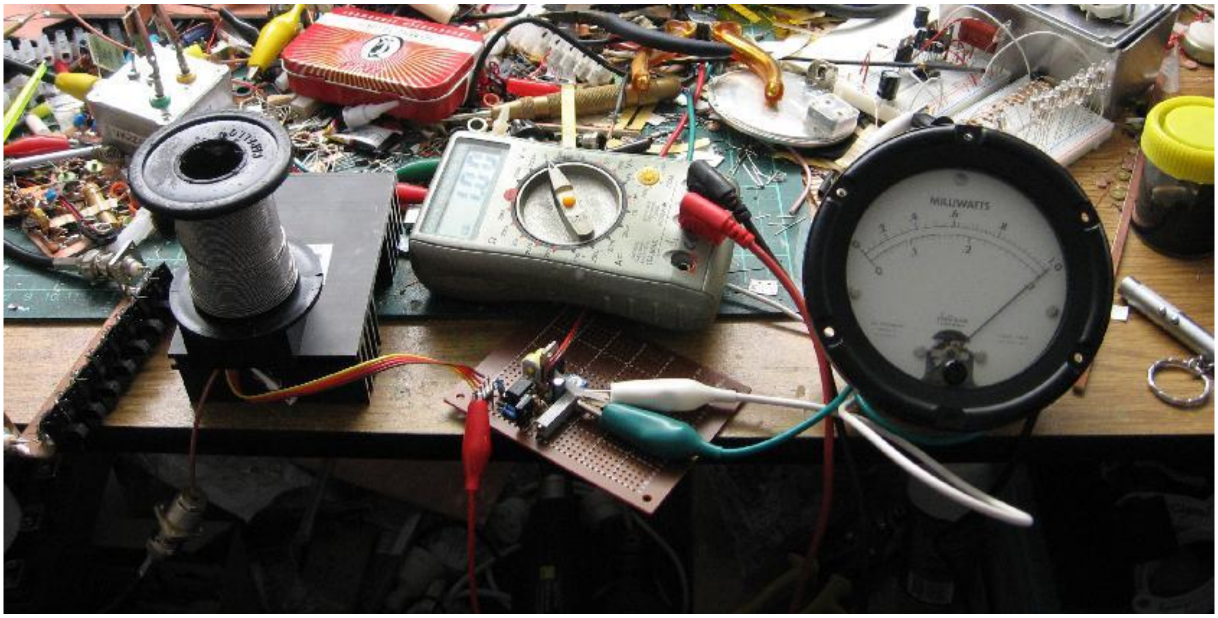

The system is configured to measure 1 mW to 100 mW, a 20 dB dynamic range. The poor dynamic range is typical of thermal sensors. Its lower limits could be improved by active cooling of the sensor head to achieve more sensitivity and more attention to noise filtering. Its upper limit is really only constrained by the temperature limit of the sensor assembly. I used 125 mW rated SMD resistors and set the bias power a little bit above that level (which is safe due to the heatsinking effect of the detector assembly). With larger resistors in the detector and a suitable power Darlington follower the circuit could be used up to kW. Efficiency is of course terrible, the sensor bias power must exceed the power to be measured. Larger RF powers are more easily measured with diode peak-voltage measurements, but as a thermal device is a natural integrator it can measure the true average power of complex waveforms containing multiple frequencies (including 0, DC biases which may or may not be a problem depending on the application). Attenuation from higher power is probably the most practical method. Amplification can be used to measure smaller signals, but the calibration of the amplifier then becomes an issue. MMICs with reasonably flat gain and compression points exceeding the detector range are available. The load offers a good return loss well into VHF and is therefore capable of absolute average power measurements from DC to several hundred MHz.

The general design can implement all kinds of radiant energy sensors. A pyrometer is simply a matter of optics and calibration. The device is already a fairly usable laser power meter.

Notes

- The heater and load resistors must be as ohmic as possible for best accuracy, although thermal variation in their resistance can be calibrated out. Obviously the 50 Ohm dummy load resistor must not vary too much or else the return loss will be compromised. For most commercial metal film devices the thermal coefficients are quite tiny.
- The "approximately" constant current biasing of the diodes is fine in this configuration, as it is thermal power balancing (not temperature sensing where two different currents would ideally be used to extract the absolute temperature). We don't really care too much about the diode coefficients, just that they hold the sensor at a constant temperature difference to the reference. The reference diode compensates the bias power for variations with ambient temperature. At constant ambient temperature a reference diode isn't strictly required. It is only required that the sensor diode temperature is held constant between zero and applied power conditions for the heater power difference to match the applied power.
- Small decoupling capacitors may be needed across the diodes to prevent RF pick-up upsetting things.
- Watch out for photovoltaic problems with glass encapsulated sensor diodes. SMD encapsulation is not susceptible, but glass cased 1N4148s needed painting black. Not all black paints are opaque at the full range of optical sensitivity of Silicon diodes.
- With a single resistor sensor for RF watch out for magnetic saturation in the RFC core which may drop its impedance dramatically shunting RF and degrading the return loss. Fortunately this is likely to happen with the smallest power inputs so mismatch damage to the DUT is unlikely, but the measurements will be compromised and perhaps go unnoticed. The RFC is problematic anyway, it must have a large impedance with respect to the 50 ohm load across the bandwidth of the instrument - a somewhat challenging requirement for a MF-UHF device. Commercial SMD RFCs designed for MMIC biasing may be useful in this service.
- Larger heating resistances allow the circuit to operate down to smaller power levels, but noise will become a problem.
- Filtering most of the resistor noise out of the bandwidth of the loop is probably a good idea if going for more sensitivity.
- The supply must be well regulated, even with the supply rejection of the op-amp the zeroing for the analogue meter will change (minor) but more importantly the slightly different dynamic resistance of the sensor and reference diodes reflects supply noise into the loop. It might be worth building an ultra stabilised supply for the sensor diode biasing, or even perhaps the entire circuit.
- The heater supply follower temperature coefficient has a small effect and can be minimised by thermal bonding to the reference heatsink.

- At constant ambient temperature the amplifier circuit temperature drift is unimportant, simply measure the change in heater power between quiescent and signal applied to find the applied power.
- The unit must come to thermal equilibrium before meaningful measurements can be made. Fortunately for the small sensor head this happens rapidly (a few seconds).
- Initial adjustment of the bias power should be done while monitoring the heater voltage. Set for about 150 mW of quiescent power and allow to stabilise. Do not set such that the op-amp is pegged out against its upper output voltage limit. Similarly in operation if the input power exceeds the bias power the op-amp will bottom out at near zero volts, the follower transistor will obviously be cut-off a bit before that happens.
- 1/f noise is going to be a major problem when trying to increase sensitivity. It might be practical with extremely small detectors to modulate the heater current and extract the AC response from the diode drop voltage variation and amplify it at that more reasonable frequency. This removes the direct-reading nature of the power balance, and linearity would be compromised based on the diode thermal coefficient linearity around the detector temperature, but baseline drift should be fairly small.

Source: http://www.vk2zay.net/article/210