A $10 Thermal Infrared Imager

Philip C. D. Hobbs

IBM T. J. Watson Research Center
PO Box 218,
Yorktown Heights NY 10598

1 Abstract

A thermal infrared imager of competitive sensitivity and very simple construction is presented. It is a pyroelectric device of 96 pixels, based on ferroelectric polyvinylidene fluoride (PVDF). It uses a novel charge-dispensing multiplexer based on ordinary light emitting diodes to achieve a noise-equivalent temperature change (NEΔT) of 0.13 K at a 5 Hz frame rate (2.1 Hz BW). Design information, theory, and measured performance are presented. Achieving such a low total system cost requires the use of the very least expensive optical system, a moulded polyethylene Fresnel lens, whose advantages and limitations are discussed. Several possible improvements, aggregating approximately 30 dB in sensitivity are also discussed, leading to the interesting possibility of few-millikelvin NEΔT values with an uncooled pyroelectric device of extremely low cost.

2 Introduction

Infrared imaging sensors are expensive, insensitive, or (usually) some of each. This limits their usefulness, though there are many applications in which thermal IR images are very useful, e.g. looking at traffic patterns in indoor public spaces, where the self-luminosity of people makes them easy to discriminate from inanimate objects such as shopping carts. Achieving high sensitivity at high resolution requires cooled indium antimonide (InSb) and mercury cadmium telluride (HgCdTe) imagers cost tens of thousands of dollars, while even uncooled devices with few pixels such as 256-pixel lead zirconate-titanate (PZT) sensors, cost at least a few thousand dollars. Sensitivities for these devices are typically 0.1 K noise equivalent temperature change (NEΔT), which is adequate for most purposes.

Uncooled sensors made from pyroelectric PVDF films have been used for many years for things such as automatic porch lights and intrusion sensors, but their very low sensitivity and requirement of one amplifier per pixel have seriously limited the use of PVDF in imaging sensors. Although its intrinsic sensitivity is lower than that of crystalline and ceramic pyroelectrics such as PZT, triglycine sulphate, and triglycine fluoroberylate, its very low cost, excellent electrical properties, and physical robustness make it an attractive material. The primary reason for its unpopularity is the complexity of the multiplexing schemes currently used with pyroelectrics, usually involving laminating or bump-bonding a pyroelectric material to a silicon readout circuit such as a multiplexer or CCD. If one is going to all that trouble, of course one will use the most sensitive material affordable, which is not PVDF.

from Sensors and Controls for Intelligent Manufacturing II, Proc. SPIE Vol. 4563, p. 42-51 (2001)

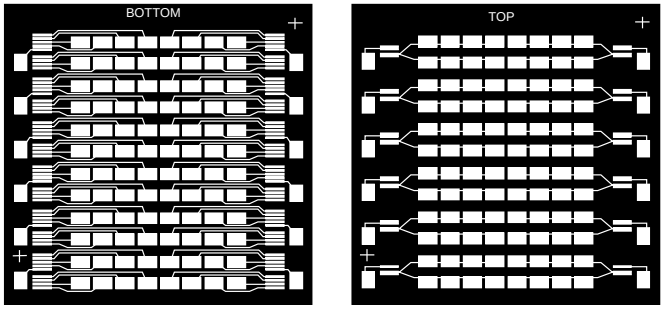

Figure 1: Mask pattern for 96-pixel pyroelectric PVDF imager. Actual wiring is screen-printed carbon ink.

3 PVDF Sensor Film

PVDF has a lot of advantages besides low cost. One is that the wiring can be put on via screen printing, like the writing on T-shirts. Screen printing of carbon-loaded ink produces a very robust film of reasonable conductivity and somewhat better thermal emissivity than metal, but at the cost of the added thermal mass of 4 μm of ink on both sides of each pixel, nearly doubling the thermal mass of the 9 μm film used here. The ink is so strong that it will take a hard crease without delaminating or opening. Compared with the extreme fragility of metal films on thin PVDF, the speed tradeoff is very worthwhile[†].

Another advantage of PVDF is that it is mechanically robust enough to be used free-standing, suspended in air, so that the thermal mass M_{th} and thermal conductance G_{th} are minimized. Air suspension increases the thermal time constant, which slows the sensor down; this might seem like a disadvantage, but really it helps the SNR at all frequencies. Far down on the slope of the thermal transfer function, the response must decay as $\varepsilon P_{opt}/M_{th}f$, where P_{opt} is the optical power incident on the film, ε is the thermal emissivity, and f is the modulation frequency. Reducing G_{th} by air suspension reduces M_{th}, since only the film itself is heated, which automatically increases high frequency SNR. Furthermore, since the low-frequency limit of the response is $\varepsilon P_{opt}/G_{th}$, the low frequency SNR is enormously enhanced by allowing the total temperature excursion of the film to be as large as possible. Therefore, insulating the film with gas or vacuum (whose thermal mass is small) improves SNR at all frequencies. Some postprocessing is required to recover decent speed of response, of course.

[†] The sensors used here were made by Measurement Specialties Inc., Norristown PA, http://www.msiusa.com.

Here the film is glued to a frame that holds it in registration. The film frame fits tightly over one of a pair of spacers made of polyethylene terephthalate (the plastic used in soda bottles). All the plastic parts can be made by injection moulding. Electrical connections to the circuit board are made via a conductive elastomer (Zebra) connector of the same sort used with liquid crystal displays. Figure 1 shows the mask pattern used in this work. The pixels are 3 × 5 mm, on a 6 × 6 mm pitch. Optically they are laid out as 12 rows by 8 columns, but electrically they are read out as 6 rows of 16 columns, with each "row" being a 4 × 4 pixel block on the sensor surface (e.g. row 1 is the 16 pixels at top left). The pixels arranged outside the main array are dark pixels, intended for first-order compensation of charge injection due to switching.

PVDF is cheap and easy to handle. If a good multiplexer were available at low cost, it would transform the economics of uncooled infrared imaging. Such a multiplexer is the principal subject of this paper. To motivate its design, it is necessary to understand the difficulty.

4 The Multiplexer Problem

The reason for the difficulty of the multiplexer problem is the extremely high impedance levels at which pyroelectric sensors work. As shown in Figure 2, the ferroelectric polymer has a frozen-in electric polarization *P*, which in a planar geometry behaves identically to a sheet of bound charge at its top and bottom surfaces. To achieve overall electrical neutrality, free charges from the environment flow to neutralize these bound charges, leading to a quiescent condition with zero volts across the ferroelectric. The remanent polarization is a strong function of temperature, however, decresing from its room-temperature value to zero at the Curie point of the polymer, about 120 °C. Its temperature coefficient at 20 °C is approximately -1%/K (PVDF properties are summarized in Table 1).

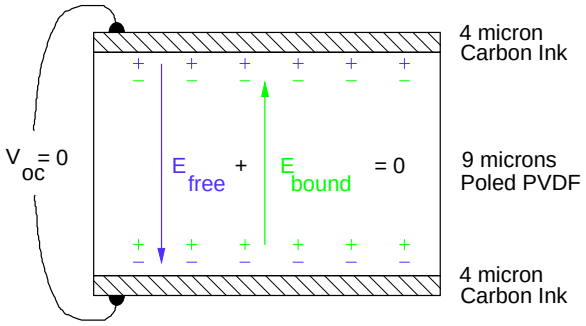

Figure 2: Pyroelectric Detection: Frozen-in field E_{bound}

For the 9 µm thick sensors used here, this corresponds to a terminal voltage change of 2 V/K. While this may seem an enormous change, there is less there than meets the eye; this is its sensitivity to *sensor* temperature changes, not *scene* temperature changes. Typical ratios $\Delta T_{sensor}/\Delta T_{scene}$ are 10^{-4} to 5×10^{-3}, so a sensitivity NEΔT of 0.1 K in the scene temperature requires a noise equivalent to 10^{-5} to 5×10^{-4} K uncertainty in T_{sensor}.

Achieving a NEΔT of 0.1 K requires an rms charge readout noise of approximately 0.1 K /($dq/dT \cdot dT_{pixel}/dT_{scene}$)≈ 0.18 pC for a bandwidth of 0.1 Hz, and about 0.06 pC for a bandwidth of 2.1 Hz (the signal and noise are not flat with frequency). While this is not in the CCD class, it is stringent for a built-up circuit. Even allowing the multiplexer to dominate the total noise, the noise current must be below 0.6 pA in the measurement bandwidth. Furthermore, since the readout circuit must be replicated 96 times, it must be very simple or very highly integrated if low cost is to be achieved.

The usual method of achieving high sensitivity in pyroelectric arrays is to bond them to a silicon CCD array or MOS multiplexer. Such devices have very low readout noise, only a few tens of electrons at room temperature, and their high degree of integration allows the use of many pixels. On the other hand, the very high thermal conductivity of the silicon substrate reduces the pyroelectric charge obtainable, which limits their sensitivities to approximately 0.1 K NEΔT. Furthermore, due to complexity, high startup costs, and low manufacturing volumes, such devices are expensive, in the $4000 range at present even for a 256-pixel device. A built-up circuit made of commonly-available parts is clearly desirable. Unfortunately, conventional approaches based on discrete silicon MOSFETs or commodity CMOS multiplexer ICs are unsuitable due to large charge injection noise, leakage, and cost.

5 DIODE SWITCHING

The simplest semiconductor switch is a diode. If its forward voltage drop can be tolerated, a diode is a good current switch, conducting heavily in the forward direction and very little in the reverse. In the present regime, though, silicon diodes are unsuitable, due to their relatively low zero-bias resistances of 100 MΩ to 10 GΩ, which make sub-picoamp leakage currents unattainable. In addition, since the pyroelectric sensors are inherently bipolar, an unbiased diode cannot be used since the negative signal swings would be lost.

On the other hand, ordinary red LEDs have leakages many orders of magnitude better. The diode equation predicts that the forward current of a diode is

$$I_F \approx I_S \left(e^{\frac{V_F}{V_\gamma}} - 1 \right) \qquad \textbf{(1)}$$

where V_γ is ideally kT/e (25.8 mV at 300K) but in real devices is 35-70 mV. The zero-bias resistance of a diode is

$$R_0 \equiv \left. \frac{\partial V_F}{\partial I_F} \right|_{V_F=0} = \frac{V_\gamma}{I_S}$$

A silicon diode with I_F = 1 mA might have V_F = 0.6V, whereas a red LED would have $V_F \approx$1.6V. Because of the exponential factor, a 1-V change in V_F corresponds to a decrease in I_S of exp(-14) to exp(-28), and an increase of the same factor in R_0, a matter of 6 to 12 orders of magnitude. Ordinary display LEDs (Chicago Miniature Lamp) conduct heavily for forward voltages of 1.3V, but leak less than 100 fA from -5 V to +0.5 V bias voltage, making them extremely good current switches.

e

Biasing is a serious problem. In our application, the maximum readout charge in a 200 ms integration time is about 1 pC, so the bias current should be around 5 pA. With a reverse bias of a few volts on the switches in the off state, the bias resistor would have to be about 1 TΩ, an unattainable value for a low-cost design. Fortunately, there is an easier way.

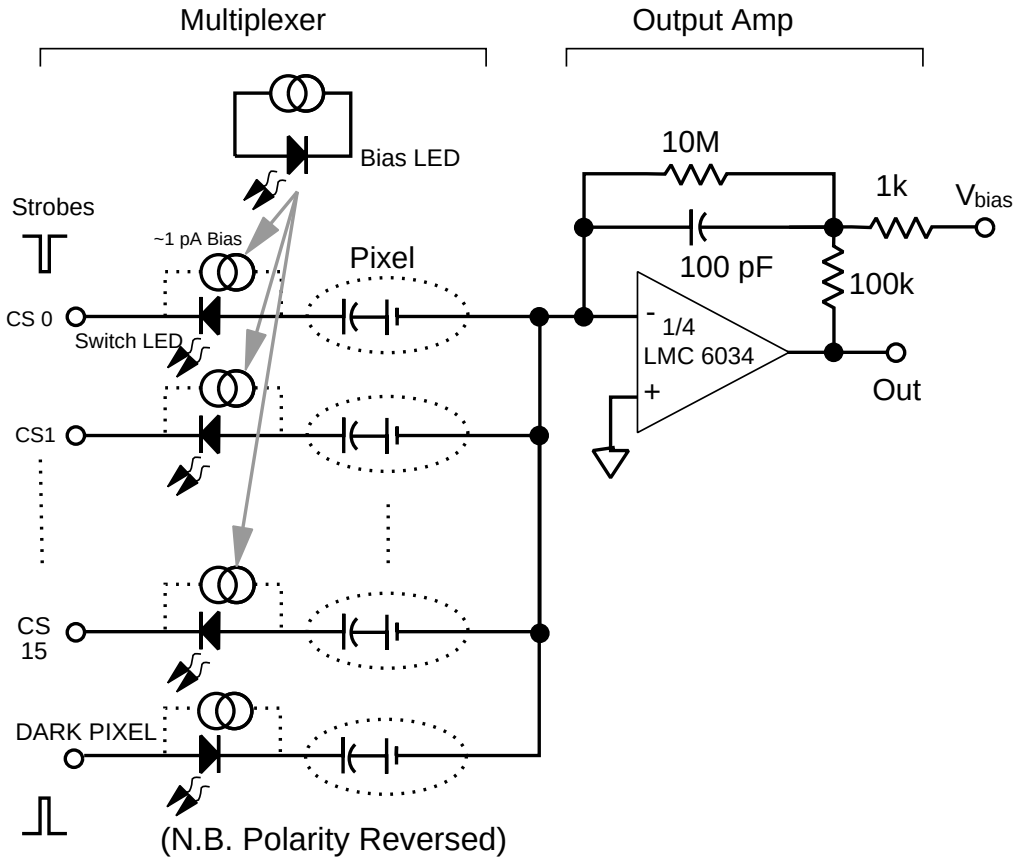

Figure 3: Simplified schematic of the multiplexer, showing the sensor elements, switching LEDs, dark pixels, and charge-sensitive readout amplifiers for one row.

5.1 Photocurrent Biasing

Besides emitting light, LEDs are also low-sensitivity photodiodes, so each pixel can be biased by shining a small amount of light on its switch LED. To get adequate uniformity and prevent ambient light from interfering, an array of four green LEDs distributed among the switch LEDs and covered with an opaque white diffusing cover. The green LEDs were controlled by a pulse-width modulation output of the processor, and conveniently a 1 mA drive current produced about 1 pA photocurrent in each switch LED, uniform to within a factor of two. Filtering ensured that the photocurrent noise was dominated by shot noise, contributing $(TeI_{bias})^{1/2} \approx 0.4$ fC charge noise per pixel per frame.

The multiplexer circuit is shown in Figure 3, a photo of the disassembled multiplexer and opaque diffuser is shown in Figure 4, and the assembly in Figure 5 and Figure 6. It achieves a fundamentally minimal component count of one per pixel, with two buffered shift registers providing the 16 column strobe signals and two quad op amps providing the 6 row amplifiers and two LED drivers. Each row has a dark pixel and associated switch LED to allow first-order cancellation of the charge injection due to the capacitance of the LEDs, which works adequately.

5.2 Readout Amplifiers

The readout amplifiers themselves are simple. When a column strobe goes active, the charge on each pixel in that column is dumped into the summing junction of its readout amplifier, which forces it to go into C_f. Within 500 µs, the resulting voltage is digitized, and after a 5 ms wait for the integrator to be reset by the 10 MΩ R_f, the next column is read out. After all columns have been read, the sensor returns to integration mode until the next 200 ms integration time has expired.

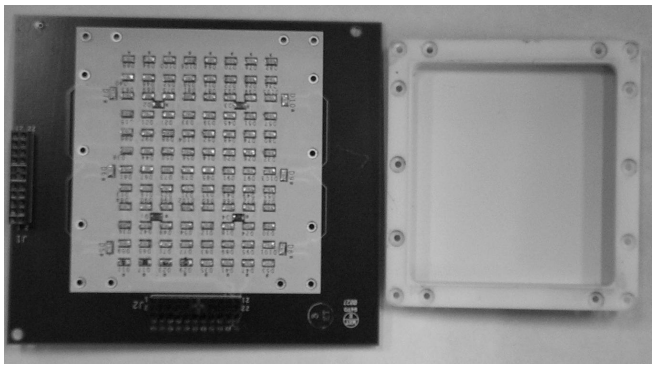

Figure 4: Disassembled multiplexer, showing switching LEDs in array, bias LEDs in each quadrant, and the diffuser made of the white silkscreen and opaque cover.

5.3 Noise Sources

The dominant noise source is the $(kTC_{\text{pixel}})^{1/2} \approx$ 0.86 fC from dumping the charge off the pixel capacitance, and the Johnson noise of the feedback resistor, which contributes an uncertainty of

$$q_{nR} = C_F \sqrt{4kTRB} = \sqrt{kTC_F}$$

or 1.1 fC, since the noise bandwidth of the parallel R_FC_F is π/2 times its 3 dB bandwidth, or $B = 1/(4R_FC_F)$. This makes the point that the uncertainty of resetting a capacitor to zero is the same regardless of how it is done. Along with the 0.4 fC shot noise, the RSS charge noise is 1.29 fC rms noise per pixel, which with a 100-pF C_F amounts to 12.9 µV at the amplifier input. Amplifier noise is comparatively insignificant. The typical total input-referred noise of the LMC6034 from 5 Hz to 4 kHz is 1.7 µV, so the total expected noise is 13.0 µV at the amplifier input. The amplifier gain is 100, so the total output noise is expected to be 1.3 mV, less than the (5 V/1024 ≈ 5 mV) analogue-to-digital converter unit (ADU).

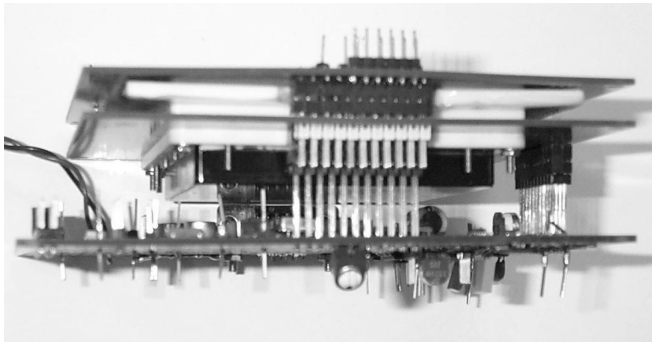

Figure 5: Side view of assembled sensor electronics. Row amplifiers are on top board, sensor film clamped between it and the multiplexer board with the opaque lid, and the CPU is at bottom.

Actual noise performance is limited by small amounts of convection and electrical pickup to 10 mV rms at the ADC input. Even so, a 15 K change in the scene temperature produces about a 1.25 V signal at the ADC.

6 Postprocessing

The transfer function of this device is somewhat unusual. If we take the impulse response g(t) at the output corresponding to an instantaneous change in scene temperature, we get a step function followed by decaying exponential of 3.1-s thermal time constant as the film cools back down. This surface charge is dumped and sampled every 200 ms, so the impulse response is convolved with an odd impulse pair (finite difference operator) of 0.2 s width. The first is approximately an integration, and the second approximately a differentiation, so that the temporal transfer function has a flat top. This is inconvenient, since passing people give rise to bipolar blobs, as shown in Figure 7. One person can disappear into the divot left by someone else, which complicates the tracking problem and makes the data hard to interpret.

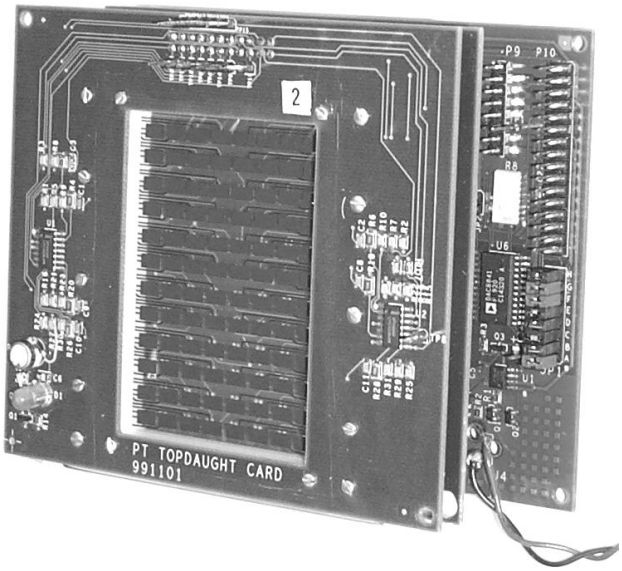

Figure 6: Front view of assembled sensor electronics, showing sensor film suspended between readout amplifier and multiplexer cards, and the CPU underneath. Not shown is a 80 μm polyethylene cover taped over the sensor to control convection.

Although the effective filter is many samples wide, it turns out to be easy to invert this transfer function by factoring. After digital baseline restoration to get the data centered on zero, a 9-sample FIR filter boosts the high frequencies, turning transfer function into that of a band-limited differentiator. A simple

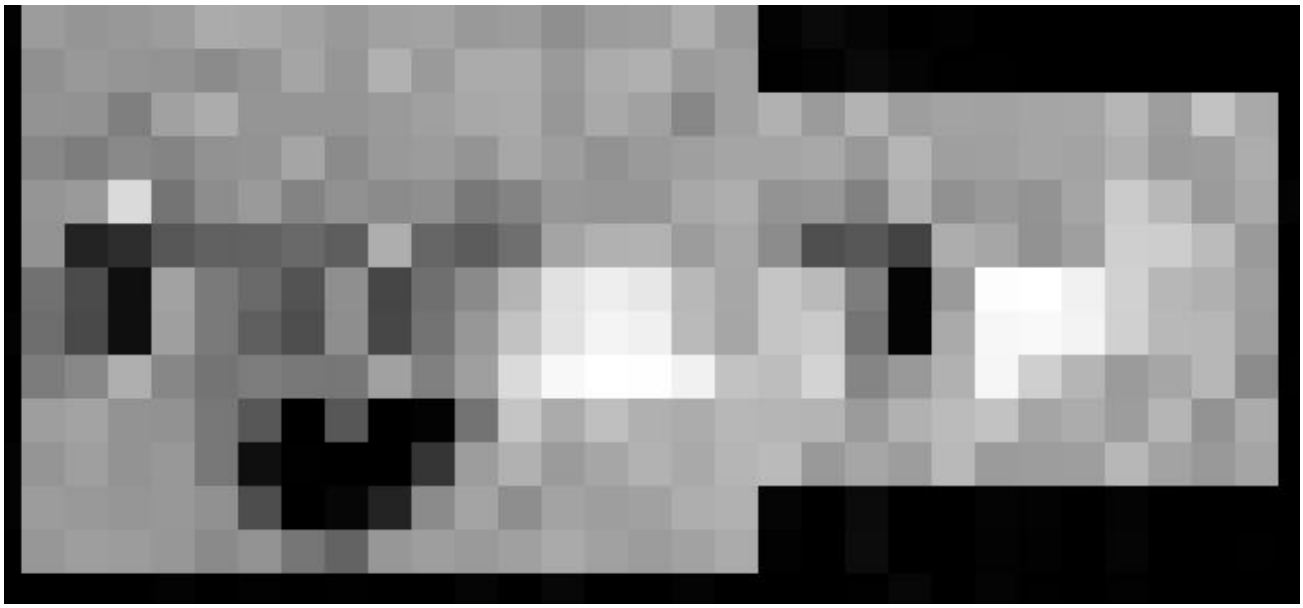

Figure 7: Mosaic image from 3 sensors based on an early 28-μm, metal-clad film, showing the bipolar response to two people walking around in the scene. Scale is 30 cm per pixel.

running sum then produces a flat transfer function up to nearly the Nyquist frequency (2.2 Hz out of 2.5 Hz).

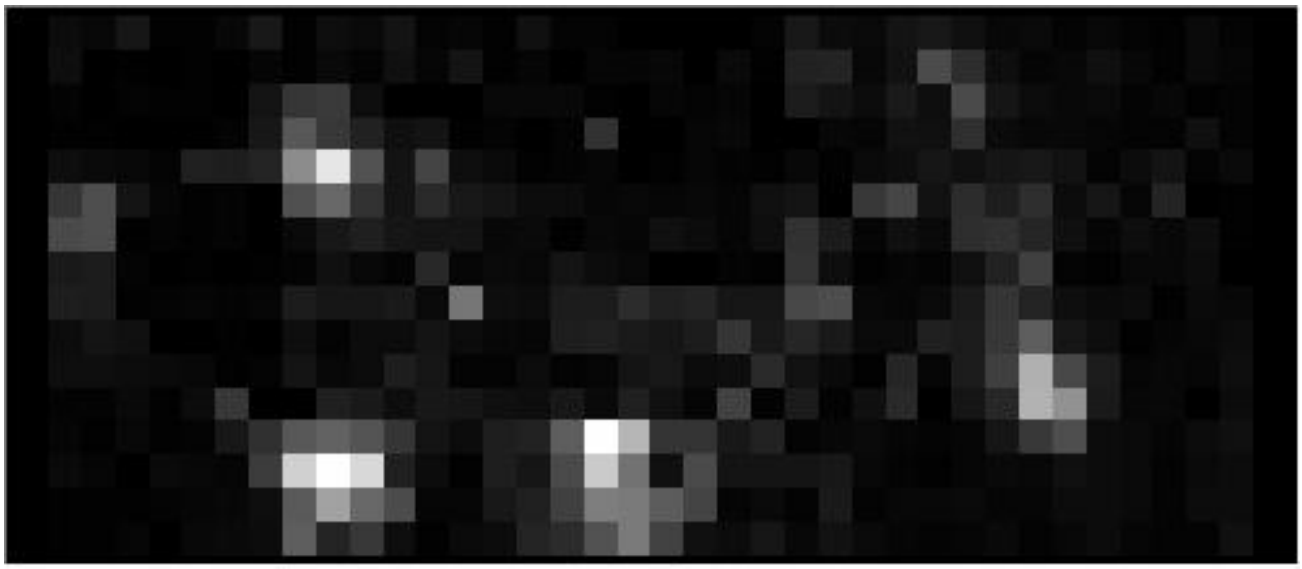

Figure 8: Pseudo-integrated data from 9-µm PVDF sensors with carbon ink. As before, the subject is a top view of people moving around, at 30 cm/pixel. Note that the divots left behind have disappeared, and that the SNR is dramatically better.

The difficulty with this is that since there is no actual signal at dc, and an integrator has infinite gain there, we wind up with a signal to noise ratio (SNR) of 0 with a linear integrator. This manifests itself as a random chequerboard pattern that grows with time, which makes this postprocessing strategy pretty useless. We're saved by a positivity constraint: people are warmer than floors. By applying a slight gain enhancement to the negative going peaks, we ensure that the integrated signal corresponding to a person crossing a pixel will go negative after he passes. Because we know that the true integral cannot be negative, we chop off the negative values, which forces the signal always to be near the baseline. A long (100 second) exponential decay is also applied to the clipped integral, and the combination ensures a stable baseline and people that look like unipolar blobs. The result is to enhance the SNR of the result, since the noise gain is principally at low frequency where the signal is strongest (since people are relatively slow-moving)[†]. A sample data frame after postprocessing is shown in Figure 8 (Figure courtesy of R. H. Wolfe).

[†] P. C. D. Hobbs, S. Pankanti, and R. H. Wolfe, in preparation.

7 Noise-Equivalent Temperature Difference *NE*Δ*T*

As shown in Table 2, the *NE*Δ*T* of this sensor is approximately 0.13 K, which is a reasonably competitive value, especially for a sensor of such low cost and simple construction. Matters can be improved by using a specially tuned absorber film made of very thin metal, whose average emittance can be as high as 0.5 without the thermal mass penalty of the carbon ink[†]. Carbon ink in an open lattice over the pixel could be used to stitch the metal film together into something electrically robust while maintaining the emissivity improvement. Such film could provide about 14 dB improved sensitivity, which would get this very simple sensor into the 0.025 K *NE*Δ*T* range.

Signal level improvements can be taken to the bank, as long as there are no noise consequences. Since the dominant noise in this sensor is a couple of ADUs' worth of pickup, switching artifacts, and thermal convection noise, emissivity improvements translate directly into SNR improvements. Noise reductions are a bit more problematical, because there may be some subdominant noise source that leaps into prominance when a dominant one is fixed. Accordingly, the factor of 5 in the noise floor that seems to be available may not be so easy to get. If it is, however, the total *NE*Δ*T* for an optimized sensor of this sort could be as low as 0.005 K in a 2 Hz bandwidth.

Table 1: PVDF Pyroelectric Element Properties

PVDF Film			
Mass Density	ρ	1.79	g/cm^3
Thermal Conductivity	α	0.19	W/(m·K)
Specific Heat	c_p	1440	J/(kg·K)
Dielectric Constant	ε	12.5	
Pyroelectric Coefficient (25°C)	p	-25	μC/(m^2·K)
Film Thickness	t	9	μm
Pixel Dimensions		3 × 5	mm
Pixel Pitch		6 x 6	mm
Pixel Capacitance	C	185	pF
Charge for 1 K ΔT_{sensor}	dq/dT	380	pC/K
Pixel Voltage for 1 K ΔT_{sensor}	dV/dT	2.0	V/K
Carbon Ink			
Thermal Emissivity (Ink)	ε_T	≈ 0.1	
Mass Density	ρ_{ink}	≈ 1.6	g/cm^3
Specific Heat	c_p	≈ 1500	J/(kg·K)
Thickness (Each Side)		≈ 4	μm
Fresnel Lens			
Transmittance of Lens	η_{lens}	0.5	
Solid Angle of IFOV at pixel	Ω_i	0.78	sr

[†] S. Bauer, S. Bauer-Gogonea, W. Becker, R. Fettig, B. Ploss, W. Ruppel, and W. von Munch, "Thin metal films as absorbers for infrared sensors", *Sensors and Actuators A*, v 37-38, pp. 497-501 (1993).

Calculated Thermal Properties (300K, still air)

Thermal Conductance:			
Air	G_{air}	$1.8 \cdot 10^{-4}$	W/K
Film	G_{film}	$9.0 \cdot 10^{-6}$	W/K
Radiative, total	G_{rad}	$1.8 \cdot 10^{-5}$	W/K
Radiative, to IFOV	G_{radI}	$1.1 \cdot 10^{-6}$	W/K
Radiative, to ambient	G_{radA}	$1.2 \cdot 10^{-5}$	W/K
Total Thermal Conductance	G_{tot}	$2.1 \cdot 10^{-4}$	W/K
DC Thermal Coupling Ratio	$\partial T_{pixel}/\partial T_{IFOV}$	0.0052	
Thermal Mass of Pixel	M_{th}	$6.6 \cdot 10^{-4}$	J/K
Thermal Time Constant of Pixel	τ_{pixel}	3.1	s
Thermal Coupling Coefficient	$\Delta T_{film}/\Delta T_{IFOV}$	≈0.005	
Input Voltage due to 1K ΔT_{IFOV} in 1 frame time (0.2 s)	$t_f \partial V/\partial T_{IFOV}$	10.	mV·s

Table 2: Noise Performance

Amplifier Bandwidth	B	4	kHz
$R_F C_F$ Bandwidth	$B_{RC} = 1/(2\pi R_F C_F)$	53	Hz
Amplifier Input Voltage Noise	V_{Namp}	1.7	μV rms
C_F Reset Noise	$\Delta q_{Cf} = \sqrt{(kTC_F)}$	0.87	fC
C_{pixel} Reset Noise	$\Delta q_{Cp} = \sqrt{(kTC_{pixel})}$	0.64	fC
Bias Current Shot Noise	$\Delta q_{shot} = \sqrt{(eq_{bias})}$	0.40	fC
RSS Charge Noise	Δq_{tot}	1.15	fC rms
Charge Noise at ADC Input	V_{noise}	1.15	mV
Amplifier Noise at ADC Input	V_{Naout}	0.17	mV
Total Front-End Noise	V_{FE}	1.16	mV
Quantization Noise	$V_{Nquant} = 5\ \mathrm{mV}/\sqrt{12}$	1.4	mV
Total Noise at ADC Input	V_N	1.8	mV
Actual RMS Noise At ADC Output	$V_{N\mathrm{meas}}$	10	mV
Actual Thermal Response for 15 K ΔT_{IFOV}	V_s(15 K)	1.25	V
Actual RMS Noise After Pseudo-Integration	$V_{N\mathrm{int}}$	40	mV
Actual Pseudo-Integral Response for 15 K ΔT_{IFOV}	V_{Sint}	19.7	V
ΔT for SNR = 1 (0.1 Hz, integral)	$NE\Delta T$(0.1 Hz)	0.03	K
ΔT for SNR = 1 (2 Hz, filtered)	$NE\Delta T$(2 Hz)	0.13	K