أكاديمية كاوست
KAUST ACADEMY

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

# Automated Pixel Selection for J–V Characterisation of 8-Pixel Perovskite Solar Cells

**Muath Alodhaibi, Sam Stranks**

moathmohhd@gmail.com, sds65@cam.ac.uk

## Abstract

Current-voltage characterisation of multi-pixel perovskite solar cells is normally performed one pixel at a time, with the operator turning a manual rotary selector switch between measurements. This is slow, tedious, and impossible to run unattended. We present an automated pixel-selection system that replaces the manual switch with a microcontroller-driven electronic switch matrix placed in series with a Keithley 2450 SourceMeter, controlled from the same Python environment that already drives the instrument.

Three architectures were designed and evaluated. The first used CMOS analogue multiplexers (CD74HC4067) in a dual-bank four-wire Kelvin configuration, needed to cancel their 70 ohm on-resistance. The second used a precision dual multiplexer (ADG1607), which exhibited an apparent leakage current that was diagnosed as conduction through the device's internal protection diodes when the signal node was driven below the negative supply rail. The final design uses eight Omron G3VM-31BR photo-MOSFET solid-state relays, whose optical isolation removes the supply-rail constraint entirely and whose on-resistance of 20 milliohms typical (40 milliohms maximum) makes remote sensing unnecessary, halving the required hardware.

The system was validated against a direct, unswitched connection using a three-trial protocol including a no-switch control. The control comparison, two direct measurements separated only by elapsed time, showed a larger discrepancy (500 microamps) than the direct-versus-switched comparison (237 microamps), establishing that run-to-run differences are dominated by device drift rather than by the switching hardware. The worst-case series error contributed by the relay is 0.2 millivolts at maximum sweep current, an order of magnitude below the source-measure unit's own accuracy specification.

**Date:** August 23, 2026

# Contents


**1 Introduction** **2**

1.1 Motivation . . . 2

1.2 Problem statement . . . 2

1.3 Assumptions and constraints . . . 2

1.4 Proposed system . . . 2

**2 Related Work** **3**

2.1 Prior approaches to multi-pixel switching . . . 3

2.2 Switching technologies and their trade-offs . . . 3

2.3 Gap addressed by this work . . . 4

**3 Methodology** **4**

3.1 System overview . . . 4

3.2 Formulation: the error budget . . . 4

3.3 Analysis of the component being replaced . . . 5

3.4 Architecture I: dual CMOS multiplexer with four-wire sensing . . . 5

3.5 Architecture II: ADG1607 dual multiplexer . . . 6

3.6 Architecture III: photo-MOSFET relay matrix (final design) . . . 6

3.7 Evaluation procedure . . . 11

**4 Results** **12**

4.1 Functional verification . . . 12

4.2 Illuminated $J$–$V$ response . . . 12

4.3 Baseline comparison: direct versus switched . . . 12

4.4 Behaviour of the relative-error metric . . . 13

**5 Discussion** **13**

5.1 Do the results meet expectation? . . . 13

5.2 Attribution: the control experiment . . . 13

5.3 Why the final architecture succeeded . . . 14

5.4 Unexpected findings . . . 14

**6 Limitations and Future Work** **15**

6.1 Limitations and their remedies . . . 15

6.2 Extensions . . . 17

**7 Conclusion** **17**

# 1 Introduction

## 1.1 Motivation

Perovskite solar cells are fabricated as substrates carrying several nominally identical devices, or *pixels*, each defined by the overlap of a patterned top electrode with a continuous bottom electrode. The substrates used in this work carry eight pixels sharing one common electrode. This multiplicity is deliberate: device-to-device variation in solution-processed perovskites is large, so conclusions about a fabrication recipe require statistics across pixels rather than a single measurement.

Each pixel is characterised by sweeping the applied bias and recording the resulting current, giving a current-voltage ($J$–$V$) curve from which the short-circuit current density $J_{sc}$, open-circuit voltage $V_{oc}$, fill factor $FF$ and power conversion efficiency $PCE$ are extracted. A Keithley 2450 SourceMeter (SMU) performs both the sourcing and the measurement.

## 1.2 Problem statement

In the existing setup, the choice of which pixel is connected to the SMU is made with a manual rotary selector switch: a mechanical wafer switch whose wiper routes one of the eight pixel contacts to the instrument. Measuring a full device therefore requires an operator to stand at the bench and turn a knob between every sweep, repeated for each measurement condition (forward and reverse scans, light and dark, and any stabilisation protocol). This imposes four costs:

- **Throughput.** Measurement time is bounded by operator availability rather than by instrument speed.
- **Reproducibility.** The interval between pixels — and hence the light-soaking history of each pixel — varies from run to run.
- **No unattended operation.** Long or repeated protocols, such as stability sequences over hours, cannot be left to run.
- **Bookkeeping risk.** The mapping from knob position to pixel identity is maintained by the operator, not recorded in the data file.

## 1.3 Assumptions and constraints

Two properties of the measurement constrain any solution. First, all eight pixels share a single common electrode, so only the eight top contacts require switching; the common electrode remains permanently connected to the instrument's LO terminal. Second, and more demanding: *the switching system must not perturb the measurement*. Any element inserted between the SMU and the pixel contributes series resistance, and any voltage dropped across that resistance is a voltage the pixel does not receive.

The central question of this work is therefore not whether pixel selection can be automated — it plainly can — but whether it can be automated without degrading the data, and how that claim can be demonstrated experimentally rather than merely asserted.

## 1.4 Proposed system

The manual switch is replaced by eight optically isolated photo-MOSFET solid-state relays, one per pixel, driven by an Arduino Nano microcontroller that receives pixel-selection commands over USB serial from the same Python environment that already controls the SMU. Selecting a pixel becomes a single function call; a complete eight-pixel characterisation becomes a loop.

The relays are chosen so that their contribution to the measurement is negligible by construction rather than by correction: their on-resistance of 20 mΩ produces a worst-case voltage error of 0.2 mV at the maximum sweep current, below the instrument's own accuracy specification. This claim is verified experimentally in Section 4 against a direct, unswitched connection, using a control experiment that separates the effect of the hardware from the drift of the device itself.

## 2 Related Work

### 2.1 Prior approaches to multi-pixel switching

Automated pixel selection is standard practice in commercial photovoltaic test equipment. Push-fit test boards and similar products integrate a switching stage so that all pixels on a substrate can be measured in sequence, and Keithley's own application literature for photovoltaic characterisation describes multi-device switching using dedicated switch mainframes with reed-relay or solid-state cards. These systems are effective but are sold as complete instruments; retrofitting one to an existing bench-built setup means replacing the whole measurement chain rather than the one component that is manual.

Within research laboratories, the common alternatives are either the manual rotary switch used here, or a bank of electromechanical relays driven by a microcontroller. Both are electrically excellent — metal contacts contribute milliohms — but the former is not automatable and the latter is bulky, audible, and mechanically wearing.

The approach taken in this work sits between these: a microcontroller-driven switch matrix built from discrete components, plug-compatible with the existing harness so that the instrument, sample holder and measurement software are unchanged.

### 2.2 Switching technologies and their trade-offs

Three families of electronically controlled switch were considered.

**CMOS analogue multiplexers..** Devices such as the CD74HC4067 connect one of several signal pins to a common pin under binary address control. Each channel is a CMOS transmission gate: an NMOS and a PMOS transistor in parallel, biased into inversion to form a conductive channel. The complementary pair is necessary because the gate overdrive of a single device collapses near one supply rail. Two consequences follow directly from this construction and dominate the design trade-off. First, an inversion layer is a resistor, so the on-state is not an ideal contact: $R_{\mathrm{on}} \approx 70\ \Omega$ for the CD74HC4067 at a 5 V supply. Second, the signal must remain between the supply rails, because outside that window the gate overdrive is insufficient and the internal protection structures conduct.

**Precision analogue multiplexers..** Parts such as the ADG1408 and the dual ADG1607 use the same principle with substantially lower on-resistance ($\approx 4\ \Omega$ and below) and support bipolar supplies. The supply-rail constraint on the signal remains.

**Photo-MOSFET solid-state relays..** Devices such as the Omron G3VM place an infrared LED and a photodiode array on opposite sides of an optically transparent, electrically insulating gap. Current through the LED produces light; the photodiode array — effectively a series stack of microscopic photovoltaic cells — converts that light into a voltage that charges the gates of two back-to-back MOSFETs, closing the switch. The switch side therefore requires no supply rails of its own, is galvanically isolated from the control side, and achieves on-resistance in the tens of milliohms.

### 2.3 Gap addressed by this work

Commercial switch matrices solve the problem at the cost of replacing the measurement chain; laboratory relay banks solve it at the cost of size and mechanical wear; naive semiconductor multiplexers are compact and cheap but introduce series resistance comparable to the features being measured, and impose a supply-rail constraint that is easy to violate in a floating measurement circuit. This work identifies the photo-MOSFET relay as the configuration that satisfies all three requirements simultaneously — automatable, compact, and electrically transparent — and, equally importantly, demonstrates the transparency experimentally rather than assuming it from the datasheet.

## 3 Methodology

### 3.1 System overview

The system consists of four blocks in series between the host computer and the sample:

1. A **host PC** running Python (Jupyter), which orchestrates the measurement.
2. An **Arduino Nano**, which receives pixel-selection commands over USB serial and drives the switching elements.
3. A **switch matrix**, which routes the instrument's force line to exactly one of the eight pixel contacts.
4. The **Keithley 2450 SMU**, which sources the voltage sweep and measures the resulting current.

The common electrode of the sample is permanently wired to the instrument's LO terminal and is never switched. A complete measurement is therefore: select pixel $n$, run the sweep, record, repeat.

### 3.2 Formulation: the error budget

An SMU sources a voltage and measures current at its own terminals. In a two-wire (local sense) configuration, everything between the instrument and the device — cables, connectors, switch contacts, holder pins — appears as a series resistance $R_{\text{series}}$ in the measurement loop, so the voltage actually reaching the device is

$$V_{\text{cell}} = V_{\text{set}} - I\,R_{\text{series}}, \tag{1}$$

while the instrument records $V_{\text{set}}$. The measured *current* is unaffected: the loop is a series circuit, so the current through the internal ammeter is exactly the current through the device. Series resistance corrupts the voltage attribution, not the current.

In a four-wire (Kelvin, remote sense) configuration the force and sense paths are separated. The force pair carries the current; the sense pair terminates in a voltmeter of very high input impedance and therefore carries almost none. The instrument regulates the *sensed* voltage to the set point, raising its output as required to compensate the drop:

$$V_{\text{out}} = V_{\text{set}} + I\,R_{\text{series}}. \tag{2}$$

Four-wire sensing does not remove the drop; it removes the drop from the *measurement*. A corollary that governs switch design is that the sense connection must be made at or beyond the point of interest: a sense line tapped upstream of a switching element leaves that element's

resistance uncorrected, so a sense line reaching the pixel must pass through a switch of its own, moving in lockstep with the force switch.

The design criterion follows immediately. The voltage error introduced by a switching element at the operating current is

$$\Delta V = I\, R_{\mathrm{on}}, \tag{3}$$

which must be judged against the span of the $J$–$V$ curve. For the devices measured here, currents reach approximately 5 mA over a sweep spanning 1.2 V. Table 1 evaluates Eq. (3) for each candidate technology and shows that the choice of switching element, not the sophistication of the sensing scheme, is the decisive factor.

**Table 1** Series-resistance error at $I = 5$ mA on a 1.2 V sweep, for each switching technology considered. The original mechanical switch is the baseline the replacement must match.

| Switching element | $R_{on}$ | ΔV | Relative |
|---|---|---|---|
| CD74HC4067 (CMOS mux) | 70 Ω | 350 mV | 29% |
| ADG1408/1607 (precision mux) | 4 Ω | 20 mV | 1.7% |
| G3VM-31BR (photo-MOSFET, typ.) | 0.02 Ω | 0.1 mV | 0.008% |
| G3VM-31BR (photo-MOSFET, max.) | 0.04 Ω | 0.2 mV | 0.017% |
| Rotary switch (baseline) | ∼ mΩ | negligible | — |

### 3.3 Analysis of the component being replaced

The existing switch was characterised before designing a replacement. Continuity mapping established that it is a single-pole, eight-position mechanical wafer switch: a rotating wiper contacting one of eight fixed terminals, each wired to one pixel's top contact, with the wiper wired to the instrument. This gives the functional specification for any replacement: an eight-way single-pole switch, exclusive (exactly one channel closed at a time), break-before-make, with series resistance comparable to the metal contacts it replaces.

### 3.4 Architecture I: dual CMOS multiplexer with four-wire sensing

The first design used two CD74HC4067 sixteen-channel analogue multiplexers. Because $R_{\mathrm{on}} \approx 70\ \Omega$ would introduce a 350 mV error in two-wire operation (Table 1), four-wire sensing was adopted — which, by the corollary in Section 3.2, required a second multiplexer:

- **Mux A** routes Force HI to the selected pixel and carries the measurement current.
- **Mux B** routes Sense HI to the *same* pixel and carries essentially no current.
- Both share the address lines $S_0$–$S_2$ and the enable line, so they cannot select different channels.
- Force LO and Sense LO connect to the common electrode.

Joining Mux A's and Mux B's outputs at each pixel places the sense tap downstream of Mux A's resistance, so both on-resistances vanish from the measurement: Mux B's because it carries no current, Mux A's because the instrument's feedback loop compensates it per Eq. (2). The fourth address bit $S_3$ was tied to ground to restrict the devices to channels 0–7.

**Bench validation on resistive loads..** Before connecting a device, the scheme was tested against two resistive loads chosen to bracket the current a pixel would draw: 200 Ω and 10.4 Ω. The instrument was stepped through a series of set voltages and the voltage actually appearing

across the load was measured directly, in both two-wire and four-wire configurations (Figure 1). Two-wire operation delivers markedly less than the set voltage in both cases, and the shortfall is severe at 10.4 Ω, where the larger current produces a correspondingly larger drop across the multiplexer channel. Four-wire sensing recovers the set voltage almost exactly at 200 Ω, confirming Eq. (2) in practice.

At 10.4 Ω the correction therefore holds only over part of the range. The reason is shown in Figure 2: to place roughly 0.9 V across the load, the instrument must raise its own output beyond 7 V, almost all of which is dropped across the multiplexer channel rather than the load. Past that point the output saturates and the delivered voltage stops following the set point. The feedback loop does not fail — it runs out of output range.

This is the practical consequence of a 70 Ω switching element. Four-wire sensing removes the drop from the *measurement*, but it cannot manufacture compliance headroom, and the burden it places on the source grows with the load current — that is, precisely in the high-current region of the *J–V* curve that matters most.

### 3.5 Architecture II: ADG1607 dual multiplexer

To reduce on-resistance, the CD74HC4067 pair was replaced by a single ADG1607 dual 8:1 multiplexer, whose two internal banks map naturally onto the force and sense paths. During commissioning two coupled faults appeared: an apparent leakage current that persisted *with the pixel disconnected*, and heating of the device when several control lines were active.

The first fault appeared as an implausibly large current in an ordinary sweep (Figure 3a): the reading rises steeply beyond 0.6 V and clamps at the instrument's 100 mA compliance limit. Repeating the sweep with nothing at all connected to the switched output produced the same knee at the same voltage and reached 25 mA by 0.8 V (Figure 3b). Since no device was present, none of that current can have flowed through one.

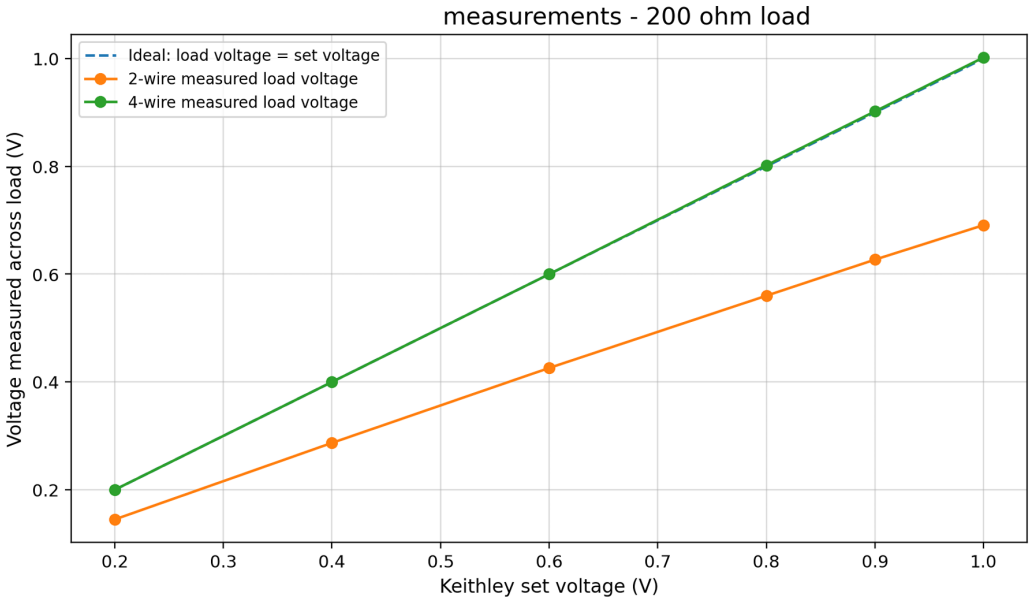
The first observation identified the mechanism. A current that flows while the pixel is disconnected cannot be flowing through the pixel, so its path must be internal to the multiplexer. With $V_{SS}$ tied to 0 V in the single-supply configuration adopted, and with the chosen polarity placing Force LO on the multiplexer signal path, the signal node is driven *below* $V_{SS}$ during a positive sweep. Every analogue pin carries an internal protection diode to the supply rail; once the node falls roughly a diode drop below $V_{SS}$, that diode conducts, providing a current path from the signal pin into the supply rail and back to the instrument that bypasses the device under test entirely. The same conduction dissipates power in the package, explaining the heating.

The available remedies were to add a negative supply ($V_{SS} = -5$ V) so that the signal stays inside the valid analogue window, or to invert the polarity so that the switched node never goes negative. Rather than introduce a negative rail into a floating measurement circuit, the architecture was reconsidered.

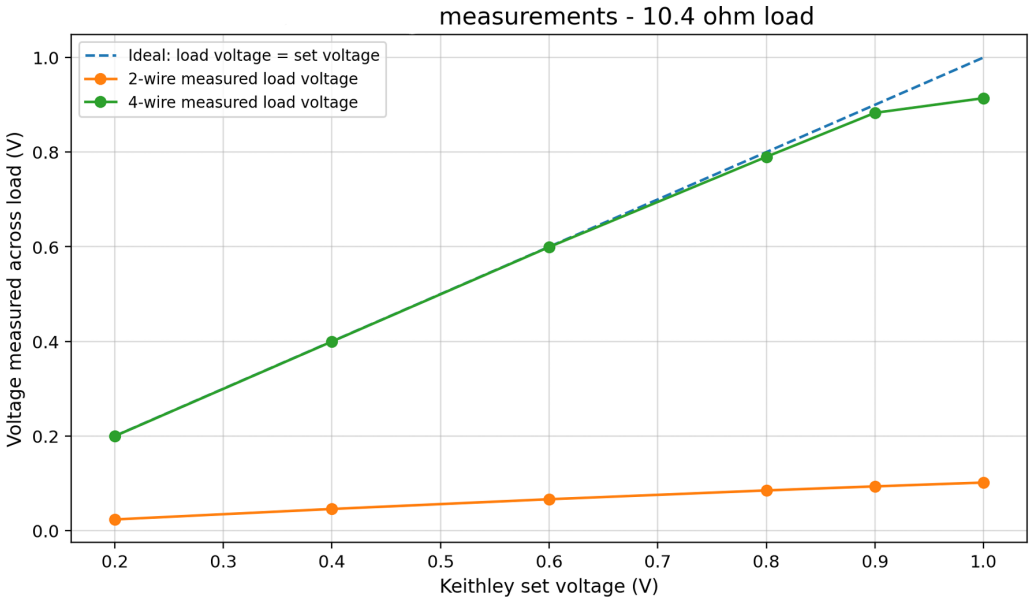
### 3.6 Architecture III: photo-MOSFET relay matrix (final design)

The final design uses eight Omron G3VM-31BR photo-MOSFET relays, one per pixel. It eliminates the failure modes of both earlier architectures by construction:

- **No supply rails on the switch side.** The MOSFET gates are driven by the internal photodiode array, powered by the LED's light. There is no $V_{SS}$ for the signal to fall below, hence no protection-diode conduction path and no polarity restriction — the switch is bidirectional.
- **Galvanic isolation.** Control and measurement sides share no electrical connection (specified dielectric strength 2500 $V_{rms}$), so their grounds stay independent and no ground loop is formed.




**(a)** 200 Ω load.




**(b)** 10.4 Ω load.

**Figure 1** Voltage delivered to a resistive load through the CD74HC4067, in two-wire and four-wire configurations, against the ideal case in which the load receives the set voltage. At 200 Ω the four-wire trace lies on the ideal line across the whole range. At 10.4 Ω it tracks the ideal line only up to about 0.6 V and then falls away.

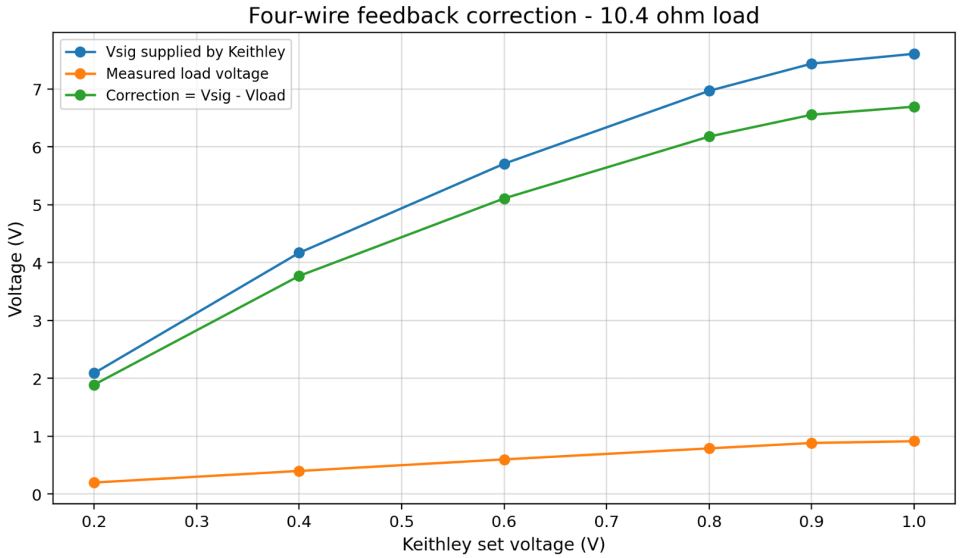



**Figure 2** The compensation demanded of the source in four-wire operation with the 10.4 Ω load. The correction $V_{\text{sig}} - V_{\text{load}}$ is the voltage the instrument must add to overcome the multiplexer's on-resistance; it exceeds 6 V while under 1 V reaches the load, and flattens as the source approaches its output limit.

- **Negligible on-resistance.** 20 mΩ typical, 40 mΩ maximum, giving $\Delta V \leq 0.2$ mV at 5 mA.
- **Negligible off-state leakage.** 0.01 $\mu$A typical.

Because the on-resistance is three orders of magnitude below that of the CD74HC4067, the entire justification for four-wire sensing disappears. The final system is therefore **two-wire, force-only**: eight relays instead of sixteen, half the control lines, and no sense path — which had been the source of several commissioning faults — at all.

### 3.6.1 Wiring

Each relay is wired in the datasheet's Connection A configuration. The control side is driven from one Arduino digital pin per relay through a series resistor sized for the datasheet's typical operating current:

$$R = \frac{V_{\text{pin}} - V_F}{I_F} = \frac{5 - 1.64}{10 \times 10^{-3}} \approx 336\ \Omega \longrightarrow 330\ \Omega \quad (I_F \approx 10.2\ \text{mA}), \tag{4}$$

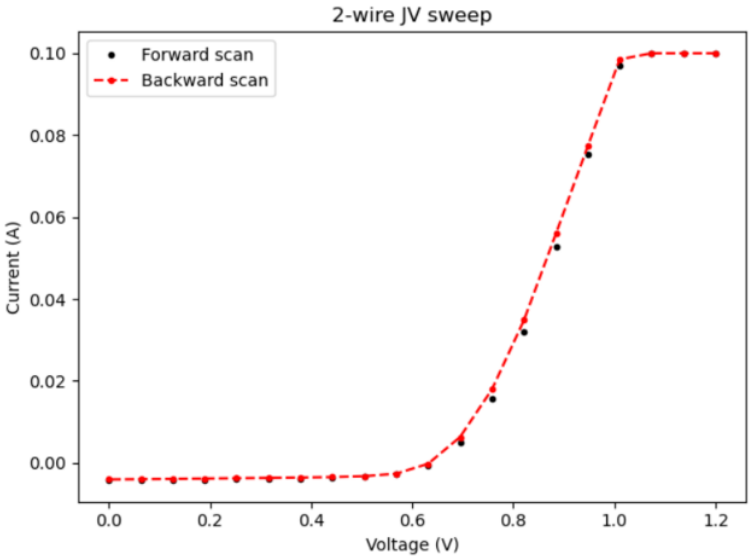
where $V_F = 1.64$ V is the LED forward voltage of the G3VM-31BR at $I_F = 10$ mA. This sits well above the 0.2 mA trigger threshold and well below the 25 mA maximum. The complete connection list is given in Table 2 and the corresponding schematic in Figure 4.

Two invariants are enforced throughout: the control ground and the measurement circuit share no connection, preserving the optical isolation; and at most one relay is closed at any time.

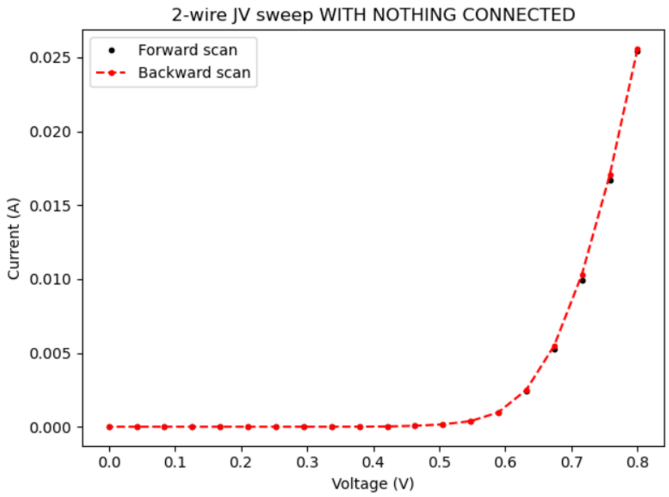
### 3.6.2 Firmware

Firmware on the Arduino Nano exposes a line-oriented serial protocol at 9600 baud: `P0`–`P7` select a pixel, `OFF` opens all channels, and query commands report state and identity. Every selection performs break-before-make switching with delays sized to the device's specified switching times ($t_{\text{ON}} \leq 5$ ms, $t_{\text{OFF}} \leq 1$ ms), and all channels are opened on power-up and on reset, so that a microcontroller reset cannot leave a pixel connected to a biased instrument.

**Listing 1** Break-before-make pixel selection. All relays are opened and allowed to settle before the target relay is closed, so two pixels are never bridged.




**(a)** Sweep with the pixel connected.




**(b)** The same sweep with nothing connected.

**Figure 3** The ADG1607 fault and the null test that identified it. The two traces share the same exponential turn-on near 0.6 V, so the current in (a) is dominated by a path that does not include the device under test. A diode-shaped characteristic measured across an open circuit can only be internal to the multiplexer.

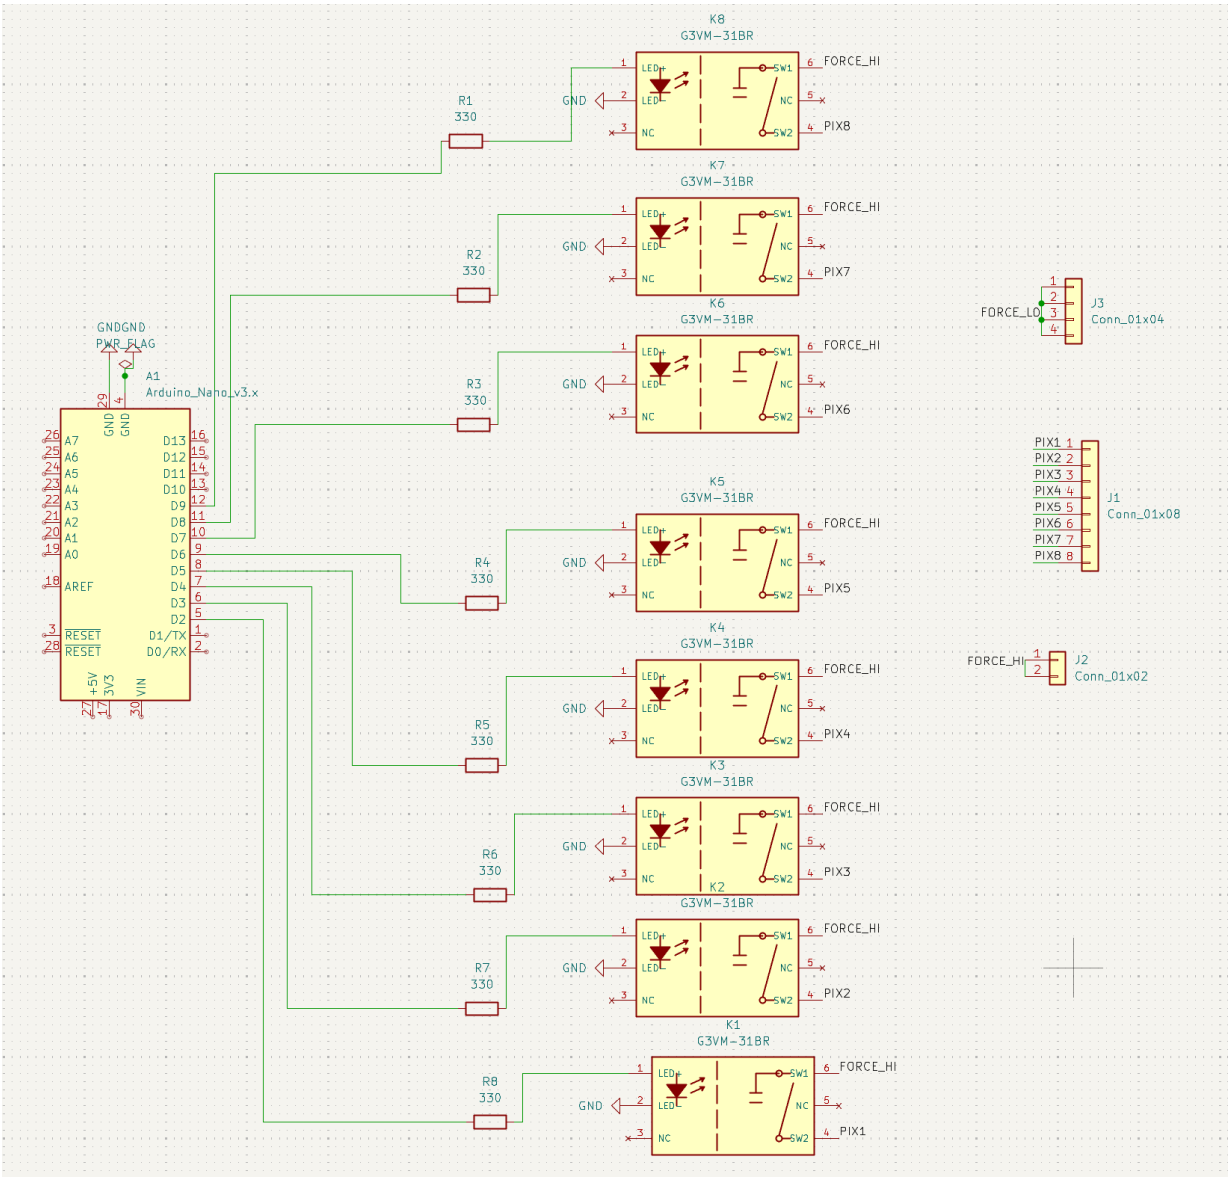



**Figure 4** Schematic of the eight-relay pixel selector. The vertical dashed line inside each G3VM-31BR symbol is the optical isolation barrier: everything to its left — the Arduino, the eight LED series resistors and the control ground — shares no electrical connection with the measurement side to its right, which carries only Force HI, the eight pixel nets and Force LO. Relay pins 3 and 5 are unconnected, as required by the Connection A configuration. Note that the LED resistors are numbered independently of the channel map: R8 feeds K1. The net names `PIX1`–`PIX8` are one higher than the firmware pixel indices, so `PIX1` is selected by the command `P0`.

**Table 2** Connection list for the eight-relay pixel selector. Relay pin numbers follow the G3VM Connection A configuration.

| From | To | Net |
|---|---|---|
| Arduino D2–D9 (via 330 Ω) | Relays K1–K8 pin 1 (LED anode) | LED drive |
| Arduino GND | All relays pin 2 (LED cathode) | Control ground |
| Keithley Force HI | Shared bus → all relays pin 6 | Force |
| Relay K$n$ pin 4 | Pixel $n-1$ contact | Pixel |
| Keithley Force LO | Cell common electrode (unswitched) | Return |
| Arduino USB, Keithley USB | Host PC | Data |

*Relay pins 3 and 5 are left unconnected in Connection A.*
*The LED series resistors R1–R8 are numbered independently of the channel map.*

```
void selectPixel(int idx) {            // idx = 0..7
  allOff();                            // break...
  delay(T_OFF_MS);                     // let the previous relay open
  digitalWrite(RELAY_PINS[idx], HIGH); // ...before make
  delay(T_ON_MS);                      // let the MOSFET fully close
  currentPixel = idx;
}
```

The protocol was deliberately made identical to that of the earlier multiplexer firmware, so that the existing Python measurement notebook required no modification when the hardware was replaced.

### 3.6.3 Host-side control

On the host, pixel selection is a single function call alongside the existing SMU control, so a full characterisation is a loop over the eight channels.

**Listing 2** Host-side pixel selection over USB serial. The reply is checked, so a failed switch raises rather than silently mislabelling the data.

```
def select_pixel(pixel):               # pixel = 0..7
    arduino.write(f"P{pixel}\n".encode())
    reply = arduino.readline().decode().strip()
    assert reply == f"OK P{pixel}", f"switch error: {reply}"
```

## 3.7 Evaluation procedure

**Tools..** Measurements were controlled from Python (Jupyter) using `pyserial` for the microcontroller and a laboratory wrapper over `pyvisa`/SCPI for the Keithley 2450. Analysis used `numpy`, `pandas` and `matplotlib`. Firmware was written in C++ using the Arduino toolchain.

**Bench verification..** Before any connection to a device, each channel was commanded in turn and continuity was checked across the switch terminals with a digital multimeter, confirming that exactly one channel closes per command, that all channels open on the `OFF` command and after a reset, and that no continuity exists between the control and measurement sides.

**Test input..** Rather than a synthetic dataset, the test input is a real illuminated perovskite pixel under a solar simulator — the same measurement the system is intended to perform. Sweep

settings were held identical across all trials: $0 \rightarrow 1.2\ \mathrm{V} \rightarrow 0$, 20 points per direction, two-wire, forward and reverse scans.

**Baseline and protocol..** The baseline is the direct, unswitched connection: Force HI wired straight to the pixel contact. Three sweeps were performed on the same pixel in the following order:

1. **Trial 1 — direct** (baseline).
2. **Trial 2 — relay:** Force HI routed through the relay board with the corresponding channel closed.
3. **Trial 3 — direct:** the baseline connection restored and the sweep repeated.

Force LO remained on the common electrode throughout. Trial 3 is the control: it differs from Trial 1 only by elapsed time, not by wiring, and therefore quantifies the device drift against which any switch-induced difference must be judged. This is essential, because the direct-versus-relay comparison alone confounds two variables — the wiring change and the passage of time — and cannot by itself attribute a difference to the hardware.

**Metrics..** Two metrics were computed pointwise across the sweep,

$$\Delta I = I_{\mathrm{direct}} - I_{\mathrm{relay}}, \qquad \varepsilon = \frac{I_{\mathrm{direct}} - I_{\mathrm{relay}}}{I_{\mathrm{direct}}} \times 100\%, \tag{5}$$

together with the maximum and mean of $|\Delta I|$. The absolute difference is the primary metric because it remains well defined across the whole sweep and is directly comparable with the prediction $IR_{\mathrm{on}}$ of Eq. (3); the relative error is reported for conventional comparison but is ill-conditioned near $V_{\mathrm{oc}}$, as discussed in Section 4.

# 4 Results

## 4.1 Functional verification

Bench verification (Section 3.7) passed on all counts: exactly one channel closed per command across all eight channels, all channels open on the `OFF` command and after a microcontroller reset, and no continuity between the control and measurement sides, confirming the optical isolation. The switching system is therefore functionally correct before any question of measurement accuracy arises.

## 4.2 Illuminated $J$–$V$ response

With the sample under the solar simulator, the relay-switched measurement produced the expected photovoltaic response: negative current at zero bias ($J_{\mathrm{sc}} \approx -3.1$ mA), a zero crossing at $V_{\mathrm{oc}} \approx 1.0$ V, and forward turn-on beyond it, with the mild forward/reverse hysteresis characteristic of perovskite devices.

An earlier measurement, taken before the common-electrode connection had been correctly identified, produced a purely rectifying curve with no photocurrent — the dark diode response of the same device. It is recorded here because it was initially mistaken for a fault in the switching hardware. A solar cell is a large-area diode; the absence of a fourth-quadrant excursion indicates a missing photocurrent path, not a defective switch.

## 4.3 Baseline comparison: direct versus switched

The three trials produce curves that are closely similar in shape and in extracted parameters, with the relay trial lying within the spread of the two direct trials — the first indication that the switch is not the dominant source of variation.

Table 3 collects the headline numbers against the datasheet prediction. The essential observation is that *the control comparison, which contains no switch at all, shows a larger discrepancy than the comparison that does*.

**Table 3** Validation results. The no-switch control (Trial 1 vs Trial 3) is the baseline against which the switched comparison must be judged; it shows roughly twice the discrepancy.

| Comparison | Switch present? | Wiring changed? | max $\lvert\Delta I\rvert$ |
|---|---|---|---|
| Trial 1 (direct) vs Trial 2 (relay) | yes | yes | 237 $\mu$A |
| Trial 1 (direct) vs Trial 3 (direct) | **no** | **no** | 500 $\mu$A |
| Predicted from $R_{\mathrm{on,max}} = 40$ m$\Omega$ at 5 mA | | | $\Delta V = 0.2$ mV |

### 4.4 Behaviour of the relative-error metric

The relative error of Eq. (5) shows excursions exceeding ±30% in the vicinity of $V_{\mathrm{oc}}$. These are artefacts of the metric rather than features of the measurement: the denominator $I_{\mathrm{direct}}$ passes through zero at $V_{\mathrm{oc}}$, so a fixed absolute difference of a few microamperes yields an unbounded percentage. The diagnostic signature is the sign reversal across the crossing, which a genuine series-resistance effect cannot produce. Away from the crossing, at biases below about 0.9 V, the relative differences are of order 1–3% — still three orders of magnitude larger than the switch could contribute, which is the subject of the following section.

## 5 Discussion

### 5.1 Do the results meet expectation?

The error budget predicted, from the datasheet alone, that the relay would contribute a worst-case error of

$$\Delta V = I\, R_{\mathrm{on,max}} = 5\ \mathrm{mA} \times 0.04\ \Omega = 0.2\ \mathrm{mV}, \tag{6}$$

i.e. 0.017% of the 1.2 V sweep, and half that for the typical 20 m$\Omega$. The measured discrepancies are of order hundreds of microamperes, which at the local slope of the $J$–$V$ curve corresponds to tens of millivolts — two to three orders of magnitude larger than Eq. (6) allows.

At face value this looks like a failure to meet expectation. The control experiment shows that it is not, and this is the central result of the work.

### 5.2 Attribution: the control experiment

Trial 1 and Trial 3 used identical wiring. The only variable between them was elapsed time, and they differ by *more* than either differs from the switched trial. A difference that is larger when the switch is absent cannot be caused by the switch. Three independent lines of evidence support this attribution.

**Magnitude..** Eq. (6) bounds the switch's contribution at 0.2 mV. Nothing in the measured data can be attributed to a 20 m$\Omega$ resistance.

**Shape..** A series resistance produces a difference that grows monotonically with current and does not change sign within the sweep, since $\Delta V = IR$ inherits the sign of $I$. The observed difference curves instead cross zero near $V_{\mathrm{oc}}$ and grow in both directions away from it. That is the signature of a change in the *shape* of the curve — a shift of the knee and of $V_{\mathrm{oc}}$ — not of a resistive offset.

**Ordering..** The trials ran direct → relay → direct, so Trials 1 and 3 span the longest interval and, correspondingly, show the largest difference. The discrepancy scales with elapsed time, not with the presence of the switch.

All three are consistent with the well-known behaviour of perovskite solar cells, whose $J$–$V$ response evolves under illumination and bias on timescales of minutes through ion migration and light soaking.

### 5.3 Why the final architecture succeeded

Both abandoned architectures shared a root cause: a semiconductor switch whose signal path is referenced to supply rails.

For the CD74HC4067 the consequence was resistive. The conducting channel is an inversion layer, and an inversion layer is a resistor; 70 Ω in series with a milliampere-scale photocurrent produces an error comparable to the features of the curve being measured. Four-wire sensing was a sound response, and the bench tests of Figure 1 show it working exactly as intended while the current stayed low. What they also show is its cost: the correction the source must supply scales with the load current, so the scheme degrades in the same region of the sweep where accuracy matters most, and eventually runs into the instrument's output range. Beyond that, four-wire sensing doubled the switch count and introduced a sense path whose correctness had to be separately established — and which was, in practice, the source of most commissioning faults.

For the ADG1607 the consequence was worse than resistive. The requirement that the signal stay between the supply rails is not a recommendation but a physical constraint: below $V_{SS}$ the internal protection diodes conduct. The resulting current bypassed the device under test entirely — which is precisely why it persisted with the pixel disconnected, the observation that identified the mechanism.

The photo-MOSFET relay removes the premise rather than mitigating the symptom. Its switch side has no supply rails, because the gate drive is generated optically inside the package; there is no rail to exceed and no protection structure to conduct. The same isolation keeps the control and measurement grounds independent. That its on-resistance is also three orders of magnitude lower is almost a bonus: it renders four-wire sensing unnecessary and returns the system to the electrical simplicity of the mechanical switch it replaces, while adding programmability. The four-wire architecture was correct, but it was sophistication deployed to compensate for a poor choice of switching element; choosing a better element dissolved the problem the sophistication existed to solve, and halved the hardware.

### 5.4 Unexpected findings

**The diagnostic value of a null condition..** The ADG1607 fault was resolved not by more careful measurement but by a single observation: current continued to flow with the device under test *disconnected* (Figure 3). That immediately excluded every explanation involving the measurement loop and located the fault inside the multiplexer. Testing what happens when nothing should happen proved more informative than refining the measurement itself.

**An ill-conditioned metric can look like a catastrophic result..** The relative error reached ±40% near $V_{\mathrm{oc}}$ — a figure that, quoted alone, would suggest the system was unusable. It is entirely an artefact of dividing by a current that passes through zero. For switching hardware evaluated on a solar cell, the absolute difference $\Delta I$ is the appropriate metric because it stays well conditioned across the sweep and is directly comparable with $IR_{\mathrm{on}}$; relative error should be quoted only where the current is well away from zero, or replaced by differences in the extracted parameters $J_{\mathrm{sc}}$, $V_{\mathrm{oc}}$, $FF$ and $PCE$.

**Automation improves more than throughput..** Because device drift dominates run-to-run variation, the timing of measurements is itself an experimental variable. Manual operation leaves that variable uncontrolled; the automated system fixes it. The automation therefore improves reproducibility, not merely speed.

# 6 Limitations and Future Work

Each limitation below is stated together with the concrete step that would address it.

## 6.1 Limitations and their remedies

**Validation precision is bounded by device drift, not by the switch..** Because the device changes measurably over the interval needed to rewire between trials, the experiment establishes only that the switch contributes *less than the drift* ($\ll 237\ \mu$A), not what the switch contributes. **Next step:** repeat the comparison on a stable dummy load — a precision resistor of similar impedance to the pixel — where no drift exists. This would resolve the switch's contribution directly and should agree with the 0.2 mV bound of Eq. (6). A second option is a relay that switches between the direct and routed paths, eliminating the manual rewiring interval entirely.

**Validation on a single channel..** The direct-versus-switched comparison was performed on one pixel. Although all eight relays are the same part with the same specification, per-channel behaviour has not been verified against a baseline. **Next step:** repeat the dummy-load comparison on each of the eight channels and record the results as a commissioning record for the instrument.

**Prototype construction..** The measurements reported here were taken on a protoboard build with hand-soldered connections. It is functional and verified, but not robust to handling by other users, and its point-to-point wiring makes the isolation between the control and measurement sides a property of how carefully it was assembled rather than of the design. **Next step:** a dedicated printed circuit board. This has since been designed (Figure 5) and passes design-rule checking with no violations and no unconnected pads, but it has not yet been fabricated, so none of the results in this report were taken on it. The layout enforces the isolation geometrically: the relay column is a physical barrier, with the Arduino, the resistors and the control ground confined to one side and the pixel, force and return connectors to the other, and no copper crossing between them. Commissioning it will require repeating the bench verification of Section 3.7 on the assembled board, together with the dummy-load comparison described above, before it replaces the protoboard for measurement work.

**Two-wire operation leaves the rest of the path uncorrected..** The choice of two-wire measurement is justified by the relay's on-resistance, but the cabling, connectors and holder contacts remain uncorrected — as they were in the original manual setup. This was not the object of the present work, but it sets the accuracy floor of the measurement as a whole. **Next step:** if higher accuracy is later required, a second relay bank switching Sense HI in lockstep would restore four-wire operation; the firmware and control software already support the required simultaneous switching.

**No automated fault detection..** A relay that fails open would produce a plausible-looking but wrong dataset. **Next step:** a compliance check after each switch — confirming the instrument sees a device rather than an open circuit before starting the sweep — would catch this in software.

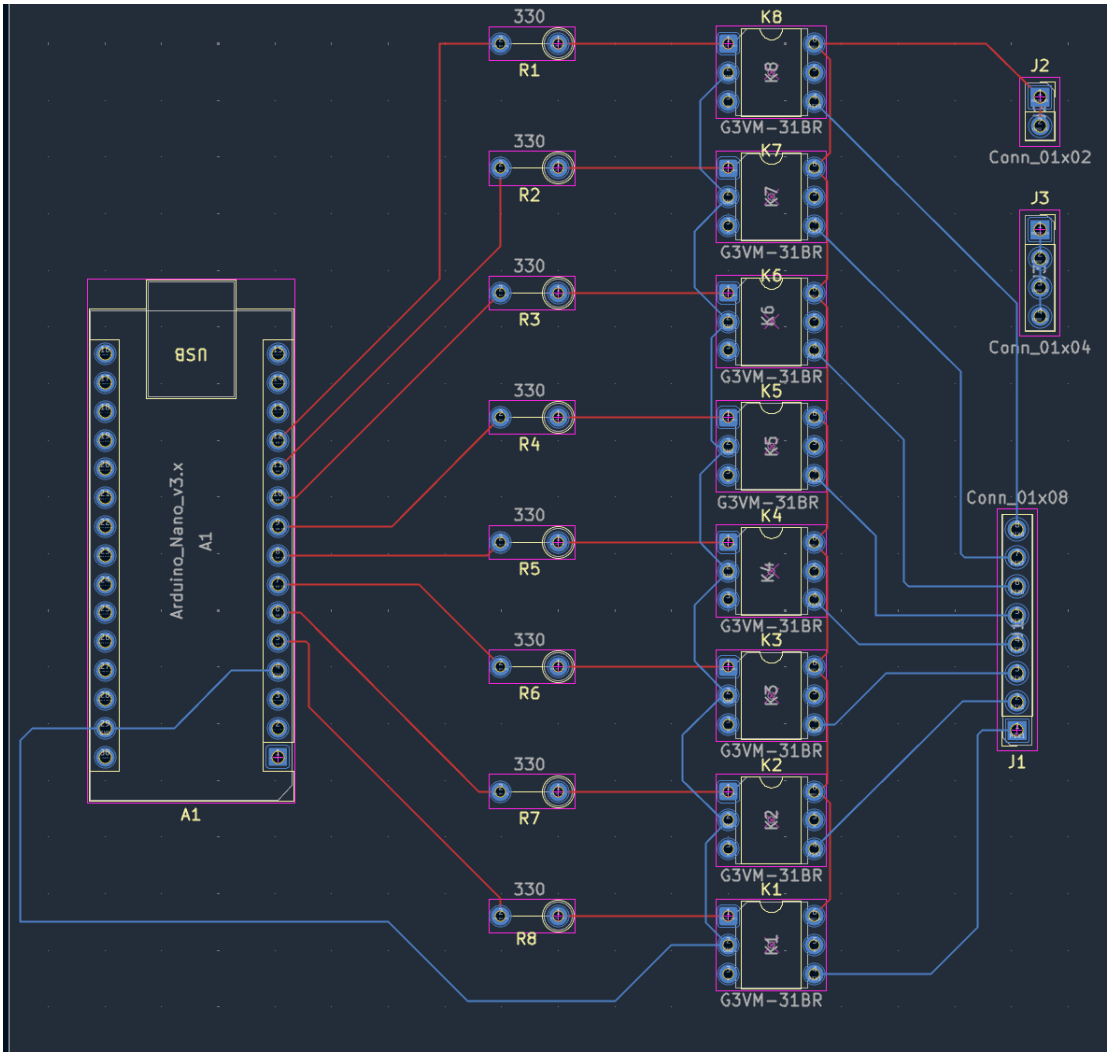



**Figure 5** Routed two-layer board for the eight-relay selector, 97.5 × 93.3 mm, all through-hole. The column of eight G3VM-31BR relays down the centre separates the control side — the Arduino Nano and the eight LED series resistors, left — from the measurement side — the eight-way pixel harness J1, Force HI on J2 and Force LO on J3, right. Two layers are required because the ground bus and the force bus each run past every relay, so the resistor and pixel traces must cross them; the crossings are free once each net can change layer.

### 6.2 Extensions

**Unattended stability protocols..** The principal benefit of automation is the ability to run without an operator. Natural extensions are maximum-power-point tracking cycled across pixels, repeated $J$–$V$ sweeps over hours to characterise light soaking and degradation, and automatic extraction of $J_{sc}$, $V_{oc}$, $FF$ and $PCE$ into a per-device summary table.

**Drift as a measurand rather than a nuisance..** This work found device drift to dominate run-to-run variation. Because the system now controls the interval between measurements precisely, that drift becomes quantifiable: a scripted sequence of identical sweeps at defined intervals would measure it directly, turning the limiting factor of this study into an experiment of its own.

**Extension to other measurements on the same bench..** The switch matrix is agnostic to what is connected to its common line. The same hardware could route pixels to a source-measure unit for $J$–$V$, to a photodetector chain for electroluminescence, or to a signal generator and oscilloscope for modulation measurements on perovskite LEDs — allowing several characterisation modes to share one sample holder without rewiring.

## 7 Conclusion

A manual rotary pixel selector used for $J$–$V$ characterisation of eight-pixel perovskite solar cells has been replaced by a microcontroller-controlled electronic switch matrix, so that a complete device can now be characterised without operator intervention from the same Python environment that already drives the Keithley 2450 SourceMeter. Three architectures were designed and evaluated: CMOS analogue multiplexers, which required four-wire sensing and a second multiplexer bank to overcome a 70 Ω on-resistance; a precision dual multiplexer, whose apparent leakage current was diagnosed — from the observation that it persisted with the device under test disconnected — as conduction through internal protection diodes when the signal node was driven below the negative supply rail; and finally eight photo-MOSFET solid-state relays, whose optically generated gate drive removes the supply-rail constraint altogether and whose 20 mΩ on-resistance (40 mΩ maximum) bounds the worst-case series error at 0.2 mV — below the instrument's own accuracy specification. This made four-wire sensing unnecessary and halved the required hardware. Validation against a direct, unswitched baseline used a three-trial protocol including a no-switch control, and the control comparison showed a *larger* discrepancy (500 $\mu$A) than the direct-versus-switched comparison (237 $\mu$A), establishing that run-to-run differences are dominated by device drift rather than by the switching hardware — a conclusion supported independently by the magnitude, the shape and the time-ordering of the difference curves.

The main takeaway is twofold. Technically, the right way to make a switch invisible to a measurement is to choose an element whose series resistance is negligible by construction, rather than to build a sensing scheme that compensates for a poor one. Methodologically, a claim that added hardware does not perturb a measurement is only meaningful when tested against a control that holds the hardware fixed — without which the drift of the device under test is indistinguishable from an instrumentation artefact.