

APPLICATION NOTE

Operating the PVX-4110, PVX-4130, PVX-4141 and PVX-4151 with common-polarity pulse and bias supplies

Why the lower-voltage supply needs its own current path to ground when both high voltage inputs carry the same polarity, and how to size the shunt resistor that provides it.

Two rails, one output

How a half-bridge pulser forms the pulse

The mechanism

Why the bias node drifts when both rails share a polarity

The shunt resistor

The published four-step sizing procedure

Bench confirmation

Reading the bias level at the voltage monitor

Operating the PVX-4110, PVX-4130, PVX-4141 and PVX-4151 with common-polarity pulse and bias supplies

Why the lower-voltage supply needs its own current path to ground when both high voltage inputs carry the same polarity, and how to size the shunt resistor that provides it.

ABSTRACT

The PVX-4110, PVX-4130, PVX-4141 and PVX-4151 pulse generators switch their output between two user-supplied high voltage rails, so a pulse that starts and ends away from ground is set up with two DC supplies of the same polarity, for example a bias supply at +500 V and a pulse supply at +1,500 V. Most high voltage DC supplies pass current in one direction only, and the lower-voltage supply in such a pair cannot absorb the current that leakage and the returning load charge push into its node. Without a path to ground the output drifts from the bias level toward the pulse level while the generator is not gated. This note explains the mechanism, gives the published procedure for sizing a shunt resistor from the bias input to ground, works the procedure through at 1 W of switching power and a 100 V bias (a 5 kohm, 5 W resistor), states what each DC supply must then deliver, and describes how to confirm the fix at the voltage monitor output.

WHO THIS PAPER IS FOR

Design and test engineers integrating a PVX-4110, PVX-4130, PVX-4141 or PVX-4151 into a time-of-flight mass spectrometer, an electro-optic drive, a detector gate or a beam kicker where the pulse baseline is offset from ground.

Technicians who see the output baseline climb toward the pulse level and need to know why.

David Brown, President, Berkeley Nucleonics Corp.

KEY RESULTS

5 kΩ, 5 W

Shunt resistor from the bias input to ground for 1 W of switching power at a 100 V bias, by the published procedure.

10 mA

Discharge current the shunt must carry at 1 W and 100 V, before the resistor value is halved to keep the bias supply sourcing current.

2 W

Continuous power the bias supply must deliver into the 5 kΩ shunt at 100 V; the pulse supply delivers the $C V^2 F$ switching power.

2000:1, 1000:1

Voltage monitor scaling into 1 MΩ on the PVX-4110 and PVX-4130, and on the PVX-4141 and PVX-4151, for checking the bias level without a high voltage probe.


CONTENTS

- | | |
|--|--|
| <ul style="list-style-type: none"> 1 Two rails and one output 2 What a DC supply can and cannot do 3 How the bias node drifts toward the pulse level 4 The remedy: a current path to ground 5 Sizing the shunt resistor 6 Worked example 7 What each DC supply must deliver | <ul style="list-style-type: none"> 8 When the problem does not arise 9 Confirming the fix on the bench 10 Applicability across the four models 11 Limitations 12 Summary and selection 13 Definitions 14 Further reading and one route onward |
|--|--|

1 Two rails and one output

A direct-coupled half-bridge high voltage pulse generator places two solid-state switches in series between a high rail and a low rail that the user supplies; the output is their

junction. A logic-level gate decides which switch conducts: gate high connects the high rail to the output, gate low connects the low rail. The pulse amplitude is therefore the difference between the two rail voltages, and the pulse width and repetition rate follow the gate. Figure 1 draws the arrangement.



SCHEMATIC Figure 1. A half-bridge output stage. Two switches in series between the user-supplied high rail (V_{HIGH} , the +HV In input) and low rail (V_{LOW} , the -HV In input); the gate selects which rail reaches the capacitive load. The pulse amplitude is V_{HIGH} minus V_{LOW} .

The two rails give the topology its polarity options. With the low rail at ground the output pulses from ground up to the high rail; with the high rail at ground it pulses from ground down to the low rail. The third case is the subject of this note: a pulse whose baseline is not ground. A detector gate that rests at a bias and pulses further in the same direction, or a deflection plate that pulses between two potentials of one sign, is set up with both rails on supplies of the same polarity, for example -HV In at +500 V and +HV In at +1,500 V, or -HV In at -1,000 V and +HV In at -500 V. The output then rests at the lower voltage and pulses to the higher one. The output stage accepts that; its only requirement is that the high rail is never below the low rail. The difficulty is on the supply side.

2 What a DC supply can and cannot do

A regulated DC supply holds its output voltage by controlling the current it passes, and most high voltage supplies pass that current in one direction only. A positive supply draws electrons from ground; in conventional terms it sources current out of its positive terminal and takes it back at ground. A negative supply does the opposite. Neither can absorb current pushed back into its output terminal, because


the regulating element conducts one way: if something external raises a positive supply's output above its setpoint, the supply's current falls to zero and its terminal voltage is set by whatever is pushing. With supplies of opposite polarity, or with one rail grounded, this limit is never tested, because every current returns through ground in each supply's natural direction. With two supplies of the same polarity the lower-voltage supply is connected to a node into which the pulse generator delivers current, and it has no way to take that current away.

3 How the bias node drifts toward the pulse level

Two things deliver current into the low rail node. The first is leakage inside the pulse generator from the high rail toward the output while the generator is not gated: small, continuous, and with no return path it can only flow into the low rail's supply terminal. The second is the load. A capacitive load is charged to the high rail during each pulse and discharged to the low rail at each falling edge, and the charge that went in at the top comes out through the lower switch into the low rail node, more often as the repetition rate rises.

If the low rail's supply cannot sink that current, the node charges. The output baseline, which should rest at the -HV In potential whenever the generator is not gated, drifts toward the +HV In potential (Figure 2). The symptom is most visible when the load draws no DC current of its own and as the

pulse repetition frequency is increased. At low rates with small loads it may not be seen, because the DC supply's own internal paths to ground absorb a small current; that is a property of the particular supply, not something to design around.




SCHEMATIC Figure 2. Output voltage against time with both rails at the same polarity and no return path for the low rail. The baseline that should hold at V_{LOW} between pulses climbs toward V_{HIGH} as leakage and returned load charge accumulate on the low rail node; the pulse amplitude shrinks by the same amount. With a shunt path fitted the baseline holds.

The result is a pulse of the wrong amplitude from the wrong baseline, and a bias supply whose current reads zero while its terminal voltage reads high. Neither instrument is faulty; the circuit is missing a return path.

4 The remedy: a current path to ground

The low rail needs a path to ground that carries, in the supply's natural direction, at least as much current as the generator pushes into the node. In low power applications a

resistor from the -HV In input to ground does this: it draws a steady current from the bias supply, which stays in regulation because it is always sourcing, and the current the generator delivers into the node flows out through the resistor instead of charging the node (Figure 3). In high power applications, where the resistor's dissipation would be large, an active shunt regulator from the bias input to ground does the same job with less heat, conducting only as much as it must to hold the node at the bias voltage.



■ **SCHEMATIC** **Figure 3.** The same output stage with a shunt resistor R_{SHUNT} from the -HV In input to ground. The bias supply sources a steady current through the resistor; leakage from the high rail and the charge returned from the load at each falling edge flow out through the same resistor to ground instead of accumulating on the low rail node.

The next two sections size the resistor by the method published with these pulse generators. It is deliberately conservative: a shunt too large in value leaves the node uncontrolled, while one too small costs only heat and supply capacity.

5 Sizing the shunt resistor

Three quantities are needed, all of them from the application rather than from the pulse generator:

- the maximum pulse repetition frequency, F ;
- the maximum differential voltage, $V = V_{HIGH}$ minus V_{LOW} , which is the pulse amplitude;
- the total load capacitance, C , including the load itself, the interconnect cable and strays.

Because the load is capacitive, the charge that goes into it must also come out, and the power associated with switching it is the quantity the whole procedure is built on. The minimum power that the resistor or shunt regulator must be able to dissipate from switching the capacitive load is

$$P = C \cdot V^2 \cdot F \quad (1)$$

where

P the switching power to be dissipated, in W

C the total load capacitance, load plus cable plus strays, in F

V the differential voltage V_{HIGH} minus V_{LOW} , in V

F the pulse repetition frequency, in Hz

From that power the discharge current through the shunt is

$$I = P / V_{LOW} \quad (2)$$

where

I the discharge current the shunt must carry, in A

V_{LOW} the bias voltage at the -HV In input, in V

and the shunt resistor value follows from the bias voltage and that current:

$$R_{SHUNT} = V_{LOW} / I \quad (3)$$

where

R_{SHUNT} the resistance from the -HV In input to ground, in Ω

A regulated supply holds the bias, and it stays in control only while it is sourcing current. If the current pushed into the node ever exceeds the current of equation (2), the supply's output current falls to zero and the bias is no longer controlled. To keep a margin, the value from equation (3) is halved, so the shunt always draws more current than the generator can deliver into the node and the supply is always sourcing.

DC current from the bias supply flows continuously through the shunt, and the resistor dissipates

$$P_R = V_{\text{LOW}}^2 / R \quad (4)$$

where

P_R the continuous power in the shunt resistor, in W

R the halved shunt resistance, in Ω

For reliability the resistor is chosen with at least twice that power rating. It also stands at the full bias voltage continuously, so its voltage rating must cover V_{LOW} with

margin: above a few hundred volts that means a high-voltage rated part or several resistors in series with the creepage between them respected.

6 Worked example

The example is the one published with the original note.


Suppose the load, the pulse amplitude and the repetition rate give a switching power of 1 W from equation (1), and the bias supply at -HV In is set to 100 V. Table 1 carries the steps: equation (2) gives 10 mA, equation (3) gives 10 k Ω , halved to 5 k Ω , and equation (4) gives 2 W, so a 5 k Ω resistor rated 5 W is chosen.

Table 1. The published sizing procedure worked at 1 W of switching power and a 100 V bias. The result is a 5 k Ω resistor rated 5 W from the -HV In input to ground.

Step	Relation	Value in the example
Switching power	$P = C V^2 F$	1 W (given)
Discharge current	$I = P / V_{\text{LOW}}$	1 W / 100 V = 10 mA
Shunt resistance	$R = V_{\text{LOW}} / I$	100 V / 10 mA = 10 k Ω
Regulation margin	$R / 2$	5 k Ω
Shunt dissipation	$P_R = V_{\text{LOW}}^2 / R$	10,000 / 5,000 = 2 W
Rating margin	$2 P_R$	5 W part

Figure 4 plots the procedure over a range of bias voltages. The halved resistance is $V_{\text{LOW}}^2 / 2P$, so it rises with the square of the bias voltage, while the continuous dissipation is

$2P$ at any bias: at 1 W the shunt dissipates 2 W whether the bias is 50 V (1.25 k Ω) or 500 V (125 k Ω).



DERIVED Figure 4. Halved shunt resistance against bias voltage for switching powers of 0.5 W, 1 W and 2 W, from $R = V_{LOW}^2 / 2P$. The 1 W curve passes through 5 k Ω at 100 V, the worked example. The continuous dissipation in the shunt is 2P at every bias voltage, so the resistor's power rating is set by the switching power, not by the bias.

7 What each DC supply must deliver

Once the shunt is fitted the two supplies carry different loads. The bias supply at -HV In must deliver the shunt power, 2 W in the example, which at 100 V is 20 mA of continuous current. The pulse supply at +HV In must deliver the $C V^2 F$ switching power of equation (1), 1 W in the example, plus the generator's quiescent losses. Both must be regulated, and the bias supply's current rating must exceed the shunt current with margin, because a bias supply in current limit has lost control of its voltage.

Where the published method is conservative

The procedure sizes the shunt from the full switching power divided by the bias voltage, so the resistor draws more than the generator pushes into the node at most operating points. That margin is intended: it costs heat and bias-supply capacity, and it buys a bias that stays regulated at every repetition rate the application will reach.

8 When the problem does not arise

Three common configurations never meet it, and need no shunt:

- Supplies of opposite polarity, for example -HV In at -1,000 V and +HV In at +500 V. Each supply passes current in its natural direction and the return path exists through ground.
- One input grounded, for example -HV In at ground and +HV In at +1,500 V, or the reverse. Ground absorbs whatever the generator returns.
- Small loads at low repetition rates, where the current delivered into the bias node is below what the supply's own internal paths to ground absorb. This depends on the supply, and raising the rate or the load capacitance later can cross the threshold: a setup that works at 100 Hz into a test load is not evidence that it will work at 10 kHz into the real one. If both rails share a polarity, fit the shunt.

9 Confirming the fix on the bench

The front-panel voltage monitor is the safe way to watch the baseline: it reproduces the output waveform at a fixed ratio into a 1 M Ω oscilloscope input and needs no high voltage probe. The ratio is 2000:1 on the PVX-4110 and PVX-4130 and 1000:1 on the PVX-4141 and PVX-4151, from the published datasheets, so a 100 V bias reads 50 mV on the PVX-4110 monitor and 100 mV on the PVX-4141 monitor.

1. Set both supplies to their operating voltages with the gate held low and the generator enabled. Confirm at the

- voltage monitor that the output rests at the -HV In voltage.
- Gate the generator at the lowest repetition rate the application uses, watch the baseline between pulses for several seconds, then raise the rate in steps to the maximum. A baseline that holds at every rate confirms the return path; one that climbs with rate, or over time at a fixed rate, means the path is missing or too small.
 - Read the bias supply's output current. With the shunt fitted it should show a steady V_{Low} / R , 20 mA at 100 V into 5 k Ω , and never fall to zero while the generator runs. Zero current with the terminal voltage above the setpoint is the signature of section §§mechanism§§.
 - After several minutes at full rate, check the shunt resistor's temperature with the supplies off. A resistor dissipating 2 W is warm; one that is hot points to a switching power above the one assumed, and the sizing should be repeated with the measured load capacitance.

Table 2. Published values that enter the shunt calculation, from the current datasheets on berkeleynucleonics.com. Internal capacitance is part of the total C in equation (1); the maximum average power is the generator's ceiling on $C V^2 F$. The rail limits are symmetric, so a common-polarity pair may be placed anywhere inside $\pm V_{\text{max}}$ as long as the differential does not exceed the model's maximum.

Model	V_{HIGH} minus V_{Low} , max.	Max. average power	Internal capacitance	PRF, continuous	Voltage monitor
PVX-4110	$\pm 10,000$ V	100 W	about 50 pF	single-shot to 10 kHz	2000:1 into 1 M Ω
PVX-4130	$\pm 6,000$ V	100 W, derated 2 W/ $^{\circ}\text{C}$ above 25 $^{\circ}\text{C}$	about 100 pF	single-shot to 10 kHz at 6,000 V	2000:1 into 1 M Ω
PVX-4141	$\pm 3,500$ V	100 W	about 120 pF	single-shot to 30 kHz at 3,500 V	1000:1 into 1 M Ω
PVX-4151	$\pm 1,500$ V	150 W, derated 2 W/ $^{\circ}\text{C}$ above 25 $^{\circ}\text{C}$	200 pF	single-shot to 240 kHz at 1,500 V	1000:1 into 1 M Ω

11 Limitations

The procedure is a sizing rule, not a model of the generator's leakage. It takes the full switching power as the quantity the shunt must handle, which gives a resistor that is safe at every operating point but larger in dissipation than a measurement of the actual return current would justify. Where that shunt power is inconvenient, the active shunt regulator of section §§remedy§§ is the alternative; this note does not design one.

The procedure assumes a capacitive load; the datasheets state that $C V^2 F$ does not apply to resistive or inductive loads. The worked example uses a round 1 W and 100 V; a real application computes equation (1) from its measured load capacitance, cable included. And the behavior of a bias supply pushed toward zero output current depends on that supply's design, which is why the fix is a return path.

The shunt carries the full bias voltage to ground; treat it, and any resistor string, with the spacing and creepage of any other high voltage component.

10 Applicability across the four models

The four pulse generators in scope share the half-bridge output stage and the pair of user-supplied inputs, so the mechanism and the procedure apply to all of them; they differ in voltage, power and rate, which set the switching power of equation (1). Table 2 collects the published values. The PVM-4210 is out of scope: it carries its own integrated high voltage supplies and has no +HV In or -HV In inputs, so there is no external bias supply to shunt.

12 Summary and selection

A pulse that starts and ends away from ground needs two supplies, and when both carry the same polarity the lower-voltage supply cannot sink the current the generator delivers into its node, so the baseline drifts toward the pulse level. A shunt resistor from the -HV In input to ground, sized by section §§sizing§§ and worked in section §§example§§, keeps the bias supply sourcing current and the baseline where it was set. Confirm it at the voltage monitor, and size the bias supply for the shunt current.

Choosing among the four models is a matter of Table 2: the differential voltage the load needs and the switching power at the intended rate. Within any of them the shunt is sized the same way.

13 Definitions

Half-bridge (totem-pole) output stage

Two switches in series between a high rail and a low rail, with the output at their junction; one switch conducts at a time.

Pulse supply (+HV In, VHIG)H)

The user-supplied DC input connected to the output while the gate is high. It sets the top of the pulse.

Bias supply (-HV In, VLOW)

The user-supplied DC input connected to the output while the gate is low. It sets the pulse baseline and is needed only when the output is biased away from ground.

Shunt resistor

A resistor from the -HV In input to ground: a continuous current for the bias supply to source, and a path to ground for returned charge.

Symbols used in this document.

Symbol	Meaning	Units
C	Total load capacitance: load, interconnect cable, strays and the generator's internal capacitance	F
V	Differential voltage, V_{HIGH} minus V_{LOW} (the +HV In and -HV In voltages); the pulse amplitude	V
F	Pulse repetition frequency	Hz
P	Switching power, $C V^2 F$	W
I	Discharge current the shunt must carry, P / V_{LOW}	A
R_{SHUNT}, R	Shunt resistance before and after halving	Ω
P_R	Continuous power in the shunt resistor, V_{LOW}^2 / R	W

Acronyms

DC Direct current

HV High voltage

PRF Pulse repetition frequency

SHV Safe high voltage connector

TTL Transistor-transistor logic level, the gate input standard

14 Further reading and one route onward

- *Capacitive-Load Pulsing*, Directed Energy technical note: the physics of $I = C \, dV/dt$ and the $C V^2 F$ power budget behind equation (1).
<https://www.berkeleynucleonics.com/docs/article-dei-capacitive-load-pulsing.html>

- *Selecting an HV Pulse Generator and How Do I Slow the Rise Time of My Voltage Pulser?*, Directed Energy technical notes on model selection and on resistor voltage ratings and strings, which apply to the shunt as well. Both are listed at
<https://www.berkeleynucleonics.com/docs/dei-technical-papers.html>
- The four model datasheets, the authority for Table 2:
<https://www.berkeleynucleonics.com/docs/bnc-dei-pvx-4110-datasheet.html> and the matching pages for the PVX-4130, PVX-4141 and PVX-4151.

If you are setting up a common-polarity pulse on one of these generators and would like the shunt sized against your actual load capacitance, cable and repetition rate, our applications engineers would be glad to work through it with you. Write to info@berkeleynucleonics.com or call 800-234-7858.