PRECISION TIMING REFERENCE HANDBOOK
BNC-HB-003 REV B
SEPTEMBER 2026

Digital Delay Generators in the pulsed laboratory

How a counted clock turns one trigger into an experiment, what jitter and insertion delay really cost at the bench, and the working reference for the Model 577 and Model 588B.

ABSTRACT

Every pulsed experiment is a scheduling problem: a handful of instruments must act in a fixed order, at fixed intervals, shot after shot. This handbook is about the instrument that owns that schedule. It builds the timing physics first, jitter and its references, insertion delay, the counted-clock architecture, and the accuracy arithmetic, then reviews the Berkeley Nucleonics Model 577 and Model 588B against the requirements that physics derives. Six application chapters follow, from Q-switched lasers to streak cameras, each with its worked numbers and, wherever the wiring is the lesson, its channel map. It closes with the troubleshooting table, the cable and logic-level references, and the quick-reference page a working bench actually opens a handbook for.

WHO THIS PAPER IS FOR

Experimental physicists and laser lab researchers sequencing pumps, probes, gates and cameras. **Test and integration engineers** who need repeatable multi-channel timing with numbers they can defend. **Graduate students** building a first pulsed bench. **Anyone** whose experiment produces a different answer when a cable is moved.

CONTENTS

1 How to use this handbook

This is a reference document, arranged in four parts that can be entered separately. Part I is the physics and carries no product in it. Part II is the instrument review, with every published value quoted with its condition. Part III is six applications at wiring level. Part IV is the bench material: practice, troubleshooting, reference tables, and the quick reference.

Table 1 is the router. Find the row that matches why you opened this document and start where it points.

Table 1. A router, not a summary. Each row names the shortest path through the handbook for one kind of question. Section 24 collects every governing relation and published value in one place with its condition attached, because a reader who has been through the handbook once will open it again only for numbers.

If you are	Start at	And you can skip
Learning what jitter and insertion delay actually are	Part I, sections 2 and 3 to 6	Parts II and III
Specifying a delay generator for an experiment	Section 7, then Part II, then section 24	Part III until wiring day
Wiring a laser, camera or detector bench this week	The matching recipe in Part III, then section 22	Part I on the first pass
Chasing a timing fault that appeared from nowhere	Section 22, then section 23	Everything that is not the fault
Looking for a number you already understand	Sections 23 and 24	Everything else

Conventions used throughout

- **Every amplitude names its load.** A pulse amplitude into 50 Ω and the same output into a high-impedance input differ by a factor of two, and section 10 works the arithmetic. An amplitude without a load is not a specification.
- **Every jitter figure names its reference.** Trigger-to-output and channel-to-channel jitter are different physical quantities that differ by an order of magnitude on the same instrument. Section 3 is about why.
- **Every accuracy figure names the delay it assumes.** Delay accuracy has a term proportional to the setting, so the error at 1 s is a thousand times the error at 1 ms.
- **Published, documented and stated are used deliberately.** Published means it appears on a Berkeley Nucleonics datasheet or user manual. Documented means it is worked through with evidence somewhere else. Stated means somebody said it.

PART I

Time on the bench

What a delay generator is for, what jitter and insertion delay really are, how a counted clock makes a delay, and the accuracy arithmetic. No product is named in this part.

2 Why experiments run on delayed edges

A pulsed measurement is a set of instruments that must act in one fixed order at fixed intervals: pump a laser, open its cavity, gate a camera, arm a digitizer, read out. None of those devices knows about the others. Each exposes one input, a trigger edge, and does its one thing some time after that edge arrives. The whole experiment therefore reduces to a scheduling problem: produce N electrical edges with defined separations, once per shot, forever.

The improvised answer grows one box at a time. A function generator supplies the repetition rate. A coil of coax delays one line by its length. A one-shot circuit stretches a pulse for a gate. Each addition works, and the sum does not: every analog delay drifts with temperature, every added box has its own threshold and its own latency, and by the fifth device nobody can say where time zero is anymore.

The antagonist of this handbook is not any single bad instrument. It is the accumulation of small, separately reasonable timing decisions into a bench that cannot be trusted or reproduced.

A Digital Delay Generator collapses that accumulation into one instrument: one timebase, one trigger, and many output channels, each programmed with a delay and a width. Every edge the experiment needs comes from the same counted clock, so every interval is a number in a register rather than a length of cable. The three questions that decide whether such an instrument can carry a given experiment are the subjects of the next three sections: how repeatable are its edges (jitter), how soon after a trigger can it act (insertion delay), and how true are its programmed intervals (accuracy).

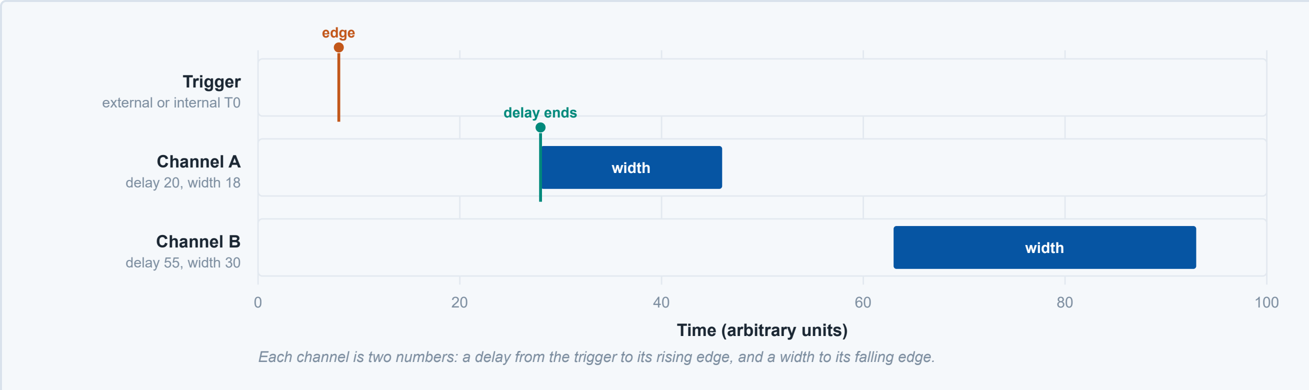

Figure 1. The whole programming model in one picture. A trigger edge starts every channel's delay counter at once; each output goes high when its own delay elapses and low after its own width. Every interval in the experiment is the difference of two programmed numbers, which is what makes the schedule reproducible: registers do not drift the way cable coils and one-shots do. SCHEMATIC

Figure 1 is worth holding through everything that follows. Every specification in this handbook attaches to one of its elements: jitter to the repeatability of the edges, insertion delay to the gap between the trigger and the earliest possible output, accuracy to the truth of the programmed numbers.

3 Jitter, and the reference that gives it meaning

Jitter is the shot-to-shot wander of an edge about its ideal position. It is quoted as an RMS value, the standard deviation of the edge-time distribution, and the first thing to understand is that no instrument has one jitter. It has one jitter per pair of points you compare. The two that matter for a delay generator are trigger-to-output jitter, an external trigger edge against a delayed output edge, and channel-to-channel jitter, one output of the instrument against another on the same shot. Section 4 shows why those two differ by an order of magnitude on the same instrument, and that one difference sets the design rule the recipes of Part III lean on.

RMS against peak-to-peak

Random jitter is Gaussian and therefore unbounded: watch long enough and some edge lands arbitrarily far out. A peak-to-peak jitter figure with no population size and no probability attached is therefore meaningless. The honest conversion is a crest factor:

$$t_{\text{pk-pk}} = k \times t_{\text{RMS}} \tag{1}$$

where

k is a crest factor chosen for a target exceedance probability: 6 covers about 99.7 percent of edges, 6.6 covers 99.9 percent, and serial-data work uses 14.1 for one edge in 10^{12} outside the window

t_{RMS} is the RMS jitter, the number datasheets publish

Worked once so the habit forms: an output with 50 ps RMS jitter has a 6-sigma peak-to-peak of 6 × 50 = 300 ps, and a 10^{-12}-grade peak-to-peak of 14.1 × 50 = 705 ps. Same instrument, same physics, both numbers correct, differing only by the convention in equation (1). Quote the convention or quote nothing. Deterministic jitter, the bounded kind that comes from supply ripple, crosstalk and the quantization of section 4, adds on top linearly rather than in quadrature, because it is not Gaussian.

Where jitter comes from: the threshold crossing

Every trigger input is a comparator, and a comparator converts voltage noise into timing noise at a rate set by how fast the signal crosses the threshold:

$$t_{\text{j}} = v_{\text{n}} \,/\, (\mathrm{d}V/\mathrm{d}t) \tag{2}$$

where

v_n is the RMS voltage noise at the comparator input, typically a millivolt or so at a bench input

dV/dt is the slew rate of the signal through the threshold

Worked at both extremes. A clean 2 V edge that rises in 4 ns slews at 0.5 V/ns; with 1 mV of input noise, equation (2) gives 1 mV / (0.5 V/ns) = 2 ps of jitter. The same 2 V arriving over 100 ns slews at 0.02 V/ns and the same millivolt becomes 50 ps. A 25 times slower edge costs 25 times the jitter, from nothing but the geometry of the crossing. Drive every trigger input with the fastest clean edge available: that is a physics statement, not a preference.

THE WRONG INTUITION, NAMED

The intuition to break is that jitter is a property of the instrument alone. It is a property of a path: source slew rate, cable, comparator, clock domain and output driver together. The same delay generator shows an order of magnitude between its best-referenced and worst-referenced numbers, and a slow trigger edge can dominate everything the datasheet says. When a jitter measurement disappoints, suspect the path before the box.

4 The counted clock, and why channels agree with each other

Nearly every Digital Delay Generator is built the same way, and knowing the architecture is what lets a reader predict its behavior instead of memorizing it. A crystal oscillator, stable but slow, is multiplied by a phase-locked loop up to a counting frequency of order 100 MHz. Each channel is a counter that counts N whole periods of that clock, followed by an analog vernier that adds the fraction of a period the counter cannot represent:

$$\text{delay} = N \times T_{\text{clk}} + t_{\text{vernier}} \tag{3}$$

where

N is the whole number of clock periods counted

T_{clk} is the clock period, 10 ns at 100 MHz, 5 ns at 200 MHz

$t_{vernier}$ is an analog fine delay, between zero and one clock period, which is where sub-period resolution comes from

The counter gives range bounded only by its width, in practice a thousand seconds and more, at crystal stability; the vernier gives resolution far below one clock period. The split explains the specification sheet: range comes from counter width, resolution from the vernier, long-delay accuracy from the crystal, and fine-scale accuracy from vernier linearity.

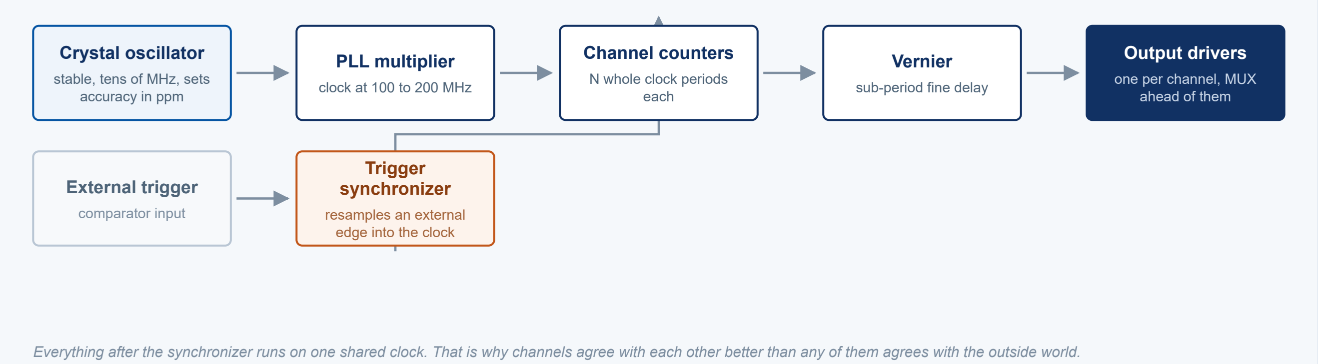

Figure 2. The canonical delay generator. An external trigger must pass through the synchronizer to enter the clock domain, and that crossing carries a cost every datasheet reflects: the trigger-referenced jitter figure is always far larger than the channel-to-channel figure, because only the trigger path pays the resampling toll. Internally triggered operation starts inside the clock domain and never pays it. SCHEMATIC

The synchronizer, quantified

An external trigger is asynchronous: it arrives anywhere within a clock period with uniform probability, and the synchronizer of Figure 2 aligns it to the next clock edge. The timing error of that alignment is uniformly distributed across one period, and the RMS of a uniform distribution of width *T* is *T*/√12:

$$t_{\text{RMS,sync}} = T_{\text{clk}} \,/\, \sqrt{12} \tag{4}$$

where

T_{clk} is the counting clock period

$\sqrt{12}$ is 3.464, the ratio between the width and the RMS of a uniform distribution

Worked at the two clock rates in this handbook. At 100 MHz the period is 10 ns and equation (4) gives 10 / 3.464 = 2.89 ns RMS. At 200 MHz, 5 ns and 1.44 ns RMS. A bare synchronizer, nothing else wrong, contributes nanoseconds. When an instrument publishes an external-trigger jitter below its own $T_{clk}/\sqrt{12}$, as both instruments in Part II do at 800 ps against bounds of 2.89 and 1.44 ns, it is telling you the trigger path measures the arrival phase within the clock period rather than merely resampling, and the published number is the residual of that finer mechanism.

Why channel-to-channel wins

Here is the fact that decides experiment architecture. Whatever error the synchronizer makes on a given shot, it makes once, before the counters start, so it moves every channel of that shot together. In the interval between channel A and channel B the synchronizer term is common mode and cancels exactly. What remains is only the differential part: vernier noise, driver threshold noise, clock distribution skew noise, tens of picoseconds. The trigger-referenced number carries the whole synchronizer residual; the channel-to-channel number carries none of it.

> **THE DESIGN RULE THIS BUYS**
>
> Put every interval that matters between two channels of one instrument, and let the external trigger position only the whole shot. An interval built as channel minus channel inherits the tens-of-picoseconds figure; the same interval built across two separately triggered devices inherits the trigger figure at both ends. The recipes of Part III lean on this rule throughout, and section 16 works the case where it is worth a factor of sixteen.

From phase noise to jitter, when the timebase is the question

For completeness, the conversion between the frequency-domain and time-domain descriptions of a timebase, needed whenever an oscillator datasheet quotes phase noise and the experiment budget needs seconds:

$$t_{RMS} = \sqrt{2 \int L(f)\, df} \,/\, (2\pi f_0) \qquad (5)$$

where

$L(f)$ is the single-sideband phase noise as a linear power ratio per hertz, integrated over the stated offset band

f_0 is the carrier frequency of the clock

Worked once, digit by digit, at a 100 MHz clock with phase noise flat at −120 dBc/Hz from 1 kHz to 1 MHz offset. The linear ratio is 10^{-12} per hertz; the band is 999,000 Hz wide; the integral is 9.99×10^{-7}; doubled for both sidebands, 2.0×10^{-6}; the square root is 1.41×10^{-3} radians; divided by $2\pi \times 10^{8}$ radians per second, the result is 2.25 ps RMS. The integration limits are part of the answer, and equation (5) without its limits is not a number.

5 Insertion delay, the latency floor

Insertion delay is the time from an external trigger edge at the connector to the earliest possible output transition with the programmed delay at zero. It is the instrument's built-in latency: the comparator that turns an analog edge into a decision, the synchronizer of section 4, the counter start logic, and the output driver, summed. No programmed setting can produce an output earlier than the trigger plus this floor, and for the instruments in Part II the floor is around a hundred nanoseconds.

Internally triggered operation has no insertion delay, and it is worth being precise about why. When the instrument free-runs on its internal rate generator, time zero is created inside the clock domain: there is no comparator and no synchronizer in the path, so there is nothing to insert. The question the floor answers, how soon after an outside event can this instrument act, simply does not arise when the instrument itself is the source of the event.

- **Advance the master event.** If the process allows it, trigger the generator earlier than the moment that matters by at least the insertion delay, and program the true delays on top. A rotating machine, a modelocker, a clock divider: anything periodic offers an earlier edge.
- **Make the generator the master.** Run the experiment from the internal rate generator and drive the laser, the camera and the digitizer as channels. The floor vanishes because nothing upstream of the clock domain exists. This is the architecture every recipe in Part III prefers, and section 18 shows the case where it is the only one that works.
- **Split one trigger through matched paths.** When several instruments must share an outside event, fan the same edge to all of them so each absorbs its own latency from the same time zero, then absorb the fixed differences into programmed offsets. Insertion delay only hurts when one path carries it and another does not.

> **A BOUND WITH ITS CAUSE**
>
> A gate cannot open sooner after an outside flash than the insertion delay allows: about 110 ns on the benchtop instrument of Part II, because a comparator, a synchronizer and a driver stand in the path. If the experiment needs a gate within nanoseconds of an event, the event must come from the generator, not the generator from the event. That single sentence decides the architecture of section 18.

6 Accuracy, which is not jitter

Jitter is shot-to-shot scatter; accuracy is where the mean lands. The two are set by different hardware, the vernier and driver for jitter, the crystal for accuracy, and an experiment can need either, both, or neither. The published form of delay accuracy for the instruments in this handbook is a fixed term plus a proportional term:

$$\text{error} \leq 1\ \text{ns} + 10^{-4} \times \text{delay} \tag{6}$$

where

$1\ ns$ is the fixed part: calibration residual and vernier zero error

$10^{-4} \times delay$ is the timebase part, 100 ppm of the programmed delay, which dominates beyond 10 µs

Equation (6) rewards being evaluated at three settings, because its two terms change places. At 1 µs the bound is 1 ns + 0.1 ns, dominated by the fixed term. At 180 µs, a laser timing value section 15 uses, it is 1 ns + 18 ns = 19 ns. At 1 s it is 100 µs, one hundred thousand times the fixed term. Same instrument, same specification line, five orders of magnitude between the absolute errors, purely because the delay grew. Any experiment quoting an accuracy figure without the delay it assumes is quoting nothing.

The repair for the proportional term is a better ruler. Locking the instrument's PLL to an external reference replaces the internal crystal's error with the reference's, and two distinct benefits follow. Absolute accuracy inherits the reference: a GPS-disciplined standard at one part in 10^{11} shrinks that 100 µs proportional error at one second to 10 ps, leaving the 1 ns fixed term as what remains of the bound. And mutual coherence, usually worth more: every instrument locked to one reference shares the same frequency error, so inter-instrument drift cancels as common mode even when the reference itself is mediocre. Two free-running instruments at 100 ppm each can disagree by 200 ppm of any interval; locked to one house clock they track together. Section 12 covers what the instruments accept.

7 Derived requirements

Part I turned into a checklist. Table 2 states what the physics demands of any delay generator asked to carry a pulsed experiment, each requirement traceable to a section rather than to a brochure. Part II reviews two instruments against it, and a reader can score any vendor with the same seven rows.

Table 2. Seven requirements from five sections of physics. R7 is the quiet one: the moment an experiment outgrows its channel count, intervals start crossing instrument boundaries, and every one that does trades a channel-to-channel number for a trigger-referenced one.

#	Requirement	Derived from
R1	Channel-to-channel jitter specified, in RMS, tens of picoseconds, because that is the number every internal interval inherits	Sections 3, 4
R2	Trigger-referenced jitter specified separately, and honestly larger	Section 4
R3	Insertion delay published, so the earliest-gate arithmetic can be done before wiring day	Section 5
R4	An internal rate generator good enough to be the master clock of the bench	Sections 4, 5
R5	Delay accuracy in the fixed-plus-proportional form, with resolution far finer than any interval being set	Section 6
R6	External clock input and output, for house-reference lock and multi-unit coherence	Section 6
R7	Enough independent channels, with per-channel modes, gating and routing, that one timebase can carry the whole experiment	Sections 2, 4

PART II

The Model 577 and Model 588B, reviewed

Two instruments measured against the seven requirements of Part I, every published value carrying its condition and its source, and every known documentation conflict named instead of papered over.

8 The family: a benchtop and a rack

Berkeley Nucleonics builds Digital Delay Generators in two working formats. The Model 577 is the benchtop: four or eight independent channels behind a display and keypad, made to be driven by hand at a laser table and by script overnight. The Model 588B is the rack system: twelve channels in a 1U chassis or twenty-four in 2U, with front-panel status LEDs and no local display, made to be installed once and driven entirely by remote interface. The two share the programming model of Figure 1, 250 ps setting resolution, and the accuracy form of equation (6), which is what lets one handbook serve both.

Table 3 is the published family at a glance; the engine and I/O detail follows in sections 9 to 12.

Table 3. The published family. Sources: the Model 577 datasheet and manual (Document Version 1.5), the Model 588B datasheet (V1.1) and manual (Document Version 1.0). The 588B's thirty-six outputs are a build-to-order configuration, routed by multiplexing a 24-channel unit's timers rather than by thirty-six independent timing channels.

	Model 577	Model 588B
Channels	4 (577-4C) or 8 (577-8C)	12 (588B-12C) or 24 (588B-24C), up to 36 outputs to order
Format	Benchtop, 10.5 × 8.25 × 5.5 in, 8 lb, color display and keypad	Rack mount, 19 × 10 × 1.75 in (1U) at twelve channels or 19 × 10 × 3.50 in (2U) at twenty-four, 8 lb, LED status, no display
Delay and width resolution	250 ps	250 ps
Delay range	0 to 1000 s	0 to 2000 s
Width range	10 ns to 1000 s	10 ns to 2000 s
Internal rate	0.001 Hz to 20.000 MHz	0.0002 Hz to 20.000 MHz
Timebase	100 MHz low-jitter PLL, 50 MHz 50 ppm crystal	200 MHz low-jitter PLL, 50 MHz 25 ppm crystal
Interfaces	USB and RS-232 standard; Ethernet and GPIB optional	USB, RS-232 and Ethernet, all standard

Figure 3. The Model 577 benchtop: eight output connectors, A through H, down the left of the panel with the gate and trigger inputs below them, a color display with soft keys, keypad and cursor knob. The display's status line lists all eight channels at once, and the panel itself marks the input pair at 30 V peak maximum, the conservative limit this handbook quotes. PHOTOGRAPH

Figure 4. The Model 588B rack unit: twenty-four channel status LEDs, a dual-purpose trigger and gate input beside a sync connector, and a Run button, in 2U of rack height. There is no local display and no controls beyond Run and power: configuration arrives over the remote interfaces of section 13. It is the shape of an instrument meant to disappear into a facility. PHOTOGRAPH

Figure 3 and Figure 4 carry the selection question most readers arrive with. The 577 is the right instrument when a person stands at the bench: the display makes delay scanning and troubleshooting immediate. The 588B is the right instrument when channel count or facility integration rules: twenty-four channels on one timebase means twenty-four intervals that all enjoy the channel-to-channel number of section 4, with no interval forced across an instrument boundary.

9 The timing engine, against R1 to R5

Table 4 is the timing specification both models publish, arranged by the requirement each row answers. Every value carries its reference and its source; the one known conflict is stated in the table rather than hidden under it.

Table 4. The timing engine. Sources: 577 datasheet and manual; 588B datasheet and manual. Channel-to-channel jitter is 50 ps RMS on both instruments. Budget every interval you build on either one against that figure.

Quantity	Model 577	Model 588B	Answers
Channel-to-channel jitter	< 50 ps RMS	50 ps RMS (output module tables)	R1
Internal rate (T0) jitter	< 50 ps RMS	< 50 ps RMS (manual; see note)	R4
External trigger jitter	< 800 ps RMS	800 ps RMS	R2
Trigger insertion delay	< 110 ns	< 160 ns	R3
Delay accuracy	1 ns + 0.0001 × delay	1 ns + 0.0001 × delay	R5
Resolution, delay and width	250 ps	250 ps	R5
Max external trigger rate	1/(200 ns + longest active pulse), 5 MHz ceiling	1/(200 ns + longest active pulse), 5 MHz ceiling	R2

Two of those rows repay a second look. The 800 ps trigger figure against the 2.89 ns bare-synchronizer bound of equation (4) at 100 MHz says the trigger path resolves arrival phase within the clock period, as section 4 predicted. And the 16:1 ratio between the trigger and channel figures is the measured form of the common-mode cancellation argument: both instruments deliver roughly the same channel-to-channel performance despite different clock rates, because the differential noise of vernier and drivers, not the clock, is what remains after the synchronizer term cancels.

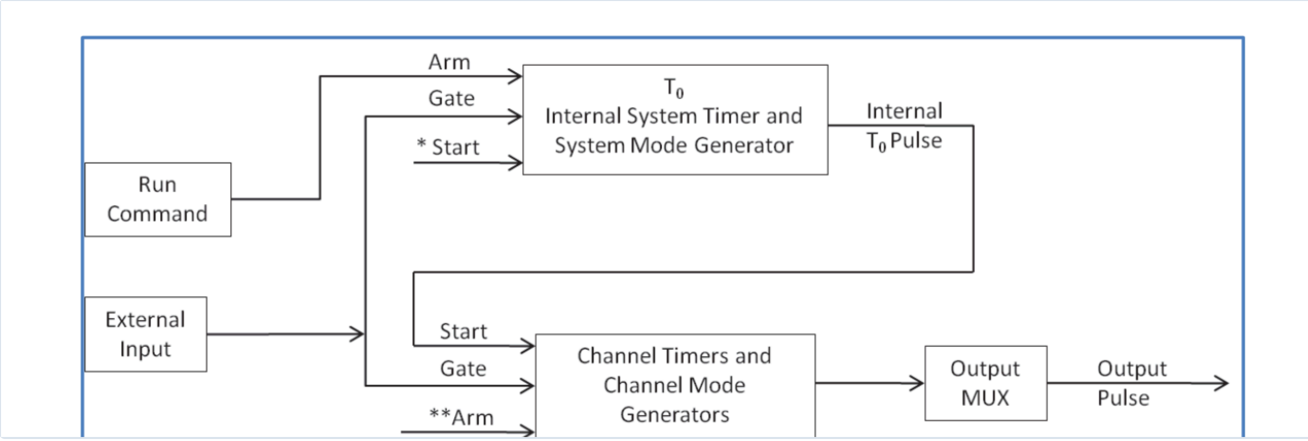

Figure 5. The instrument's own architecture drawing, from the Model 588B manual: run command and external inputs feed arm, gate and start logic into the T0 system timer, whose internal T0 pulse starts the channel timers, and an output multiplexer routes timers to output drivers. It is Figure 2 of Part I with the labels the products actually use, and it locates every specification of this section: trigger jitter enters at the left edge, channel-to-channel jitter lives between the channel timers and the output pulse they produce, and the MUX block ahead of the outputs is the routing fabric section 11 exercises. SCHEMATIC

Behind Figure 5 is the retrigger rule both instruments enforce: the system ignores incoming triggers until every channel has completed its programmed pulse. The manual's own worked example is exact: with 100 µs widths and a 50 µs delay, no trigger can be honored faster than every 150 µs, about 6.7 kHz, however far below the 5 MHz ceiling that is. The 588B adds a stated budget: delay plus width plus 75 ns of hardware reset must fit inside the T0 period, or pulses are dropped.

10 Outputs, and the arithmetic of a 50 ohm world

Every amplitude specification in this section names its load, because the load is half the number. The physics is one voltage divider: a source of impedance Z_S driving a load Z_L delivers

$$V_{load} = V_{open} \times Z_L / (Z_S + Z_L) \qquad (7)$$

where

V_{open} is the open-circuit swing the driver produces

Z_S is the output impedance, 50 Ω for the TTL mode here

Z_L is the load: 50 Ω for a terminated line, about 1 MΩ for a bare oscilloscope input

The Model 577 manual's own output table is the worked example, and it is worth reproducing exactly because it is the most-asked support question in this product class. Table 5 adds the third load the other documents publish and the divider arithmetic of equation (7) beside each row.

Table 5. The Model 577 TTL/CMOS output (50 Ω output impedance) into its three published loads. Sources: the 577 manual for the 1 kΩ and 50 Ω rows, the comparison-era datasheet material for high impedance. The halving is not a defect; it is equation (7) doing what it must. The trap it lays is a logic-level one: 2.0 V meets a TTL input threshold of 2.0 V with no margin at all and misses every 5 V CMOS input threshold outright, 3.15 V for HC parts and 3.5 V for 4000-series, so a terminated line into CMOS-gated hardware needs the high-current option below, not hope.

Load	Published TTL level	The divider says
High impedance	4.0 V typical	the full open-circuit swing
1 kΩ	4.0 V typical	1000/1050 = 0.95 of the swing, printed as 4.0 V
50 Ω	2.0 V typical	50/100 = exactly half: 4.0 V becomes 2.0 V

Adjustable mode on the same module is published at 2.0 to 20 VDC into 1 kΩ and 0.8 to 8.0 VDC into 50 Ω against a 75 Ω source impedance, and 50/(50+75) = 0.4 reproduces the second range from the first, reading the published 1 kΩ figures as the open-circuit swing they approximate at 1000/1075 = 0.93 of it. The datasheet's older 1 to 10 V into 50 Ω lineage implies a 50 Ω source. The 0.8 to 8.0 V pair is the range to design against.

The module options answer the traps the standard output leaves open. The TZ50 high-current option holds at least 4 V of TTL drive into a 50 Ω load, for terminated lines that must still clear CMOS thresholds. The AT35 and AT45 options raise the ceiling for Pockels cell drivers and other stiff loads: the AT45 delivers 4 to 45 V in 20 mV steps, with published rise times below 2 ns into low impedance and below 9 ns into high impedance, a 150 V total-amplitude budget across the system, no short-circuit protection, and a stated limit of four channels. Optical output and input modules (820 and 1300 nm, ST fiber) carry triggers across kilovolt-potential gaps, at a price the datasheet states plainly: the optical input adds its own insertion delay below 300 ns and jitter below 1.4 ns RMS. Every one of those numbers is a published module specification, quoted with its condition, and the AT45's amplitude floor is the corrected table value, 4 V, not the 5 V its own datasheet prose once said.

11 Routing, gating and modes: R7 at work

Channel count only matters if the channels can be shaped and combined. Three mechanisms do the shaping on both instruments; the oscilloscope captures in this section are from the Model 588B manual, evidence rather than illustration, and the mechanisms they record are common to both instruments.

The multiplexer

Ahead of each output driver is an OR fabric. On the Model 577, any or all channel timers may be OR'd onto any output, an 8-bit mask per output; on the Model 588B, up to five timers may be routed to each output. Only timing is multiplexed, not amplitude: the output level is the output module's. The fabric turns channel timers into an instruction set, and the two captures below, Figure 6 and Figure 7, are its two canonical programs.

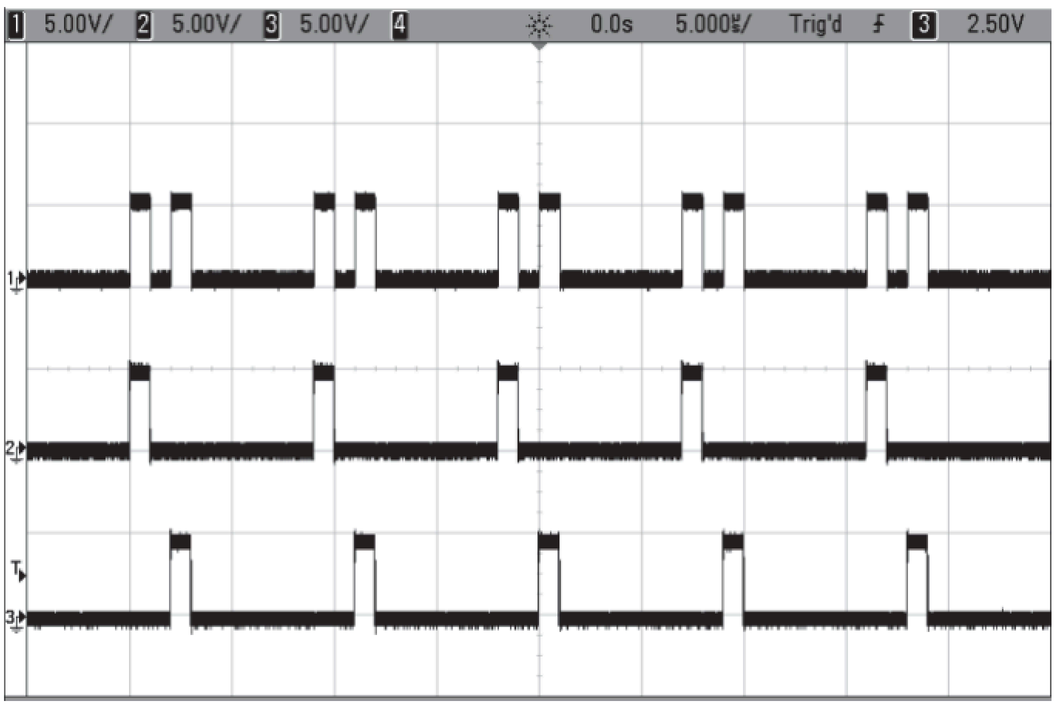

Figure 6. Double pulse from one output, no extra hardware: the lower two traces are two channel timers programmed with offset delays, and the top trace is one output carrying both, because both timers are routed to it through the multiplexer. Twice per period, one connector. This is how a pump laser gets a pre-pulse, or a camera a bracketing pair, without consuming a second cable run.

MEASURED

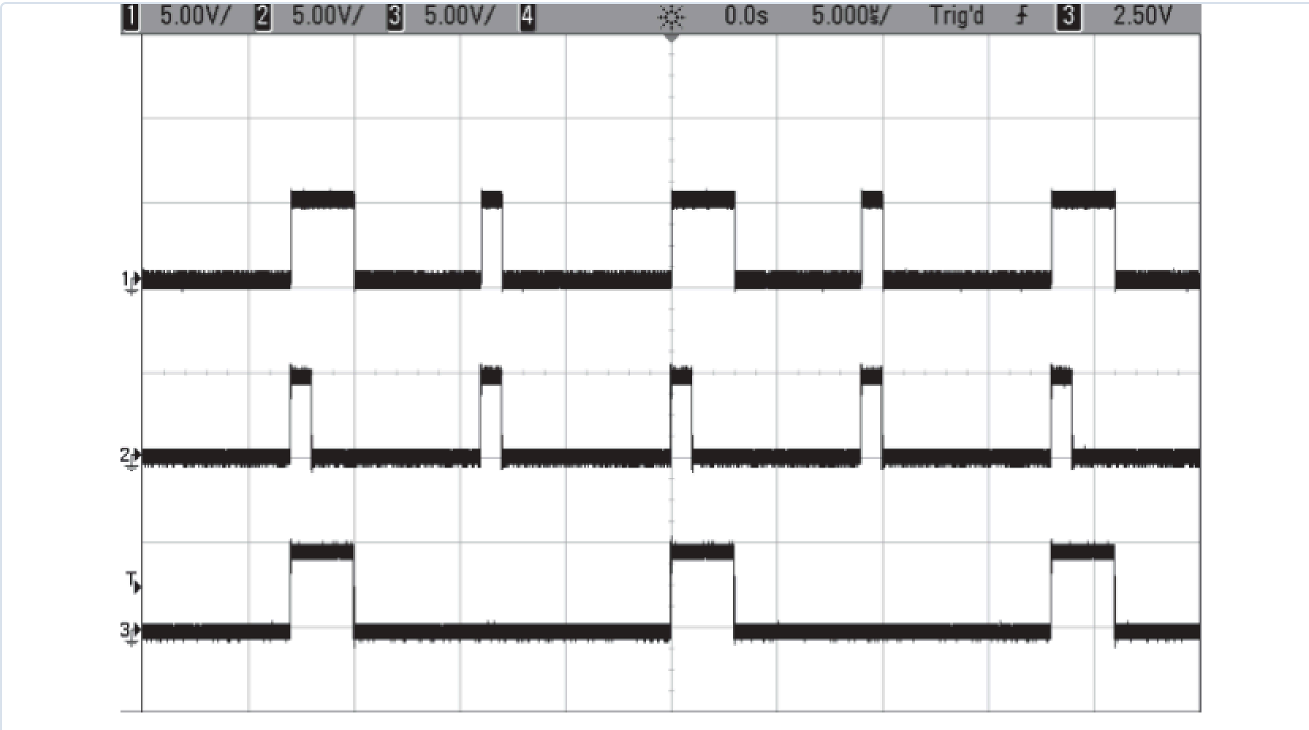

Figure 7. Alternating pulse width on a single output: the middle trace is a narrow-pulse timer firing every period, the bottom trace is a wide-pulse timer in duty-cycle mode firing every other period, and the top trace is one output carrying both. Where the wide pulse fires it covers the narrow one, so the output alternates wide, narrow, wide, narrow: a two-state program written entirely in timing registers. Substitute any shapes and any on-off pattern up to millions of pulses and the mechanism is unchanged. MEASURED

Gating

Both instruments accept a gate input that can enable or inhibit output, globally or per channel, active high or active low. The two inhibit flavors differ in one respect a laser bench cares about: pulse inhibit lets a pulse in progress finish and blocks the next one, while output inhibit cuts the output immediately. The Model 588B adds channel-timer gating, GATA as a global gate, GATB as a bank gate and INHB as a bank inhibit, so one timer channel can gate a group of others with no external wiring at all. Figure 8 is the behavior on the glass.

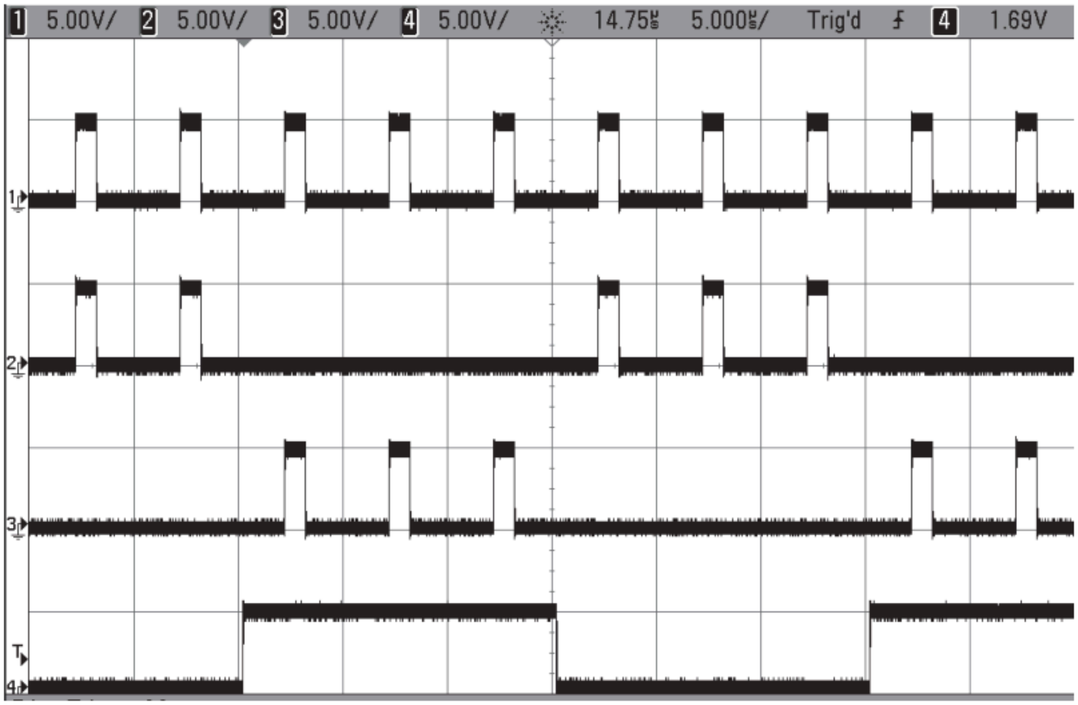

Figure 8. One gate, three verdicts: the bottom trace is the gate, the top channel ignores it, the second channel is inhibited while the gate is high, and the third runs only while the gate is high. Per-channel gate modes let one input silence one bank and enable another in complement, which is the interlock pattern of Part III: a safety shutter or ready signal holds off the hazardous channels while diagnostics keep running on the same instrument. MEASURED

Per-channel modes, and divide-by-N

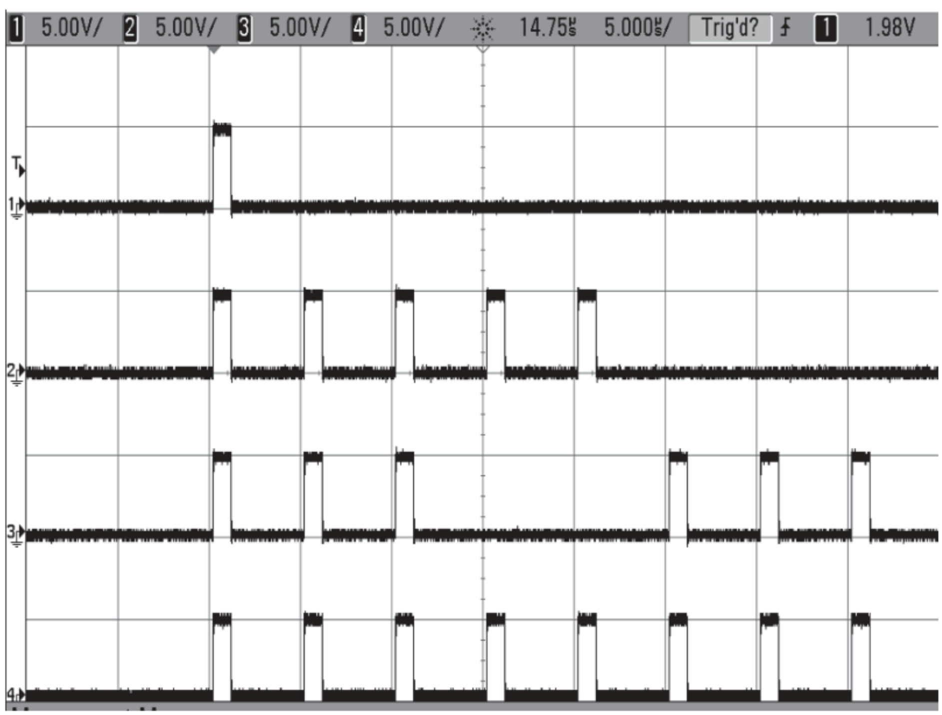

Figure 9. Four channel personalities against one trigger, one instrument, one capture: a single-shot channel that fires once, a burst channel that fires a programmed count, a duty-cycle channel that fires a pattern, and a continuous channel. Modes are per channel, so one instrument runs the one-off camera arm, the thousand-shot acquisition burst and the always-on sync from the same trigger. MEASURED

Figure 9 includes the mode that quietly replaces a rack of dividers. Duty cycle with an on-count of 1 and an off-count of N − 1 fires every Nth pulse, which is the divide-by-N function of older pulse generators as a special case; the Model 588 manual states the equivalence in exactly those terms, and on both instruments reviewed here the counters run to 10,000,000 per channel. A photodiode watching an 80 MHz oscillator, a channel dividing by 80,000, and the experiment runs at 1 kHz in lockstep with the laser's own clock. Burst counters reach 10,000,000 per channel on both models, and 4,000,000,000 at system level on the 588B; wait counts let a channel hold off for a programmed number of periods before joining.

Two referencing features complete the set. A channel may take another channel, rather than T0, as its sync source, so a chain of events can be programmed as relative offsets that survive retuning of the head of the chain, and the 577 documentation describes referencing one channel against another in positive or negative time, its negative-delay feature. Circular chains are refused by the firmware. The instruments' own captures of independent per-channel rates and of negative delay are reproduced in the datasheet set; the negative-delay capture shows one channel firing ahead of the reference channel's edge, which is only possible because everything is programmed against the same future T0.

12 Clocks in, clocks out, and staying coherent

R6 asks that the instrument accept the house ruler and offer its own. Both comply, differently enough to matter. The Model 577 accepts an external clock from 10 to 100 MHz in discrete steps and can emit T0 or a reference output over the same range. The Model 588B accepts a discrete list, 10, 20, 25, 30, 40, 50, 60 or 80 MHz, and emits T0 or 10 to 100 MHz; its manual's preliminary appendix mentions electrical acceptance to 100 MHz, but 80 MHz is the top of every published selectable list, so 80 MHz is the number to design against. The Model 588 variant documents its clock input's sensitivity separately: 50 mV to 2.5 V peak, minimum pulse width 100 ps, insertion delay of the clock path at most 10 ns, a specification written for exactly one customer, the photodiode watching a modelocked laser.

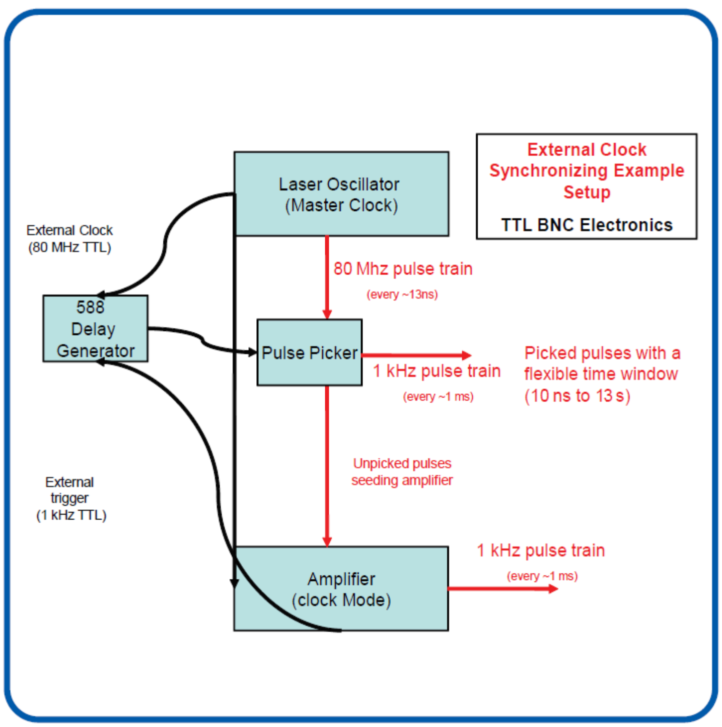

Figure 10. The canonical lock-up, from the datasheet set: a laser oscillator's own clock output drives the delay generator's clock input, and the generator's channels then drive the pulse picker and amplifier of the same laser chain. Everything downstream counts on the laser's oscillator, so the picker window and the amplifier gate cannot walk against the pulse train they select from. This one diagram is why the external clock input exists. SCHEMATIC

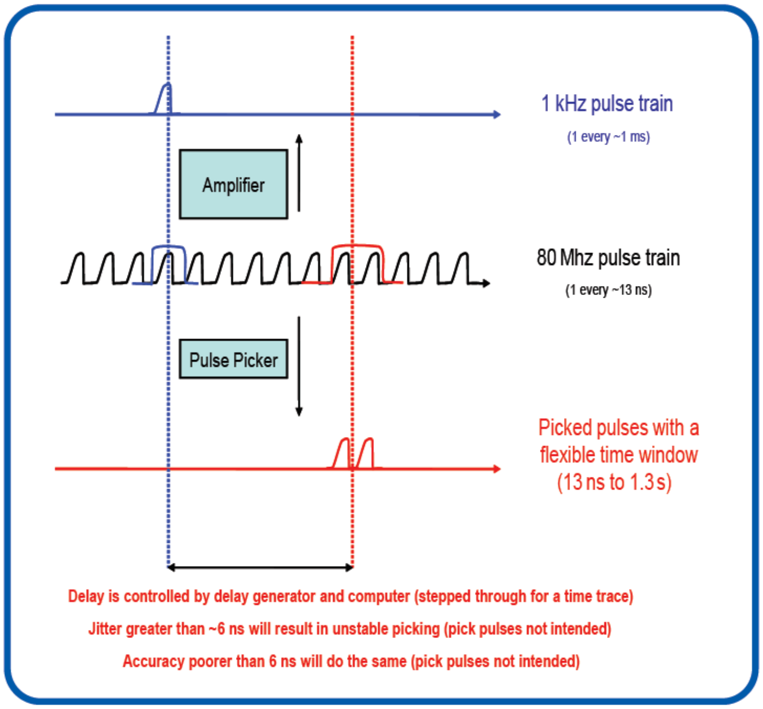

Figure 11. Pulse picking, timed: from an 80 MHz oscillator train, the picker window lifts one pulse, or a short run of them, down to a 1 kHz experiment rate, and the drawing carries its own engineering bound: picking becomes unstable when timing error approaches 6 ns, because at an 80 MHz train the pulses are 12.5 ns apart and a window wandering by half a period no longer selects the pulse it meant to. The selectable window is 13 ns to 1.3 s. SCHEMATIC

Figure 10 and Figure 11 together are the coherence argument of section 6 made concrete: lock the generator to the oscillator and the 6 ns bound of Figure 11 is met with room to spare, because a clock-locked window counts the oscillator's own periods and tracks the train by construction. The residual should be channel-class, tens of picoseconds. The Model 588B manual's external-clock appendix publishes 300 ps cycle-to-cycle and 1 ns period jitter for that path, both preliminary. No equivalent measurement exists for the Model 577, so its margin belongs on the wiring-day checklist of section 21 before an experiment depends on it. Trigger the picker from an unlocked box instead and the 800 ps trigger figure plus drift eats the margin.

13 Remote operation

Both instruments speak the same command family over their serial paths: SCPI-style ASCII, with every front-panel function reachable as a command. The Model 577 carries USB and RS-232 as standard with Ethernet and GPIB as a factory communications option; the Model 588B carries USB, RS-232 and Ethernet as standard, and since it has no front panel display, the remote path is not an option but the instrument's face. A supplied Windows control application exposes every 588B channel, the system modes and a command terminal, and requires no installation. Configuration storage is on board: 16 named setups on the 577, one per channel count on the 588B, 12 or 24, with a factory-default configuration zero that cannot be overwritten. For scripted work the practical numbers are the defaults: 115200 baud on RS-232, 38400 on the 588B's USB port, and a stated 10 ms command pacing at full rate on the 588B.

WHICH PATH TO USE FOR WHAT

Use the front panel, on the 577, for bring-up and troubleshooting: a delay scan with a knob finds an optimum faster than an edit-compile loop. Use the remote interface for anything repeated: a pump-probe scan of two hundred delay points is a for-loop, and section 17 budgets one. Use stored configurations for the handoff between them, so the state a human tuned is the state the script recalls.

14 Where these instruments are not the answer

A reference document that lists only capability is a brochure. The boundaries below are stated in the same spirit as the specifications, with numbers and sources, and Part III routes around each one where a route exists.

- **Sub-picosecond work is out of reach, and not narrowly.** Setting resolution is 250 ps and channel-to-channel jitter is 50 ps RMS. Streak-camera sweeps, femtosecond pump-probe fine axes and anything else that counts single picoseconds needs an optical delay line or the camera's own sweep electronics, with the delay generator supplying the surrounding framework. Section 20 is the worked example, including the budget that proves the point: degrading a 2 ps camera by no more than ten percent leaves a 0.9 ps trigger budget, which no instrument in this class meets.
- **The external trigger path costs 800 ps RMS and about 110 to 160 ns of latency.** Any interval built across the trigger input inherits both. The cure is architectural and free, run the bench from the internal rate generator, but an experiment that must slave to an outside, aperiodic event keeps the cost.
- **Retriggering has a floor.** Triggers arriving before every channel completes are ignored: 200 ns plus the longest active pulse, 5 MHz at absolute best. An experiment needing every event of a megahertz-class random source timestamped wants a time-to-digital converter, not a delay generator.
- **High-voltage and optical modules are Model 577 territory.** No high-voltage or optical output option is published for the 588B; its published outputs are the TTL/adjustable module and the TZ50 high-current option. A 24-channel system needing 45 V drive on some channels is a factory conversation, not an order form line.
- **Design the 588B against 50 ps.** Both the datasheet and the manual give 50 ps RMS for channel-to-channel jitter, and that is the budget every interval in this handbook assumes.

PART III

The benches, wired

Application recipes at wiring level: worked numbers for every bench, and a channel map wherever the wiring is the lesson. One design rule runs through them, the generator as the master clock of the bench, and the last chapter is honest about where the rule reaches its edge. The laser and detector physics quoted throughout, upper-state lifetimes, flashlamp delays, build-up jitter, window ranges and detector spreads, are standard literature and practice values rather than Berkeley Nucleonics specifications.

15 The Q-switched laser, and the delay that is an energy knob

Most of Part III is variations on one skeleton, so it is built once, here, on the commonest pulsed laser in the laboratory. A flashlamp-pumped, Q-switched Nd:YAG stores pump energy in its upper laser level, whose fluorescence lifetime is about 230 µs. While the cavity is held closed, the lamp pumps population in and spontaneous decay drains it out, so the stored inversion rises, peaks and falls. Opening the Q-switch at the peak dumps the maximum energy into one nanosecond-class pulse; opening off-peak yields less. The flashlamp-to-Q-switch delay is therefore not merely a setting, it is the laser's pulse-energy control, and it touches no optics.

The arithmetic that makes the knob quantitative: waiting an extra 50 µs past the optimum loses roughly $1 - e^{-50/230}$, about 20 percent, of the stored inversion to spontaneous decay; an extra 23 µs costs about 10 percent. A delay swept from 100 to 250 µs in 1 µs steps maps the whole energy curve in 150 shots, and against that sweep the instrument's contribution is invisible: at a 180 µs setting the accuracy bound of equation (6) is 19 ns, fifty times finer than the step, and the shot-to-shot scatter is 50 ps, more than four orders below the step.

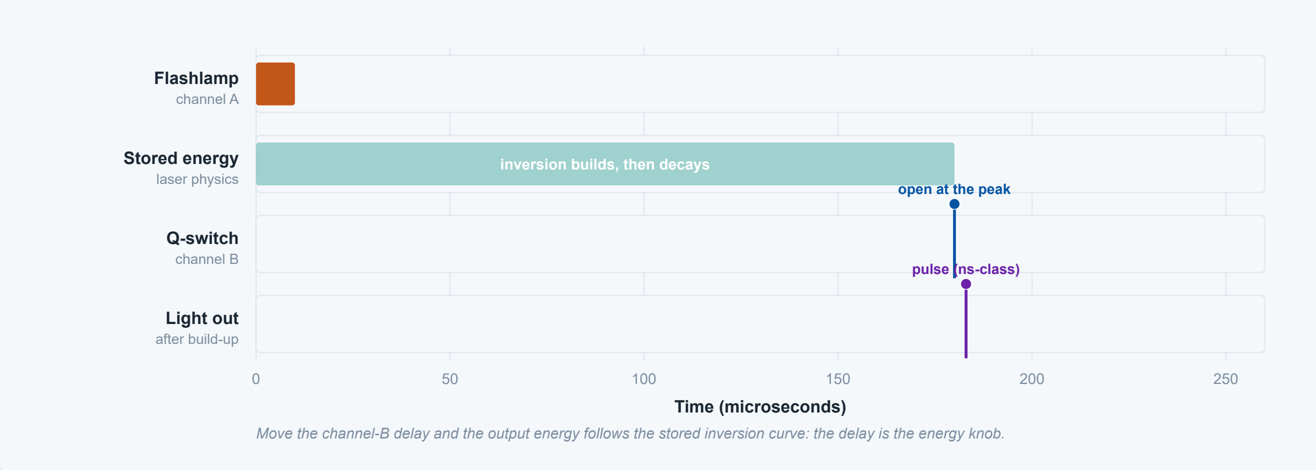

Figure 12. The skeleton of half of experimental laser physics. The lamp fires at time zero, and the shaded band marks the interval the inversion is stored, from pump to dump; the true stored-energy curve rises against its 230 microsecond decay and peaks near 180 microseconds, which is where the Q-switch channel opens the cavity here. The optical pulse emerges tens of nanoseconds later, after build-up from noise, an offset drawn far wider here than its true scale so it can be seen at all on a microsecond axis. That build-up interval jitters by tens of nanoseconds shot to shot, which is why anything needing the true emission time watches a photodiode rather than trusting the Q-switch edge. SCHEMATIC

Table 6 is the channel map, the format every following recipe reuses. The photodiode row is the honesty row: the electrical schedule is exact to picoseconds, but the laser adds its own tens-of-nanoseconds build-up jitter on top, so downstream picosecond work re-references to measured light, per Figure 12.

Table 6. The Model 577 as the master of a Q-switched bench, internal rate generator at 10 Hz. Every interval is channel-to-channel, so every interval enjoys the 50 ps figure; the trigger input is not used at all, which is the design rule of Part III applied with nothing left over for it to fix.

Ch	Connects to	Delay	Width	Why
A	Flashlamp driver trigger	0	10 µs	Starts the pump; the lamp supply shapes the current, the edge only starts it
B	Q-switch (Pockels cell) driver	180 µs, tuned 100 to 250 µs	100 ns	Opens the cavity at peak inversion; this delay is the energy knob
C	Camera or experiment gate	180.1 µs	as needed	Coarse sync; carries the laser's build-up jitter, see the photodiode row
D	Digitizer trigger	179.9 µs	1 µs	Pre-trigger so the acquisition window brackets the light
PD	Photodiode on a pick-off, to any ps-critical input	n/a	n/a	Re-references downstream timing to actual photon arrival

16 Particle image velocimetry: two pulses, one number

PIV measures a flow by photographing seeded particles twice, a known interval Δt apart, and correlating the two frames: velocity is displacement over Δt. The relative velocity error therefore contains the relative timing error directly,

$$\delta v / v = \delta(\Delta t) / \Delta t \tag{8}$$

where

Δt is the pulse separation, sized so particles move a few pixels between frames

$\delta(\Delta t)$ is the shot-to-shot error of that separation, which is a jitter figure from Part II

Sizing first: at 20 µm per pixel and a 10 m/s flow, a 5 pixel target displacement is 100 µm, so Δt = 10 µs. Now equation (8) prices the architecture. Both Q-switch channels on one instrument's internal clock: 50 ps on 10 µs is 5×10^{-6}, three orders below the correlation error of about one percent, invisible. The same two edges built across the external trigger path instead: 800 ps gives 8×10^{-5}, sixteen times worse, still survivable at this Δt. But fast flows shrink Δt: at 100 ns the external-trigger figure is 0.8 percent of the measurement, as large as the correlation error itself, while the channel-to-channel figure is still only 0.05 percent. The design rule of Part III is worth a factor of sixteen here, and the factor lands exactly where the measurement is hardest.

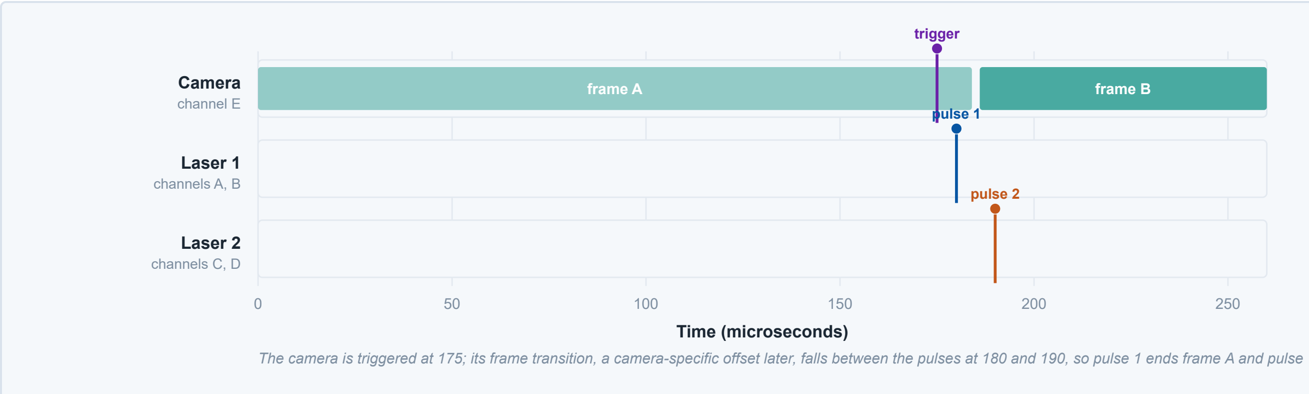

Figure 13. Frame straddling, the trick that frees the pulse separation from the camera. The frame boundary falls between the two laser pulses, so a camera whose frame period is milliseconds still delivers image pairs 10 microseconds apart, and the pulse separation can shrink until it meets the camera's one hard limit, the interframe transfer time, typically a hundred nanoseconds and more, set by the sensor. The whole scheme is five delay numbers on one instrument. SCHEMATIC

Figure 13 places the pulses; Table 7 shows why a dual-cavity PIV system uses four laser channels rather than two. Each head keeps its own flashlamp-to-Q-switch delay, its energy knob from section 15, while Δt is set by sliding head 2's pair together. Energy balance and pulse separation stay independent, which is what makes the system tunable in the field.

Table 7. Dual-cavity PIV on a Model 577. Sweeping Δt means changing two numbers, C and D, by the same amount; balancing energies means trimming A-to-B or C-to-D individually.

Ch	Connects to	Delay	Width	Why
A	Laser 1 flashlamp	0	10 μs	Pump head 1
B	Laser 1 Q-switch	180 μs	100 ns	Pulse 1; the A-to-B interval is head 1's energy knob
C	Laser 2 flashlamp	Δt = 10 μs	10 μs	Pump head 2, shifted whole
D	Laser 2 Q-switch	180 μs + Δt	100 ns	Pulse 2, exactly Δt after pulse 1; C-to-D stays head 2's energy knob
E	Camera frame trigger	175 μs	per camera	Places the frame boundary between the two pulses

The field evidence behind Table 7 is the Berkeley Nucleonics PIV technical note, the published record of this wiring in working installations: the four-signal dual-laser architecture is its recommended hookup, the direct dependence of velocity on Δt is its central caution, and the channel map above compresses its guidance into one table. The practical consequence it documents is operational: a facility that changes flow regime retunes two delay numbers and touches no optics, which is why this wiring survives operator handoffs that would break an optically trimmed bench.

17 Pump-probe: the delay axis is the data axis

Time-resolved spectroscopy launches a process with a pump pulse and samples it with a probe a delay t later; stepping t films the kinetics. The delay generator's specifications become the experiment's sampling parameters directly: resolution sets the finest time step, jitter smears each sample, range sets the slowest process reachable. The working rule is ten samples per decay constant and total timing noise well under the fastest feature.

- **Microsecond kinetics**, phosphorescence or plasma evolution: steps of 100 ns or coarser, 250 ps resolution gives four hundredfold margin, and 50 ps of jitter is five parts in 10^5 of even a 1 μs decay. Fully electronic scanning, no moving parts, nine decades of range.
- **Nanosecond kinetics**, fast fluorescence with a 10 ns lifetime: 1 ns steps, still comfortable. Jitter now matters: 50 ps of generator scatter against 30 ps of laser trigger response combines as $\sqrt{(50^2 + 30^2)} = 58$ ps, a 0.6 percent smear on the lifetime. Acceptable, and now every picosecond of the budget is accounted for.
- **Picosecond and faster**: electronics is out, per section 14. Light travels 0.3 mm per picosecond, and a

retroreflector doubles the path, so 0.15 mm of stage travel is 1 ps: the fine axis becomes optical, and the generator supplies the coarse framework of laser, chopper and camera around it.

The scan itself is a remote-interface loop: channel B's delay stepped from 0 to 2 μs in 10 ns increments is 200 points; at 10 Hz with 100 shots averaged per point, one kinetic trace costs 200 × 10 s, about 33 minutes, unattended. A chopper on a spare channel at half the repetition rate gives alternating pump-on, pump-off shots for background subtraction with no extra hardware.

18 Gated cameras: ICCD timing and the LIBS window

An intensified CCD is an electronic shutter: its photocathode conducts only during a gate pulse, with opening times in nanoseconds. Two pieces of arithmetic govern every gated-camera experiment, and both are duty-cycle arguments.

Background rejection first. Continuous background, room light, flame luminosity, plasma glow, integrates only while the gate is open. A 20 ns fluorescence signal captured with a 50 ns gate at 10 Hz accumulates 500 ns of open time per second of experiment; against an ungated one-second exposure, the continuous background is suppressed by 1 s / 500 ns, a factor of two million, while the pulsed signal inside the gate is untouched. Second, flash rejection: elastic scatter of the excitation pulse arrives during the pulse, fluorescence persists after it, so opening the gate just after the pulse ends trades a known fraction of signal, $1 - e^{-5/20}$, about 22 percent for a 5 ns wait on a 20 ns lifetime, for the removal of essentially all prompt scatter.

THE INSERTION-DELAY TRAP, SPRUNG

Trigger the generator from a photodiode watching the laser and no gate can open sooner than the insertion delay, about 110 ns on the Model 577, after the light. For a gate that must open within nanoseconds of the pulse, the sequence must run the other way: the generator, on its internal clock, fires the laser as one channel and the gate as another, and the photodiode merely verifies. Section 5 predicted this; here is where it bites.

Laser-induced breakdown spectroscopy is the canonical delayed-gate application and the forgiving one. The ablation pulse ignites a plasma whose first few hundred nanoseconds emit structureless continuum; the elemental lines outlive it. The measurement lives in a window, gate delay of roughly 0.5 to 2 μs after the pulse, width of 1 to 20 μs, standard starting ranges that each matrix and laser energy shift somewhat. Stepping the gate delay in 50 ns increments while watching the line-to-background ratio finds the optimum per element in a few dozen shots; the microsecond scale of the window makes the laser's tens-of-nanoseconds build-up jitter irrelevant, which is why this recipe tolerates any triggering scheme and makes a good first bench for a new user.

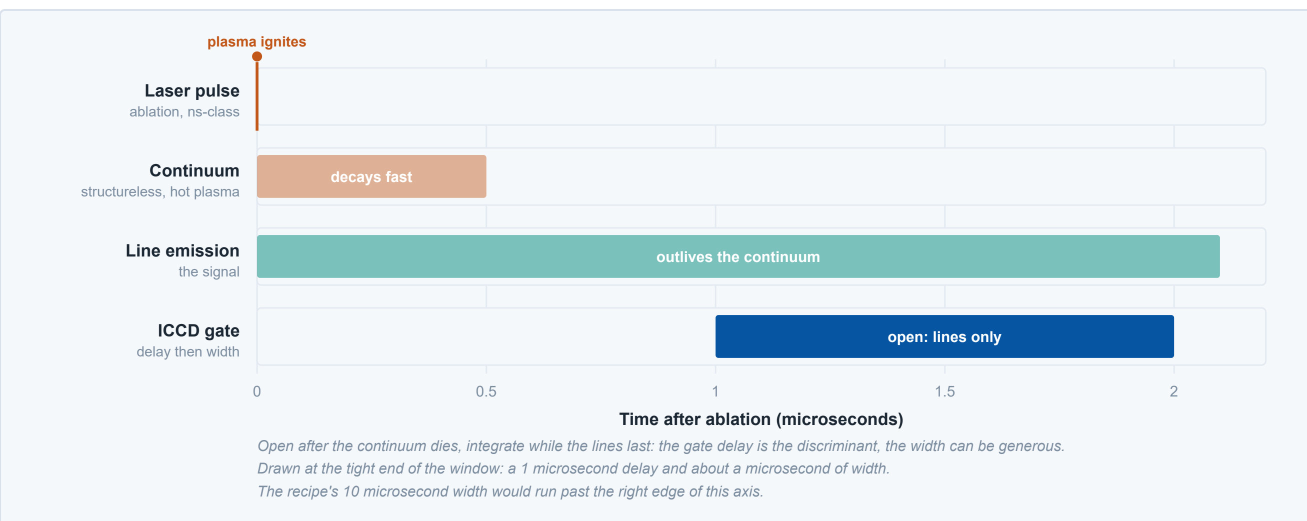

Figure 14. The LIBS timing argument in one drawing. The continuum and the lines both start at ablation; the continuum dies first; the gate opens in the gap. Delay is the critical number and width is forgiving, which is the exact inverse of a fluorescence-rejection gate, where the width must be tight and the delay generous. Knowing which of the two numbers is doing the work is the transferable skill of this section. SCHEMATIC

Figure 14 generalizes to every gated detector on the bench: the question is always which emission to catch and which to outwait. The channel map is the skeleton of section 15 with the camera promoted: lamp on A, Q-switch on B at 180 μs, spectrometer gate on C at 180 μs plus 1 μs with a 10 μs width, readout on D after the gate closes, and an optional early 300 ns gate on E to record the continuum itself for background modeling.

19 Coincidence counting: the window arithmetic

Photon-pair optics, positron annihilation, gamma-gamma correlation and muon telescopes all reduce to the same question: two detectors click, are the clicks related? A coincidence window of half-width τ accepts pairs arriving within ±τ, and uncorrelated singles at rates R_1 and R_2 fake coincidences at

$$R_{acc} = 2\tau\, R_1\, R_2 \qquad (9)$$

where

τ is the coincidence half-width: either click may lead by up to τ, so the full window is 2τ

R_1, R_2 is the singles rate of each detector, in counts per second

Worked at bench rates: both detectors at 10^5 counts per second and τ = 10 ns give, by equation (9), $2 \times 10^{-8} \times 10^{10}$ = 200 accidental coincidences per second. Tighten τ to 1 ns and the accidentals fall to 20 per second: with a true coincidence rate of 1000 per second, the signal-to-accidentals ratio improves from 5:1 to 50:1 by changing one number, provided true pairs still land inside. That proviso is the floor: the window cannot shrink below the total timing spread of the two branches, detector responses and generator jitter in quadrature. Single-photon avalanche detectors near 100 ps each, with 50 ps of channel-to-channel generator scatter, give $\sqrt{(100^2 + 100^2 + 50^2)}$, about 150 ps, so windows down to a nanosecond are safe and the generator is not the limit. Fast photomultipliers at a nanosecond each dominate the budget themselves, and sub-nanosecond windows start clipping real pairs.

Gated detector bias uses the same channels twice over: biasing an avalanche detector above breakdown only in a window around expected arrival suppresses dark counts by the duty-cycle argument of section 18, a 5 ns gate at a microsecond period being a two-hundredfold suppression. And the standard accidentals measurement is itself a timing recipe: duplicate one gate on a spare channel, offset it far outside any true correlation, 100 ns is ample, and that channel counts $2\tau R_1 R_2$ directly, the background measured live for subtraction. This is the Model 588B's natural habitat: a detector array wants a gate per detector, path-length trims as per-channel delays, a common reference for the time-digitizer, calibration pulser lines and veto channels, and twenty-four channels on one timebase supply all of it with the 50 ps figure on every pairwise interval.

20 Streak cameras, honestly

A streak camera converts time to space, sweeping photoelectrons across a sensor with a picosecond-class ramp, and it is included here because it is where this instrument class reaches its edge, and because knowing the edge exactly is worth more than pretending it is further out. In any accumulating mode, trigger jitter shifts each sweep's start and smears the summed image; the effective resolution adds in quadrature,

$$\sigma_{eff} = \sqrt{(\sigma_{intrinsic}^2 + \sigma_{trigger}^2)} \qquad (10)$$

where

$\sigma_{intrinsic}$ is the camera's own resolution, single picoseconds for a good tube

$\sigma_{trigger}$ is the RMS jitter of the sweep trigger relative to the light

Worked with a 2 ps camera and 10 ps of trigger jitter: equation (10) gives √104, 10.2 ps. The trigger owns the measurement and a five times better tube would change nothing. Run the budget backwards to see how hard the requirement is: to degrade that 2 ps camera by no more than ten percent, the trigger may contribute $\sqrt{(2.2^2 - 2^2)}$, 0.9 ps, beyond any Digital Delay Generator, this vendor's or anyone's. The honest architecture follows: the sweep trigger comes optically, a photodiode or photoconductive switch watching the actual beam feeding the camera's own sweep electronics, and the delay generator owns everything slower, lamp and Q-switch, the arm window that readies the sweep unit just before light arrives, readout, and the housekeeping of the shot. Even in single-shot work, where jitter only positions the window rather than smearing a sum, placement matters: 800 ps of external-trigger scatter on a 10 ns sweep window spends eight percent of the window on placement uncertainty, while running the arm chain from the internal clock at 50 ps spends half a percent.

The cable is part of this budget, and of every budget in Part III: signals travel about 1.5 ns per foot in solid-dielectric coax, so a forgotten meter of cable is five nanoseconds, one hundred times the channel-to-channel jitter. Section 23 tabulates the velocity factors; the working rule, one foot is a nanosecond and a half, is worth committing to memory before wiring day.

PART IV

Practice and reference

The bench habits, the failure modes, and the numbers, arranged to be found again rather than read once.

21 Measurement practice

None of the following depends on which delay generator is on the bench, and all of it decides whether the timing is worth anything.

1. **Make the generator the master.** Run the bench from the internal rate generator and drive everything as channels. Every recipe in Part III applies this rule; every interval it covers inherits the 50 ps figure instead of the 800 ps one.
2. **Terminate, or know why not.** Every 50 Ω output into an unterminated line reflects; section 23 carries the arithmetic. A through terminator at the scope is the cheapest instrument on the bench.
3. **State the load with every amplitude.** Write "2.0 V into 50 Ω" in the lab book, never "2 V". Half the mysterious amplitude changes in a pulsed lab are a termination that came or went.
4. **Fix the cables and label their lengths.** A meter of coax is five nanoseconds. A timing map that records cable lengths in nanoseconds alongside programmed delays can be rebuilt after any teardown; one that does not cannot.
5. **Trigger on the fastest clean edge available**, per section 3, and set the threshold at the steepest part of that edge, mid-transition, not near its foot where slew is low and noise is amplified.
6. **Let the timebase warm up before precision work.** Crystal frequency moves with temperature as the chassis comes to equilibrium; delays in the seconds range move with it, per the proportional term of equation (6).
7. **Record the whole timing state, not the two numbers you changed.** Both instruments store complete configurations on board and both dump state over the remote interface; a saved configuration in the run log is the difference between reproducing a run and guessing at one.
8. **Verify one interval independently after wiring day.** Two channels into one oscilloscope with matched cables confirms the schedule as built, catches the swapped cable and the forgotten attenuator, and takes five minutes.

22 Troubleshooting

Table 8 is arranged by symptom, because a symptom is what a reader has when they open this page. The third column is the one that matters: most symptoms here have two or three plausible causes, and the discriminating check is what keeps an afternoon from being spent on the wrong one.

Table 8. Thirteen symptoms, each with the check that separates its causes. The first four rows resolve most first-day calls, and every one of them is arithmetic from Parts I and II rather than folklore.

Symptom	Probable cause	The check that discriminates	Fix
No output on any channel	System not armed, or waiting on an external trigger that never comes	Is the run indicator active? Switch trigger source to internal: does it fire?	Arm the system; restore the intended trigger source and verify its level
No output on one channel	Channel disabled, gated off, or its output routed away by the multiplexer	Check the channel enable, the gate mode, and the MUX mask for that output	Re-enable; clear the gate; restore the default one-timer-per-output routing
Output at half the expected amplitude	50 Ω termination on an output specified into high impedance, or vice versa	Compare against the three-load table in section 10	Match the load to the specification row, or fit the high-current option
Levels fine on a scope, downstream device never fires	2.0 V terminated TTL against a threshold at or above 2.0 V, commonly 5 V CMOS at 3.15 V or more	Look up the receiver's actual input threshold in section 23	Drive it with the adjustable output at a healthy margin, or the high-current TTL option

Symptom	Probable cause	The check that discriminates	Fix
Doubled edges, staircase steps, or ringing on a long line	Unterminated 50 Ω line: the reflection returns after the edge	Does 2 × cable delay exceed the rise time? Section 23	Terminate at the far end; keep stubs short
Timing wanders by nanoseconds shot to shot	Triggering across the external input from a slow or noisy edge	Compare jitter with the bench run from the internal rate generator	Make the generator the master; if it must slave, speed up the trigger edge
Fixed offset of tens to hundreds of nanoseconds appears	Insertion delay in one path and not another, or an unaccounted cable	Sum the published insertion delays and 5 ns/m of cable along each path	Rebalance with programmed delays; document the offsets in the timing map
Delay of seconds lands visibly off	The proportional accuracy term: 10^{-4} × delay is 100 μs at 1 s	Does the error scale with the programmed delay?	Lock to a house reference via the external clock input, per section 12
Instrument ignores triggers at high rate	The retrigger rule: triggers during an active pulse train are discarded	Compute 200 ns + longest active pulse; compare the trigger period	Shorten widths and delays, or accept the ceiling and decimate upstream
Pulses vanish at high rate on the rack unit	The 75 ns budget: delay + width + reset exceeds the T0 period	Check delay + width + 75 ns against the period, channel by channel	Lower the rate or shorten the offending channel's schedule
Two instruments drift apart over minutes	Free-running timebases: two crystals disagree by their combined ppm	Does the drift accumulate linearly with elapsed time?	Clock both from one reference; coherence, not accuracy, is what is needed
Remote commands garble or drop	Baud mismatch, or commands sent faster than the instrument's pacing	Echo a query at the default rate; then add pacing, 10 ms on the rack unit	Set the documented defaults; pace the script; prefer query-echo handshakes
Gated channel fires once more after the gate asserts	Pulse inhibit lets a pulse in progress finish; output inhibit truncates	Scope the gate against the output: does the pulse complete or cut?	Choose the inhibit flavor the experiment needs, per section 11

23 Bench reference tables

The numbers a timing bench reaches for, collected. Values here are standard reference figures from manufacturer datasheets and physical constants; the instrument-specific numbers live in section 24.

Coaxial cable delay

A signal in coax travels at the velocity factor times c, so delay per length is 3.336 ns/m, or 1.017 ns/ft, divided by the velocity factor. Table 9 evaluates the families a timing bench actually stocks.

Table 9. Propagation delay by cable family, computed from published velocity factors. One caution the label will not give you: RG-58 names a connector fit more than a dielectric, and foam-dielectric cables carrying the RG-58 designation run near 0.76 rather than 0.66. On a ten-foot run that is a 2 ns disagreement, forty times the channel-to-channel jitter, so measure or look up the actual cable, not the family name.

Cable family	Dielectric	Velocity factor	Delay, ns/m	Delay, ns/ft
RG-58, RG-213, RG-174 (solid PE)	solid polyethylene	0.66	5.05	1.54
RG-142 (PTFE)	PTFE	0.69	4.83	1.47
RG-8X, foam types	foam polyethylene	0.78 to 0.85	3.92 to 4.28	1.20 to 1.30

Rule of thumb: one foot of solid-dielectric coax is a nanosecond and a half; one meter is five nanoseconds.

Logic thresholds

Table 10 is why the half-amplitude trap of section 10 matters: 2.0 V clears some thresholds and not others.

Table 10. Input and output thresholds from the family datasheets. The row that causes bench calls is 5 V CMOS: its 3.15 V input-high requirement is not met by a 2.4 V TTL-grade high, and certainly not by 2.0 V of terminated TTL, which is the failure the fourth row of the troubleshooting table diagnoses.

Family	V_{IL} max	V_{IH} min	V_{OL} max	V_{OH} min
5 V TTL and LS	0.8 V	2.0 V	0.4 to 0.5 V	2.4 V
LVTTL, 3.3 V	0.8 V	2.0 V	0.4 V	2.4 V
5 V CMOS (HC)	1.35 V at V_{CC} 4.5 V	3.15 V at V_{CC} 4.5 V	0.1 V lightly loaded	4.4 V lightly loaded
4000-series at 5 V	1.5 V	3.5 V	0.05 V typical	4.95 V typical
3.3 V CMOS (LVC)	0.8 V	2.0 V	0.4 V	V_{CC} − 0.4 V

Emitter-coupled families switch near −1.3 V; their exact output levels vary by family and are not tabulated here.

Reflections, and when termination is mandatory

$$\Gamma = (Z_L - Z_0) / (Z_L + Z_0) \quad (11)$$

where

Γ is the fraction of the incident wave reflected at the load: +1 at an open end, −1 at a short, 0 into a matched 50 Ω termination

Z_0 is the line impedance, 50 Ω throughout this handbook

An open end reflects at +1, so the arriving edge doubles there: a back-terminated source launches half its open-circuit swing, and the reflection restores the full swing at the open end, which is why a bare high-impedance scope input reads double the terminated value, and why the returning wave dies quietly in the matched source. The rule for when this matters is a race: terminate whenever the rise time is shorter than the round trip, twice the cable's one-way delay. Two meters of RG-58 is 10.1 ns one way, a 20.2 ns round trip, and a 2 ns edge on that line without a termination produces a doubled plateau at the far end and a 20 ns staircase at the near end, per equation (11). A 15 cm jumper with a 10 ns edge, round trip 1.5 ns, blends its reflection invisibly into the edge, which is why short unterminated stubs are tolerated for logic and never for calibrated amplitude. Connector bandwidth follows the same conservatism: the common bayonet connector is a 4 GHz part and threaded subminiature types reach 18 GHz, typical figures per manufacturer catalogs, all far above any edge this instrument class produces.

24 Quick reference

Table 11 collects the governing relations of Part I; Table 12 collects the published values of Part II. Together they are the pages to photocopy, and nothing in either is quotable without its condition column.

Table 11. Every governing relation in this handbook with a worked anchor beside it, so the arithmetic can be checked rather than trusted.

Quantity	Relation	Worked anchor
Peak-to-peak from RMS	$t_{pk\text{-}pk} = k \times t_{RMS}$	k = 6: 50 ps RMS is 300 ps pk-pk at 99.7 percent
Noise-to-jitter conversion	$t_j = v_n / (dV/dt)$	1 mV on 0.5 V/ns: 2 ps; on 0.02 V/ns: 50 ps
Synchronizer quantization	$t_{RMS} = T_{clk}/\sqrt{12}$	2.89 ns at 100 MHz; 1.44 ns at 200 MHz
Delay construction	delay = $N \times T_{clk}$ + vernier	range from the counter, resolution from the vernier
Delay accuracy	1 ns + 10^{-4} × delay	19 ns at 180 μs; 100 μs at 1 s
Loaded amplitude	$V = V_{open} Z_L/(Z_S+Z_L)$	half into 50 Ω from a 50 Ω source
Reflection	$\Gamma = (Z_L - Z_0)/(Z_L + Z_0)$	+1 open, −1 short, 0 matched
Termination rule	terminate when t_{rise} < 2 × cable delay	2 m of RG-58: round trip 20 ns
Cable delay	1.017 ns/ft divided by velocity factor	RG-58: 1.54 ns/ft, 5.05 ns/m
Accidental coincidences	$R_{acc} = 2\tau R_1 R_2$	200 per second at 10 ns and 10^5 singles
Velocity error, PIV	$\delta v/v = \delta(\Delta t)/\Delta t$	50 ps on 10 μs: 5×10^{-6}
Resolution under jitter	$\sigma_{eff} = \sqrt{\sigma_i^2 + \sigma_t^2}$	2 ps camera + 10 ps trigger: 10.2 ps

Table 12. Every published value this handbook argues from, with its condition. Sources: the Model 577 datasheet and manual, the Model 588 manual for shared-family items, and the Model 588B datasheet and manual.

Quantity	Model 577	Model 588B	Condition
Channels	4 or 8	12 or 24	per ordering configuration
Delay, width resolution	250 ps	250 ps	
Delay range	0 to 1000 s	0 to 2000 s	width 10 ns to same maximum
Delay accuracy	1 ns + 0.0001×delay	1 ns + 0.0001×delay	
Channel-to-channel jitter	< 50 ps RMS	50 ps RMS	output module tables
Trigger jitter	< 800 ps RMS	800 ps RMS	external trigger input
Insertion delay	< 110 ns	< 160 ns	external trigger to output
Internal rate	0.001 Hz to 20 MHz	0.0002 Hz to 20 MHz	internal generator
Max trigger rate	5 MHz	5 MHz	and 1/(200 ns + longest active pulse)
Timebase	100 MHz PLL, 50 ppm osc	200 MHz PLL, 25 ppm osc	
TTL output	4.0 V into 1 kΩ; 2.0 V into 50 Ω	4.0 V into 1 kΩ	50 Ω output impedance; rise < 3 ns typical
Adjustable output	2 to 20 V into 1 kΩ; 0.8 to 8 V into 50 Ω	2 to 20 V into 1 kΩ; 1 to 10 V into 50 Ω	10 mV steps
High-voltage option	AT45: 4 to 45 V, 20 mV steps, 4 channels max	none published	not short-circuit protected
External clock in	10 to 100 MHz, discrete	10 to 80 MHz, discrete	clock out to 100 MHz on both
Trigger threshold	0.2 to 15 V, 10 mV steps	0.2 to 15 V	rising or falling
Burst counters	10,000,000 per channel	10,000,000 per channel; 4×10^9 system	duty cycle the same
Interfaces	USB, RS-232; Ethernet, GPIB optional	USB, RS-232, Ethernet	SCPI-style command set
Storage	16 configurations	12 or 24 configurations	config 0 is the factory default

25 Definitions, symbols and further reading

Channel-to-channel jitter
The RMS scatter of the interval between two outputs of one instrument on the same shot. The synchronizer error is common to both and cancels, which is why this is the small number.

Crest factor
The multiplier k converting RMS jitter to a peak-to-peak figure at a stated exceedance probability. Peak-to-peak without k is not a specification.

Delay generator
An instrument producing multiple electrical edges at programmed delays and widths from a common trigger, all counted on one timebase.

Duty cycle mode
A channel mode firing a programmed on-count of pulses then skipping an off-count. With on = 1 and off = N − 1 it is divide-by-N.

Gate
An input or internal signal that enables or inhibits outputs. Pulse inhibit lets the pulse in progress finish; output inhibit cuts it immediately.

Insertion delay
The time from an external trigger edge to the earliest possible output at zero programmed delay: comparator, synchronizer and driver latency summed. Internally triggered operation has none.

Multiplexer (MUX)
The OR fabric routing channel timers to output drivers, letting one output carry several timers' pulses. Timing is multiplexed; amplitude is the module's.

Rate generator (T0)
The internal periodic source that starts every channel's counters each cycle when the instrument is its own master.

Synchronizer

The flip-flop stage resampling an asynchronous trigger into the clock domain, at a cost of up to one clock period of quantization error.

Termination
A resistor matching the line impedance at the receiving end, absorbing the wave and preventing reflection. Mandatory when the rise time is shorter than the round trip.

Timebase
The crystal oscillator and PLL supplying the counting clock. Its fractional error, in ppm, multiplies every programmed delay.

Trigger-to-output jitter
The RMS scatter of an output edge against the external trigger edge. Carries the full synchronizer residual, which is why it is the large number.

Vernier
The analog fine-delay stage adding sub-clock-period resolution to the counted delay. Its linearity sets the fine-scale accuracy floor.

Velocity factor
The ratio of signal speed in a cable to c, set by the dielectric: 0.66 for solid polyethylene, up to 0.85 for foam.

Symbols used in this paper.

Symbol	Meaning	Units
T_{clk}	counting clock period	s
N	whole clock periods counted in a delay	1
t_{RMS}	RMS jitter	s
k	crest factor for peak-to-peak conversion	1
v_n	RMS voltage noise at a threshold	V
dV/dt	slew rate through a threshold	V/s
$L(f)$	single-sideband phase noise, linear ratio	1/Hz
Δt	PIV pulse separation	s
τ	coincidence half-width	s
Γ	reflection coefficient	1
Z_0, Z_L, Z_S	line, load and source impedances	Ω

Symbol	Meaning	Units
σ_{eff}	effective resolution under trigger jitter	s

CMOS complementary metal-oxide-semiconductor logic

GPIB general purpose interface bus

ICCD intensified charge-coupled device camera

LIBS laser-induced breakdown spectroscopy

MUX multiplexer, the timer-to-output routing fabric

PIV particle image velocimetry

PLL phase-locked loop

ppm parts per million, fractional frequency error

RMS root mean square

SCPI standard commands for programmable instruments

T0 the internal rate generator's time-zero event

TTL transistor-transistor logic levels

Further reading

- The Model 577 datasheet and user manual (Document Version 1.5), and the Model 588B datasheet (V1.1) and user manual (Document Version 1.0), Berkeley Nucleonics: the instruments of record for every specification quoted here.
- The Berkeley Nucleonics PIV technical note: the four-signal dual-laser architecture of section 16, with field wiring detail.
- IEEE Std 1139, *Definitions of Physical Quantities for Fundamental Frequency and Time Metrology*: the formal footing under section 3's phase-noise and jitter definitions.
- JCGM 100 (GUM), *Evaluation of Measurement Data: Guide to the Expression of Uncertainty in Measurement*: the framework behind every quadrature budget in Part III.
- ANSI Z136.1, *Safe Use of Lasers*: binding on every bench in Part III that fires one.
- Manufacturer cable and logic-family datasheets, for the velocity factors and thresholds of section 23, which vary by part and lot.

One route onward

A timing budget that will not close, a channel count that will not fit, or a value that has to be right before an order exists: a Berkeley Nucleonics applications engineer will talk any of them through before an order exists. Bring the timing map of section 21 and the channel count you actually need. The conversation costs nothing, and it usually shortens wiring day.

Berkeley Nucleonics Corp.
Test, Measurement and Nuclear Instrumentation since 1963
For details: info@berkeleynucleonics.com or 800-234-7858

BNC-HB-003 Rev B
September 2026