
ANGLE-DEPENDENT CROSSTALK BETWEEN TWO ADJACENT VL53L1X TIME-OF-FLIGHT SENSORS

Jessica Wang
Trinity School NYC
New York, NY 10024, USA

ABSTRACT

Time-of-flight ranging sensors mounted side by side can disturb one another, but we found no published measurement of how this depends on the angle between them. We measured the extra ambient light one VL53L1X sensor sees when an identical neighbor 20 millimeters away also ranges, for one pair of sensors tested against a gray board at 0.3, 0.6 and 0.9 meters in two window sizes, with five repeats per setting, a divider test, and the whole sweep repeated after exchanging the two sensors. Three findings follow. First, the extra light fell to 10 percent of its peak at about 16 to 18 degrees of total angle between the sensors (half that per sensor) for a full sensor array and 8 to 11 degrees for a narrower window, and to 1 percent at about 20 to 22 and 11 to 13 degrees, growing with distance; a geometric model using the datasheet field of view overestimated these angles by three to seven degrees, and no single emitter beam angle fitted both windows. The extra light also arrived in bursts about every 1.9 seconds, consistent with unsynchronized sensor clocks, and after the sensors were exchanged the angles reproduced to within about 0.4 degrees, with the small difference in crossing angle between the two sensors following the unit, not its mounting position. Second, ranging itself was barely disturbed: a large rise in ambient rate at close range shifted the mean distance by half a millimeter or less. Third, near a sensor's acceptance threshold the neighbor widened the signal scatter and could change the share of accepted readings by up to about 20 percentage points, in either direction. A thin divider changed the effect by no more than about 4 percent in twelve of fourteen settings before the swap and fifteen of eighteen after it, consistent with, but not proof of, a shared-target path rather than direct coupling. Separating the sensors' pointing angles past this crossover is expected to remove the effect; a barrier was not shown to help, and ranging in turn, which avoids the interference by construction, was not tested.

Keywords Crosstalk · interference · sensor arrays · time-of-flight sensors · VL53L1X

1 Introduction

Ranging by the travel time of light is now a common way to measure distance without touching the target. Automotive and robotic systems use pulsed time-of-flight sensors for obstacle detection and navigation¹, and single-photon-avalanche-diode (SPAD) arrays of the kind inside the sensor used here are a standard building block of these systems². Two further points set the scene for this paper. First, not every such sensor scans or images a scene: a scanning lidar sweeps a beam and is comparatively expensive, and a multizone sensor such as the VL53L5CX reports a small grid of distances from a fixed mount³; but arrays of cheap, single-point sensors like the VL53L1X remain common because each one is inexpensive and can be pointed in any direction independently of the others. Second, when several such sensors are mounted close together, as in a drone's collision-avoidance ring or a mobile robot's bumper, each can potentially receive light meant for its neighbor (Fig. 1).

The VL53L1X is a compact direct time-of-flight (ToF) sensor of exactly this kind: it pairs a 940 nm laser emitter with an array of single-photon avalanche diodes (SPADs) and a programmable region of interest (ROI)⁴, and is the sensor used throughout this paper. A manufacturer community thread discusses whether neighboring units disturb one another⁵. Mutual interference of this kind has already been studied for other sensor families: a model of the distance error exists for two ToF cameras⁶ and methods for operating several of them were developed^{7,8}, and for scanning lidars, where it was

measured^{9,10} and direct and scattered interference were distinguished¹¹. Crosstalk between two multizone ToF sensors viewing through one light guide has been characterized for a tactile and proximity skin¹², and crosstalk compensation in multizone ToF sensors has been reported¹³. For small single-point sensors such as the VL53L1X, by contrast, we found only the forum thread above: a manufacturer employee replied that units generally do not interfere, while relaying a report of interference at 5 cm when two face each other⁵. The original poster, who was designing a multi-sensor drone, asked whether an angle (45° or 90°) exists at which the interference disappears; the reply gave no angle data⁵. In the sources we searched (web searches of the research literature, manufacturer documents and community posts, repeated on Oct. 15, 2026 with the terms “VL53L1X”, “time-of-flight sensor crosstalk”, “multi-sensor interference” and “angle” against Google Scholar, Scopus, IEEE Xplore and arXiv; not an exhaustive bibliographic-database search) we found no published measurement of how the interaction between two adjacent VL53L1X-class sensors depends on the angle between them. This is the gap this paper addresses.

We chose the VL53L1X for three reasons: it is inexpensive and widely used, so results for it are relevant to many existing designs; its field of view can be narrowed in software from the full 16×16 SPAD array to a 4×4 window, so the same pair of units lets us test two optical geometries without new hardware; and it reports the signal and ambient photon rates directly (Section 1.1), so the interference can be measured as a rate rather than inferred from range alone.

This paper measures that dependence for one pair of VL53L1X sensors 20 mm apart, in two ROI sizes and at three distances. We ask four questions. (i) At what angle does the extra ambient rate fall to 10% and 1% of its peak, and how precisely can that angle be pinned down? (ii) Does the response fall monotonically with angle in both ROI sizes? (iii) Does a geometric model built from the datasheet’s field-of-view figures predict the angle at which the effect disappears? (iv) Does a thin opaque divider between the sensors reduce the effect, which would point to light reaching the neighbor directly rather than only via the target? Because one campaign cannot separate the properties of the two units, of their mounting positions and of one measurement session, we repeated the sweeps after exchanging the two units (Section 2).

1.1 How the sensor works

The VL53L1X emits laser pulses and times the returning photons with an array of SPADs⁴. Photons arriving within a narrow window after each pulse, reflected from the target, cluster at one delay and form a sharp peak in the histogram of arrival times (Fig. 2a); photons arriving at other times, from room light or any other unsynchronized source, are spread roughly uniformly across the histogram and form a flat floor beneath the peak. The driver reports three quantities from this histogram, each summed over the active SPADs and given in kilocounts per second (kcps): the *range*, from the position of the peak; the *signal rate*, the height of the peak above the floor; and the *ambient rate*, the height of the floor itself. A reading is accepted or rejected according to status codes defined by the manufacturer’s driver¹⁴; in the data reported here, every accepted reading has a signal rate above 1024 kcps (the lowest accepted value across both campaigns is 1032 kcps and 1024 kcps itself is always rejected; rates are quantized in 8 kcps steps). The ultra-lite driver’s own manual documents a default signal-rate threshold of exactly 1024 kcps^{14, Sec. 7.37}, consistent with what we observe (Section 3.3).

1.2 What crosstalk means here

When a second sensor is also ranging nearby, some of its emitted light can reach the first sensor’s receiver, either directly or after reflecting off the target. Because the two sensors run on independent, unsynchronized clocks, the neighbor’s photons arrive at random delays relative to the first sensor’s own pulse, so they add to the flat floor rather than to the peak (Fig. 2a): the ambient rate rises, while the range, read from the peak, is affected only indirectly. This is what Section 2 measures as the *extra ambient rate*.

A simple geometric picture bounds the angle at which this stops: sensor B’s emitter illuminates a cone of half-angle $E/2$ and sensor A’s receiver accepts light from a cone of half-angle $F/2$ (Fig. 2b); as the two sensors are turned apart by a relative angle $\Delta\theta$, these cones’ footprints on the target slide apart and eventually stop overlapping. Equating the point where the two footprint edges meet gives

$$s + 2d \tan(\Delta\theta_g/2) = d [\tan(E/2) + \tan(F/2)] \quad (1)$$

where s is the sensor spacing, d the target distance, and $\Delta\theta_g$ the predicted end-of-overlap angle; this treats both footprints as circular cones, centered on each sensor’s nominal axis. The in-plane variant used in Section 2.5 instead converts the receive field to a value appropriate for a square array, so the two assumptions are not mutually consistent and (1) should be read as an order-of-magnitude guide rather than a precise model. Section 3.4 evaluates (1) against the measured crossover angles.




Figure 1: Concept: sensors in an array point in different directions, but adjacent ones' emitter and receive cones can still overlap on a shared target.




Figure 2: Principle. (a) Sketch of the photon-arrival-time histogram one sensor builds: its own return forms a sharp peak (the signal rate); unsynchronized light, including a neighbor's, spreads across the window and raises the floor (the ambient rate). (b) The overlap geometry of (1): sensor B's emitter cone (half-angle $E/2$) and sensor A's receive cone (half-angle $F/2$), rotated apart by a relative angle $\Delta\theta_g$, stop sharing a common footprint on the target beyond the predicted angle.

2 Methodology

To answer the four questions above, one pair of VL53L1X sensors was turned apart by a known angle and read while alone and together, at three distances and two field-of-view sizes, and the extra ambient rate this produces was measured and related back to the overlap geometry of Section 1.2. Because one campaign cannot separate the properties of the two units, of their mounting positions and of one measurement session, the whole sweep was repeated after exchanging the two units (campaign 2; Section 2.1). The rest of this section gives the apparatus, the acquisition protocol and acceptance rules, what was computed and why, how the crossover angle and its uncertainty were estimated, and the divider check.

2.1 Apparatus and geometry

Two VL53L1X sensors⁴ on Pololu 3415 carrier boards (campaign 1: unit 3 at position A, unit 1 at B) were read by an ESP32-S3 through a TCA9548A I²C multiplexer¹⁵ at 400 kHz. The STMicroelectronics ultra-lite driver¹⁴ set short distance mode, a 33 ms timing budget and a 40 ms period. Each interrupt-timestamped reading gave the range, its status, and the total signal and ambient rates (Section 1.1). Two ROIs were used, the full 16×16 array and a 4×4 window, each centered on the factory optical center. The datasheet gives their fields of view as 27° and 15° , as diagonal values, and gives no emitter beam angle⁴. Table 1 lists every setting; Fig. 3 is a schematic of the geometry it describes, and Fig. 4 shows the physical rig.

Table 1: Hardware, settings and experimental parameters

Parameter	Value	Parameter	Value
Sensors	Two VL53L1X ⁴ on Pololu 3415 carriers	Target distances d	0.3, 0.6, 0.9 m (tape; see Section 2)
Host	ESP32-S3, via TCA9548A I ² C mux ¹⁵	Target	Flat dark-gray board, 1.2 m wide
I ² C clock	400 kHz	Room lighting	Not recorded or varied
Driver mode	Short distance, ultra-lite driver ¹⁴	Repeats per cell	5 blocks (A alone, B alone, A+B)
Timing budget	33 ms; period 40 ms	Run length	30 s
ROI (zone)	16 × 16 (full array) and 4 × 4 (window)	Warm-up	600 s, both sensors, each distance
Field of view (datasheet, diagonal)	27° and 15°	Sensor spacing s	20 mm (nominal)
Angle step	1° (pin-set plate; accuracy not verified)		

What differs between the two campaigns:

Parameter	Campaign 1	Campaign 2 (unit swap)
Unit positions / ROI centers	unit 3 at A, unit 1 at B	unit 1 at A, unit 3 at B (ROI centers exchanged)
θ sweep ($\Delta\theta = 2\theta$)	$\theta = 0^\circ$ and $2\text{--}12^\circ$ in 1° steps (12 levels, $\Delta\theta = 0, 4, 6, \dots, 24^\circ$; at 0.6 m 13 levels, $\Delta\theta = 0, 2, \dots, 24^\circ$)	$0\text{--}12^\circ$ step 1° (13 levels, $\Delta\theta = 0, 2, \dots, 24^\circ$)*
Divider angles (θ)	$2^\circ, 5^\circ$ ($\Delta\theta = 4, 10^\circ$); also 1° ($\Delta\theta = 2^\circ$) at 0.6 m	$1^\circ, 2^\circ, 5^\circ$ ($\Delta\theta = 2, 4, 10^\circ$)
Session dates (UTC)	Oct. 6–17, 2026	Oct. 12–14, 2026

*At 0.6 m the sweep stopped early at $\theta = 11^\circ$ ($\Delta\theta = 22^\circ$) once an adaptive range-quality check passed, so that session has no $\Delta\theta = 24^\circ$ level.

The sensors, nominally $s = 20$ mm apart, sat on pucks turned by $-\theta$ (A) and $+\theta$ (B), so their axes diverged by $\Delta\theta = 2\theta$ (Fig. 3). Pins in a plate with holes at 1° steps set the angles (accuracy not verified). Campaign 1 covered $\theta = 0^\circ$ and $2\text{--}12^\circ$, so $\Delta\theta = 0, 4, 6, \dots, 24^\circ$ (12 levels). Campaign 2 exchanged the two units (unit 1 at A, unit 3 at B, with the ROI centers exchanged accordingly) and repeated the sweeps on different days at every 1° step of θ ($\Delta\theta = 0, 2, \dots, 24^\circ$; 13 levels). In campaign 2 each sweep was allowed to stop early once the worst-case range SD ratio fell below 1.10; this happened only at 0.6 m, after $\theta = 11^\circ$, so that sweep has no $\Delta\theta = 24^\circ$ level, while the 0.3 and 0.9 m sweeps ran to $\Delta\theta = 24^\circ$; the early stop affects only which angles were sampled at 0.6 m, not the measured values at the angles it did reach. A swap of this kind separates the properties of a unit from those of its mounting position, but because each unit still looks at the other’s emitter, it cannot separate a unit’s receiver behavior from its emitter behavior. The target was a flat dark-gray board 1.2 m wide at tape distances $d = 0.3, 0.6$ and 0.9 m (at $\Delta\theta = 0^\circ$ with the full array the sensors read 7–23 mm less than the tape distance). Before each setting a guided check required, three times in a row, the mean 16×16 range within 15 mm of a stored value and B–A within 10 mm of a per-distance offset (2.5, 2.5 and 3.0 mm, from a squared board), and a self-test to pass. The measured B–A range offsets at $\Delta\theta = 0^\circ$ were +3.0, −4.4 and +1.2 mm in campaign 1 and +3.5, +4.2 and +4.9 mm in campaign 2 at 0.3, 0.6 and 0.9 m, all within the 10 mm tolerance; the 0.9 m check in campaign 2 kept the campaign-1 offset of 3.0 mm. The sensors’ own mean ranges at $\Delta\theta = 0^\circ$ differed between the campaigns by −8, −2.5 and −1 mm at 0.3, 0.6 and 0.9 m (mean of the two units; individual units −11 to +2 mm); a change in distance of this size shifts a crossing by only about 0.03° . Room temperature, taken from the room’s digital temperature control (74 °F, about 23 °C), was not logged with the data.

2.2 Acquisition and acceptance

Each distance was one session on its own day (campaign 1: Oct. 6–17, 2026; campaign 2: Oct. 12–14, 2026), angles ascending (resumed once after an interruption at 0.3 m in campaign 1), after a 600 s two-sensor warm-up. Each cell




Figure 3: Measurement geometry (schematic, not to scale): sensors A and B are angled apart by $\Delta\theta = 2\theta$; the common footprint is where B's emitter cone and A's receive field overlap. The divider was present only in the divider tests. Cone angles are generic.




Figure 4: Physical setup, close-up: the two sensor carriers on their angled pucks, the pin-set angle-setting plate (a pin at each hole sets θ), and the wiring to the TCA9548A multiplexer above the frame. The target board stands opposite at the tape distance of Fig. 3.

(zone $\times \Delta\theta$) had five blocks of three 30 s runs, A alone, B alone and A with B, in seeded random order with the idle sensor stopped (748–767 readings per active sensor). A run was accepted if it had no abort or I²C/read error, the idle sensor stayed silent, the reading count was within 5% of expected and at most 1% of results were missed. A block was accepted if its runs were and each sensor's solo mean ambient rate was within 10% (at least 0.5 kcps) of that of the previous accepted block at the same distance, angle and zone; failures were repeated up to three times. Blocks failing three times (1 of 120 at 0.3 m, none of 130 at 0.6 m and 7 of 120 at 0.9 m in campaign 1; one 0.9 m sweep block and four 0.9 m divider blocks in campaign 2) were re-measured in separate refill sessions and replaced; excluding them moves no crossover by more than 0.01° .

2.3 What was computed, and why

For each sensor and block, the extra ambient rate is the difference of run-mean ambient rates $\bar{a}$ with both sensors ranging and with that sensor alone, pooled over the two sensors (Eqs. 2–3 below). A difference is used, not a ratio to the solo rate, because for the 4×4 window the solo rate is only about 0.1–10.9 kcps across both campaigns (including the divider sessions), comparable to its own run-to-run noise, and so a ratio would be dominated by that noise rather than by the crosstalk. Range effects are the companion measurements of whether ranging itself is disturbed: the difference of mean ranges (bias) and the ratio of range standard deviations between two-sensor and solo runs, from valid readings, pooled as in (3) (geometric mean for the ratio). Cells in which fewer than 20% of readings were valid are excluded from the range statistics, since their standard deviations then rest on a handful of readings (Section 3.2). The rejected-reading fraction is tracked separately, since it depends on the signal threshold of Section 1.1 rather than on the ambient floor.

The block, not the individual reading, is the unit of statistical analysis throughout, since readings within a 30-second run are not independent.

For sensor $s \in \{A, B\}$ in block b , the extra ambient rate is the difference of run-mean ambient rates $\bar{a}$ with both sensors ranging and with s alone, pooled over the sensors:

$$\Delta a_s(b) = \bar{a}_s^{AB}(b) - \bar{a}_s^{solo}(b) \quad (2)$$

$$\Delta a(b) = \frac{1}{2} [\Delta a_A(b) + \Delta a_B(b)] \quad (3)$$

Five blocks give a cell mean $\Delta \bar{a}$ and standard deviation.

2.4 Crossover and uncertainty

Let $\Delta \theta_{pk}$ be the level with the largest cell mean $\Delta \bar{a}_{pk}$. The crossover at threshold φ is the first angle beyond $\Delta \theta_{pk}$ at which the response falls to $\varphi \Delta \bar{a}_{pk}$ ($\varphi = 0.10$ unless stated). With levels 2° apart, the data alone fix only a *bracket* (the last level above and the first at or below the threshold); a point estimate needs an interpolant or a fit. The primary estimator is shape-preserving piecewise cubic Hermite interpolation (PCHIP)¹⁶ of the cell means (M4). Four further estimators are compared: linear interpolation (M3), the only other interpolant; and three logistic-family fits (M1, M2, M5), one of which can rise before it falls,

$$g(\Delta \theta) = \frac{a(1 + b\Delta \theta)}{1 + \exp[(\Delta \theta - x_0)/w]} \quad (4)$$

which is a plain logistic when $b = 0$. M3 and the tilted logistic M5 of (4) agree with M4 within 0.22° in every cell (M5 within 0.13°); a plain logistic fitted to all levels and read at $\Delta \theta = 0$ (M1) can lie up to 1° from M4 and falls entirely outside the model-free bracket at 0.9 m (full array) in campaign 1, so M1 and M2 are reported but not used to bound the estimate. The five are, in full:

- M1: three-parameter logistic fitted to all levels; reference is the fitted value at $\Delta \theta = 0$;
- M2: the same logistic fitted only from the observed peak level outward; reference is the fitted value there (identical to M1 whenever the peak is at $\Delta \theta = 0$);
- M3: piecewise-linear interpolation of the block means; reference is the observed peak mean; first crossing after the peak;
- M4: PCHIP interpolation of the block means, same reference and crossing rule;
- M5: the tilted logistic of (4); reference is its own maximum; first crossing after that maximum.

Sampling uncertainty is a block bootstrap¹⁷: the five blocks of every level are resampled with replacement $B = 2000$ times, with common resamples for all estimators, and percentile 95% intervals are reported. This describes within-session repeatability only. The unit-swap campaign gives an empirical between-session spread: the pooled 10% crossing moved by $0.07\text{--}0.42^\circ$ (mean 0.23°) between campaigns, and the 1% angle by at most 0.3° (Table 3), so the crossings are better quoted as good to about $\pm 0.4^\circ$ than the $0.01\text{--}0.07^\circ$ bootstrap half-width alone would suggest; at 0.3 and 0.9 m this comparison also mixes in campaign 2's denser level grid, so it is not a clean replicate there. The analytical uncertainty is the range across estimators; we also vary $\varphi \in \{0.05, 0.10, 0.20\}$ and analyze A and B separately. For one cell and one threshold the procedure is: (1) average the five blocks at each level and locate the peak level p ; (2) find the first level after p whose mean is at or below φ times the peak mean, which gives the bracket; (3) take the PCHIP root (M4) and the linear root (M3) in the bracket, and the crossing of each fitted curve (M1, M2, M5), where the references of M1, M2 and M5 follow their definitions above rather than the observed peak; (4) repeat steps 1–3 on 2000 resamples and report the bracket, the M4 estimate with its 2.5 and 97.5 percentiles, and the range across the five estimators. A cell in which no level falls to the threshold has no crossing within the sampled range; this did not occur in any cell reported here. The analysis was implemented in Python (scripts/crossover_analysis.py; the swap campaign is selected with CROSSOVER_CAMPAIGN=2).

2.5 Geometric reference

Equation (1) (Section 1.2) needs the emitter cone E and receive field F as numbers. We take $E = 27^\circ$ (the datasheet's stated field of view, with no separate emitter figure given) and F as the datasheet's diagonal field-of-view value for each zone. We also evaluate (1) with F converted to an in-plane value for a square array, $\tan \beta = \tan \alpha / \sqrt{2}$ (where $\alpha = F/2$ is the datasheet's diagonal half-angle and β the in-plane half-angle it converts to) — an assumption that, as noted in Section 1.2, is not consistent with treating the footprints themselves as circular — and separately invert the equation for the E that reproduces the measured 1% angle. This is a sensitivity analysis, not a claim that either convention is correct: the orientation of the array axes relative to the plane of rotation was not recorded. Because (1) is purely geometric, its predictions do not depend on campaign.

2.6 Divider check

The purpose of this check is to ask whether the light reaches the neighboring sensor directly, or only by reflecting off the target: if a barrier between the sensors reduced the effect, that would point to a direct path. At each distance a thin black strip stood on edge between the sensor windows; the protocol fixed what was checked — that the strip stayed out of both fields of view and within about 15 mm of the sensor chips — but not its length, height, thickness or exact placement within that envelope, none of which was measured or recorded at the time. The cells at $\theta = 2^\circ$ and 5° ($\Delta\theta = 4$ and 10° ; also $\theta = 1^\circ$, $\Delta\theta = 2^\circ$, at 0.6 m in campaign 1 and at every distance in campaign 2) were re-measured in both zones, five blocks each, and compared with the no-divider cells through the difference of cell means of (3).

With the apparatus, the acceptance rules and the quantities above in place, Section 3 applies them to the four questions of Section 1: what each sensor’s own ambient and signal rates looked like without a neighbor, what changed with one present, and how that change depended on angle, zone, distance and, now, campaign.

3 Results and Discussion

3.1 One measurement at the worst case

Before the sweeps, one setting shows what the numbers below mean: the full array, 0.3 m, $\Delta\theta = 0^\circ$. Fig. 5 plots sensor A’s ambient rate and range across one block, first running alone, then with B. The ambient rate jumps from a mean of 43 kcps to 864 kcps, about 20 times higher, the moment B starts; the mean range moves by 0.4 mm, from 287.9 mm to 287.5 mm, indistinguishable from no effect at this scale. The same block after the unit swap (Fig. 5b) gives 52 kcps alone and 1187 kcps with B (23-fold) and a range shift of 0.5 mm, from 280.0 mm to 279.5 mm.

The increase is not steady: with B ranging, the ambient rate alternates between roughly its solo level and a burst level several times higher, in cycles that repeat every 1.88 s on average (1.84–1.92 s across every AB block examined at $\Delta\theta = 0$, both zones, all three distances, and again after the unit swap), so the 864 kcps mean is an average over bursts rather than a steady rate. We located the period as the lag of the strongest peak, outside the trivial zero-lag peak, in the autocorrelation of the ambient-rate series resampled onto a uniform 25 Hz grid; the peak autocorrelation at that lag is about 0.90–0.93 in every full-array block at $\Delta\theta = 0^\circ$, so the cycle is not a noisy artifact of one run. A 1D (Otsu) threshold splits each run’s readings into a low and a high state; the high state occupies close to half of all readings (we find 49–56% across the runs examined), though because the rise between states spans several consecutive readings rather than being instantaneous, the exact split depends on where the threshold falls within that transition, so we do not treat this figure as more precise than “roughly half.” Unsynchronized sensor clocks would produce a beat of about this period: a 2% mismatch between the two sensors’ nominal 40 ms periods gives a beat of $T^2/\Delta T \approx 2$ s; we did not test this directly, for instance by comparing the two sensors’ interrupt timestamps.

The rest of this section asks why ranging survives such a large jump in ambient light, when the effect matters despite that, how far apart the sensors must point for it to disappear, and where the light is coming from.

3.2 Why ranging survives

“Barely affected” means: at $\Delta\theta = 0^\circ$ the range bias was at most 0.3 mm in magnitude at every distance and zone, and the range standard deviation was 7–10% larger with both sensors ranging for the full array and 13% larger for the window at 0.3 m; by about 16° the full-array ratio is already 1.00–1.03 at every distance, indistinguishable from no effect (it stays that way out to 24° , so 16° is the earliest point at which the effect has clearly disappeared, not a later one). Two sensor-blocks (both sensor B, 4×4 window, 0.3 m, $\Delta\theta = 22^\circ$) have a range SD ratio near 0.37, about a third of solo scatter; this is not a crosstalk effect, since there is essentially no extra ambient left at that angle: fewer than 1% of the window’s readings were valid at that angle (about 4–5 of some 750), so both standard deviations in the ratio rest on a handful of readings and are not meaningful; cells below the 20%-valid cutoff of Section 2 are excluded from Fig. 6 for this reason. No other points depart from the main trend by a comparable amount.

The reason ranging survives is that the extra ambient a neighbor adds is small next to the sensor’s own signal (Section 1.1): the histogram’s peak, which fixes the range, is barely moved by raising its floor. If this is the mechanism, the range standard deviation ratio (A+B over solo) should track the size of the extra ambient relative to that signal, regardless of distance or zone. Fig. 6 plots this for every accepted (block, sensor) in the sweep with at least 20% valid readings, before the swap (a, b) and after it (c, d): the SD ratio rises from about 1.0 toward 1.1–1.2 as extra ambient / own signal grows from 0 to about 7%, with broadly the same trend at all three distances and both zones, and follows $1 + kx$ with $k = 3.5$ before the swap and $k = 3.2$ after it (Fig. 6a,c; x is a fraction here, while the axes plot it in percent). The correlation over all plotted points is $r = 0.73$ before the swap and 0.69 after it (0.68 and 0.63 for the full-array points alone); because both quantities fall together as the sensors turn apart, it is partly a joint dependence




Figure 5: One block, full array, 0.3 m, $\Delta\theta = 0^\circ$: sensor A’s ambient rate (top) and range (bottom) against elapsed time, first alone then with B (dotted line marks the switch). (a) Before the swap (campaign 1, sensor A is unit 3): ambient rate rises about 20-fold, in bursts that repeat about every 1.9 s; range is not visibly affected. (b) After the swap (campaign 2, sensor A is unit 1): the same picture, a 23-fold rise in bursts of the same period, and a range shift of 0.5 mm.

on angle rather than a clean test of the mechanism on its own, so it supports, without on its own confirming, the SD ratio tracking extra ambient relative to signal. The range itself does not move: the mean range shift (neighbor on minus alone) is -0.01 mm before the swap and -0.01 mm after it, with a standard deviation of 0.25 mm in both, and no trend with extra ambient (Fig. 6b,d). At 0.9 m the mean shift is $+0.03$ mm before and $+0.02$ mm after the swap; the largest full-array shift is 0.76 mm (before the swap, 0.3 m) and 0.72 mm (after the swap, 0.6 m), and the largest 4×4 shift is 0.83 mm (after the swap, 0.6 m).

At $\Delta\theta = 0^\circ$, where the two beams overlap most, extra ambient / own signal is 0.056, 0.061 and 0.062 for the full array at 0.3, 0.6 and 0.9 m, and 0.073, 0.062 and 0.044 for the window. The solo signal and the extra ambient both fall with distance, but only by a few-fold: from 0.3 to 0.9 m the full-array solo signal falls 3.6-fold and the extra ambient 3.3-fold, so their ratio stays within about 10% across that range; for the window the solo signal falls 8.2-fold and the extra ambient 14-fold, so the ratio changes by a factor of about 1.7. The full-array result — that the ratio barely moves even though the two quantities individually fall several-fold — is consistent with both the neighbor’s light and the sensor’s own light reaching the sensor by the same route, reflected off the same target, so that whatever reduces one with distance reduces the other by nearly the same factor; Section 3.5 returns to this point. The window’s larger swing in the ratio weakens, without ruling out, the same reading there.

3.3 When it does matter: near the detection limit

The SD-ratio trend above is small because the sensor’s own signal is normally much stronger than the extra background; it follows that the effect should matter more exactly when the signal is weak, for a far or dark target, the narrower 4×4 window, or a brighter room. The 4×4 window across the three distances already spans this range: pooled over the two sensors, its solo signal at $\Delta\theta = 0$ is far above the 1024 kcps acceptance threshold (Section 1.1) at 0.3 m (about 3800 kcps), close to it at 0.6 m (about 1190 kcps), and below it at 0.9 m (about 470 kcps). The acceptance threshold, though, applies per sensor and per reading, and pooling the two sensors hides a real split at 0.6 m: unit 3 sits higher, at 1310–1370 kcps, and is (almost) always accepted at $\Delta\theta = 0$, while unit 1 sits lower, at 1010–1070 kcps, close to the threshold, and is accepted in 81% of its readings at B (campaign 1) but only 18% at A (campaign 2) (Table 5), so unit 1’s validity depends on its mounting position as well as on the unit. Only in the middle, mixed regime can the small extra ambient from a neighbor plausibly tip a reading from one side of the threshold to the other, since far above it nearly everything is accepted regardless and far below it nearly everything is rejected regardless.

Fig. 7 demonstrates this where it is clearest: the 4×4 window at 0.6 m for unit 3, which sat at position A before the swap ($\Delta\theta = 4^\circ$, top row) and at B after it ($\Delta\theta = 8^\circ$, bottom row). Before the swap, alone, this sensor’s per-reading signal averages 1179 kcps with a standard deviation of 48 kcps, close above the threshold, and 0.1% of its readings are




Figure 6: Range SD ratio (A+B over solo) (a, c) and range shift (b, d) against extra ambient / own signal, every accepted (block, sensor) with at least 20% valid readings, both zones and the three distances; 4×4 -window points are open triangles (none at 0.9 m, where the window returns no valid range). (a, b) Before the swap (campaign 1); (c, d) after the swap (campaign 2). The curve is $1 + kx$ fitted to the points.

rejected; with B also running the average barely moves (1177 kcps) but the standard deviation widens to 81 kcps and 4.9% are rejected (Fig. 7a). After the swap, at a similar mean signal (1106 kcps), the standard deviation widens from 46 to 63 kcps and the rejected share rises from 5.1% to 12.1% (Fig. 7d). The acceptance rule itself is the same with or without the neighbor — a reading is accepted only if its own signal exceeds the threshold, a step at 1024 kcps (Fig. 7b,e) — so the extra rejections appear only near the threshold, at the angles where the mean signal is close to it (Fig. 7c,f); what changes is how much of the reading-by-reading spread of signal values falls on each side of that step. Because widening the spread moves the accepted share toward 50% from either side, the direction of the change depends on which side of 50% the sensor started on: across the fifteen partly accepted 4×4 cells of both campaigns, the neighbor lowered the accepted share in ten and raised it in five; the raised cells all start from a solo share of 18% or less, and all but one of the lowered cells start from 65% or more, by 0.6 to 19.9 points (the exception changes by only -0.3 points); Table 2 lists all fifteen, with a further four near-degenerate cells (changes under 0.2 points) omitted. This was seen in these settings and summarized over fifteen cells, not as a general measurement: it demonstrates the mechanism by which a detection threshold turns a small shift in ambient-light spread into a change in how many readings are returned at all, not a general figure for how often it happens.

3.4 How far apart the sensors need to point

Having shown where the effect does and does not matter, we return to question (i) of Section 1: how it falls with angle, and where it stops. The extra ambient rate peaked at 1006, 697 and 302 kcps for the 16×16 zone (the full array) in campaign 1 (1069, 710 and 319 kcps in campaign 2) and at 273, 75 and 24 kcps for the 4×4 window at 0.3, 0.6 and 0.9 m (309, 70 and 24 kcps in campaign 2) — within -7 to $+13\%$ of campaign 1 (higher in campaign 2 in five of the six cells). For the full array that is 23, 8 and 3.5 times the solo rate; no ratio is given for the window, whose solo rate is near the noise floor. The response fell with increasing angular separation at every distance for the full array and at 0.3 m for the window (Fig. 8a,b). At 0.9 m the window instead peaked at $\Delta\theta = 4^\circ$, above its value at $\Delta\theta = 0$ by 23% (Welch




Figure 7: 4×4 window, 0.6 m, unit 3. (a, d) Per-reading signal histogram, neighbor off and on, with the acceptance threshold (dashed; shaded: rejected) and the share of readings rejected. (b, e) Readings accepted against own signal, binned: the rule is the same with and without the neighbor. (c, f) Readings rejected against the angle between the sensors; the mean signal (kcps) at the labelled angles is printed below the axis. Top row: before the swap, unit 3 at A, histogram at $\Delta\theta = 4^\circ$. Bottom row: after the swap, unit 3 at B, histogram at $\Delta\theta = 8^\circ$.

Table 2: The fifteen partly accepted 4×4 cells: solo vs. A+B accepted share

d (m)	$\Delta\theta$ (deg), unit (campaign)	Solo %	A+B %	Δ (pts)
0.3	20, unit 1 (1)	65.1	61.8	-3.3
0.6	0, unit 3 (1)	100.0	99.4	-0.6
0.6	0, unit 1 (1)	80.9	61.0	-19.9
0.6	0, unit 1 (2)	18.2	20.1	+1.9
0.6	2, unit 3 (1)	100.0	98.4	-1.6
0.6	2, unit 1 (1)	22.2	21.8	-0.3
0.6	2, unit 1 (2)	0.0	0.5	+0.5
0.6	4, unit 3 (1)*	99.9	95.1	-4.8
0.6	4, unit 1 (1)	0.0	1.2	+1.2
0.6	4, unit 3 (2)	100.0	99.2	-0.8
0.6	6, unit 3 (1)	97.8	87.9	-9.9
0.6	6, unit 3 (2)	100.0	96.8	-3.2
0.6	8, unit 3 (1)	4.7	6.6	+2.0
0.6	8, unit 3 (2)**	94.9	87.9	-7.1
0.6	10, unit 3 (2)	5.6	7.3	+1.6

*, **The settings of the histograms in Fig. 7a (before the swap) and 7d (after the swap). The neighbor lowers the accepted share in ten cells and raises it in five; the raised cells all start from a solo share of 18% or less, and all but one of the lowered cells start from 65% or more, by 0.6 to 19.9 points (the exception, unit 1 at $\Delta\theta = 2^\circ$ in campaign 1, starts at 22.2% and changes by only -0.3 points). A further four cells (solo shares of 99.9%, 0.8%, 100.0% and 100.0%) change by less than 0.2 points and are omitted as neither meaningfully partial nor informative.

$p < 0.0001$) in campaign 1; at 0.6 m the maximum was at $\Delta\theta = 2^\circ$, only 2.2% above $\Delta\theta = 0$ ($p = 0.09$), with the response 10% lower at $\Delta\theta = 4^\circ$ ($p < 10^{-4}$); at 0.3 m it was 25% lower at $\Delta\theta = 4^\circ$ — answering question (ii): the response is not always monotonic. Distance and session coincide within a campaign, so a single campaign cannot tell a distance effect from a one-off session artifact; but the same pattern reappeared in campaign 2, with the units exchanged and measured on different days, growing with distance in the same way (absent, within noise, at 0.3 m; weak at 0.6 m; clear at 0.9 m), so it is unlikely to be a session or unit artifact and more likely tied to distance itself. Campaign 2 also sampled $\theta = 1^\circ$ ($\Delta\theta = 2^\circ$) at every distance (campaign 1 did so only at 0.6 m): at 0.9 m the extra ambient there was already 12% above its $\Delta\theta = 0$ value (Welch $p < 10^{-5}$), much as at $\Delta\theta = 4$; at 0.6 m it was 3% above at $\Delta\theta = 2$ ($p = 0.06$) and essentially level at $\Delta\theta = 4$ in campaign 2; at 0.3 m it was already 4% lower at $\Delta\theta = 2$ ($p < 10^{-4}$), as at $\Delta\theta = 4$. So the maximum, where it exists, lies at or before $\Delta\theta = 4$ in both campaigns.




Figure 8: (a, b, d, e) Extra ambient rate against the angle between the sensors, as a percentage of each cell's observed peak, for the 16×16 and 4×4 zones at three distances. Markers are cell means of five blocks (error bars, one standard deviation, are mostly hidden); lines are PCHIP interpolants; filled triangles mark the 10% crossing and open triangles the geometric end-of-overlap angle. (c, f) The 10% crossing (bar: model-free bracket) and the angle at 1% of the peak against distance, with the model (solid) and with the receive field converted to an in-plane value (dashed). Top row: before the swap (campaign 1); bottom row: after the swap (campaign 2).

The model-free bracket of the 10% crossing (Table 3) was $16\text{--}18^\circ$ for the full array at all three distances and $8\text{--}10$, $8\text{--}10$ and $10\text{--}12^\circ$ for the window in campaign 1, and $16\text{--}18$, $16\text{--}18$ and $18\text{--}20^\circ$ (full array) and $8\text{--}10$, $10\text{--}12$ and $10\text{--}12^\circ$ (window) in campaign 2. PCHIP (M4) puts the crossing at 16.15 , 17.29 and 17.90° (full array) and 8.22 , 9.72 and 10.80° (window) in campaign 1, and 16.32 , 17.36 and 18.18° (full array) and 8.51 , 10.14 and 10.96° (window) in campaign 2 — $0.07\text{--}0.42^\circ$ higher (mean 0.23°), rising with distance in most estimators within each campaign (full precision in Table 3). We quote these to 0.1° : the five estimators already spread by up to 1° within a cell, and the swap moved the pooled crossing by up to 0.42° , so the practical precision is about $\pm 0.4^\circ$. Fig. 9 compares the block means, three of the estimators and their 10% crossings in every cell, before and after the swap; in the 4×4 zone at 0.6 and 0.9 m, where the response peaks at $\Delta\theta = 2\text{--}4^\circ$ rather than at 0° , the three curves differ most at small angles, near the peak, while their crossings stay close. The bootstrap half-width of the PCHIP value is $0.01\text{--}0.07^\circ$, but the five estimators span up to about 1° within a cell (0.13° at the least), several to tens of times as much. The linear interpolant (M3) and the tilted logistic (M5) agree with PCHIP (M4) within 0.22° and 0.13° respectively in every cell; a plain logistic fitted to all levels and read at $\Delta\theta = 0$ (M1) lies $0.9\text{--}1.0^\circ$ above PCHIP for the full array, outside the campaign-1




Figure 9: Extra ambient rate against $\Delta\theta$ in each cell, as a percentage of the largest cell mean. Points are block means $\pm$ SD ($n = 5$). Curves: M1 logistic, M5 tilted logistic and M4 PCHIP interpolant. Triangles mark where each of these three crosses 10% of the peak (color identifies the estimator). The dash-dotted vertical line is the geometric end-of-overlap angle of the model of (1), which is the same in both campaigns. Upper two rows: before the swap (campaign 1); lower two rows: after the swap (campaign 2).

Table 3: Angle between the sensors (degrees) at which the extra ambient rate falls to 10% (crossover) and 1% of its peak, and the model’s end-of-overlap angle, both campaigns

Zone	d (m)	Campaign	Bracket	Crossover (M4, PCHIP)	Range, M3–M5	1% angle	Model	Model, <i>F</i> in-plane
16 × 16	0.3	1	16–18	16.15	16.15–16.20	19.9	23.4	19.5
16 × 16	0.3	2	16–18	16.32	16.17–16.38	19.8	23.4	19.5
16 × 16	0.6	1	16–18	17.29	17.29–17.42	20.9	25.2	21.3
16 × 16	0.6	2	16–18	17.36	17.36–17.50	21.1	25.2	21.3
16 × 16	0.9	1	16–18	17.90	17.90–17.94	21.6	25.8	21.9
16 × 16	0.9	2	18–20	18.18	18.18–18.22	21.5	25.8	21.9
4 × 4	0.3	1	8–10	8.22	8.20–8.33	11.4	17.3	15.2
4 × 4	0.3	2	8–10	8.51	8.51–8.70	11.4	17.3	15.2
4 × 4	0.6	1	8–10	9.72	9.65–9.82	12.7	19.2	17.1
4 × 4	0.6	2	10–12	10.14	10.14–10.28	13.0	19.2	17.1
4 × 4	0.9	1	10–12	10.80	10.80–11.02	13.1	19.8	17.7
4 × 4	0.9	2	10–12	10.96	10.96–11.19	13.3	19.8	17.7

Campaign 1: unit 3 at A, unit 1 at B. Campaign 2: units exchanged (Table 1). Bracket is model-free; range is across M3 (linear), M4 (PCHIP) and M5 (tilted logistic), which together bound the 10% crossing reliably (M1, M2, the logistics read at or from $\Delta\theta = 0$, can lie up to 1° from M4). Model and model-with-in-plane-*F* are purely geometric and so do not depend on campaign. Bootstrap 95% half-widths of M4 are 0.01–0.07° within either campaign; the crossover moved by 0.07–0.42° (mean 0.23°) and the 1% angle by at most 0.3° between campaigns, so between-session reproducibility is closer to $\pm 0.4^\circ$ than the within-session bootstrap suggests.

bracket at 0.9 m (18.84° against 16–18°), and fits worse than the tilted logistic (Akaike difference above 140 in five of six cells). Doubling or halving the threshold moves the crossing by 0.7–1.7°, so the 10% value is a convention.

Analyzed separately, the crossing-angle gap between sensors follows the unit, not the position, in all six comparisons (three distances × two zones). In campaign 1 (unit 3 at A, unit 1 at B), sensor A crossed over about 1.1° earlier than B for the full array and 0.6–0.7° later for the window; after the swap (unit 1 at A, unit 3 at B), the full-array gap reversed sign (A 1.2–1.7° later than B) and so did the window gap (A 0.4–0.7° earlier than B) — in both campaigns, unit 3 crosses earlier in the full array and later in the window, and unit 1 the opposite, regardless of which position it occupies. A unit swap cannot, however, separate a unit’s receiver properties from its emitter properties, since each unit still looks at the other’s emitter.

Table 4 sets the two campaigns side by side, unit by unit. The crossing angle of each unit moves by at most 0.5° between campaigns (unit 3: -0.2 – 0.4° ; unit 1: 0.1 – 0.5° , later after the swap), less than or comparable to the gap between the units, which keeps its sign. The peak extra ambient moves more than the crossing does: in the full array unit 3’s peak changes by 2–21% after the swap (when it sits at B) while unit 1’s changes by -4 to $+1\%$; in the window the changes range from -7 to $+15\%$.

Table 4: Unit-swap results: 10% crossing angle (M4, degrees) and peak extra ambient rate (kcps) of each unit, in campaign 1 (unit 3 at A, unit 1 at B) and campaign 2 (units exchanged), each sensor analyzed separately

Zone	d (m)	Crossing unit 3 camp. 1 / 2	Crossing unit 1 camp. 1 / 2	Peak unit 3 (kcps) camp. 1 / 2	Peak unit 1 (kcps) camp. 1 / 2	Gap unit 3 — unit 1 (°) camp. 1 / 2
16 × 16	0.3	15.42 / 15.25	16.50 / 16.92	825 / 996	1186 / 1142	−1.08 / −1.68
16 × 16	0.6	16.66 / 16.68	17.80 / 17.93	674 / 690	719 / 730	−1.14 / −1.25
16 × 16	0.9	17.30 / 17.24	18.36 / 18.87	272 / 318	332 / 320	−1.06 / −1.63
4 × 4	0.3	8.51 / 8.67	7.78 / 8.28	312 / 360	234 / 258	+0.73 / +0.39
4 × 4	0.6	10.05 / 10.44	9.30 / 9.72	83 / 79	68 / 63	+0.74 / +0.72
4 × 4	0.9	11.03 / 11.19	10.43 / 10.66	29 / 29	19 / 20	+0.60 / +0.53

Campaign 1: unit 3 at A, unit 1 at B. Campaign 2: units exchanged. Crossing angle is the 10% crossing (M4, degrees) of each sensor analyzed separately; the peak is the largest cell mean of the sensor’s own extra ambient rate. Unit 3 crossed earlier than unit 1 in the full array and later in the window in both campaigns, whichever position (A or B) it occupied: the gap keeps its sign after the swap. A swap cannot separate a unit’s receiver from its emitter (Section 3.4).

The extra ambient fell to 1% of its peak at 19.9, 20.9 and 21.6° (full array) and 11.4, 12.7 and 13.1° (window) in campaign 1, and at 19.8, 21.1 and 21.5° (full array) and 11.4, 13.0 and 13.3° (window) in campaign 2 — within 0.3° of campaign 1 in every cell. At those angles the full array’s extra ambient was still 0.19, 0.09 and 0.04 times the solo rate. Per sensor and per angle, the window’s valid-range share depends on the unit and, for unit 1, on its mounting position

(Table 5): at 0.3 m unit 1 stays valid to $\Delta\theta = 24^\circ$ at position A but falls from 20° at position B, while unit 3 is valid throughout at every sampled angle in both campaigns. At 0.6 m, $\Delta\theta = 0$, unit 3 is valid in both campaigns, while unit 1’s share there is 81% in campaign 1 and 18% in campaign 2 — high in one campaign and low in the other, so it is governed by more than the unit’s identity alone — and none are valid at 0.9 m for either unit, so range effects are not available there (Section 3.3 explains why).

Table 5: Valid-range share (%) of the 4×4 window by angle, per unit and mounting position

d	Unit (position, campaign)	0	2	4	6	8	10	12	14	16	18	20	22	24
0.3 m	unit 1 (B, campaign 1)	100	--	100	100	100	100	100	100	100	100	65	1	0
0.3 m	unit 1 (A, campaign 2)	100	100	100	100	100	100	100	100	100	100	100	100	100
0.6 m	unit 3 (A, campaign 1)	100	100	100	98	5	0	0	0	0	0	0	0	0
0.6 m	unit 3 (B, campaign 2)	100	100	100	100	95	6	0	0	0	0	0	0	--
0.6 m	unit 1 (B, campaign 1)	81	22	0	0	0	0	0	0	0	0	0	0	0
0.6 m	unit 1 (A, campaign 2)	18	0	0	0	0	0	0	0	0	0	0	0	--

Full array (16×16) range is valid throughout for both units at every angle and distance, and is omitted. Unit 3’s 4×4 window is valid at every sampled angle at 0.3 m (both campaigns) and is omitted. Neither unit returns a valid 4×4 window reading at any angle at 0.9 m. A dash marks an angle not sampled in that campaign; at 0.6 m, campaign 2 stopped at $\Delta\theta = 22^\circ$.

Turning to question (iii), whether the geometric model predicts this angle: the overlap model of (1), which is purely geometric and so identical in both campaigns, puts the end of overlap at $23.4\text{--}25.8^\circ$ (full array) and $17.3\text{--}19.8^\circ$ (window), which is $3.4\text{--}4.3^\circ$ and $5.9\text{--}6.7^\circ$ beyond the 1% angles: the model overshoots in both zones, so the agreement we expected does not fully hold. Converting the datasheet’s diagonal field of view to an in-plane value moves the full-array values to $19.5\text{--}21.9^\circ$, within 0.5° of the 1% angles, but leaves the window $3.8\text{--}4.6^\circ$ too late (Fig. 8c). Inverting the model for the emitter cone gives $26\text{--}28^\circ$ for the full array and $18\text{--}20^\circ$ for the window: nearly constant with distance within each zone (spread at most 1.8°) but not between zones. So the model reproduces how the required angle grows with distance, which follows directly from (1)’s geometry, but not its dependence on zone, under either convention: candidate explanations we did not test are soft beam edges, a window field that is not simply a narrower copy of the full cone, and array axes outside the plane of rotation. The window’s non-monotonic response, present at 0.9 m, weakly at 0.6 m, but not at 0.3 m, may reflect near-field saturation at 0.3 m (273 kcps); session differences alone are unlikely, since the same pattern reappeared after the unit swap.

3.5 Where the interfering light comes from, and how to avoid it

Section 3.4 located where the effect stops; question (iv) asks where it comes from in the first place, which the divider check was designed to test (Section 2). In campaign 1, with the strip between the sensors the extra ambient changed by -7.7 to $+2.5$ % across the fourteen cells (both zones, $\Delta\theta = 4$ and 10° at all three distances, plus $\Delta\theta = 2^\circ$ at 0.6 m), within ± 4.2 % in twelve of them; the exceptions are the window cells at $\Delta\theta = 10^\circ$ at 0.3 and 0.6 m (-7.7 % and -4.3 %; absolute changes -0.5 and -0.3 kcps). The sign did not follow distance or zone in any clear pattern: all four 0.3 m cells fell (-0.7 to -7.7 %), the 0.6 m cells were small and mixed (-4.3 to $+2.3$ %), and they changed by -0.1 to $+2.5$ % at 0.9 m. The 95% bootstrap intervals excluded zero in seven of the fourteen cells (all four 0.3 m cells, the 0.6 m window at $\Delta\theta = 10^\circ$, and the 0.9 m full array), while the sensors’ full-array solo ambient rate differed by $+0.9$ to $+11.3$ % between the divider and no-divider sessions of campaign 1, so session differences are as large as or larger than the divider effects. Campaign 2 repeated this check at all three distances, with the units exchanged and $\Delta\theta = 2^\circ$ added: the full-array changes were -1.0 to $+0.8$ %, similar to campaign 1’s (window changes spanned -5.6 to $+11.2$ %, at most about 10 kcps absolute). The small, similar spread in both campaigns supports reading these signed changes mainly as session noise and weighs against a strong direct path; it does not remove the confound, since campaign 2’s own solo ambient still differed by up to 18% (full array, 0.9 m, sensor A/unit 1 specifically; 4%–5% for sensor B/unit 3 at the same cell) between its divider and no-divider sessions — as large as, or larger than, the divider effect it is meant to isolate. A thin barrier, in other words, did not show a measurable benefit in either campaign, which answers question (iv) only partially: it does not rule out a direct path, and a small one is not excluded, but it does not demonstrate one either (the strip’s length, height, thickness and exact placement within the protocol’s envelope were not recorded, and the divider cells came from separate sessions in both campaigns). Table 6 gives each cell, both campaigns, with its bootstrap interval.

Two further observations point the same way, toward light that reaches the neighbor via the target rather than directly. First, extra ambient / own signal at $\Delta\theta = 0^\circ$ changed little with distance for the full array (within about 10%, across a 3.6-fold fall in signal; Section 3.2), as expected if both the neighbor’s light and the sensor’s own light are reflected by the same target and fall off with distance together. Second, the response fell smoothly with increasing angular separation

Table 6: Change in the extra ambient rate with the divider, % of the no-divider cell mean (95% bootstrap interval), both campaigns

Zone	d (m)	$\Delta\theta = 2^\circ$ (camp. 2; 0.6 m in camp. 1)	$\Delta\theta = 4^\circ$	$\Delta\theta = 10^\circ$
Campaign 1 (unit 3 at A, unit 1 at B)				
16 × 16	0.3	--	-1.3 (-1.8, -0.8)	-0.7 (-1.3, -0.1)
16 × 16	0.6	-0.5 (-1.1, +0.1)	-0.1 (-0.6, +0.5)	+0.1 (-0.8, +1.1)
16 × 16	0.9	--	+1.6 (+1.1, +2.2)	+2.5 (+1.1, +4.0)
4 × 4	0.3	--	-3.9 (-4.5, -3.3)	-7.7 (-11.8, -4.0)
4 × 4	0.6	+0.3 (-2.5, +3.1)	+2.3 (-0.3, +4.8)	-4.3 (-6.6, -2.0)
4 × 4	0.9	--	+0.2 (-1.1, +1.5)	-0.1 (-2.8, +2.7)
Campaign 2 (unit 1 at A, unit 3 at B)				
16 × 16	0.3	+0.4 (-0.1, +0.9)	+0.2 (-0.2, +0.7)	-0.4 (-1.3, +0.5)
16 × 16	0.6	-1.0 (-1.9, -0.1)	-0.7 (-1.5, +0.2)	-0.8 (-1.5, -0.1)
16 × 16	0.9	+0.7 (-0.1, +1.4)	+0.6 (-0.5, +1.8)	+0.8 (-0.3, +1.7)
4 × 4	0.3	-0.0 (-0.6, +0.8)	-4.2 (-4.8, -3.6)	+11.2 (+7.6, +14.4)
4 × 4	0.6	-1.0 (-2.2, +0.3)	-3.8 (-5.0, -2.8)	-5.6 (-8.1, -3.0)
4 × 4	0.9	-0.3 (-1.3, +0.8)	+1.2 (+0.1, +2.4)	+0.7 (-3.2, +4.5)

$\Delta\theta = 2^\circ$ ($\theta = 1^\circ$) was measured at every distance in campaign 2 but in campaign 1 only at 0.6 m. Full-array changes span -1.3 to $+2.5\%$ in campaign 1 and -1.0 to $+0.8\%$ in campaign 2; the window spans -7.7 to $+2.3\%$ and -5.6 to $+11.2\%$, the latter corresponding to at most about 10 kcps absolute change in either campaign.

at every distance for the full array (Fig. 8a): if the dominant path were direct emitter-to-receiver coupling rather than a shared target, the footprint argument of Section 1.2 would not obviously apply, and the divider, sitting between the sensor windows, would be expected to block it; it did not show a reduction. Neither observation proves the path is exclusively via the target, but both are what that explanation predicts and neither is what a dominant direct path would predict; taken with the divider result, we read the evidence as consistent with, but not proof of, a shared-target path.

Practically, for sensors spaced and mounted like this pair, a barrier between them was not shown to help. Two approaches should work: separate the sensors' pointing angles by more than the crossover found in Section 3.4, or let them take turns so that only one ranges at a time; we expect the first to remove the effect and the second to avoid it by construction, but we tested neither as an operating mode, only as the solo runs that make up two of the three runs in every block. The solo runs throughout this paper are exactly what each sensor sees when the other is idle, and they show no crosstalk at all, by construction.

3.6 Limitations and open questions

One sensor pair (two units); one session per distance in each campaign; angle and time confounded within a campaign; a unit's receiver and emitter properties confounded by the swap; $\theta = 1^\circ$ sampled at every distance only in campaign 2 (and at 0.6 m in campaign 1); plate accuracy, board reflectance, room light and strip geometry unrecorded. Only one timing budget (33 ms, short mode) was used; ST staff state that the sensors do not interfere at longer timing budgets but might when run very fast^{18;19}, so the angles found here may be larger at shorter budgets. Re-measured blocks are at most 6% of the accepted blocks in any sweep session (7 of 120) and 13% (4 of 30) in the 0.9 m divider session of campaign 2, mostly in the tails. The 4×4 window's peak at $\Delta\theta = 2^\circ$ to 4° rather than 0° (clear at 0.9 m in both campaigns, weak at 0.6 m, absent at 0.3 m) is unexplained (Section 3.4). At 0.3 m only, the full array's own (solo) ambient rate rose at large angles: by about 50% for each unit between $\Delta\theta = 14^\circ$ and $\Delta\theta = 24^\circ$ in campaign 1, and by 21% (unit 1) and 51% (unit 3) in campaign 2 (35% pooled), while the window's solo ambient did not rise, so it is not unique to one unit or one session; since this appears in the solo runs, it is not a crosstalk effect, but it adds noise, and conceivably a small bias, to the extra-ambient estimate at the largest angles, where the crosstalk signal itself is already small. The detection-limit effect of Section 3.3 was shown in two settings (Fig. 7) and summarized over fifteen cells (Table 2), not as a general measurement of how often it occurs.

4 Conclusion

For one pair of VL53L1X sensors 20 mm apart, repeated after exchanging the two units, three things follow. First, the extra ambient rate from running both fell to 10% of its peak at about $16\text{--}18^\circ$ of total angular separation between the sensors (half that per sensor) for the full array and $8\text{--}11^\circ$ for the 4×4 window, reaching 1% at about $20\text{--}22^\circ$ and $11\text{--}13^\circ$, growing with distance from 0.3 to 0.9 m. The choice of estimator, not sampling, limits the precision of these angles within a session, and the unit swap shows a between-session spread of about 0.4° ; a geometric model overshoots

the angles by 3–7°. The window’s non-monotonic bump reproduced at 0.9 m after the swap, and the crossing-angle gap between the two sensors followed the unit, not its mounting position, in every comparison. Second, ranging itself holds up well even under heavy interference: a 20-fold jump in ambient rate at close range, arriving in bursts about every 1.9 s, moved the mean distance by half a millimeter or less. Third, near a sensor’s acceptance threshold the neighbor widens the signal scatter and can change the share of accepted readings by up to about 20 percentage points in either direction; we demonstrated this in fifteen cells, not as a general rate. A thin divider changed the effect by no more than about 4% in twelve of fourteen settings before the swap and fifteen of eighteen after it, which does not rule out direct coupling; the evidence instead is consistent with, but does not prove, light reaching the neighbor mainly via the target. For someone building an array of these sensors, the practical guidance is: a barrier between adjacent units was not shown to help; separating their pointing angles past the crossover found here (about 11° each for the full array, 7° each for the window) is expected to remove the effect, and running them in turn should avoid it by construction, but neither was tested as an operating mode. Other units and targets, a third unit (to separate a unit’s emitter properties from its receiver properties), bright room light, and a direct timing analysis of the interrupt timestamps are the next steps we would take before treating these angles as a general design rule.

Acknowledgment

The experiments were self-funded. Python was used to write and run the analysis scripts, to generate the figures and tables, and to draft and edit the text. The author designed and ran the experiments, reviewed the analyses and text, and takes responsibility for the content.

Data and Code Availability

The logged sessions of both campaigns (12 consolidated session folders), the data and code are available at https://github.com/jscwang/lidar_crosstalk.

References

- [1] A. Incoronato, I. Cusini, K. Pasquinelli, and F. Zappa. Single-shot pulsed-LiDAR SPAD sensor with on-chip peak detection for background rejection. *IEEE J. Sel. Top. Quantum Electron.*, 28(5), Sep./Oct. 2022. doi:10.1109/JSTQE.2022.3180687. art. no. 3804210.
- [2] P. Padmanabhan, C. Zhang, and E. Charbon. Modeling and analysis of a direct time-of-flight sensor architecture for LiDAR applications. *Sensors*, 19(24):5464, Dec. 2019.
- [3] STMicroelectronics. VL53L5CX: Time-of-flight 8×8 multizone ranging sensor with wide field of view. datasheet DS13754, Rev. 13, Sep. 2024. URL <https://www.st.com/resource/en/datasheet/vl53l5cx.pdf>. Accessed: Oct. 17, 2026.
- [4] STMicroelectronics. VL53L1X: A new generation, long distance ranging time-of-flight sensor based on ST’s FlightSense technology. datasheet DS12385, Rev. 8, Aug. 2024. URL <https://www.st.com/resource/en/datasheet/vl53l1x.pdf>. Accessed: Oct. 17, 2026.
- [5] J. E. Kvam. Reply to ‘VCSEL signals interferences for closed sensors (VL53L1X)’. STMicroelectronics Community forum, Apr. 2021. URL <https://community.st.com/imaging-sensors-49/vcsl-signals-interferences-for-closed-sensors-vl53l1x-25089>. accessed: Oct. 17, 2026.
- [6] F. Wermke and B. Meffert. Interference model of two time-of-flight cameras. In *Proc. IEEE SENSORS*, Montreal, QC, Canada, Oct. 2019. doi:10.1109/SENSORS43011.2019.8956892.
- [7] R. Z. Whyte, A. D. Payne, A. A. Dorrington, and M. J. Cree. Multiple range imaging camera operation with minimal performance impact. In *Proc. SPIE, Image Processing: Machine Vision Applications III*, volume 7538, page 75380I, Feb. 2010.
- [8] L. Li, S. Xiang, Y. Yang, and L. Yu. Multi-camera interference cancellation of time-of-flight (TOF) cameras. In *Proc. IEEE Int. Conf. Image Process. (ICIP)*, pages 556–560, Sep. 2015.
- [9] G. Kim, J. Eom, and Y. Park. An experiment of mutual interference between automotive LIDAR scanners. In *Proc. 12th Int. Conf. Inf. Technol.—New Gener. (ITNG)*, pages 680–685, Apr. 2015.
- [10] G. Kim, J. Eom, and Y. Park. Investigation on the occurrence of mutual interference between pulsed terrestrial LIDAR scanners. In *Proc. IEEE Intell. Veh. Symp. (IV)*, pages 437–442, Jun. 2015.
- [11] G. B. Popko, T. K. Gaylord, and C. R. Valenta. Interference measurements between single-beam, mechanical scanning, time-of-flight lidars. *Opt. Eng.*, 59(5), May 2020. art. no. 053106.

- [12] A. Ogurcovs, I. Aulika, S. Cartiel, J. Garcia-Pueyo, and A. Muñoz. Crosstalk effects in a dual ToF-based tactile-proximity sensing platform integrated in a flat PMMA light guide. *Sensors*, 25(23):7319, Dec. 2025.
- [13] A. Fasolino, L. E. Andrada, R. Liguori, L. Di Benedetto, A. Rubino, D. Pau, and G. D. Licciardo. Real-time crosstalk compensation in multizone time-of-flight sensors. In *Proc. SPIE, Real-time Processing of Image, Depth, and Video Information 2025*, volume 13526, page 135260D, May 2025.
- [14] STMicroelectronics. A guide to using the VL53L1X ultra lite driver. user manual UM2510, Rev. 7, Dec. 2025. URL https://www.st.com/resource/en/user_manual/um2510-a-guide-to-using-the-vl53l1x-ultra-lite-driver-stmicroelectronics.pdf. Accessed: Oct. 17, 2026.
- [15] Texas Instruments. TCA9548A low-voltage 8-channel I2C switch with reset. datasheet SCPS207G, Nov. 2019. URL <https://ti.com/lit/pdf/SCPS207>. Accessed: Oct. 17, 2026.
- [16] F. N. Fritsch and R. E. Carlson. Monotone piecewise cubic interpolation. *SIAM J. Numer. Anal.*, 17(2):238–246, Apr. 1980.
- [17] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall, New York, NY, USA, 1993.
- [18] J. E. Kvam. Reply to ‘Two VL53L1X sensors influencing their measurements’. STMicroelectronics Community forum, May 2020. URL <https://community.st.com/mems-sensors-48/two-vl53l1x-sensors-influencing-their-measurements-52421>. accessed: Oct. 17, 2026.
- [19] J. E. Kvam. Reply to ‘Multiple VL53L1X sensors and FOV interference [sic]’. STMicroelectronics Community forum, Apr. 2020. URL <https://community.st.com/mems-sensors-48/multiple-vl53l1x-sensors-and-fov-interference-53290>. accessed: Oct. 17, 2026.