Opto-Electronics Review

journal homepage: publishinghouse.pan.pl/index.php/opelre



Bulk-optics architectures for a stochastic optical AND gate: Nominal behaviour, non-idealities, and robustness comparison

Oleg Angelsky[1,2], Myroslav Strynadko[2]*, Claudia Zenkova[2], Bogdan Tymochko[4], Roman Zaiats[2], Zhang Xinzheng[3], Jun Zheng[1]**, Cai Jingxian[1]

[1] Taizhou Institute of Zhejiang University, Taizhou, China
[2] Yuriy Fedkovych Chernivtsi National University, Chernivtsi, Ukraine
[3] Nankai University, Tianjin, China
[4] Bukovinian State Medical University, Chernivtsi, Ukraine

Article info

Article history:
Received 27 Mar. 2026
Received in revised form 23 Aug. 2026
Accepted 14 Sep. 2026
Available online 16 Nov. 2026

Keywords:
stochastic optical logic;
optical AND gate;
bulk optics;
polarisation encoding;
robustness analysis.

Abstract

This paper compares three bulk-optics realisations of a stochastic optical gate performing the logical conjunction operation within a unified analytical framework based on probability accuracy, logical correctness, feasibility margins, and robustness under parameter perturbations. The considered designs share the same nominal logical target but differ in the physical mechanism by which the output decision is formed: a metrology-oriented architecture with ratio-based readout, an electro-optically controlled architecture formulated at the non-coherent intensity level, and an architecture based on thresholded nonlinear optical switching. Under nominal conditions, all three realisations reproduce the stochastic conjunction law and exhibit equivalent global performance on the adopted dense probability grid. Their practical behaviour diverges once realistic non-idealities are introduced. The metrology-oriented architecture retains the broadest robust operating region and the highest overall tolerance. The electro-optically controlled architecture preserves a structurally clean polarisation-coded output and is therefore attractive for cascaded stochastic logic. The nonlinear switching architecture provides the strongest regenerative output formation but operates within a narrower admissible window. The results show that nominal logical correctness alone is insufficient for comparing stochastic optical gates and that practical design choice should be guided by robustness of the underlying decision mechanism.

1. Introduction

Stochastic computing represents numerical values and logical variables through the statistics of bit streams rather than through conventional deterministic multibit words. In this paradigm, the value of a signal is associated with the probability of observing logic 1 within a sequence, which makes simple logic elements capable of performing probability-domain operations and has motivated renewed interest in stochastic architectures as compact and noise-tolerant computational systems [1–3]. For optical implementations, this viewpoint is especially attractive because optical states can be naturally encoded, transported, separated, and measured through physically distinct channels such as polarisation, intensity, or switching state, enabling direct mapping between optical signal statistics and stochastic logic variables.

Optical logic itself is not a new idea. Polarisation-based logic schemes and broader optical logic concepts have long shown that logical states can be embedded into controllable properties of light and processed by passive or active optical elements [4, 6]. More recently, all-optical processing roadmaps and electro-optical logic demonstrations have reinforced the relevance of optical logic as a platform for high-speed and functionally rich information processing [5, 8]. Within this broader landscape, polarisation-coded and intensity-level optical representations remain particularly useful because they provide physically interpretable decision channels and a practical route toward experimental

Corresponding author at: *m.strinadko@chnu.edu.ua
**dbzj@netease.com

https://doi.org/10.24425/opelre.2026.7670
1896-3757/ Association of Polish Electrical Engineers (SEP) and Polish Academy of Sciences (PAS). Published by PAS
© 2026 The Author(s). This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

realisation with bulk-optics or hybrid optoelectronic components [4–8].

At the same time, the existence of a nominally correct optical truth table is not sufficient for a practically meaningful stochastic optical gate. In stochastic logic, the function is defined not only by individual state transitions but also by the stability of the output probability under realistic perturbations. For optical hardware, this requirement translates into sensitivity to leakage, imperfect extinction, finite threshold margins, timing mismatch, imperfect electro-optic control, and nonlinear switching constraints. The importance of these issues is implicit in both stochastic-computing literature, which emphasises error behaviour and implementation constraints, and in optical-logic literature, where physical realisations often depend strongly on device non-idealities and operating windows [2, 3, 5].

This problem becomes especially clear for the optical AND operation. AND is the most natural logical primitive for stochastic computation because it directly corresponds to probability multiplication under suitable statistical assumptions [1–3]. In optics, however, the same Boolean target can be realised through quite different physical mechanisms. A polarisation- and measurement-oriented architecture may derive the output from analogue intensity ratios and threshold separation; an electro-optically controlled architecture may generate the required output stream through guided intensity-level state transformation; and a nonlinear all-optical-switch architecture may produce strong regenerative contrast through thresholded switching behaviour. Prior work has shown the feasibility of optical AND/logic elements in nonlinear and electro-optical settings [7–11], but these implementations are usually analysed within their own physical context rather than within a unified comparison framework that asks which realisation remains most usable once realistic non-idealities are introduced.

A further motivation for such comparison comes from the measurement and physical-optics perspective. In a bulk-optics stochastic gate, the experimentally relevant quantity is often not merely the presence of a logical output, but the degree to which the output can be separated, measured, calibrated, and trusted under non-ideal conditions. In this sense, correlation-optics and optical-diagnostics viewpoints are useful as a broader conceptual backdrop: modern treatments of coherence, structured optical fields, and optical diagnostics emphasise that the informative content of an optical field is inseparable from the way it is analysed and measured [12–14]. For the present work, this context is most directly relevant to the metrology-oriented architecture, where the output decision is explicitly tied to measurable analogue ratios and threshold margins, while the broader singular/coherence context remains supportive rather than central [12–14].

Against this background, the goal of this paper is to compare three bulk-optics realisations of a stochastic optical AND gate within a single harmonised analytical framework. The compared designs share the same Boolean target and stochastic interpretation but differ in their internal decision physics: (i) a metrology-oriented architecture with ratio-based readout, (ii) an electro-optic (EO)-controlled architecture described in the final model at the non-coherent intensity level, and (iii) an all-optical switching (AOS)-based architecture with strong nominal contrast but a narrower feasible operating window. Rather than treating these schemes as unrelated implementations, we evaluate them by common probability-level and decision-level metrics, feasible-region analysis, and Monte Carlo robustness criteria, while preserving the architecture-specific physics of each model.

The main contribution of the paper is therefore not only the demonstration that all three architectures reproduce the stochastic AND function under nominal conditions, but also the clarification that nominal equivalence does not imply practical equivalence. We show that the metrology-oriented architecture remains the most robust after non-ideal modelling, the EO-controlled architecture is attractive for cascaded logic propagation but more sensitive to control and timing imperfections, and the AOS-based architecture offers the strongest nominal contrast at the cost of a narrower robust operating region. In this way, the paper converts three qualitatively different bulk-optics realisations into a common comparison space and provides design-oriented guidance for selecting a stochastic optical AND architecture according to robustness priorities rather than nominal truth-table behaviour alone.

All references in the main text to Supplements S1–S9 refer to the corresponding sections of the supplementary package provided with this submission.

2. Common analytical framework and comparison protocol

2.1. Common stochastic representation

All three architectures are evaluated within the same stochastic-logic representation. The two input streams, A and B are polarisation-coded, with horizontal polarisation H corresponding to logical 1 and vertical polarisation V corresponding to logical 0.

The analysis is source-agnostic at the level of the adopted normalised intensity-and-polarisation model. The logical operation does not require a particular emitter or mutual optical coherence because no phase-sensitive interference term is used in the decision laws. The source power is absorbed into the reference intensities, while the spectral and angular dependences of the optical components and detector are represented by effective transmission, reflection, and detection coefficients. Therefore, the model supports coherent, incoherent, and broadband illumination, provided the required input polarisation states are prepared and the effective coefficients remain calibrated over the employed spectrum. For a broadband source, these coefficients should be interpreted as spectrally averaged quantities. This compatibility is a theoretical consequence of the model and was not separately validated experimentally with incoherent or white-light sources.

For stochastic inputs with probabilities p_A and p_B, the ideal AND target is

$$p_{\mathrm{out,ideal}} = p_A p_B. \tag{1}$$

At the state level, the four joint input cases are

$$HH,\ HV,\ VH,\ VV, \tag{2}$$

with occurrence probabilities:

$$P_{HH} = p_A p_B, \quad P_{HV} = p_A(1-p_B),$$
$$P_{VH} = (1-p_A)p_B, \quad P_{VV} = (1-p_A)(1-p_B). \tag{3}$$

The ideal truth table is therefore defined by the rule that only the *HH* state produces logical 1, while *HV*, *VH*, and *VV* produce logical 0.

To compare physically different implementations on a common basis, each architecture is reduced to the state-conditioned output decision probabilities:

$$p_1(HH) = P(C=1\,|HH),$$
$$p_1(HV) = P(C=1\,|HV),$$
$$p_1(VH) = P(C=1\,|VH),$$
$$p_1(VV) = P(C=1\,|VV), \tag{4}$$

where *C* denotes the output logic state. Under this notation, the predicted output probability becomes

$$\hat{p}_{\text{out}} = P_{HH}\,p_1(HH) + P_{HV}\,p_1(HV) + P_{VH}\,p_1(VH) + P_{VV}\,p_1(VV). \tag{5}$$

Here, $p_1(s) = P(C=1\,|s)$ for $s \in \{HH, HV, VH, VV\}$.

This representation is central to the final notebook because it allows the three architectures to be compared through the same probability-domain and decision-domain metrics while preserving their different internal physics.

Real non-polarising beam splitters may exhibit polarisation-dependent amplitude transmission/reflection and a differential phase shift between the *H* and *V* components. The latter can convert a linearly polarised mixed-component state into an elliptically polarised state, particularly in a reflected branch. We therefore added a vendor-independent Jones-model sensitivity analysis for Figs. 1 and 2. The analysis separates diattenuation from retardance: pure differential phase changes the polarisation ellipse, whereas an immediately resolved *H*/*V* intensity ratio in aligned analyser axes is affected primarily by polarisation-dependent amplitude imbalance. If a substantial downstream ellipticity is present, a clean-up linear polariser should be placed after the affected a non-polarising beam splitter (NPBS) branch and before the polarisation-sensitive decision element; an input polariser alone cannot remove ellipticity generated later in the optical train. The associated projection loss must be included in the corresponding effective branch coefficient.

2.2. Common probability-level and decision-level metrics

The primary probability-domain deviation at a grid point (p_A, p_B) is defined as:

$$\Delta p(p_A, p_B) = \hat{p}_{\text{out}} - p_A p_B,$$
$$|\Delta p(p_A, p_B)| = |\hat{p}_{\text{out}} - p_A p_B|. \tag{6}$$

The principal global probability metric is the maximum absolute deviation over the test grid,

$$\Delta p_{\max} = \max_{(p_A, p_B)} |\Delta p(p_A, p_B)|. \tag{7}$$

Following the metric convention used in the final notebook, the aggregated quantities $\text{RMSE}_{\text{mean}}$ and MAE_{mean} are computed from the gridwise probability deviations rather than from analogue intensity traces:

$$\text{RMSE}_{\text{mean}} = \sqrt{\frac{1}{K}\sum_{k=1}^{K} \Delta p_k^2}, \tag{8}$$

$$\text{MAE}_{\text{mean}} = \frac{1}{K}\sum_{k=1}^{K} |\Delta p_k|, \tag{9}$$

where *K* is the number of grid points.

In parallel, decision-level performance is evaluated through the mean BER over the probability grid. For a given pair (p_A, p_B), the analytical local BER is written as:

$$\text{BER}(p_A, p_B) = P_{HH}(1 - p_1(HH), + P_{HV}\,p_1(HV) + P_{VH}\,p_1(VH) + P_{VV}\,p_1(VV). \tag{10}$$

The grid-averaged BER is then

$$\text{BER}_{\text{mean}} = \frac{1}{K}\sum_{k=1}^{K} \text{BER}_k. \tag{11}$$

To distinguish the type of logical failure, false-positive and false-negative indicators are also tracked. The local false-positive rate is defined as the probability of generating output 1 when the ideal output should be 0,

$$\text{FP}(p_A, p_B) = \frac{P_{HV}\,p_1(HV) + P_{VH}\,p_1(VH) + P_{VV}\,p_1(VV)}{1 - P_{HH}}, \tag{12}$$

for $1 - P_{HH} > 0$, whereas the false-negative rate is

$$\text{FN}(p_A, p_B) = 1 - p_1(HH). \tag{13}$$

Their grid-averaged values are reported as FP_mean and FN_mean. These quantities are identical to FPR_mean and FNR_mean, respectively, in the supplementary material.

This common metric layer is used for all three architectures and forms the basis of the cross-architecture comparison in section 6.

2.3. Architecture-specific feasibility metrics

Although the global metrics are shared, each architecture requires one additional feasibility descriptor reflecting its internal decision mechanism.

For the metrology-oriented architecture (Fig. 1), the output is determined from the measured intensity ratio:

$$q = \frac{I_H}{I_H + I_V}, \tag{14}$$

with the decision threshold T_{th}. The state-level threshold margins are

$$M_1 = q_{HH} - T_{\text{th}}, \quad M_0 = T_{\text{th}} - \max(q_{HV}, q_{VH}, q_{VV}), \tag{15}$$

and the effective threshold-separation margin is

$$M = \min(M_1, M_0). \tag{16}$$

A nominal threshold setting is feasible when $M > 0$. In the noisy version of the model, an additional robustness quantity is used:

$$M_{\mathrm{rob}} = q_{HH}^{(5\%)} - q_{0,\max}^{(95\%)}, \tag{17}$$

where $q_{HH}^{(5\%)}$ is the lower quantile of the *HH* decision ratio and $q_{0,\max}^{(95\%)}$ is the upper quantile of the most dangerous zero-output state. This quantity directly measures whether threshold separation survives detector noise and threshold uncertainty.

For the EO-controlled architecture (Fig. 2), the final model is intentionally restricted to a non-coherent intensity-level description. The decisive auxiliary quantities are the horizontal-output purity for the desired *HH* state, denoted $\Pi_{H,HH}$, and the maximum residual horizontal component among the zero-output states. For each input state $s \in \{HH, HV, VH, VV\}$, the normalised horizontal-output purity is defined as $\Pi_{H,s} = H_{\mathrm{out},s}/(H_{\mathrm{out},s} + V_{\mathrm{out},s})$. The quantity $\Pi_{H,0,\max}$ denotes the largest residual horizontal purity among the three states whose ideal AND output is zero:

$$\Pi_{H,0,\max} = \max\left(\Pi_{H,HV}, \Pi_{H,VH}, \Pi_{H,VV}\right). \tag{18}$$

The corresponding state-separation margin is

$$M_{\Pi} = \Pi_{H,HH} - \Pi_{H,0,\max}. \tag{19}$$

The EO-controlled architecture is considered state-feasible when $M_{\Pi} > 0$. This formulation is fully consistent with the adopted final model, in which coherence and interference terms are excluded from the main manuscript. For the AOS-based architecture (Fig. 3), the dominant auxiliary quantities are the effective threshold margins of the nonlinear switching stages. In the notation of the final notebook, these are written as:

$$\begin{aligned} M_H &= \min\left(I_{\mathrm{th},2} - I_{1H}^{\mathrm{eff}}, I_{2H}^{\mathrm{eff}} - I_{\mathrm{th},2}\right), \\ M_V &= I_{1V}^{\mathrm{eff}} - I_{\mathrm{th},1}. \end{aligned} \tag{20}$$

Here, M_H measures whether the second switching threshold can separate the undesired and desired horizontal-control conditions, whereas M_V measures whether the first switching threshold is crossed reliably in the required branch, I_{1H}^{eff} and I_{2H}^{eff} are the effective horizontal control intensities generated by one and two horizontal input contributions, respectively; I_{1V}^{eff} is the effective vertical control intensity generated by one vertical contribution; and $I_{\mathrm{th},1}$ and $I_{\mathrm{th},2}$ are the switching thresholds of AOS1 and AOS2, respectively. The superscript "eff" indicates that loss, leakage, overlap, and other non-ideal factors have already been applied.

Threshold feasibility is achieved when

$$M_H > 0, \quad M_V > 0\,. \tag{21}$$

In addition, the AOS architecture must satisfy dynamic feasibility conditions related to switching and recovery times, as well as sufficient temporal overlap of the control pulses.

For clarity, Table 1 summarises the common notation, global comparison metrics, and architecture-specific feasibility criteria used throughout the manuscript.

2.4. Test grids, nominal evaluation, and robustness protocol

The comparison protocol is organised in two layers. First, nominal behaviour is evaluated on a dense probability grid covering

$$p_A, |p_B \in \{0, 0.1, 0.2, \ldots, 1.0\}, \tag{22}$$

which yields $11 \times 11 = 121$ grid points. This dense grid is used for the nominal output surfaces and for the summary nominal metrics reported from the S2, S4, and S6 packages. Second, robustness analysis is performed on a coarse grid

$$p_A, |p_B \in \{0, 0.25, 0.5, 0.75, 1.0\}, \tag{23}$$

which yields $5 \times 5 = 25$ grid points. This reduced grid is used in the non-ideal and Monte Carlo packages S3, S5, and S7, where repeated evaluation over perturbed parameter sets is required. In the final notebook, each architecture is also assessed by a Monte Carlo ensemble of 250 realisations.

To keep the comparison uniform, the same acceptance thresholds are used across architectures:

$$\Delta p_{\max} \le 0.05, \quad \mathrm{BER}_{\mathrm{mean}} \le 0.02. \tag{24}$$

These two conditions are necessary but not sufficient. Each architecture must also satisfy its own internal feasibility requirement:

- Figure 1: $M > 0$ and positive noisy separation margin $M_{\mathrm{rob}} > 0$;
- Figure 2: $M_{\Pi} > 0$;
- Figure 3: $M_H > 0$, $M_V > 0$ and dynamic feasibility of the nonlinear switching sequence.

Under this protocol, a realisation is classified as robust only if it simultaneously satisfies the common probability/error criteria and the architecture-specific feasibility conditions. This makes it possible to compare the three designs fairly without forcing them into a single artificial physical model.

2.5. Position of the supplementary material

The supplementary package is used as the supporting analytical layer rather than as a duplicate of the main paper. Supplement S1 provides the unified notation, ideal truth table, parameter dictionary, and common metrics. Supplements S2, S4, and S6 contain the architecture-specific nominal packages for Fig. 1–3. Supplements S3, S5, and S7 contain the corresponding non-ideal, sensitivity, and Monte Carlo analyses. Supplement S8 (Figs. S8.1–S8.5) provides a harmonised cross-architecture comparison used later in the manuscript, while S9 remains a reproducibility and package-consistency layer.

Accordingly, the main manuscript retains only the common framework, the key nominal and robustness figures, and the final comparative interpretation, whereas extended sweeps, additional maps, histograms, manifests, and reproducibility artifacts are delegated to the supplementary material.

Table 1.
Common notation, probability-level metrics, decision-level metrics, and architecture-specific feasibility criteria used for the comparison of the three stochastic optical AND-gate architectures.

Block	Symbol / term	Definition	Role in the manuscript	Applies to
Stochastic inputs	p_A, p_B	Probabilities of logical 1 in the input stochastic streams (A) and (B)	Define the stochastic operating point	Figs. 1–3
Ideal AND target	$p_{\text{out}}^{\text{ideal}}$	$p_{\text{out}}^{\text{ideal}} = p_A\, p_B$	Reference probability for all comparisons	Figs. 1–3
Joint input-state probabilities	P_{HH}, P_{HV}, P_{VH}, P_{VV}	$P_{HH} = p_A p_B$, $P_{HV} = p_A(1 - p_B)$, $P_{VH} = (1 - p_A)p_B$, $P_{VV} = (1 - p_A)(1 - p_B)$	Weight state-conditioned output behaviour	Figs. 1–3
State-conditioned output probabilities	$p_1(HH)$, $p_1(HV)$, $p_1(VH)$, $p_1(VV)$	Conditional probabilities of generating output 1 for each input state	Unified state-level representation of all architectures	Figs. 1–3
Predicted output probability	$\hat{p}_{\text{out}}$	$\hat{p}_{\text{out}} = P_{HH}p_1(HH) + P_{HV}p_1(HV) + P_{VH}p_1(VH) + P_{VV}p_1(VV)$.	Main model output compared against the ideal AND law	Figs. 1–3
Probability deviation	Δp	$\Delta p = \hat{p}_{\text{out}} - p_A p_B$	Local probability-level error	Figs. 1–3
Maximum absolute deviation	$\Delta p_{\max}$	Maximum of absolute $\lvert\Delta p\rvert$ over the probability grid	Main global probability-accuracy metric	Figs. 1–3
Mean RMSE	$\text{RMSE}_{\text{mean}}$	Root-mean-square error over all grid points	Aggregated probability-domain performance metric	Figs. 1–3
Mean MAE	MAE_{mean}	Mean absolute error over all grid points	Aggregated probability-domain performance metric	Figs. 1–3
Local BER	BER (p_A, p_B)	Analytical bit-error probability at a given stochastic operating point	Local decision-level accuracy measure	Figs. 1–3
Mean BER	BER_{mean}	Mean BER over the probability grid	Main global logical-correctness metric	Figs. 1–3
False-positive rate	FP, FP_mean (≡ FPR, FPR_mean)	Probability of output 1 when the ideal output should be 0	Characterises spurious logical activation	Figs. 1–3
False-negative rate	FN, FN_mean (≡ FNR, FNR_mean)	Probability of output 0 when the ideal output should be 1	Characterises loss of the desired AND response	Figs. 1–3
Decision ratio	q	$q = I_H/(I_H + I_V)$	Observable analogue decision variable for thresholding	Fig. 1
Threshold separation margin	M	$M = \min(M_1, M_0)$, where $M_1 = q_{HH} - T_{\text{th}}$, $M_0 = T_{\text{th}} - \max(q_{HV}, q_{VH}, q_{VV})$	Nominal feasibility criterion for state separation	Fig. 1
Noisy robustness margin	M_{rob}	$M_{rob} = q_{HH}^{(5\%)} - q_{0,max}^{(95\%)}$	Verifies separation survival under noise and threshold uncertainty	Fig. 1
Horizontal-output purity margin	M_Π	$M_\Pi = M_{H,HH} - M_{H,0,\max}$	State-feasibility criterion in the non-coherent EO model	Fig. 2
Horizontal-output purity for target state	$\Pi_{H,HH}$	Horizontal output purity in the desired HH state	Measures desired-state fidelity	Fig. 2
Maximum residual horizontal component	$\Pi_{H,0,\max}$	Maximum horizontal leakage among (HV), (VH), (VV) states	Measures undesired-state contamination	Fig. 2
Horizontal switching margin	M_H	$M_H = \min\left(I_{\text{th},2} - I_{1H}^{\text{eff}}, I_{2H}^{\text{eff}} - I_{\text{th},2}\right)$	Feasibility of threshold placement in the horizontal AOS branch	Fig. 3
Vertical switching margin	M_V	$M_V = I_{1V}^{\text{eff}} - I_{\text{th},1}$	Feasibility of threshold crossing in the vertical AOS branch	Fig. 3
Dynamic feasibility	–	Satisfaction of switching-time, recovery-time, and pulse-overlap conditions	Ensures realisability of the nonlinear switching sequence	Fig. 3
Acceptance threshold	$\Delta p_{\max} \le 0.05$	Common robustness requirement for probability accuracy	Cross-architecture acceptance criterion	Figs. 1–3
Acceptance threshold	$\text{BER}_{\text{mean}} \le 0.02$	Common robustness requirement for logical accuracy	Cross-architecture acceptance criterion	Figs. 1–3
Robust realisation	–	Simultaneous satisfaction of common acceptance criteria and architecture-specific feasibility conditions	Final classification used in robustness comparison	Figs. 1–3

3. Metrology-oriented architecture (Fig. 1)

3.1. Operating principle and nominal state response

The first implementation is a metrology-oriented realisation of the stochastic optical AND function (Fig. 1). Its defining feature is that the output decision is not treated as an implicit internal state, but is derived from an explicitly measurable analogue quantity, namely the ratio of the detected horizontal component to the total detected output intensity. This makes the architecture especially suitable for calibration-aware analysis, uncertainty evaluation, and experimentally transparent verification.

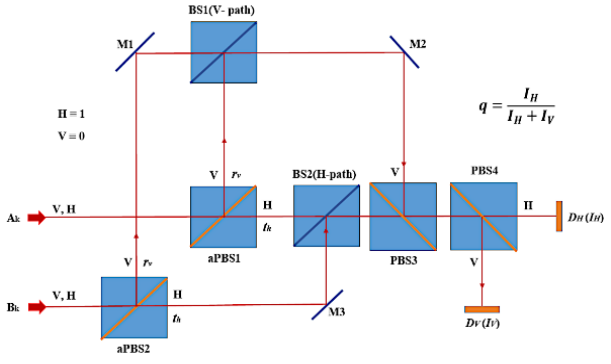


Fig. 1. Metrology-oriented bulk-optics architecture for the stochastic optical AND gate with ratio-based decision readout $q = I_H/(I_H + I_V)$.

In the adopted model, the two polarisation-coded inputs A and B are processed by a bulk-optics network that produces two detected output components, I_H and I_V .

The output logic state is then determined from the ratio:

$$q = \frac{I_H}{I_H + I_V}, \tag{25}$$

with the scalar threshold T_{th}. A value $q > T_{\text{th}}$ is interpreted as output 1, whereas $q \le T_{\text{th}}$ is interpreted as output 0.

For the nominal parameter set used in the final notebook, the deterministic state-conditioned intensities are

$$\begin{aligned} I_H &= I_0 t_H \left(A_H T_H + B_H n_B R_H \right), \\ I_V &= I_0 r_V \left(A_V m_A T_V + B_V m_B R_V \right). \end{aligned} \tag{26}$$

Here, I_0 is the reference input intensity; A_H, A_V, B_H, $B_V \in \{0,1\}$ are the state indicators of the two polarisation-coded inputs; t_H and r_V are the effective routing coefficients of the horizontal and vertical analyser channels shown in Fig. 1; T_H, R_H, T_V, R_V are the effective transmission and reflection factors of the corresponding horizontal and vertical optical paths; and m_A, m_B, n_B are the branch-specific balance or loss factors. These effective coefficients refer to the paths labelled in Fig. 1 and include the net response of the relevant beam splitters, polarisation elements, mirrors, and propagation losses. The complete nominal parameter set, truth table, and calculation workflow are given in Supplement S2 (Figs. S2.1–S2.3).

For the nominal baseline ($I_0 = 1$, $t_H = 0.2$, $r_V = 0.8$, $T_H = R_H = T_V = R_V = 0.5$, $m_A = m_B = n_B = 1$), the four input states produce the following deterministic outputs:

- HH: $I_H = 0.2$, $I_V = 0$, $q = 1.0$
- HV: $I_H = 0.1$, $I_V = 0.4$, $q = 0.2$
- VH: $I_H = 0.1$, $I_V = 0.4$, $q = 0.2$
- VV: $I_H = 0$, $I_V = 0.8$, $q = 0$.

Thus, the nominal state response is fully ordered: the desired HH state is separated from all zero-output states by a large ratio gap, while the two mixed states HV and VH remain symmetric, as expected from the AND logic. This gives the architecture a particularly transparent physical interpretation: the logical output is determined by measurable analogue contrast rather than by an internal black-box decision mechanism.

3.2. Threshold window and nominal stochastic behaviour

Since the nominal state ratios satisfy

$$q_{HH} = 1.0, \quad q_{HV} = q_{VH} = 0.2, \quad q_{VV} = 0,$$

the threshold remains feasible over the interval

$$0.2 < T_{th} < 1.0 \, .$$

For the baseline choice $T_{\text{th}} = 0.6$, the threshold margins are

$$\begin{aligned} M_1 &= q_{HH} - T_{th} = 0.4, \\ M_0 &= T_{th} - \max\left(q_{HV}, q_{VH}, q_{VV} \right) = 0.4, \end{aligned}$$

so that the effective nominal margin is

$$M = \min\left(M_1, M_0 \right) = 0.4.$$

This symmetric margin is one of the key reasons why Fig. 1 remains the most robust architecture after the introduction of non-idealities.

At the stochastic level, the nominal output reproduces the ideal AND law over the dense 11×11 probability grid used in Supplement S2 (Figs. S2.1–S2.3). The dense-grid comparison gives $\Delta p_{\max} = 0.0109$, $\text{RMSE}_{\text{mean}} = 0.003255$, $\text{MAE}_{\text{mean}} = 0.002298$, $\text{BER}_{\text{mean}} = 0$. In practical terms, the architecture is exact at the binary truth-table level and closely reproduces the ideal probability surface; the small residual probability-level deviations arise from the finite-stream stochastic evaluation used in the notebook workflow rather than from a structural error of the architecture itself.

The detailed nominal state-response plot, threshold-window visualisation, and dense-grid output surface are provided in Supplement S2 (Figs. S2.1–S2.3).

This nominal behaviour is important for two reasons. First, it confirms that the ratio-based readout does not distort the underlying AND law. Second, it establishes a broad threshold window before any non-ideal modelling is introduced, which later translates into unusually stable robustness statistics compared with the EO-controlled and AOS-based alternatives.

3.3. Non-ideal model and robustness under perturbations

Let $I_H^{(0)}$ and $I_V^{(0)}$ denote the nominal intensities from (26). With ε_{HV} defined as the leakage from H to V and ε_{VH} as the leakage from V to H, the effective channel intensities are

$$I_H^{\text{eff}} = (1-\varepsilon_{HV})I_H^{(0)} + \varepsilon_{VH}I_V^{(0)},$$
$$I_V^{\text{eff}} = (1-\varepsilon_{VH})I_V^{(0)} + \varepsilon_{HV}I_H^{(0)}. \quad (27)$$

Detector gain, offset, and additive noise are introduced as

$$\tilde{I}_H = \max\left[0, g_H I_H^{\text{eff}} + d_H + \xi_H\right],$$
$$\tilde{I}_V = \max\left[0, g_V I_V^{\text{eff}} + d_V + \xi_V\right], \quad (28)$$

where g_H and g_V are the detector gains, d_H and d_V are the dark/background offsets, and $\xi_H, \xi_V \sim N(0, \sigma_{\text{det}}^2)$. The noisy decision ratio and threshold are $\tilde{q} = \tilde{I}_H/(\tilde{I}_H + \tilde{I}_V)$ and $\tilde{T}_{\text{th}} = T_{\text{th}} + \xi_T$, with $\xi_T \sim N(0, \sigma_T^2)$. The output bit is one when $\tilde{q} > \tilde{T}_{\text{th}}$. The ideal reference uses $\varepsilon_{HV} = \varepsilon_{VH} = 0$, $g_H = g_V = 1$, $d_H = d_V = 0$, and $T_{\text{th}} = 0.6$; the noise sweeps cover $0 \le \sigma_{\text{det}} \le 0.08$ and $0 \le \sigma_T \le 0.04$. Full parameter sweeps and Monte Carlo distributions are provided in Supplement S3 (Figs. S3.1–S3.5).

This formulation has an important methodological advantage: all major degradations remain expressible in directly measurable quantities. As a result, robustness can be discussed in terms of threshold separation rather than only in terms of abstract logic failure, which is fully consistent with the metrology-oriented motivation of this architecture.

The noise-aware state analysis confirms that Fig. 1 retains substantial separation even after perturbation. In the dedicated threshold/noise sweep of Supplement S3 (Figs. S3.1–S3.5), the quantile-based robust margin decreases monotonically from 0.800 at the ideal point ($\sigma_{\text{det}} = 0$, $\sigma_T = 0$) to 0.129 at ($\sigma_{\text{det}} = 0.08$, $\sigma_T = 0.04$) yet remains positive throughout the tested range. At the same time, the common accuracy criteria reveal the practical operating boundary more strictly than the margin alone. Up to (σ_{det}, σ_T) = (0.06, 0.03), the exported metrics remain comfortably within the common acceptance limits, with $\Delta p_{\max} = 0.0366$, $\text{BER}_{\text{mean}} = 0.00915$. When the noise level is increased to (0.07, 0.035), the separation margin is still positive, but the probability-accuracy criterion is no longer satisfied because $\Delta p_{\max} = 0.0608 > 0.05$. This is a useful practical result: in Fig. 1, the analogue state separation survives somewhat beyond the point where the unified robustness target is violated. In other words, the architecture degrades gracefully rather than catastrophically.

The one parameter that most strongly affects deterministic logical correctness is the threshold itself. The sensitivity summary in Supplement S3 (Figs. S3.1–S3.5) shows that threshold mis-setting dominates all other tested single-parameter perturbations. However, this sensitivity is highly structured rather than fragile. Over the interval $0.30 \le T_{\text{th}} \le 0.80$, the exported mean BER and $\Delta p_{\max}$ remain zero to reported precision, while threshold feasibility is preserved throughout the broader interval $0.25 \le T_{\text{th}} \le 0.90$. The baseline choice T_{th}=0.6 therefore lies well inside the stable plateau rather than near a critical edge. By contrast, the remaining tested deterministic parameters have much smaller influence over the explored ranges: for example, sweeping t_H changes $\Delta p_{\max}$ only up to 0.0025 and BER_{mean} only up to 0.000625, while the tested detector offsets and leakage terms do not produce comparable dominant errors in the exported summary.

The broad robustness of Fig. 1 is also evident in the feasible-region analysis. In the two-parameter sweep reported in Supplement S3 (Figs. S3.1–S3.5), where symmetric leakage $\varepsilon_{HV} = \varepsilon_{VH}$ is varied over $0 \le \varepsilon \le 0.10$ and the threshold is varied over $0.4 \le T_{\text{th}} \le 0.8$, the feasible fraction remains 1.0 throughout the entire tested domain under the adopted moderate noise setting ($\sigma_{\text{det}} = 0.01$, $\sigma_T = 0.02$). This means that the architecture does not merely possess a single robust operating point; it preserves robustness over a broad region of practical parameter choices.

Extended sensitivity sweeps, the representative Δp heatmap, the BER-*vs*.-threshold curve, and the full feasible-region map are provided in Supplement S3 (Figs. S3.1–S3.5).

Finally, the Monte Carlo analysis provides the strongest condensed summary of Fig. 1 behaviour. Across 250 realisations, the mean performance remains extremely close to the nominal ideal:

- $\Delta p_{\max,\text{mean}} = 2.52 \times 10^{-4}$
- $\text{RMSE}_{\text{mean}} = 9.4 \times 10^{-5}$
- $\text{MAE}_{\text{mean}} = 6.3 \times 10^{-5}$
- $\text{BER}_{\text{mean}} = 6.3 \times 10^{-5}$.

The false-positive component remains zero in the exported Monte Carlo summary, while the false-negative component is very small, $\text{FN}_{\text{mean}} = 2.52 \times 10^{-4}$. At the same time, the median threshold margin remains high ($M_{\text{median}} = 0.352$), the median robust margin is even larger (0.613), and the robust fraction is exactly 1.0 over all 250 Monte Carlo realisations. Taken together, these results justify the conclusion that Fig. 1 is not only nominally correct, but also the most stable and interpretable architecture in the full three-way comparison.

In the generic *NPBS* sweep reported in Supplement S3 (Fig. S3.5), the intensity diattenuation ratio was varied from 0.8 to 1.2 and the differential phase from 0° to 180°. The deterministic Fig. 1 margin remained positive over the full adopted range ($0.362 \le M \le 0.400$).

3.4. *Interpretation within the comparison framework*

Within the unified comparison adopted in this paper, Fig. 1 occupies a distinctive position. It does not aim to maximise regenerative behaviour or architectural compactness. Instead, it maximises observability of the decision variable and preserves a broad threshold-separation reserve. This is precisely why it remains the most robust architecture after non-ideal modelling and why it is especially well suited for metrology-driven validation, uncertainty analysis, and teaching-oriented demonstration of stochastic optical logic.

At the same time, this strength also defines its main limitation. Because the logical output is determined through analogue ratio separation, the architecture is less naturally interpreted as a clean cascaded logic stage than the EO-controlled realisation discussed next. In other words, Fig. 1 is the strongest architecture for understanding, validating, and calibrating the stochastic AND operation, but not necessarily the most elegant one for direct multi-stage propagation.

4. EO-controlled architecture (Fig. 2)

4.1. Operating principle and nominal behaviour

In contrast to the metrology-oriented architecture, Fig. 2 is organised as a logic-propagation stage: its main design advantage is that the output remains explicitly polarisation-coded and is therefore naturally compatible with cascaded stochastic logic.

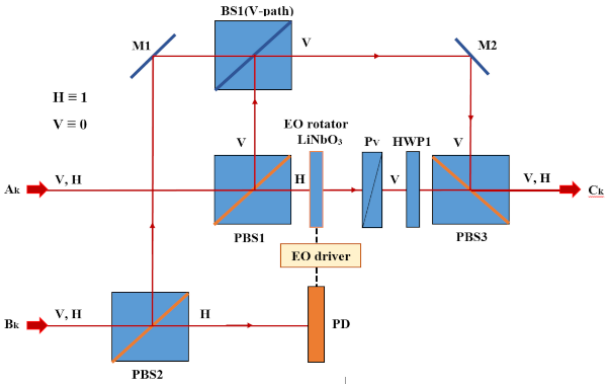


Fig. 2. EO-controlled bulk-optics architecture for the stochastic optical AND gate in the final non-coherent intensity-level formulation.

In the nominal structural model, the control logic is simple and fully state dependent. The *B*-input activates the EO element only when $B = H$, while the *A*-input contributes to the upper branch only when $A = H$. As a result, the upper branch can generate an *H*-polarised output only for the *HH* input state. For all remaining states, the lower branch provides a *V*-polarised output. In the notation of the final notebook, this nominal logic is expressed through three structural flags: EO activation, upper-path activity, and upper-branch *H*-generation. Their combined action yields the nominal truth flow summarised in Supplement S4 (Figs. S4.1–S4.3).

The nominal state table confirms exact state-level AND behaviour. For the *HH* case, EO activation is on, the upper path is active, and the upper branch generates an *H*-polarised output. For *HV*, EO activation is absent, so the upper branch does not produce the desired converted *H* output, and the resulting state is *V*. For *VH*, EO activation is present, but the *A*-signal does not enter the upper path, so the output also remains *V*. For *VV*, neither upper branch, *H* generation nor a valid control event occurs, and the output is again *V*.

The nominal state purity is therefore ideal:

$$\Pi_{H,HH} = 1, \quad \Pi_{H,HV} = \Pi_{H,VH} = \Pi_{H,VV} = 0, \tag{29}$$

which gives

$$\Pi_{H,0,\max} = 0, \quad M_{\Pi} = 1. \tag{30}$$

At the stochastic level, the nominal dense-grid evaluation in Supplement S4 (Figs. S4.1–S4.3) yields the same global nominal result as the other two architectures:

$$\begin{aligned} \Delta p_{\max} &= 0.0109, \qquad \mathrm{RMSE}_{\mathrm{mean}} = 0.003255, \\ \mathrm{MAE}_{\mathrm{mean}} &= 0.002298, \quad \mathrm{BER}_{\mathrm{mean}} = 0, \end{aligned} \tag{31}$$

to the reported precision. Thus, as long as the EO-controlled stage is treated ideally, it reproduces the stochastic AND law without any logical ambiguity and with a perfectly clean nominal polarisation-coded output. This nominal behaviour is important because it establishes Fig. 2 as the most natural candidate for cascaded functional propagation among the three architectures. The logical result is already encoded in the output polarisation state itself, rather than being reconstructed through an analogue readout variable.

The detailed nominal truth-flow diagram, nominal Δp map, and dense-grid output surface are provided in Supplement S4 (Figs. S4.1–S4.3).

4.2. Final non-coherent intensity-level model

For the main manuscript, Fig. 2 is described by the final non-ideal model adopted in Supplement S5 (Figs. S5.1–S5.6).

The relations in this subsection constitute an author-proposed normalised phenomenological model developed for comparative robustness analysis. They were not derived by fitting measurements from a specific EO device and are not claimed as a device-specific transfer function. The logistic activation, Gaussian timing-overlap factor, normalised intensity balance, and purity margin are used as effective descriptors of control activation, temporal overlap, and output-state separation.

Importantly, this is a non-coherent intensity-level model: coherence and interference effects are intentionally excluded from the main analysis. The goal is not to describe phase-sensitive propagation, but to evaluate whether the EO-controlled architecture can preserve correct stochastic logic under physically meaningful intensity-level imperfections.

The model begins with a soft EO activation rather than a hard binary switch. If the effective control signal is denoted by

$$u_{\mathrm{ctrl}} = \eta_{PD} B_{\mathrm{ctrl}} + u_{\mathrm{leak}}, \tag{32}$$

then the EO activation is

$$a_{EO} = \frac{1}{1 + \exp\left[-\left(u_{\mathrm{ctrl}} - u_{\mathrm{th}}\right)/u_{\sigma}\right]}. \tag{33}$$

Here, η_{PD} is the control-path detection efficiency, u_{th} is the control threshold, u_{σ} sets the softness of the transition, and u_{leak} represents the spurious control activation.

Timing mismatch is included through a normalised overlap factor:

$$\kappa_{ov} = \exp\left[-\left(\frac{\tau_{\mathrm{ctrl,norm}}}{0.40}\right)^2\right], \tag{34}$$

so that the effective EO conversion efficiency becomes

$$\eta_{\mathrm{eff}} = \eta_{\mathrm{EO}} a_{\mathrm{EO}} \kappa_{\mathrm{ov}}. \tag{35}$$

In (32), B_{ctrl} is the normalised control-branch intensity (or binary control-state indicator, according to the implemented notebook), and η_{EO} is the nominal EO conversion efficiency. The value 0.40 in (34) is the adopted

characteristic normalised timing scale of the phenomenological overlap model.

This expression captures the central physical limitation of Fig. 2: even if the EO device is intrinsically efficient, insufficient temporal overlap between the control event and the optical slot can sharply reduce the usable conversion. The upper branch carries the potentially converted A-signal, while the lower branch provides the background V-state support for the zero-output cases. In the final notebook, the upper branch includes finite recovery efficiency after the half-wave plate (HWP), finite vertical-polariser extinction ratio, and branch losses. The lower branch includes the residual V-support channels. Additional penalty terms are introduced at the virtual readout layer to account for timing-induced and EO-efficiency-induced loss of usable H state formation. Accordingly, the model outputs effective intensity-level quantities H_{out} and V_{out} from which the virtual horizontal purity is computed as:

$$\Pi_H = \frac{H_{\mathrm{out}}}{H_{\mathrm{out}} + V_{\mathrm{out}}}. \tag{36}$$

The quantity $\Pi_{H,s}$ is the analogue state-purity metric. It is not itself the probability of logical output 1. For a deterministic parameter realisation, the output bit is $y_s = 1[\Pi_{H,s} > T_\Pi]$; the state-conditioned probability $p_1(s) = P(y_s = 1|s)$ is obtained only over repeated noisy or parameter-perturbed realisations.

The logic decision is then defined by a fixed purity threshold T_Π with the baseline value:

$$T_\Pi = 0.5. \tag{37}$$

Thus, output 1 corresponds to $\Pi_H > T_\Pi$ while output 0 corresponds to $\Pi_H \leq T_\Pi$. This formulation preserves the logic-level interpretation of the EO-controlled stage while remaining fully compatible with the adopted non-coherent intensity-level framework.

The key architecture-specific feasibility quantities are therefore

$$\Pi_{H,HH}, \quad \Pi_{H,0,\max} = \max\left(\Pi_{H,HV}, \Pi_{H,VH}, \Pi_{H,VV}\right), \qquad M_\Pi = \Pi_{H,HH} - \Pi_{H,0,\max}. \tag{38}$$

A parameter set is state-feasible whenever $M_\Pi > 0$. Unlike Fig. 1, the decisive issue here is not analogue threshold spacing of a measured ratio, but preservation of the desired H-state purity relative to all zero-output states under imperfect EO actuation and imperfect timing.

Extended sensitivity sweeps, the Π_H-purity trends, the BER-*vs*.-timing plot, the feasible-region map, and the Monte Carlo histograms are provided in Supplement S5 (Figs. S5.1–S5.6).

4.3. *Sensitivity and robustness under non-idealities*

The non-ideal results show that Fig. 2 remains highly attractive as a cascaded logic stage, but its robustness is governed by a much more specific bottleneck than in Fig. 1. In the exported sensitivity ranking of Supplement S8 (Figs. S8.1–S8.5), the dominant parameter is clearly the normalized control-path mismatch $\tau_{\mathrm{ctrl,norm}}$ whereas all other explored parameters remain secondary in terms of logical correctness.

This dominance is already visible in the dedicated timing sweep of Supplement S5 (Figs. S5.1–S5.6). At the ideal point, $\tau_{\mathrm{ctrl,norm}} = 0$, the architecture retains $\Pi_{H,HH} = 0.97084$, $\Pi_{H,0,\max} = 0.001904$, $M_\Pi = 0.96894$ with zero probability and BER error on the coarse grid. The same remains true up to $\tau_{\mathrm{ctrl,norm}} = 0.4$ for which the state margin is still substantial: $\Pi_{H,HH} = 0.66969$, $\Pi_{H,0,\max} = 0.001497$, $M_\Pi = 0.0.66819$. However, at $\tau_{\mathrm{ctrl,norm}} = 0.5$, the architecture crosses a sharp transition. The desired HH state purity drops below the decision threshold, $\Pi_{H,HH} = 0.48913 < T_\Pi = 0.5$, and the logical behaviour collapses discretely rather than gradually. As a result, $\Delta p_{\max} = 1.0$, $\mathrm{RMSE}_{\mathrm{mean}} = 0.375$, $\mathrm{MAE}_{\mathrm{mean}} = 0.25$, $\mathrm{BER}_{\mathrm{mean}} = 0.25$, $\mathrm{FN}_{\mathrm{mean}} = 1.0$. This failure mode is physically transparent: once the effective timing overlap becomes insufficient, the HH state no longer produces a usable H-dominant output, so the AND response is lost entirely. Importantly, the error is almost purely false-negative rather than false-positive, which is consistent with the state logic of the architecture: the zero-output states do not spuriously generate H, but the desired HH state can fail to convert strongly enough.

The remaining deterministic sensitivity sweeps are much less severe in the explored ranges. Varying the intrinsic EO efficiency over $0.4 \leq \eta_{\mathrm{EO}} \leq 1.0$ does not produce any logical error in the exported summary, while the HH-state purity changes from approximately 0.797 to 0.971 and the purity margin remains positive throughout. Similarly, sweeping the control threshold u_{th}, HWP recovery efficiency, symmetric polarising beam splitter (PBS) leakage, and control leakage does not trigger decision-level failure within the tested domain. Even the finite extinction ratio of the vertical polariser primarily affects the zero-state purity reserve rather than causing outright logical breakdown. In other words, Fig. 2 is not broadly fragile; it is selectively sensitive, with timing overlap acting as the dominant critical parameter.

This selectivity is also reflected in the feasible-region analysis. The two-parameter robustness map in Supplement S5 (Figs. S5.1–S5.6) uses EO efficiency η_{EO} and normalised timing mismatch $\tau_{\mathrm{ctrl,norm}}$ as the comparison axes. Over the tested grid, the architecture is fully feasible for low timing mismatch, but the feasible region contracts sharply once $\tau_{\mathrm{ctrl,norm}}$ exceeds roughly 0.40. For $\eta_{\mathrm{EO}} = 0.4$ and 0.50, full feasibility is retained only up to $\tau_{\mathrm{ctrl,norm}} = 0.2$, whereas for $\eta_{\mathrm{EO}} \geq 0.6$, the feasible plateau extends to $\tau_{\mathrm{ctrl,norm}} = 0.4$. Beyond that boundary, the feasible fraction drops to zero in the exported map. This confirms that Fig. 2 has a meaningful robust window, but it is a window defined primarily by synchronisation quality.

The Monte Carlo statistics provide a compact summary of this behaviour. Across 250 realisations, the average metrics remain acceptable, $\Delta p_{\max,\mathrm{mean}} = 0.0600$, $\mathrm{RMSE}_{\mathrm{mean}} = 0.225$, $\mathrm{MAE}_{\mathrm{mean}} = 0.0150$, $\mathrm{BER}_{\mathrm{mean}} = 0.0150$, $\mathrm{FP}_{\mathrm{mean}} = 0$, $\mathrm{FN}_{\mathrm{mean}} = 0.0600$. The state-purity statistics remain strong on average, $\Pi_{H,HH,\mathrm{mean}} = 0.8017$, $\Pi_{H,0,\max,\mathrm{mean}} = 0.0303$, $M_{\Pi,\mathrm{mean}} = 0.7713$ and the robust fraction is 0.94. At the same time, the distribution is not uniformly smooth: the median logical metrics remain zero, but the upper quantiles include discrete failure cases associated with the timing-overlap collapse described

above. This is why Fig. 2 remains highly usable in practice yet still ranks below Fig. 1 in robustness.

The corresponding generic NPBS sweep for Fig. 2 is reported in Supplement S5 (Fig. S5.6). Over the same diattenuation and retardance ranges, the purity margin remained positive ($0.9636 \le M_H \le 0.9742$). These values represent a model-based sensitivity study rather than characterisation of a specific commercial beam splitter.

4.4. *Interpretation within the comparison framework*

Within the three-way comparison, Fig. 2 occupies a distinct middle position. It is more naturally suited to cascaded functional logic than Fig. 1 because the output state is already polarisation-coded and structurally clean. It is also less dependent on finely tuned nonlinear threshold placement than Fig. 3. However, it trades these advantages for a pronounced dependence on EO timing fidelity and control-path synchronisation.

This result is important for the practical interpretation of the architecture. Figure 2 should not be described as broadly unstable; the deterministic sensitivity sweeps show that many non-ideal parameters leave the logic function intact over surprisingly wide ranges. The real issue is narrower and more physically specific: if the control event is synchronised well enough, the architecture performs cleanly and predictably, but once the overlap falls below a certain level, the desired *HH* response disappears abruptly. Accordingly, Fig. 2 is best viewed as the preferred architecture for structurally clean cascaded stochastic logic, provided that control fidelity and slot synchronisation can be maintained at an adequate level.

5. AOS-based architecture (Fig. 3)

5.1. *Operating principle and nominal behaviour*

In contrast to Fig. 1, where the output decision is extracted from an analogue ratio, and Fig. 2, where the output is formed through EO-controlled polarisation transfer, Fig. 3 relies on threshold-selective activation of two optical switching stages. This gives the architecture its characteristic regenerative behaviour: once the control conditions are satisfied, the output state is reconstructed from calibrated reference channels rather than inherited directly from the input amplitude balance.

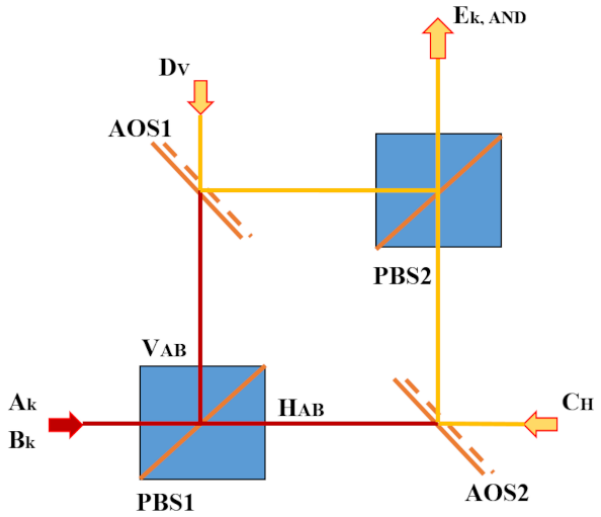


Fig. 3. AOS-based bulk-optics architecture for the stochastic optical AND gate with thresholded nonlinear switching stages.

In the nominal model, the two inputs generate horizontal and vertical control intensities applied to the second and first switching stages, respectively. For a given input slot, the nominal control intensities are determined by the number of horizontally and vertically polarised inputs. With the baseline parameter set, the resulting state-conditioned controls are:

- *HH*: $I_{H,\mathrm{ctrl}} = 2.0$, $I_{V,\mathrm{ctrl}} = 0.0$
- *HV*: $I_{H,\mathrm{ctrl}} = 1.0$, $I_{V,\mathrm{ctrl}} = 1.0$
- *VH*: $I_{H,\mathrm{ctrl}} = 1.0$, $I_{V,\mathrm{ctrl}} = 1.0$
- *VV*: $I_{H,\mathrm{ctrl}} = 0.0$, $I_{V,\mathrm{ctrl}} = 2.0$.

The switching logic is then defined by two thresholds, $I_{\mathrm{th},2}$ or the horizontal AOS stage and $I_{\mathrm{th},1}$ for the vertical AOS stage. In the nominal baseline, $I_{\mathrm{th},2} = 1.5$, $I_{\mathrm{th},1} = 0.5$. Accordingly, only the *HH* state crosses the horizontal threshold strongly enough to activate the *H*-output branch, whereas all non-*HH* states activate only the vertical branch. The nominal truth table exported in Supplement S6 (Figs. S6.1–S6.3) is therefore exact:

- $HH \rightarrow H \rightarrow 1$
- $HV \rightarrow V \rightarrow 0$
- $VH \rightarrow V \rightarrow 0$
- $VV \rightarrow V \rightarrow 0$.

At the output level, the reconstructed reference intensities are $H_{\mathrm{out}} = 1$, $V_{\mathrm{out}} = 1$ for *HH* and $H_{\mathrm{out}} = 0$, $V_{\mathrm{out}} = 1$ for *HV*, *VH*, *VV*. This gives the architecture an ideal nominal contrast between the desired and undesired logical states.

5.2. *Threshold feasibility and nominal stochastic behaviour*

The nominal threshold feasibility of Fig. 3 is governed by two effective margins. Using the notation fixed in section 2, the horizontal switching margin is

$$M_H = \min\left(I_{th,2} - I_{1H},\, I_{2H} - I_{th,2}\right), \tag{39}$$

while the vertical switching margin is

$$M_V = I_{1V} - I_{th,1}. \tag{40}$$

For the nominal baseline, $I_{1H} = 1.0$, $I_{2H} = 2.0$, $I_{1V} = 1.0$, so that $M_H = \min(1.5 - 1.0,\, 2.0 - 1.5) = 0.5$, $M_V = 1.0 - 0.5 = 0.5$. Thus, the baseline thresholds are feasible with equal reserve in the horizontal and vertical control channels.

The nominal threshold window is correspondingly narrower and more structured than in Fig. 1. From the dedicated sensitivity export, the vertical threshold remains feasible only while $I_{\mathrm{th},1} < 1.0$ whereas the horizontal threshold must lie strictly between the single-*H* and double-*H* control levels, $1.0 < I_{\mathrm{th},2} < 2.0$. Within the tested discrete range, this yields feasibility for $0.10 \le I_{\mathrm{th},1} \le 0.8857$ and for $1.2 \le I_{\mathrm{th},2} \le 1.8$, with the widest horizontal margin around the central part of that interval. This immediately shows why Fig. 3 has a narrower robust operating window than Fig. 1: its correctness depends on calibrated placement of two thresholds rather than on a single broad analogue separation window.

At the stochastic level, the nominal result remains practically ideal. Over the dense 11×11 grid used in Supplement S6 (Figs. S6.1–S6.3), the architecture yields $\Delta p_{max} = 0.0109$, $RMSE_{mean} = 0.003255$, $MAE_{mean} = 0.002298$, $BER_{mean} = 0$. Thus, as in Figs. 1 and 2, the nominal stochastic AND law is reproduced correctly; the differences appear only once threshold drift, overlap degradation, and non-ideal switching are introduced.

The detailed nominal truth-flow diagram, threshold-feasibility visualisation, and dense-grid output surface are provided in Supplement S6 (Figs. S6.1–S6.3).

5.3. *Non-ideal model: soft switching, overlap factor, and dynamic constraints*

For the non-ideal analysis, Fig. 3 is modelled by soft threshold switching rather than by an abrupt ideal comparator. If the control intensity at a switching stage is I_{ctrl}, the effective reflectance is written as:

$$R_{eff} = R_{off} + \frac{R_{on} - R_{off}}{1 + \exp\left[-k\left(I_{ctrl} - I_{th}\right)\right]}, \tag{41}$$

where k is the soft-switch steepness parameter, R_{off} and R_{on} are the asymptotic reflectances of the switch in its off and on states, respectively. This formulation preserves the nominal threshold logic while allowing smooth degradation when the control level approaches the switching threshold.

The effective horizontal and vertical output intensities are further modulated by dynamic factors representing finite switching and recovery. In the final notebook, these dynamic penalties are applied multiplicatively to the reconstructed reference outputs. The two channels are then mixed by a small output cross-coupling term ε_{out} after which the final decision variable is the normalised horizontal fraction:

$$q_{AOS} = \frac{H_{out}}{H_{out} + V_{out}}. \tag{42}$$

The output bit is determined by the rule $q_{AOS} > 0.5 \Rightarrow 1$, $q_{AOS} \le 0.5 \Rightarrow 0$. A central non-ideal quantity in Fig. 3 is the overlap factor κ_{AB}, which scales the effective horizontal control level for the *HH* state. In the final notebook, the non-ideal double-*H* control becomes $I_{2H}^{eff} = I_{2H}(1 - \varepsilon_{pbs})\kappa_{AB}$ (where ε_{pbs} is the effective PBS leakage coefficient) so the horizontal threshold margin is directly compressed when the overlap between the two *H*-contributions is reduced. The threshold-feasibility descriptor is therefore:

$$\begin{aligned} M_H &= \min\left(I_{th,2} - I_{1H}^{eff},\ I_{2H}^{eff} - I_{th,2}\right), \\ M_V &= I_{1V}^{eff} - I_{th,1}. \end{aligned} \tag{43}$$

The realisation is threshold-feasible only when both margins are positive.

In addition to threshold feasibility, the robust-comparison protocol also imposes a dynamic acceptability condition. In the implemented comparison logic, the nonlinear sequence is treated as dynamically feasible only when the switching and recovery times remain below the normalised slot duration and when the overlap factor is not too small. Thus, Fig. 3 is evaluated not only as a static threshold gate, but as a threshold gate with an explicitly constrained operating sequence.

5.4. *Sensitivity and robustness under non-idealities*

The exported sensitivity ranking in Supplement S8 (Figs. S8.1–S8.5) identifies κ_{AB} as the dominant non-ideal parameter of Fig. 3. This is the exact analogue of the timing-overlap bottleneck in Fig. 2, but here it acts through the threshold-feasibility reserve of the nonlinear horizontal branch.

The dedicated κ_{AB} sweep in Supplement S7 (Figs. S7.1–S7.5) shows a sharp transition. At $\kappa_{AB} = 1.0$ the baseline margins are preserved, $I_{2H}^{eff} = 2.0$, $M_H = 0.5$, $M_V = 0.5$, and the logical metrics remain ideal. When the overlap decreases to $\kappa_{AB} = 0.8$, the architecture is still threshold-feasible, $I_{2H}^{eff} = 1.6$, $M_H = 0.1$, and the logical metrics remain zero on the tested coarse grid. At $\kappa_{AB} = 0.7$, the coarse-grid logical metrics are still zero, but the threshold reserve has already vanished, $M_H = -0.1$ so the operating point is no longer robust even though no decision-level error is yet observed. This is an important feature of Fig. 3: the loss of robust threshold reserve precedes the loss of nominal logical correctness.

Below that level, the failure becomes explicit. At $\kappa_{AB} = 0.5$, the architecture yields $\Delta p_{max} = 1.0$, $RMSE_{mean} = 0.375$, $MAE_{mean} = 0.25$, $BER_{mean} = 0.25$, $FN_{mean} = 1.0$ with $M_H = -0.5$. The same complete failure persists at $\kappa_{AB} = 0.4$. As in Fig. 2, the dominant breakdown mode is primarily false-negative: the desired *HH* case no longer generates a sufficiently strong regenerated *H*-output once the overlap collapses. The threshold sweeps reveal the second characteristic feature of Fig. 3. Varying $I_{th,1}$ and $I_{th,2}$ over the tested ranges changes threshold feasibility much earlier than it changes the coarse-grid logical metrics. For example, $I_{th,1}$ remains feasible up to approximately 0.8857, after which $M_V < 0$, even though the exported probability and BER metrics remain zero in that sweep. Similarly, $I_{th,2}$ is feasible only over the interval 1.2 to 1.8, while the logical metrics remain zero even at some infeasible edge settings. This means that Fig. 3 can appear nominally correct at isolated points that already lack threshold reserve. For a robustness-based manuscript, this distinction is important and should be stated explicitly. By contrast, the remaining tested non-idealities are comparatively mild in the explored ranges. Scaling R_{on}, scaling R_{off}, output mixing ε_{out}, and PBS leakage ε_{pbs} do not produce logical failure in the exported sweeps. The main effect of ε_{pbs} is gradual reduction of the threshold margins, for example from $M_H = 0.5$, $M_V = 0.5$ at $\varepsilon_{pbs} = 0$ to $M_H = 0.3$, $M_V = 0.4$ at $\varepsilon_{pbs} = 0.1$, while the decision-level metrics remain zero over the tested interval.

Thus, Fig. 3 is not broadly unstable, but it is distinctly more dependent on calibrated threshold placement and overlap preservation than Fig. 1.

5.5. *Feasible-region and Monte Carlo interpretation*

The two-parameter feasible-region analysis in Supplement S7 (Figs. S7.1–S7.5) makes this dependence especially clear. In the exported robustness map, the parameters are

the horizontal threshold $I_{th,2}$ and the overlap factor κ_{AB}. The feasible region forms a narrow wedge-like domain: full feasibility is retained for:

- $I_{th,2}$ =1.2 when $\kappa_{AB} \geq 0.7$
- $I_{th,2}$ =1.4 when $\kappa_{AB} \geq 0.8$
- $I_{th,2}$ =1.6 when $\kappa_{AB} \geq 0.9$
- $I_{th,2}$ =1.8 when $\kappa_{AB} = 1.0$

whereas no robust operating region remains for $I_{th,2} \geq 2.0$. This is a compact quantitative expression of the main conclusion: Figure 3 can operate very cleanly, but only inside a relatively narrow calibrated window.

Extended threshold sweeps, overlap-sensitivity curves, the threshold-margin map, the feasible-region plot, and the Monte Carlo histograms are provided in Supplement S7 (Figs. S7.1–S7.5).

The Monte Carlo summary confirms the same picture. Across 250 realisations, the mean metrics remain reasonably good, $\Delta p_{max} = 0.0120$, $RMSE_{mean} = 0.005236$, $MAE_{mean} = 0.0040$, $BER_{mean} = 0.0040$ with $FP_{mean} = 0.002753$, $FN_{mean} = 0.0080$.

The threshold margins also remain positive on average, $M_H = 0.2318$, $M_V = 0.4805$, but the lower tail of M_H crosses below zero, indicating that a non-negligible fraction of realisations lose horizontal threshold reserve. This is reflected directly in the classification statistics: threshold-feasible fraction = 0.864, dynamic-feasible fraction = 1.0, robust fraction = 0.

An important statistical detail is that the median logical metrics remain zero, as do their upper quartiles. The failures are therefore not gradual broad degradation events, but relatively sparse discrete realisations associated mainly with loss of threshold feasibility in the nonlinear branch. This places Fig. 3 between Fig. 1 and Fig. 2 in practical behaviour: it is more regenerative than both, but less robust than Fig. 1 and somewhat less forgiving than Fig. 2 in its admissible operating window.

5.6. *Interpretation within the comparison framework*

Within the three-way comparison, Fig. 3 is the architecture with the strongest nominal regenerative character. Once the threshold conditions are satisfied, the output is reconstructed from calibrated reference channels and therefore exhibits very clear logical contrast. This makes the architecture conceptually attractive for regenerative nonlinear logic.

Its limitation is equally clear. Robust operation depends on the simultaneous satisfaction of threshold placement and overlaps conditions, and these conditions define a distinctly narrower window than in the metrology-oriented architecture. Moreover, the architecture can remain nominally correct even after its threshold reserve has effectively disappeared, meaning raw truth-table correctness alone is not a sufficient indicator of practical robustness.

For this reason, Fig. 3 should be interpreted not as a generally fragile design, but as a calibrated nonlinear logic stage with strong regenerative output and a relatively narrow admissible operating region. In the present manuscript, temperature sensitivity of the nonlinear switching elements is not included in the main model; it is more appropriately treated as an additional practical limitation and future extension. Under the current comparison protocol, the principal conclusion remains that Fig. 3 is highly attractive when regenerative behaviour is prioritised, but it is less tolerant to threshold and overlap deviations than the metrology-oriented alternative.

6. Cross-architecture comparison

6.1. *Nominal equivalence and architecture-dependent auxiliary descriptors*

Under nominal conditions, all three architectures reproduce the stochastic AND function with the same global probability-level accuracy on the dense test grid. In the regenerated comparative summary, Figs. 1–3 yield identical nominal values: $\Delta p_{max} = 0.0109$, $RMSE_{mean} = 0.003255$, $MAE_{mean} = 0.002298$, $BER_{mean} = 0$. The false-positive and false-negative rates are also zero. Thus, at the level of the nominal stochastic truth law, the three realisations remain equivalent.

This result is important because it prevents a misleading interpretation of the architecture ranking. None of the three designs is preferred because of superior nominal logic fidelity; rather they differ in how that same nominal function is physically realised and in how strongly the realisation degrades once non-idealities are introduced.

The differences are already visible in the auxiliary descriptors associated with each nominal design, as summarised in Table 2. Figure 1 is characterised by a broad threshold-separation reserve, $M = 0.4$, $T_{th,opt} = 0.6$, which reflects the balanced ratio-based state separation of the metrology-oriented scheme. Figure 2 exhibits perfect nominal state purity, $\Pi_{H,HH} = 1.0$, $\Pi_{H,0,max} = 0$, which confirms that the EO-controlled architecture delivers a structurally clean polarisation-coded output under ideal actuation. Figure 3 is characterised by symmetric nominal threshold margins, $M_H = 0.5$, $M_V = 0.5$, which quantify the calibrated reserve of the nonlinear switching thresholds. These auxiliary descriptors are not directly interchangeable, but they play the same conceptual role: each one expresses how far the nominal operating point is from the local feasibility boundary of the corresponding architecture.

Accordingly, the nominal comparison does not establish a winner. Instead, it provides a common starting point for interpreting robustness differences without ambiguity. The nominal cross-architecture agreement is visualised in Fig. 4.

6.2. *Robustness ranking under non-idealities*

Once non-ideal modelling is introduced, the nominal equivalence breaks clearly. The three architectures separate into distinct robustness classes. Figure 1 remains the strongest design by every global robustness indicator. In the Monte Carlo summary, it yields $\Delta p_{max,mean} = 2.52 \times 10^{-4}$, $BER_{mean} = 6.3 \times 10^{-5}$, robust fraction = 1.00, while also retaining an average feasible fraction of 1.0 across the tested feasible-region analysis. This confirms that the metrology-oriented architecture is not merely locally stable at its nominal point, but broadly tolerant to the explored perturbations. Figure 2 occupies the intermediate position. Its average error metrics remain acceptable, $\Delta p_{max,mean} = 0.0600$, $BER_{mean} = 0.0150$, robust fraction = 0.94, yet its average feasible fraction is much smaller, 0.387755,

Table 2.
Nominal cross-architecture comparison of the three stochastic optical AND-gate realisations.

Architecture	Δp_{max}	$RMSE_{mean}$	MAE_{mean}	BER_{mean}	False-positive rate	False-negative rate	Auxiliary descriptor 1	Auxiliary descriptor 2
Figure 1: metrology-oriented AND	0.0109	0.003255	0.002298	0	0	0	$M = 0.4$	$T_{th,opt} = 0.6$
Figure 2: EO-controlled AND	0.0109	0.003255	0.002298	0	0	0	$\Pi_{H,HH} = 1.0$	$\Pi_{H,0,max} = 0$
Figure 3: AOS-based AND	0.0109	0.003255	0.002298	0	0	0	$M_H = 0.5$	$M_V = 0.5$

While the global nominal metrics are identical, the auxiliary descriptors already indicate that the three designs rely on different physical mechanisms and different notions of operating margin.

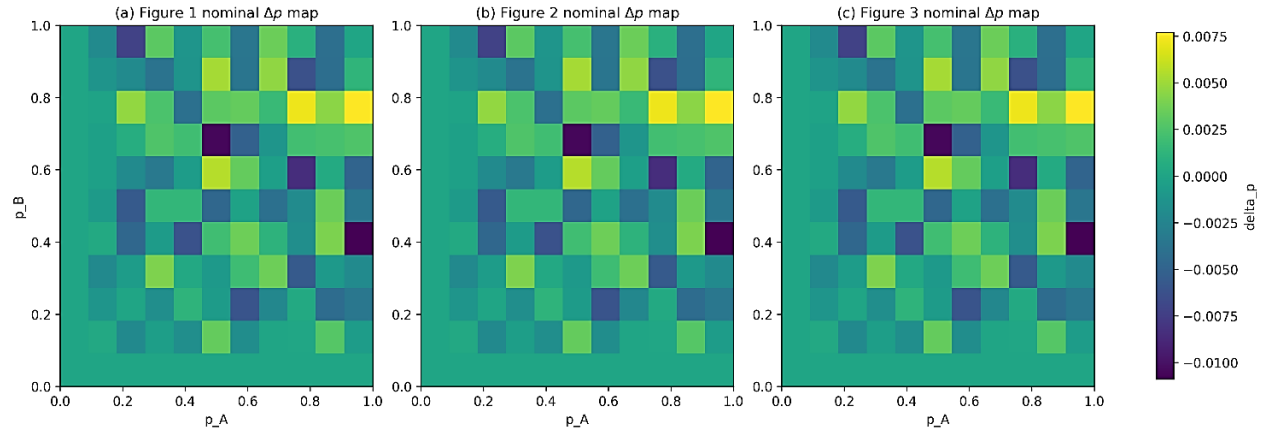


Fig. 4. Nominal cross-architecture comparison of the three stochastic optical AND-gate realisations on the dense probability grid: (a) metrology-oriented architecture, (b) EO-controlled architecture, and (c) AOS-based architecture. The identical Δp maps confirm that all three implementations reproduce the same nominal stochastic AND behaviour.

because its usable parameter domain contracts sharply once timing mismatch becomes appreciable. The exported sensitivity ranking identifies the dominant controlling parameter unambiguously as $\tau_{ctrl,norm}$, while the remaining explored non-idealities do not induce comparable logical degradation within the tested ranges. Thus, Fig. 2 is not broadly unstable, but selectively vulnerable to synchronisation failure. Figure 3 follows next. Its mean logical metrics are better than those of Fig. 2 and yield $\Delta p_{max,mean} = 0.0120$, $BER_{mean} = 0.0040$, but its robust fraction is lower, 0.86, and its average feasible fraction is the smallest among the three architectures, 0.178571. This combination is revealing. It means that once Fig. 3 operates inside its admissible region, it performs very cleanly, but the admissible region itself is narrower than for the other two realisations. The sensitivity ranking identifies κ_{AB}, as the dominant parameter, showing that the principal robustness bottleneck is preservation of the overlap-dependent horizontal threshold reserve. Taken together, the three-way comparison establishes the following robustness ordering: Figure 1 > Figure 2 > Figure 3 with respect to overall tolerance and interpretability of the operating region. At the same time, the ranking is not one-dimensional. Figure 3 can outperform Fig. 2 in mean error once it remains inside its calibrated switching window, but it is less forgiving in terms of feasible parameter space. Therefore, the comparison must be read in terms of both accuracy and admissible operating domain, not by a single scalar metric alone. The reported feasible fractions are highly informative for comparative purposes, but they should not be interpreted as universal physical yield values, because the two-parameter sweep axes differ between architectures.

Their role is comparative and design-oriented: they show how much of the tested architecture-specific parameter plane remains acceptable under the same global accuracy criteria. The corresponding robust operating windows under non-idealities are compared in Fig. 5.

6.3. *Dominant physical trade-offs*

The comparison also reveals that each architecture is controlled by a different physical trade-off. For Fig. 1, the dominant issue is threshold placement, but this sensitivity is embedded in a wide analogue separation plateau. As a result, the threshold is influential yet not intrinsically fragile. This is why the architecture combines high observability with high robustness. For Fig. 2, the dominant issue is EO timing fidelity. The architecture retains clean polarisation-coded logic provided that control activation and signal-slot overlap are sufficiently synchronised. Once this condition is violated, the desired *HH* response collapses rapidly, typically through false-negative failure rather than false-positive contamination.

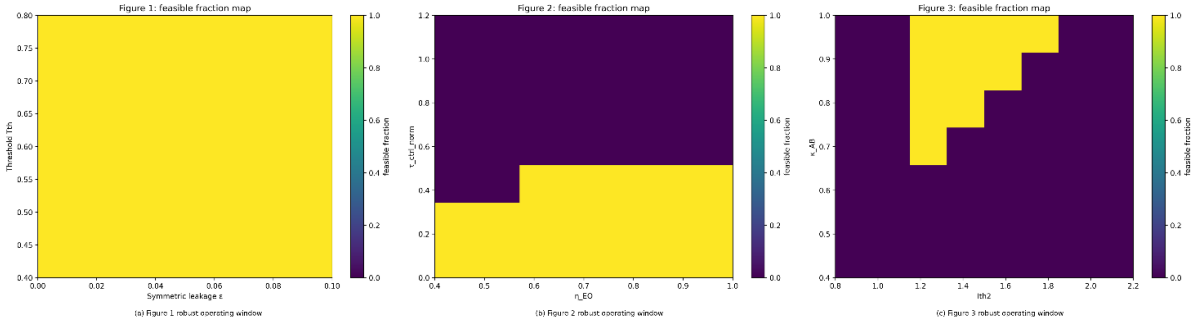


Fig. 5. Cross-architecture comparison of robust operating windows under non-idealities: (a) metrology-oriented architecture, (b) EO-controlled architecture, and (c) AOS-based architecture. The three panels show that the metrology-oriented design retains the broadest robust region, whereas the EO-controlled and AOS-based architectures operate within narrower architecture-specific admissible windows.

Table 3.
Monte Carlo robustness summary, dominant sensitivity, and recommended primary use case for the three stochastic optical AND-gate architectures.

Architecture	$\Delta p_{max,mean}$	BER_{mean}	Robust fraction	Average feasible fraction	Dominant sensitivity	Dominant strength	Dominant limitation	Recommended primary use case
Figure 1: metrology-oriented AND	0.000252	0.00006	1.00	1.000000	Threshold setting	Explicit output observability and threshold analysis	Output decision depends on analogue ratio separation	Metrology, uncertainty analysis, teaching demonstration
Figure 2: EO-controlled AND	0.0600	0.0150	0.94	0.387755	$\tau_{ctrl,norm}$	Structurally clean cascaded polarisation-coded output	Sensitive to EO fidelity, timing, and polariser quality	Functional cascaded logic module
Figure 3: AOS-based AND	0.0120	0.0040	0.86	0.178571	κ_{AB}	Regenerative output from calibrated reference channels	Sensitive to threshold placement and switching dynamics	Regenerative nonlinear logic

For Fig. 3, the dominant issue is maintenance of nonlinear threshold reserve under overlap degradation. Here, the architecture can remain nominally correct even after the reserve has vanished, making threshold-feasibility analysis indispensable. In this sense, Fig. 3 most clearly demonstrates why truth-table correctness alone is an insufficient robustness criterion for stochastic optical logic. These trade-offs explain why the three designs are complementary rather than redundant. Figure 1 is the strongest where calibration transparency and confidence in decision separation matter most. Figure 2 is the strongest where structurally clean cascaded propagation is the priority. Figure 3 is the strongest where regenerative nonlinear output formation is valued more than broad tolerance margins. The main robustness outcomes, dominant sensitivity parameters, and practical design roles are summarised in Table 3. Table 3 makes clear that the preferred architecture depends on the intended design objective: broad tolerance and interpretability favour Fig. 1, cascaded propagation favours Fig. 2, and regenerative nonlinear operation favours Fig. 3.

6.4. Practical use-case positioning

From a practical design perspective, the comparative assessment can be summarised as follows.

The metrology-oriented architecture is the preferred choice for calibration-aware experiments, uncertainty analysis, and teaching-oriented demonstrations, because its decision variable is directly observable and its operating window remains the broadest under non-idealities.

The EO-controlled architecture is the preferred choice for functional cascaded stochastic logic, because the output state remains structurally polarisation-coded and therefore ready for propagation into subsequent stages. Its main implementation requirement is adequate EO control fidelity and timing alignment. The AOS-based architecture is the preferred choice for regenerative nonlinear logic, because it reconstructs a strong output state from calibrated switching channels.

The regenerative advantage of Fig. 3 is obtained at the cost of narrower threshold and overlap tolerances, so this architecture is best suited for designs in which careful operating-point calibration is acceptable.

In this way, the comparison does not merely rank the three architectures; it assigns each of them a distinct design role.

7. Discussion

The comparative results presented above allow the three architectures to be interpreted not only as alternative

realisations of the same stochastic AND law, but as designs optimised for different physical priorities.

This comparison shows that nominal truth-table correctness is necessary but insufficient for evaluating stochastic optical logic. All three architectures considered here reproduce the same stochastic AND law under nominal conditions, and their global nominal metrics are identical within the adopted dense-grid evaluation. If we were to conclude the analysis at that point, the three realisations would appear practically equivalent. However, the non-ideal and Monte Carlo results show that this conclusion would be misleading. Once realistic perturbations are introduced, the architectures separate clearly according to the physical mechanism by which the logical decision is formed and maintained.

The most important general outcome is that practical equivalence in stochastic optical logic is determined by robustness of the decision mechanism, not only by nominal correctness of the Boolean mapping. In Fig. 1, the decision mechanism is based on the analogue ratio separation with a broad threshold reserve. In Fig. 2, it is based on the preservation of a sufficiently pure *H*-dominant output under EO actuation and timing overlap. In Fig. 3, it is based on the calibrated nonlinear threshold crossing and preservation of overlap-dependent switching reserve. These are fundamentally different forms of physical decision formation, and the robustness ranking obtained in this work is a direct consequence of that difference rather than of any difference in nominal logic functionality.

This distinction is summarised quantitatively in Fig. 6 and Table 3.

From this perspective, the superiority of Fig. 1 in the robustness analysis is physically intuitive. The metrology-oriented architecture does not depend on a narrowly tuned conversion event or on simultaneously satisfying multiple switching thresholds. Instead, it relies on an observable analogue variable with a comparatively wide acceptable interval. This produces two advantages at once: the logical state remains experimentally interpretable, and the degradation under non-idealities is gradual rather than abrupt. The Monte Carlo statistics confirm that this architecture remains deep inside its acceptable region for the tested perturbations, which explains why it is the preferred solution when calibration, uncertainty evaluation, and confidence in state separation are central priorities.

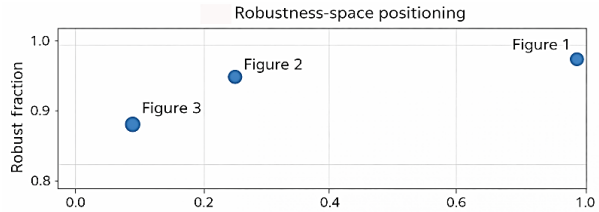


Fig. 6. Final robustness-space positioning of the three stochastic optical AND-gate architectures in terms of average feasible fraction and robust fraction. Metrology-oriented architecture (Fig. 1) occupies the most favourable region and is therefore the preferred choice for metrology, uncertainty analysis, and teaching-oriented validation; its dominant sensitivity is threshold setting. The EO-controlled architecture (Fig. 2) remains attractive for functional cascaded logic modules, with robustness governed primarily by the normalised control-path mismatch $\tau_{ctrl,norm}$. The AOS-based architecture (Fig. 3) is best suited to regenerative nonlinear logic, with robustness dominated by the overlap factor κ_{AB}.

Figure 2 illustrates a different design philosophy. Its main strength is the structural cleanliness of the output state: under correct operation, the output is already polarisation-coded and therefore naturally suited for cascading. This is an important architectural advantage if the gate is to be used not only as an isolated demonstrator, but as a building block of a larger stochastic optical chain. At the same time, this advantage depends specifically on control fidelity. The analysis shows that the dominant failure mechanism is not a broad sensitivity to all imperfections, but a comparatively sharp loss of the desired *HH* response once the timing overlap between the control event and the optical slot becomes insufficient. Thus, Fig. 2 should be understood as a design with a clean functional output and a selective synchronisation bottleneck, rather than as a generally fragile architecture.

Figure 3, in turn, is the architecture that most clearly exhibits regenerative logic behaviour. Its output is reconstructed through thresholded switching channels rather than by a direct analogue separation or direct EO transfer, and this gives it strong nominal contrast once the switching conditions are satisfied. This is conceptually attractive for nonlinear optical logic, especially when output regeneration is valued. However, the results also show why regenerative behaviour should not be confused with broad robustness. The admissible region of Fig. 3 is narrower because correct operation depends simultaneously on threshold placement, effective overlap, and dynamic feasibility. Moreover, the architecture can remain nominally correct even after its threshold reserve has already been lost, which means that truth-table verification alone can overestimate practical stability. For this architecture, threshold-feasibility analysis is therefore not an auxiliary refinement, but an essential part of the validity assessment.

A broader methodological implication of the study is that stochastic optical gates should be compared in a two-layer framework. The first layer is nominal correctness: whether the gate reproduces the intended stochastic truth law. The second layer is operational validity: whether the gate preserves that behaviour inside a non-negligible and physically interpretable operating region. The present work shows that these two layers can lead to different conclusions. Figures 2 and 3 both remain highly attractive architectures within their intended use cases, but their admissible windows are narrower and more mechanism-dependent than that of Fig. 1. Accordingly, future work on stochastic optical logic should avoid ranking designs solely by nominal truth tables or isolated output examples.

Another useful outcome is that the comparison clarifies the proper role of supplementary material in a study of this kind. The main manuscript can remain focused on the architecture definitions, unified comparison protocol, key robustness figures, and final design interpretation. Extended threshold sweeps, detailed sensitivity maps, parameter manifests, Monte Carlo histograms, and reproducibility artifacts are more appropriately placed in the supplementary package. This division is especially important here because the scientific value of the paper lies not in the sheer

volume of parameter scans, but in the conversion of those scans into a coherent cross-architecture design interpretation. The study should also state several limitations explicitly.

One component-level limitation concerns the polarisation response of real non-polarising beam splitters used in Fig. 1 and Fig. 2. Such elements may exhibit both polarisation-dependent amplitude imbalance and differential phase shift between the H and V components, particularly in a reflected branch.

To quantify the resulting effect without assuming a specific commercial component, we performed a generic fully polarised Jones-model sweep over an intensity diattenuation ratio $0.8 \le D = \frac{|a_V|^2}{|a_H|^2} \le 1.2$ and a differential phase $0° \le \delta \le 180°$. The results are summarised in Fig. 7, while the complete numerical datasets are provided in Supplements S3 and S5.

As shown in Figs. 7(a) and 7(c), a differential phase can generate appreciable ellipticity when both polarisation components are present: the maximum ellipticity angle is approximately 25.6° for the representative Fig. 1 *HV* state and 9.8° for the Fig. 2 *HH* output state.

However, pure differential phase does not by itself change the immediately resolved H/V intensities in analyser axes aligned with the Jones basis. Figures 7(b) and 7(d) show that the corresponding decision quantities are governed primarily by NPBS diattenuation: over the investigated range $0.8 \le D \le 1.2$, Fig. 1 decision margin remains positive ($0.362 \le M \le 0.400$) and the Fig. 2 purity margin also remains positive ($0.9636 \le M_{\Pi} \le 0.9742$).

If a larger component-specific ellipticity is encountered experimentally, a clean-up linear polariser may be placed after the affected NPBS branch and before the polarisation-sensitive decision element; its projection loss must then be incorporated into the relevant effective branch coefficient.

Three broader limitations should also be noted. First, the comparison is intentionally model-based and architecture-aware rather than device-specific. The purpose is to identify how different physical decision principles behave under structured perturbations, not to optimise a single laboratory implementation for maximum experimental performance. Second, the EO-controlled architecture is described in the final manuscript by a non-coherent intensity-level model. This deliberate choice isolates the logic-relevant performance of the stage without introducing phase-sensitive interference effects that are not required for the present comparison. Third, the AOS-based analysis does not include temperature sensitivity of the nonlinear elements in the main model. This omission defines a boundary of the current scope rather than implying that temperature is irrelevant. In practical nonlinear implementations, thermal stabilisation may become an important additional requirement and should be addressed in future work.

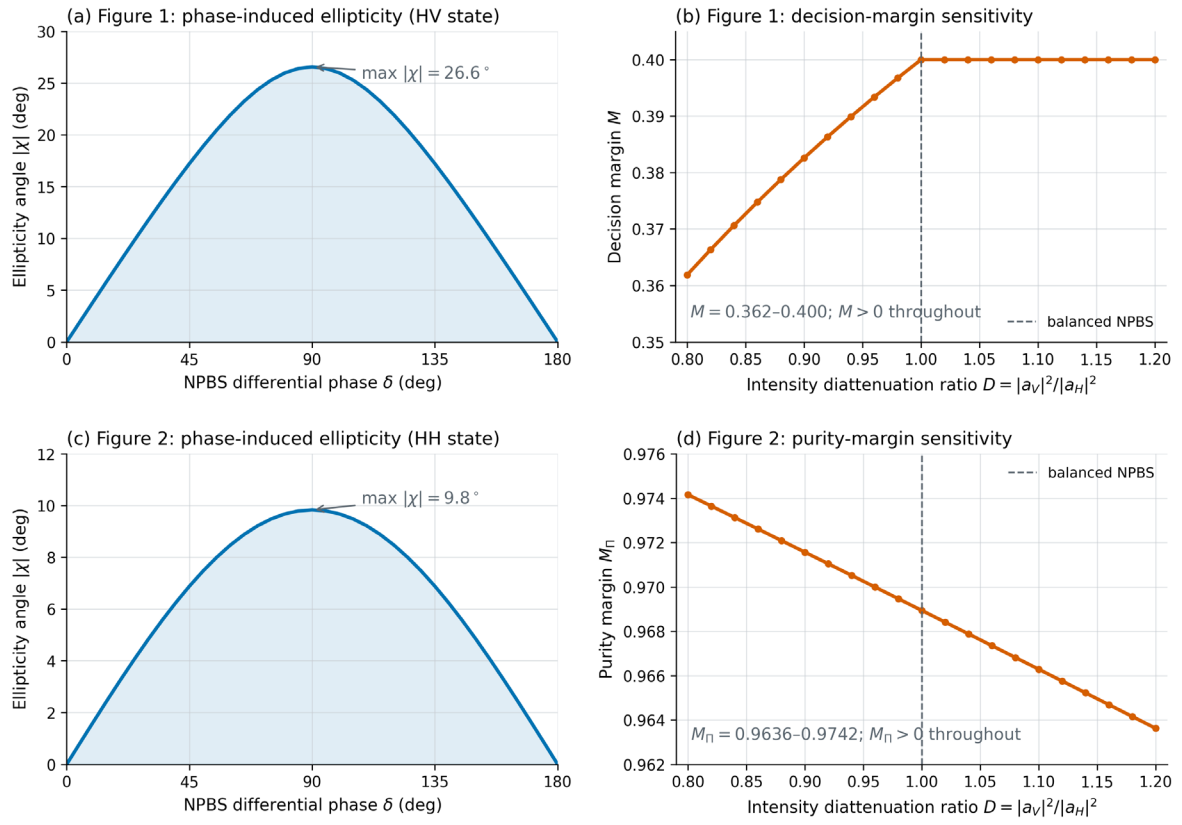


Fig. 7. Generic NPBS polarisation-sensitivity analysis for metrology-oriented and EO-controlled architectures. (a) Phase-induced ellipticity of the mixed *HV* state in Fig. 1 for a balanced NPBS ($D = 1$). (b) Figure 1 decision margin M *vs.* the intensity diattenuation ratio $D = \frac{|a_V|^2}{|a_H|^2}$. (c) Phase-induced ellipticity of the desired *HH* output state in Fig. 2 for $D = 1$. (d) Figure 2 horizontal-purity margin M_{Π} *vs. D.* The differential-phase sweep covers $0° \le \delta \le 180°$ and the diattenuation sweep covers $0.8 \le D \le 1.2$. The calculation is vendor-independent and represents a generic fully polarised limiting case rather than characterisation of a specific commercial beam splitter.

The comparison also points directly to possible optimisation directions. For Fig. 2, one natural extension is deliberate asymmetric strengthening of the *H*-branch in order to enlarge the purity margin of the desired *HH* state and thereby relax the effective timing bottleneck. In the present work, this possibility is not included in the core model, because the goal is to compare the three architectures in their fixed final forms. Nevertheless, it represents a plausible engineering route for improving the robustness of the EO-controlled design. For Fig. 3, a natural continuation is inclusion of explicit temperature-dependent switching variation and possibly broader dynamic non-idealities, so that the robustness window of the nonlinear architecture can be evaluated under more realistic environmental conditions.

Taken together, these observations define the practical meaning of the three-way comparison. Figure 1 is the architecture of choice when robustness, interpretability, and calibration transparency are prioritised. Figure 2 is the architecture of choice when a structurally clean cascaded stochastic logic stage is needed and synchronisation quality can be maintained. Figure 3 is the architecture of choice when regenerative nonlinear output formation is the main objective and operation inside a narrower calibrated window is acceptable. The main contribution of the paper is therefore not a universal declaration of one "best" architecture, but a physically grounded explanation of why different architectures become best under different design priorities.

These design-oriented distinctions are condensed into the final conclusions below.

8. Conclusions

This paper compared three bulk-optics realisations of a stochastic optical AND gate within a unified analytical framework that combines probability-level accuracy, decision-level correctness, feasibility margins, and Monte Carlo robustness. Although the three architectures are physically different, they were shown to be nominally equivalent at the level of the stochastic AND law.

In all three cases, the nominal truth table is reproduced correctly, and the dense-grid nominal metrics confirm essentially ideal stochastic behaviour.

The main result of the study is that this nominal equivalence does not survive robustness analysis. Once realistic non-idealities are introduced, the three architectures separate clearly according to the physical mechanism by which the logical decision is formed. The metrology-oriented architecture remains the most robust because it relies on an observable analogue decision variable with a broad threshold-separation reserve. The EO-controlled architecture preserves a structurally clean polarisation-coded output and is therefore the most suitable for cascaded stochastic logic, but its performance is limited mainly by control-path synchronisation and timing overlap. The AOS-based architecture provides the strongest regenerative nonlinear output formation, yet its admissible operating window is narrower because correct behaviour depends more critically on threshold placement and overlap-dependent switching reserve.

These results establish a design-oriented interpretation of the three architectures. Figure 1 is the preferred solution when calibration transparency, uncertainty-aware validation, and broad tolerance margins are the main priorities. Figure 2 is the preferred solution when the gate must function as a clean cascaded logic stage with preserved polarisation-coded output. Figure 3 is the preferred solution when regenerative nonlinear behaviour is desired and operation inside a more narrowly calibrated window is acceptable.

A broader methodological conclusion is that stochastic optical gates should not be compared only by nominal truth-table correctness. For practical implementation, the decisive issue is whether correct stochastic behaviour is preserved inside a physically meaningful and sufficiently wide operating region. The unified framework used here provides such a comparison space and makes it possible to evaluate architectures with different internal physics on a common basis without forcing them into an artificial common mechanism.

The present study also defines several natural directions for further work. For the EO-controlled architecture, robustness may be improved by engineered enhancement of the desired *H*-branch response. For the AOS-based architecture, future extensions should include temperature-dependent switching variation and broader environmental sensitivity of the nonlinear elements. More generally, the comparison framework developed here can be extended to other stochastic optical logic functions, including OR, NOT, and larger cascaded stochastic photonic modules.

In summary, the contribution of this work is twofold: first, it demonstrates that three distinct bulk-optics architectures can realise the same stochastic optical AND function under nominal conditions; second, it shows that their practical value is determined by fundamentally different robustness profiles. This turns the comparison from a purely functional exercise into a design guide for selecting the most suitable stochastic optical AND architecture according to the intended experimental or system-level objective.

The supplementary package provides the extended sensitivity maps, Monte Carlo distributions, and reproducibility artifacts supporting these conclusions.

Authors' statement

All authors contributed substantially to the conception, design, and completion of this work: research concept and design; collection and assembly of data; data analysis and interpretation; writing the article, M.S.; critical revision of the article and final approval of the article, O.A., C.Z., R.Z., B.T., X.Z., J.Z., C.J.
All authors discussed the results, revised the manuscript collaboratively, and approved the final submitted version.
The corresponding author served as the main contact with the journal during the submission and review process.

Funding

This research received no external funding.

Ethics statement

This study did not involve human participants, human data, or animal subjects, and therefore did not require ethical approval.

Supplementary material

The supplementary package includes the reproducible Colab notebook, an executable Python mirror, parameter-sweep datasets, Monte Carlo outputs, boundary-analysis results, and practical alignment/calibration protocols used to generate the figures and tables in this work. It contains the workflows corresponding to Supplements S1–S9 and the associated CSV/JSON data files. The complete package is provided as Supplementary.zip.

Data and code availability statement

All data and code supporting the results reported in this study are provided in the supplementary material accompanying this submission. The package includes the Colab notebooks used for simulation, exported datasets (CSV/JSON), and configuration manifests enabling full reproduction of the reported results.

Reproducibility statement

All numerical results reported in this study can be reproduced using the provided Colab notebooks and configuration files included in the supplementary materials. The workflows use fixed random seeds and documented parameter settings, ensuring that the figures and tables can be regenerated from the distributed datasets and scripts.

References

[1] Gaines, B. R. Stochastic Computing Systems. in *Advances in Information Systems Science* **2**, 37–172 (ed. Tou, J. T.) (Springer, 1969). https://doi.org/10.1007/978-1-4899-5841-9_2

[2] Alaghi, A. & Hayes, J. P. Survey of stochastic computing. *ACM Trans. Embed. Comput. Syst.* **12**, 92 (2013). https://doi.org/10.1145/2465787.2465794

[3] Alaghi, A., Qian, W. & Hayes, J. P. The promise and challenge of stochastic computing. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst*. **37**, 1515–1531 (2018). https://doi.org/10.1109/TCAD.2017.2778107

[4] Lohmann, A. W. Polarization and optical logic. *Appl. Opt.* **25**, 1594–1597 (1986). https://doi.org/10.1364/AO.25.001594

[5] Minzioni, P. *et al*. Roadmap on all-optical processing. *J. Opt.* **21**, 063001 (2019). https://doi.org/10.1088/2040-8986/ab0e66

[6] Zaghloul, Y. A. & Zaghloul, A. R. M. Unforced polarization-based optical implementation of Binary logic. *Opt. Express* **14**, 7252–7269 (2006). https://doi.org/10.1364/OE.14.007252

[7] Guo, L.-Q. *et al*. All-optical AND gate with improved extinction ratio using signal-induced nonlinearities in a bulk semiconductor optical amplifier. *Opt. Express* **14**, 2938–2943 (2006). https://doi.org/10.1364/OE.14.002938

[8] Lin, C.-E., Lu, Y.-H., Zhou, M.-T. & Chen, C.-C. Reconfigurable electro-optical logic gates using a 2-layer multilayer perceptron. *Sci. Rep.* **12**, 14203 (2022). https://doi.org/10.1038/s41598-022-18408-0

[9] Yanik, M. F., Fan, S., Soljačić, M. & Joannopoulos, J. D. All-optical transistor action with bistable switching in a photonic crystal cross-waveguide geometry. *Opt. Lett.* **28**, 2506–2508 (2003). https://doi.org/10.1364/OL.28.002506

[10] Zamora-Munt, J. & Masoller, C. Numerical implementation of a VCSEL-based stochastic logic gate via polarization bistability. *Opt. Express* **18**, 16418–16429 (2010). https://doi.org/10.1364/OE.18.016418

[11] Salvide, M. F., Masoller, C. & Torre, M. S. All-optical stochastic logic gate based on a VCSEL with tunable optical injection. *IEEE J. Quantum Electron.* **49**, 886–893 (2013). https://doi.org/10.1109/JQE.2013.2276122

[12] Angelsky, O. V., Bekshaev, A. Y., Zenkova, C. Y., Ivansky, D. I. & Zheng, J. Correlation optics, coherence and optical singularities: Basic concepts and practical applications. *Front. Phys.* **10**, 924508 (2022). https://doi.org/10.3389/fphy.2022.924508

[13] Angelsky, O. V. *et al*. Optical phase singularities: Physical nature, manifestations and applications. *Front. Phys.* **10**, 1060787 (2022). https://doi.org/10.3389/fphy.2022.1060787

[14] Angelsky, O. V. *et al*. Random object optical field diagnostics by using carbon nanoparticles. *Opt. Express* **29**, 916–928 (2021). https://doi.org/10.1364/OE.411118