

Risk-constrained guard bands for flexible-grid optical links

Keshav Kumar,^{*,1} Mohit Kumar Srivastava¹ and Man Mohan Shukla²

¹Department of ECE, Pranveer Singh Institute of Technology, Kanpur-209305, UP, India

²Pranveer Singh Institute of Technology, Kanpur-209305, UP, India

Received August 29, 2026; accepted October 10, 2026; published October 17, 2026

Abstract—Guard bands in flexible-grid links are usually fixed or selected from hop count or OSNR, although spectral overlap is ultimately driven by uncertain transmitter and filtering frequency errors. We formulate a risk-constrained rule that converts monitored laser and wavelength-selective-switch error variances into the smallest 6.25-GHz-grid margin satisfying a prescribed overlap probability. For 32-GBd, 0.1-roll-off channels over 1–20 switching stages, the closed-form rule matches 10^7 -trial Monte Carlo simulation within 3.1%. At a 10^{-3} risk target it reduces the mean occupied spectral span by 13.1% and 11.1% for 0.5- and 1-GHz laser-error standard deviations, relative to a fixed 12.5-GHz guard.

Elastic optical networks (EONs) replace rigid wavelength channels with demand-matched frequency slots [1, 2]. Their packing gain depends on the granularity of the spectrum: smaller slots reduce stranded bandwidth but increase switching and control complexity [3]. A guard band must nevertheless separate adjacent allocations. An excessive guard consumes capacity, whereas an insufficient one allows transmitter drift, finite spectral roll-off, and wavelength-selective-switch (WSS) edge errors to close the nominal gap.

Optimized guard-band planning has traded spectrum against regeneration [4]. Adaptive schemes have used hop count with spectral equalization [5], network occupancy, or OSNR margin [6]. Cascaded WSS filtering is known to narrow passbands [7], and measurements on an installed multi-vendor network showed crosstalk when a small channel-frequency offset coincided with weak isolation at a WSS edge [8]. Laser-frequency drift is time-varying and is therefore monitored in coherent receivers [9]. These results motivate adaptation, but the cited selectors do not map monitored frequency-error statistics onto an explicit probability of spectral overlap.

This Letter introduces a risk-constrained guard-band (RCGB) rule with three features: the risk target is interpretable, the computation is closed-form, and the result is quantized directly to the flexible-grid frequency granularity. The rule is a spectral-safety test performed before a conventional modulation- and OSNR-aware quality-of-transmission (QoT) check. Figure 1 shows where it sits in the admission path, and Algorithm 1 states it in executable form.




Fig. 1. Flowchart of the proposed risk-constrained guard-band rule within the connection-admission path. The rule executes before the OSNR-aware QoT check.

Consider two adjacent root-raised-cosine channels with symbol rate R_s and roll-off α , so that the occupied bandwidth is $B = (1 + \alpha)R_s$. Let g be the nominal empty separation between their spectral supports, and let the positive random variable X_h be the amount by which that separation is closed along a route containing h WSS stages. If the two residual laser offsets are independent with standard deviation σ_L , and each stage contributes an independent equivalent spectral-edge error σ_W , the aggregate variance is

$$\sigma_h^2 = 2\sigma_L^2 + h\sigma_W^2 \quad (1)$$

For unequal lasers, the first term in Equation (1) becomes $\sigma_{L,1}^2 + \sigma_{L,2}^2$. Under a zero-mean Gaussian approximation a spectral-overlap event occurs when $X_h > g$, hence

$$p_{ov}(g | h) = \Pr\{X_h > g\} = Q\left(\frac{g}{\sigma_h}\right) \quad (2)$$

Let ε be the permitted one-sided overlap probability and q the guard-band quantum. The smallest grid-compatible guard satisfying $p_{ov} \leq \varepsilon$ is

*keshav@gyancity.com



$$g_h^* = q \left\lceil \frac{Q^{-1}(\varepsilon)\sigma_h}{q} \right\rceil \quad (3)$$

Equation (3) gives the guarantee by construction. A controller needs only h and rolling variance estimates from transceiver and WSS telemetry, so the per-request cost is $O(1)$, as set out in Algorithm 1. If the errors are correlated or heavy-tailed, $Q^{-1}(\varepsilon)\sigma_h$ may be replaced by the measured empirical $(1 - \varepsilon)$ quantile without changing the quantization step.

For reporting, the fractional reduction in occupied spectral span relative to a fixed guard g_{fix} is, in the many-channel limit where the per-channel span is $B + g$,

$$\Delta S = \frac{g_{\text{fix}} - \bar{g}}{B + g_{\text{fix}}} \quad (4)$$

where $\bar{g}$ with an overbar is the mean guard selected by Equation (3) over the path-length mix of interest.

Algorithm 1. Risk-constrained guard-band (RCGB) selection.

Input: $h, \sigma_L, \sigma_W, q, \varepsilon, g_{\text{max}}$
Output: guard band g_h^* or a blocking decision
1: $\sigma_h \leftarrow (2\sigma_L^2 + h\sigma_W^2)^{1/2}$ $\triangleright$ Equation (1)
2: $n \leftarrow \lceil Q^{-1}(\varepsilon)\sigma_h / q \rceil$ $\triangleright$ quantize to grid
3: $g_h^* \leftarrow nq$ $\triangleright$ Equation (3)
4: if $g_h^* > g_{\text{max}}$ then return BLOCK
5: if no free slot of width $B + g_h^*$ then return REROUTE
6: return g_h^* to the OSNR-aware QoT check

Table 1 lists the simulation parameters. The 6.25-GHz quantum follows the nominal-central-frequency granularity of the flexible grid [1]; a fixed 12.5-GHz guard is the conservative baseline and a fixed 6.25-GHz guard the aggressive one. The value of σ_W is an equivalent sensitivity parameter rather than a specification for a particular WSS. For every h , 10^7 independent Gaussian trials were drawn from a single seeded stream and compared with Equation (2).

Table 1. Simulation parameters.

Parameter	Symbol	Value
Symbol rate	R_s	32 GBd
Roll-off factor	α	0.10
Occupied bandwidth	B	35.2 GHz
Guard-band quantum	q	6.25 GHz
WSS stages	h	1–20
Per-stage WSS edge error	σ_W	0.35 GHz
Laser frequency error	σ_L	0.5, 1.0, 1.5 GHz
Overlap risk target	ε	10^{-3}
Trials per stage count	N	10^7
Random seed	—	20260828

Figure 2(a) shows the quantized decision. For $\sigma_L = 0.5$ GHz the rule selects 6.25 GHz on every path. For $\sigma_L = 1$ GHz it holds 6.25 GHz through $h = 17$ and steps to 12.5

GHz at $h = 18$. At $\sigma_L = 1.5$ GHz, 12.5 GHz is required throughout, so the rule reverts to the conservative allocation once the measured uncertainty no longer allows for safe packing gain.




Fig. 2. (a) Grid-quantized guard band selected by Eq. (3) against the number of WSS stages. (b) Spectral-overlap probability for $\sigma_L = 1$ GHz; lines are analytical and open circles are Monte Carlo estimates (10^7 trials per h).

Figure 2(b) explains why a single small guard is unsafe at $\sigma_L = 1$ GHz. The analytical overlap probability for a 6.25-GHz guard rises from 8.93×10^{-6} at $h = 1$ to 1.52×10^{-3} at $h = 20$ and crosses the target after 17 stages. RCGB tracks that curve only while feasible; its 12.5-GHz decision then reduces the risk to at most 1.56×10^{-9} , more than five orders of magnitude below the target. Analysis and simulation agree to within 3.1% wherever $p_{\text{ov}} \geq 10^{-4}$, the residual deviation being consistent with the sampling error of 10^7 trials.

Table 2 summarizes the three laser-error cases. Assuming a uniform path-length mix only to summarize packing, Equation (4) gives mean-span reductions of 13.1%, 11.1% and 0% against a fixed 12.5-GHz guard, while the worst-case risk stays at or below the target in every case. The allocation for each actual path is independent of this reporting assumption.

Table 2. Worst-case spectral-overlap probability over $h = 1$ –20 and mean spectral-span reduction against a fixed 12.5-GHz guard.

σ_L (GHz)	Fixed 6.25 GHz	RCGB	Mean g_h^* (GHz)	Span reduction
0.5	1.37×10^{-4}	1.37×10^{-4}	6.25	13.1%
1.0	1.52×10^{-3}	9.90×10^{-4}	7.19	11.1%
1.5	8.88×10^{-3}	1.06×10^{-6}	12.50	0%

In conclusion, a fixed guard cannot simultaneously exploit stable links and protect uncertain ones. The proposed RCGB rule converts measured frequency uncertainty into the minimum grid-compatible margin for a chosen overlap risk, and, in the modeled 32-GBd links, it preserves a 10⁻³ constraint while reducing the mean spectral span by up to 13.1%. The rule is a transparent safety layer, not a complete BER model: spectral overlap is a conservative trigger, and ASE noise, Kerr nonlinearity, passband concatenation, modulation format, and receiver DSP still require a QoT estimator. Independence and Gaussianity may fail under common frequency locking or heavy-tailed disturbances; σ_w must be calibrated for the deployed WSS population, and laboratory validation is required before operational use. Future work should combine the rule with impairment-aware RMLSA and test non-Gaussian telemetry on a reconfigurable optical testbed.

References

- [1] ITU-T Recommendation G.694.1, *Spectral grids for WDM applications: DWDM frequency grid* (ITU-T, Geneva 2020). <https://handle.itu.int/11.1002/1000/14498>.
- [2] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, S. Matsuoka, *IEEE Commun. Mag.* **47**, 66 (2009).
- [3] V. Lohani, A. Sharma, Y.N. Singh, arXiv:2112.07831 (2021). <https://doi.org/10.48550/arXiv.2112.07831>.
- [4] A. Eira, J. Pedro, J. Pires, *Proc. OFC/NFOEC 2013, OTu2A.5* (2013).
- [5] H. Takeshita, H. Noguchi, J. Abe, S. Fujisawa, A. Tajima, *Proc. ECOC 2016, 1* (2016).
- [6] N.S. Monteiro, A.C. Fontinele, D.R. Campelo, A. Soares, *J. Internet Serv. Appl.* **11**, 5 (2020).
- [7] J.M. Fàbrega, M. Svaluto Moreolo, L. Martín *et al.*, *J. Opt. Comm. Netw.* **8**, A23 (2016).
- [8] D.J. Ives, S. Yan, L. Galdino *et al.*, *Sci. Rep.* **11**, 10750 (2021).
- [9] J. Lu, C. Lu, *Opt. Lett.* **44**, 3753 (2019).