Supplementary Information

What the line delivers: a physics-guided design platform for time-modulated transmission lines

Aiganym Abzaliyeva, Constantinos Valagiannopoulos, Bakhtiyar Orazbayev

Contents

S1 Transmission-line model: charge form, ABCD blocks, Bloch dispersion, and nodal analysis **2**
S1.1 Telegrapher equations with time-varying capacitance . . . 2
S1.2 ABCD blocks and the multi-section stub . . . 2
S1.3 Bloch–Floquet dispersion . . . 3
S1.4 Nodal-admittance analysis and per-cell power . . . 3

S2 Parametric physics of the modulated branch **3**
S2.1 Mathieu picture: operating condition and threshold depth . . . 3
S2.2 Coupled-mode picture: phase matching, coherence length and the degenerate phase 4
S2.3 Gauge identity and the real-source pair . . . 5
S2.4 Capacitance-linear and elastance-linear modulation . . . 5
S2.5 Harmonic cascade and eigenvalue classifiers . . . 6
S2.6 Observables: folding family, pump-on/off ratio and absolute out/in, and the regime detector . . . 7

S3 The delivery rule: quantities, thresholds, and closed-form contraction **7**
S3.1 Branch transfer with the full stub impedance . . . 7
S3.2 Thresholds . . . 8
S3.3 Closed-form scaling . . . 9
S3.4 Criterion statistics on the as-built part . . . 9

S4 FDTD companion: engine, checks and results **11**
S4.1 Engine . . . 11
S4.2 Charge form and legacy form . . . 11
S4.3 Prescribed and delivered pump, and the Bloch escape . . . 13

S5 Design campaigns: criteria, certification and records **13**
S5.1 Screening criteria and buildability . . . 13
S5.2 Redesign campaign records . . . 14
S5.3 Buildability sweep (main-text Fig. 4c) and design maps (main-text Fig. 4d) . . . 15

S6 Measurement reference and near-field characterization **17**
S6.1 Two-tone controls: signal-added test . . . 18
S6.2 Near-field scans . . . 19

S7 Component models and unit-cell geometry **20**
S7.1 Two published models of one varactor . . . 20
S7.2 Varactor C–V characteristic . . . 21

S8 Degenerate phase-sensitive certification 21

S9 Evidence tiers, reproducibility and released records 21

S1 Transmission-line model: charge form, ABCD blocks, Bloch dispersion, and nodal analysis

S1.1 Telegrapher equations with time-varying capacitance

For a line element of length Δz with per-unit-length parameters L, R, C and G, Kirchhoff's voltage law gives the standard series relation $-\partial_z v = R\,i + L\,\partial_t i$. The time variation enters through the node relation. When the capacitance varies in time, the shunt current is the leakage current plus the time derivative of the *charge* $q = C(z,t)\,v$, rather than the capacitance multiplied by the time derivative of the voltage,

$$-\frac{\partial i}{\partial z} \;=\; Gv + \frac{\partial}{\partial t}\big[C(z,t)v\big] \;=\; C\frac{\partial v}{\partial t} \;+\; v\frac{\partial C}{\partial t} \;+\; Gv. \tag{S1}$$

The term $v\,\partial_t C$ is required for the energy balance of the modulated element. The power entering the capacitance splits into a change of stored energy and a term exchanged with the modulation, $v\,i_C = \partial_t(\tfrac{1}{2}Cv^2) + \tfrac{1}{2}\dot{C}v^2$, and over the whole line the balance reads $\partial_z(vi) + \partial_t(\tfrac{1}{2}Li^2 + \tfrac{1}{2}Cv^2) = -Ri^2 - Gv^2 - \tfrac{1}{2}\dot{C}v^2$. For $C(t) = C_0[1 + m\sin\omega_m t]$, this term is of order $m\,\omega_m/\omega_s$ relative to $C\,\partial_t v$. A scheme that omits it therefore does not merely lose a small correction but computes a different parametric response. This form applies to a capacitance prescribed in time. For a differential capacitance $C_{\mathrm{diff}}(V)$, the charge is instead $Q(V) = \int_0^V C_{\mathrm{diff}}(u)\,du$, not $C_{\mathrm{diff}}(V)\,V$. On the staggered (Yee) grid, the charge-conserving discretization is

$$q_k^n = C_k^n v_k^n, \qquad q_k^{n+1} = q_k^n - \Delta t\left[\frac{i_{k+1/2}^{n+1/2} - i_{k-1/2}^{n+1/2}}{\Delta z} + G v_k^n\right], \qquad v_k^{n+1} = \frac{q_k^{n+1}}{C_k^{n+1}}, \tag{S2}$$

together with the usual current update $i_{k+1/2}^{n+1} = i_{k+1/2}^{n} - \frac{\Delta t}{L}\big[(v_{k+1}^n - v_k^n)/\Delta z + R\,i_{k+1/2}^n\big]$. Equation (S2) contains the term $v\,\partial_t C$ automatically and conserves charge in every cell. Under modulation, the Courant bound must be evaluated at the *minimum* capacitance, $\Delta t \le \Delta z\sqrt{LC_{\min}}$. This minimum is $C_{\min} = C_0(1-m)$ for a sinusoidal capacitance and $C_{\min} = C_0/(1+m)$ for the sinusoidal elastance used in the companion model. The implemented time step is $\Delta z/(2c_0)$. A larger step could make the scheme unstable at the modulation peaks, where the resulting numerical growth could be mistaken for amplification. The effect of omitting the term $v\,\partial_t C$ is quantified in Section S4.

S1.2 ABCD blocks and the multi-section stub

With the convention $\binom{V_1}{I_1} = M\binom{V_2}{-I_2}$, a uniform section of characteristic impedance Z_0 and propagation constant $\gamma = \alpha_c + j\beta$ is described by the matrix of Eq. (1) of the main text. Losses are retained throughout, so γ is used rather than $j\beta$. The quantities Z_0, $\varepsilon_{\mathrm{eff}}(f)$ and $\alpha_c(f)$ follow from the Hammerstad–Jensen quasi-static and Kirschning–Jansen dispersive microstrip models (main text), and the tee-junction discontinuities are represented by calibrated lumped equivalents. A shunt branch of input admittance Y_{in} contributes the matrix $M_{\mathrm{shunt}} = \left(\begin{smallmatrix} 1 & 0 \\ Y_{\mathrm{in}} & 1 \end{smallmatrix}\right)$, and the input impedance of the multi-section stub follows the recursion

$$Z_{\mathrm{in},k} \;=\; Z_{0,k}\,\frac{Z_{\mathrm{in},k+1} + Z_{0,k}\tanh(\gamma_k l_k)}{Z_{0,k} + Z_{\mathrm{in},k+1}\tanh(\gamma_k l_k)}, \tag{S3}$$

which reduces to the familiar form with $\tanh \gamma_k l_k = j \tan \beta_k l_k$ for a lossless section. The recursion starts from the termination of the stub. An open end is modelled together with its fringing capacitance (ADS `MLEF`), and a grounded via is modelled as a physical via (`VIAGND`) rather than an ideal short. This distinction matters, because a physical via adds inductance and loss. The unit cell is the T-network $M_{\text{cell}} = M_{\text{TL}}(p/2)\, M_{\text{TEE}}\, M_{\text{shunt}}\, M_{\text{TEE}}\, M_{\text{TL}}(p/2)$, the complete line is $M_{\text{total}} = M_{\text{cell}}^{N}$, and the result is converted to S-parameters with a reference impedance $Z_{\text{ref}} = 50\,\Omega$.

S1.3 Bloch–Floquet dispersion

Imposing $V_{n+1} = V_n e^{-(\alpha+j\theta)}$ and $I_{n+1} = I_n e^{-(\alpha+j\theta)}$ on the cell relation and using reciprocity, $AD - BC = 1$, gives

$$e^{2(\alpha+j\theta)} - (A+D)\, e^{\alpha+j\theta} + 1 = 0 \quad \Longrightarrow \quad \cosh(\alpha + j\theta) = \frac{A+D}{2}, \qquad Z_B^{\pm} = \pm \frac{B}{\sqrt{A^2-1}} \;\; (A = D), \tag{S4}$$

where α is the attenuation per cell (Np/cell), θ the phase shift per cell and $Z_B^{\pm}$ the Bloch impedance. For a lossless cell, $(A + D)/2$ is real. Passbands ($\alpha = 0$) then correspond to $|A + D|/2 \leq 1$ and stopbands to $|A + D|/2 > 1$, where α describes the *spatial decay of an evanescent wave*. This property later explains why stability classifiers based on eigenvalues fail (Section S2). With loss, $(A + D)/2$ becomes complex and $\alpha > 0$ at every frequency. The lossy case is therefore classified by the complex Bloch exponent and by the attenuation limits of criteria (B) and (D), not by $|A+D|/2$. The idealized 2.0 GHz design used in the length study of the main text (Varactor A, Table S6, Fig. S4b,c) operates close to a slow-wave band edge, with a phase shift of about 133–156° per cell. There, the field concentration strengthens the parametric interaction but also makes the response more sensitive to detuning.

S1.4 Nodal-admittance analysis and per-cell power

To obtain the distributions of voltage, current and power along the line, the cascade of cells is assembled into a single nodal system. For a *symmetric, reciprocal* two-port, with $A = D$ and $AD - BC = 1$, the admittance matrix is

$$Y = \frac{1}{B} \begin{pmatrix} A & -1 \\ -1 & A \end{pmatrix}, \tag{S5}$$

and each two-port adds its blocks to the global matrix F. Solving $F\mathbf{V} = \mathbf{V}_S$, with the source vector $\mathbf{V}_S = (V_s/R_s, 0, \ldots)^T$, gives all node voltages. The branch currents follow from the voltage differences, the varactor voltage from the voltage division inside the stub, and the real power entering node k from $P_k = \frac{1}{2}\Re(V_k I_k^*)$. This is the procedure required for a power-handling assessment, in which the maximum must be located along the actual node profile and compared with the ratings of the soldered component at a specified excitation. Without a documented input level, rating and verified sweep, however, the calculation remains a method rather than an established safety margin. Evaluated at $2f_s$, the same procedure is the linear-response precursor of the delivery quantities in Section S3. The pump delivered in the real device is, however, self-consistent and large-signal, which is why delivery is certified in harmonic balance rather than predicted from Eq. (S5).

S2 Parametric physics of the modulated branch

S2.1 Mathieu picture: operating condition and threshold depth

Consider a lossless uniform line with $C(t) = C_0[1 + m\cos\omega_m t]$. Eliminating the current from the telegrapher equations and expanding in spatial modes, $u = Cv \propto e^{-jkz}$, gives to first order in m

the Mathieu equation $\ddot{u} + \omega_k^2[1 - m\cos\omega_m t]\,u = 0$, with $\omega_k^2 = k^2/(LC_0)$. Its principal instability region lies at

$$\omega_m = 2\omega_k, \tag{S6}$$

which is exactly the degenerate operating condition $f_p = 2f_s$. This condition therefore follows from the equation rather than being assumed. When a loss rate δ is included, the averaged growth exponent of the free, undriven resonator in this first region becomes $s = m\omega_0/4 - \delta$. The free resonator thus becomes unstable, and oscillates parametrically by itself, once

$$m \;>\; m_{\text{inst}} \;=\; \frac{4\delta}{\omega_0} \;=\; \frac{2}{Q}. \tag{S7}$$

This is a threshold for instability in time, not a threshold for amplification. Below it, a resonantly driven branch already amplifies one quadrature of a real signal by the amplitude factor $1/(1-r)$ and attenuates the other by $1/(1+r)$, where $r = m/m_{\text{inst}}$ (Section S2.3). For the isolated branch models of Table S6, m_{inst} equals 0.03 for the idealized 0.3 pF design, 0.19 for the series RLC of the COMSOL folder note and 0.41 for the SPICE model of the as-built part. These values describe the resonance of an isolated RLC branch only. They are not thresholds for finite amplification in the loaded lines, because the stub, the loading of the line and the terminations all shift both the resonance and the effective damping.

S2.2 Coupled-mode picture: phase matching, coherence length and the degenerate phase

In a distributed description, the signal and idler amplitudes coupled by an undepleted pump obey $dA_s/dz = -\alpha A_s + j\kappa A_i^* e^{j\Delta\beta z}$ and its conjugate equation. Here z is counted in cells, the coupling κ is proportional to m, and the phase mismatch is $\Delta\beta = \theta(2f_s) - 2\theta(f_s)$, all per cell. The growth exponent, taken as the real part of

$$g \;=\; \sqrt{|\kappa|^2 - (\Delta\beta/2)^2} \;-\; \alpha, \tag{S8}$$

separates two spatial regimes of this two-mode model. For $|\kappa| < |\Delta\beta|/2$, the square root is imaginary, and the signal and idler exchange energy back and forth along the line. Beyond the exceptional point, $|\kappa| > |\Delta\beta|/2$, one solution grows exponentially, and with equal loss on both modes a net growth over long distances requires $|\kappa|^2 > (\Delta\beta/2)^2 + \alpha^2$. Neither boundary is a condition for finite phase-sensitive amplification. For example, with $\kappa = 0.5$, $\Delta\beta = 2$, $\alpha = 0.1$ and unit length, the system lies below the exceptional point and both exponents have a real part of -0.1, yet the best input quadrature is amplified by $+2.84$ dB in absolute out/in [ANALYTICAL]. The mismatch also sets the coherence length, $N_{\text{coh}} = \pi/|\theta(2f_s) - 2\theta(f_s)|$, measured in cells.

For the degenerate pair, the idler is generated at the signal frequency with the phase $\varphi_i = \varphi_p - \varphi_s$. The interference between signal and idler at the output is therefore governed by

$$\Delta \;=\; \varphi_p - 2\varphi_s, \tag{S9}$$

which predicts exactly one maximum and one minimum of the amplification per 2π of pump phase, separated by 180°. This is the signature of a phase-sensitive amplifier (PSA), which is used for certification (Section S8). Two caveats apply. First, the screening cell, which is the cell model used in the delivery rule, operates at a slow-wave band edge, with $\theta(f_s) = 9.4°$ per cell at 3.1 GHz. There, the coherence length evaluated from the unwrapped phase is about 3.2 cells, since $|\theta(2f_s) - 2\theta(f_s)| = 55.5°$ per cell. For the analytical cell and for the dispersion derived from the near-field ridge, the mismatches are 26.1° and 24.8° per cell (main-text Fig. 3c), giving coherence lengths of 6.9 and 7.3 cells. All three values are well below the 20 cells of the line. A quantitative theory of the coherence length for band-edge cells is left for future work, and the phase-matching relation is not used here to claim an operating range. Second, the gauge identity derived below holds only for a one-sided complex drive. For the physical, real source, it does not remove the pump phase (Section S2.3).

S2.3 Gauge identity and the real-source pair

Consider first the tridiagonal system of main-text Eq. (9) with only its $n = 0$ source term, that is, a one-sided complex drive $V_0\delta_{n0}$. The substitution $Q_n = \tilde{Q}_n e^{+jn\varphi}$ then removes the modulation phase from the system. The transformed system has the coupling $\Delta S/2$ on both off-diagonals and the same source as before, so $|Q_n| = |\tilde{Q}_n|$ for all n, and the sum

$$P \;=\; \frac{R}{2}\sum_n \omega_n^2\,|Q_n|^2 \tag{S10}$$

does not depend on φ. This identity, however, holds only for the one-sided formulation. A real signal, $v(t) = \mathrm{Re}[Ve^{j\omega_s t}]$, contains both $+\omega_s$ and $-\omega_s$. When $\omega_m = 2\omega_s$, the negative component is the $n = -1$ member of the same Floquet ladder, since $\omega_{-1} = \omega_s - \omega_m = -\omega_s$. The substitution that removes φ from the matrix then moves it into the relative phase of the two source terms. Keeping only the two resonant components, with $D = S_0 - L\omega_s^2 + jR\omega_s$ and $h = \Delta S/2$, gives

$$Dq_+ + he^{j\varphi}q_- = \tfrac{1}{2}V, \qquad he^{-j\varphi}q_+ + D^*q_- = \tfrac{1}{2}V^* \;\Longrightarrow\; q_+ = \frac{D^*V - he^{j\varphi}V^*}{2\,(|D|^2 - h^2)}. \tag{S11}$$

A single driven stub therefore already responds as $A_{\mathrm{out}} = H_s A_{\mathrm{in}} + H_i e^{j\varphi}A_{\mathrm{in}}^*$, which is the phase-sensitive form of a degenerate parametric amplifier. At resonance, the two quadratures are amplified by the factors $1/(1 \mp r)$, with $r = h/|D| = m/m_{\mathrm{inst}}$, and the condition $|D| = h$ is the instability threshold of Eq. (S7). As a numerical example, take $L = S_0 = \omega_s = 1$, $R = 0.1$, $\Delta S = 0.1$ and $\omega_m = 2$, with the ladder extending to $|\omega| = 21\,\omega_s$. The power dissipated in the branch, relative to the unmodulated branch, is then $+6.02\,\mathrm{dB}$ at $\varphi = 90°$ and $-3.52\,\mathrm{dB}$ at $270°$. At the same time, both Floquet multipliers of the free branch over one pump period have a modulus below one (0.790 and 0.924), so the branch is stable, and the one-sided sum stays constant. Insensitivity to the pump phase is therefore a property of the one-sided formulation, or of an observable averaged over the pump phase, and not of the physical stub. The single bounded lobe of the certified design, with its sampled maximum and minimum 180° apart at 270° and 90° on the 45° grid (Section S8), has exactly this two-term form at the output of the line.

S2.4 Capacitance-linear and elastance-linear modulation

The two natural ways of describing the modulation, linear in the capacitance or linear in the elastance, agree only to first order in the depth. Expanding $S(t) = S_0/(1 + m\cos\omega_m t)$ gives

$$S(t) = S_0\Big[\Big(1 + \tfrac{m^2}{2}\Big) - m\Big(1 + \tfrac{3m^2}{4}\Big)\cos\omega_m t + \tfrac{m^2}{2}\cos 2\omega_m t - O(m^3)\Big]. \tag{S12}$$

To first order in m, the two descriptions are equivalent, with $m_S = -m$, which corresponds to a shift of the modulation phase by π. To second order, however, they describe different devices. The mean elastance changes, which shifts the resonance of the branch, and an additional modulation at the second pump harmonic appears, with amplitude $m^2/2$, which couples harmonics two steps apart. This correction amounts to 2 % at $m = 0.2$ and to 18 % at $m = 0.6$. Figure S1 shows how the best-phase on/off ratios of the two descriptions diverge as the depth increases. For the idealized line of this figure, the best-phase on/off ratio is the ratio of the signal plus the degenerate idler at the optimum pump phase to the pump-off signal. The elastance-linear law is computed at $m_S = m$, since the best-phase ratio does not depend on the sign of the depth. The two laws differ by $7 \times 10^{-5}\,\mathrm{dB}$ at $m = 0.01$, by $0.17\,\mathrm{dB}$ at $m = 0.2$ and by $7.8\,\mathrm{dB}$ at $m = 0.5$, and their curves cross between $m = 0.5$ and 0.7. In contrast, the two numerical solutions of the capacitance-linear law agree to machine precision, below $10^{-14}\,\mathrm{dB}$, up to $m = 0.15$, and differ by only $5 \times 10^{-5}\,\mathrm{dB}$ at $m = 0.7$. Circuit-level SDD models integrate the charge $Q(v)$ and are therefore based on the capacitance description. The analytical core uses the capacitance-linear law, and its elastance form is the same law to first order. Solving through the exact relation $S = 1/C$ serves only as a cross-check of the solver.

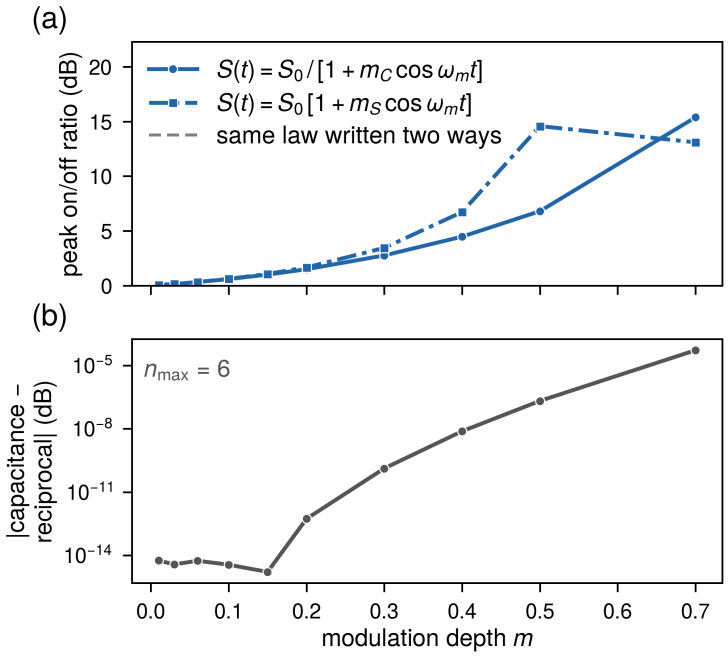


Figure S1: Capacitance-linear and elastance-linear modulation. Results for an idealized 20 cell stub loaded line with a 0.3 pF modulated capacitor, not the fabricated board, at $f_s = 3.1$ GHz and $f_m = 6.2$ GHz, with the harmonics truncated at $n_{\max} = 6$. **(a)** Peak on/off ratio at the best pump phase as a function of modulation depth, for the capacitance-linear law $S(t) = S_0/[1 + m_C \cos \omega_m t]$ (solid line, circles), where m_C is the depth m of modulation, and for the elastance-linear law $S(t) = S_0[1 + m_S \cos \omega_m t]$ (dash-dotted line, squares). The grey dashed curve, which lies under the solid one, is the capacitance-linear law solved through the exact relation $S = 1/C$, so the same law is written in two ways. **(b)** Absolute difference between these two solutions of the capacitance-linear law, on a logarithmic scale. The two laws agree to first order in the depth, with $m_S = -m_C$, but diverge as the depth increases. This comparison illustrates a single model line and does not set limits on the thresholds used in the main text.

S2.5 Harmonic cascade and eigenvalue classifiers

In a symmetric harmonic basis with $|n| \leq N_h$, the cell becomes a $2(2N_h{+}1) \times 2(2N_h{+}1)$ block matrix $\mathbf{M}_{\text{cell}}$. A line of identical cells is described by $\mathbf{M}_{\text{cell}}^N$, and a line with unequal delivered depths and phases by the ordered product of the individual cell matrices. The eigenvalues λ_k of the cell matrix are the harmonic Bloch multipliers. A natural candidate for a stability margin is $1 - \max_k |\lambda_k|$, but it does not reliably distinguish an amplifier from an oscillator in finite varactor cascades. The eigenvalues split into a propagating family, with $|\lambda| \approx 1$, and an evanescent family, and the largest magnitude $|\lambda_{\max}|$ is set by the *evanescent decay in the passive stopbands* rather than by the parametric interaction. This mechanism is reported here because eigenvalue screening is the most obvious first test a designer would try. In a resonant small-signal model with a well-defined loop gain $\mathcal{L}$, the gain is $G = G_0/(1 - \mathcal{L})$, or $|1 - \mathcal{L}|^{-2}$ in power, so a large gain lies close to the onset of oscillation. This is not, however, a general property of lossy lines. A phase-matched traveling-wave amplifier, for example, amplifies its best quadrature by $\exp[(|\kappa| - \alpha)L]$ without any closed feedback loop. The certification tests described below therefore address the proximity to oscillation directly, through operational checks. When harmonic balance fails to converge under drive, the point is recorded as unresolved rather than as an oscillation boundary, because the simulation cannot distinguish a solver failure from a genuine instability.

S2.6 Observables: folding family, pump-on/off ratio and absolute out/in, and the regime detector

With $f_p = 2f_s$, a mixing product mf_s+nf_p falls on the signal frequency whenever $m+2n = 1$. The signal-band observable is therefore defined by the truncated folding family $\{(1,0),(-1,1),(3,-1),(-3,2)\}$. In the single-fundamental harmonic-balance setup of the R17b record, all members of this family fall in the f_s bin, so harmonic balance returns their sum. On the certified R17b lobe, the two-term model of Section S8 reproduces this sum at all eight pump phases to within 9×10^{-6} dB.

The amplification bookkeeping follows strict rules. The *pump-on/off ratio* compares two paired runs, with and without the pump, whereas the *absolute out/in* is the transmission the load actually sees. A pump-on/off ratio is never quoted without its absolute transmission. A recorded Bloch design with a transmission zero illustrates this: its pump-off output is only about 10^{-11} V. Dividing by such a near-zero reference produces pump-on/off ratios above 130 dB, while the absolute out/in remains near -35 dB [two-port source].

A regime detector separates bounded amplification from oscillation by means of two tests that do not depend on loss. The first is the off-state ratio, $r_A = |\text{out(signal off)}|/|\text{out(signal on)}|$, measured with the pump on in both cases. The second is the input-scaling slope, $d\log|\text{out}|/d\log|\text{in}|$. Both tests use a fixed threshold of 0.5, which is not adjusted for individual designs. The oscillator side of the detector is illustrated by the MAVR worked example of the release (MAVR-000350, 24 cells, $f_s = 3.0$ GHz, 12 V), which is not a certified design. With a 50 Ω source, a +29 dBm pump and a pump phase of 210°, its output remains at -4.0 dBm on the signal tone $(1,0)$ and at -11.0 dBm on the family sum even when the signal is switched off (-200 dBm). On the signal tone, both tests indicate oscillation, with $r_A = 0.58 > 0.5$ and a slope of $0.15 < 0.5$. On the family sum, the two tests disagree, with $r_A = 0.171 \le 0.5$ but a slope of $0.470 < 0.5$. The verdict therefore rests on the output that persists without the signal, which identifies the response as self-oscillation rather than amplification [ADS-REALIZED].

S3 The delivery rule: quantities, thresholds, and closed-form contraction

S3.1 Branch transfer with the full stub impedance

The modulation depth at the junction is set by the pump voltage that appears across its capacitance. With the package impedance $Z_{\text{pkg}} = R_s + j\omega L_{\text{pkg}}$ and the stub impedance $Z_{\text{stub}}(f)$ calculated analytically, the branch transfer is

$$T(f) \;=\; \left| \frac{Z_C}{Z_{\text{pkg}} + Z_C + Z_{\text{stub}}(f)} \right|, \qquad Z_C = \frac{1}{j\omega C_0}. \tag{S13}$$

Two results show that Eq. (S13) is a reliable design quantity rather than a mere illustration. First, it agrees with direct large-signal readings from the same harmonic-balance simulation. Multiplying $T(6.2\,\text{GHz})$ by the pump voltage at the top of the branch, as obtained in harmonic balance, reproduces the pump voltage read directly at the junction to about 1 %, which is the precision of the rounded input values. The linear transfer with the *full* stub impedance is therefore the correct map between the two reference planes. In contrast, a simple two-element divider formed by the package alone, without Z_{stub}, misestimates the same quantity by a factor of 56 and even reverses the order of frequencies, so it must not be used. Second, the transfer explains why the pump of the fabricated line is not delivered. At the pump frequency of the board, the stub is near an antiresonance, with $|Z_{\text{stub}}(6.2\,\text{GHz})| = 1313\,\Omega$, which gives $T = 0.0026$. The branch has its overall delivery minimum at 6.42 GHz, where $T = 1.2 \times 10^{-4}$. The pump frequency of the fabricated line therefore falls almost exactly on the delivery minimum of its own branch (main-text Fig. 2c).

Figure S2 applies the same rule to a plain shunt branch on the same line, consisting of the varactor pad and a grounding via without a stub. The map is computed with the released tool,

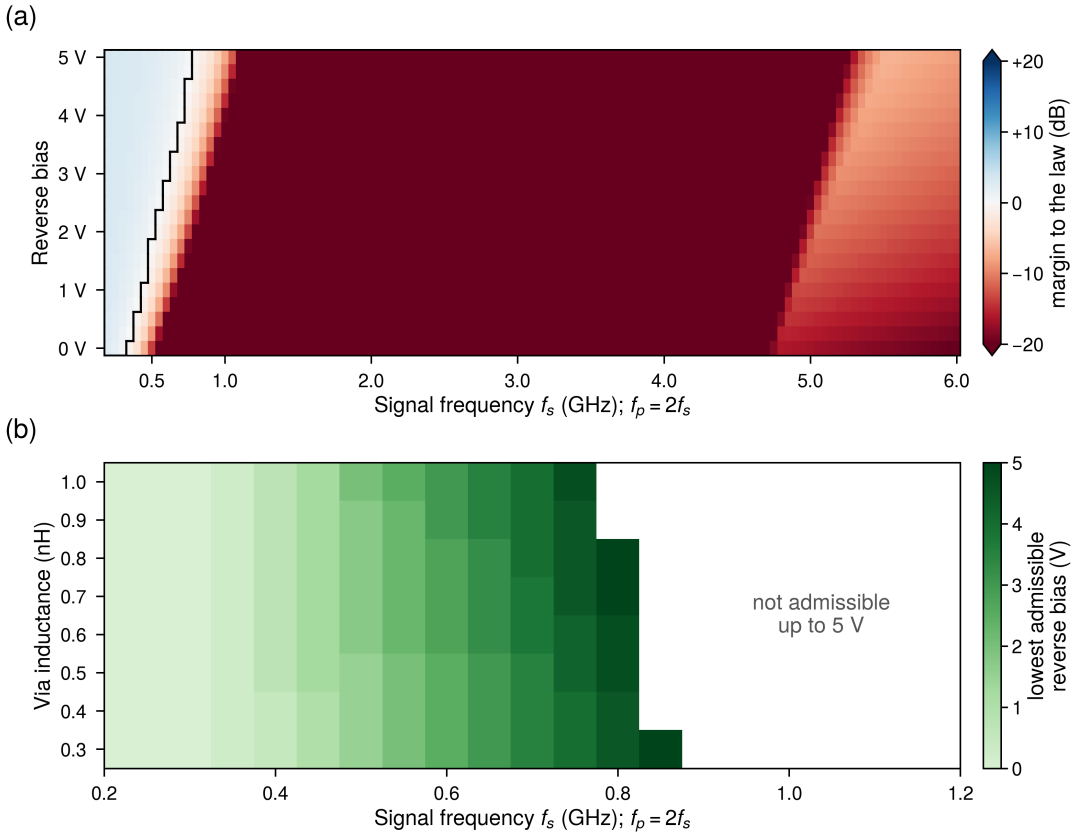


Figure S2: Plain shunt branch on the model line. Screening of the 20-cell SMV1245 model with a branch consisting only of the varactor pad and a grounding via, without a stub. **(a)** Margin to the delivery rule as a function of signal frequency and reverse bias, for a via inductance of 1.0 nH. Blue indicates a positive margin, where the rule is met, and red a negative margin, where it is not. The black contour marks zero margin. **(b)** Lowest reverse bias at which the rule is met, as a function of signal frequency and via inductance. Darker green indicates a higher required bias, and white marks the region where no bias up to 5 V is sufficient. Both panels assume zero via resistance. The maps use the earlier screening threshold $T_{\mathrm{req}} = 0.1438$, whereas main-text Fig. 4d uses the adopted value of 0.17378. Since the adopted threshold is stricter, it can only reduce the region where the rule is met. The footprint and the bias network are not assessed.

so every value in it is a prediction rather than a harmonic-balance record. Cells that satisfy the rule exist only for $f_s = 0.20$–0.85 GHz, as shown in panel (b), which is zoomed to $f_s \leq 1.2$ GHz and marks in white the cells that do not pass at any bias up to 5 V. Both panels assume a via resistance of $0\,\Omega$, and the data file also contains a slice for $0.3\,\Omega$.

S3.2 Thresholds

The delivery threshold $T_{\mathrm{req}} = 0.174$ is a chosen screening value. It is anchored on the harmonic cascade of the as-built cell, and it corresponds to a +1 dB pump-on/off ratio for one model core at one frequency, with one varactor and one drive level. The threshold follows from the requirement that the delivered junction depth $m_j = k\,T\,V_{\mathrm{bt}}$ reaches the depth $m_D^* = 0.236$ at which the core predicts +1 dB, which gives $T_{\mathrm{req}} = m_D^*/(k\,V_{\mathrm{bt}})$. For the SMV1245 at 0 V, the relative slope of the capacitance, including the package capacitance $C_P = 0.47$ pF, is $k = |dC/dV|/C = [M/(V_J + V)]\,C_j/(C_j + C_P) = (1.7/3.5)(6.9/7.37) = 0.4547\,\mathrm{V}^{-1}$. The pump amplitude at the top of the branch, $V_{\mathrm{bt}} = 2.98$ V, is the largest value over the $N = 20$ cells of a linear pump calculation at +19.5 dBm from a $50\,\Omega$ source at the board plane. The depth m_D^* is evaluated at $f_s = 3.3$ GHz, the first calibrated signal frequency that meets criterion D. Expressed in decibels, the delivery

threshold drawn in main-text Fig. 2b,c is −15.2 dB. The earlier screening threshold of 0.1438 was based on a branch-top amplitude obtained with a 25 Ω source.

The transport limit $\alpha_{\max} = 0.02\,\mathrm{Np/cell}$ corresponds to a total pump attenuation of $0.02 \times N \times 8.686\,\mathrm{dB}$, which is 3.47 dB over 20 cells and 6.95 dB over 40 cells. Conditions (B) and (D) are therefore *quantitative limits on the attenuation*, not passband checks. A simple passband criterion was tested directly and failed. A frequency window that this criterion classifies as "passband-open" gives a pump-on/off ratio of only +0.002 dB in harmonic balance [ADS-REALIZED, one 50 Ω source port]. An attenuation limit and a passband condition are therefore not interchangeable.

S3.3 Closed-form scaling

To obtain a closed form, the cell is reduced to a transmission line of impedance Z_0 loaded by a lossless branch of reactance X. The Bloch trace and the transfer then take the forms

$$\frac{A+D}{2} \;=\; \cos\vartheta \;+\; \frac{Z_0 \sin\vartheta}{2X}, \qquad T \;=\; \frac{|Z_C|}{|X|}, \tag{S14}$$

where $\vartheta = \beta_0 p$ is the electrical length of the full unloaded period p, so that each half-period section contributes $\beta_0 p/2$. This angle should not be confused with the loaded Bloch phase θ of Eq. (S4). Transport requires $|A+D|/2 \le 1$, which sets an upper bound on $|Z_0/2X|$ and hence a lower bound on $|X|$. Delivery requires $|X| \le |Z_C|/T_{\mathrm{req}} = 1/(\omega_p C_0 T_{\mathrm{req}})$, which sets an upper bound on $|X|$. The two conditions thus constrain the same reactance from opposite sides, and the fraction of angles ϑ that satisfy both shrinks with the single dimensionless parameter

$$u_{\min} \;=\; \tfrac{1}{2}\, Z_0\, T_{\mathrm{req}}\, \omega_p C_0. \tag{S15}$$

For the as-built geometry, with stub scales $s = 0.4$–3.0 and pump frequencies of 2–16 GHz, the share of stub lengths that meet both (A) and (B) at $T_{\mathrm{req}} = 0.174$ and $\alpha_{\max} = 0.02\,\mathrm{Np/cell}$, averaged over pump frequency, is 53.7, 21.7, 2.1 and 0 % for $C_0 = 0.3$, 1, 3 and 7.37 pF. For the 7.37 pF commercial part, the share is zero at every sampled point [ANALYTICAL]. The pump loading $\omega_p C_0 Z_0$ (with $Z_0 = 50\,\Omega$) is 14.4 for the board and 0.37 for the certified design at its 8.7 V bias, which is 39 times smaller.

The closed-form curve of main-text Fig. 2d is $1 - 2\arctan(u_{\min})/\pi$. It gives the fraction of angles $\vartheta \in [0, \pi]$ for which the lossless half-trace of Eq. (S14), $|\cos\vartheta + u\sin\vartheta|$ with $u = Z_0/2X$, remains below one at the delivery limit $|u| = u_{\min}$. In other words, it is the fraction of phase angles that lie in a passband when delivery is just sufficient, and it does not include the attenuation limit. The numerical scan, by contrast, counts stub lengths against both the delivery and the attenuation limits. The two curves therefore show the same trend but are not two estimates of the same quantity.

S3.4 Criterion statistics on the as-built part

The screening step covers 2,034 evaluable stub geometries, out of 3,245 candidates, for the SMV1245 at 0 V, with $N_{\mathrm{model}} = 15$ and $f_s = 3.1$ GHz ($f_p = 6.2$ GHz), and with criterion (C) applied to the 150×30 mm design envelope [ANALYTICAL]. Each of conditions (A) and (B) is met by some geometries on its own, but no geometry in the scanned family meets both at once. Across the sample, the two quantities are strongly correlated, with $\mathrm{corr}(\log T, \log\alpha) = +0.883$. The scan was run at the earlier screening threshold of 0.1438. At that threshold, statistically independent criteria (A), (B) and (C) would leave about 125 designs that pass all three ($n_{\mathrm{A}} n_{\mathrm{B}} n_{\mathrm{C}}/n^2$), and the qualitative passband criterion would admit 61 designs, or 28 when buildability is included. Since the adopted threshold $T_{\mathrm{req}} = 0.174$ is stricter, the set of designs meeting both (A) and (B) remains empty. The result is the same for every N from 9 to 21, which gives 13 identical rows, because the criteria are evaluated per cell and the length of the line never becomes the limiting factor at $f_s = 3.1$ GHz. It also holds in a separate sweep of the unchanged as-built geometry over $f_s = 2$–6 GHz, in which none of the 1,601 points shows any admissible window.

Figure S3 brings together, on a common frequency axis, the dispersion of the exported analytical cell, the attenuation per cell compared with the transport limit $\alpha_{\max}$, and the available passive measurements. The exported cell differs from the analytical tier of main-text Fig. 3b. Vertical lines mark the two signal–pump pairs, $f_s = 3.1\,\mathrm{GHz}$ with $f_p = 6.2\,\mathrm{GHz}$, and $f_s = 5.3\,\mathrm{GHz}$ with $f_p = 10.6\,\mathrm{GHz}$. Panels (a) and (b) use this export, which combines the framework cell with an ADS-tabulated stub and a grounding-via model. It gives a small pump attenuation at 6.2 GHz, but $\alpha = 1.80\,\mathrm{Np/cell}$ at 10.6 GHz. At the signal frequency of 3.1 GHz, it gives $\alpha = 0.038\,\mathrm{Np/cell}$, which exceeds $\alpha_{\max}$. Above 11.64 GHz, the tee reference-plane model is no longer valid, so the cell is evaluated with an ideal tee, which causes the step visible in these panels. Panel (c) shows Mem7, the VNA trace drawn in main-text Fig. 3b, with $S_{21}(3.1\,\mathrm{GHz}) = -1.99\,\mathrm{dB}$. The trace Trc2 in the same file gives $-2.45\,\mathrm{dB}$. The board and its state at the time of measurement are not documented. A newer trace of the bare board was supplied together with the fixture comparison (Section S6), whereas the older traces cannot be assigned to a specific board with certainty. A terminal S_{21} value alone does not establish whether the per-cell transport limit of the model is met.

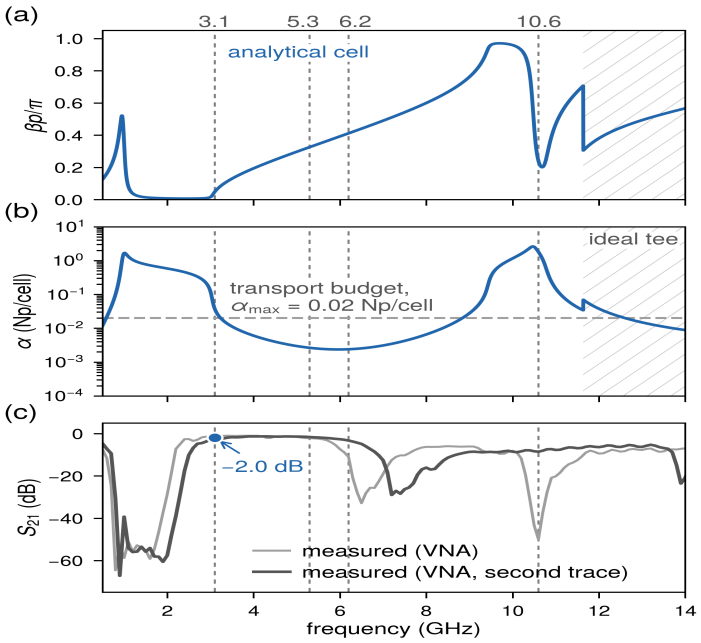


Figure S3: Dispersion, attenuation and passive transmission of the as-built line. The vertical dashed lines in all panels mark the two signal–pump pairs, $f_s = 3.1\,\mathrm{GHz}$ with $f_p = 6.2\,\mathrm{GHz}$, and $f_s = 5.3\,\mathrm{GHz}$ with $f_p = 10.6\,\mathrm{GHz}$. **(a)** Normalized Bloch phase per cell, $\beta p/\pi$, and **(b)** attenuation per cell, α, of the as-built cell, computed from the exported analytical cell with an ADS-tabulated stub [ANALYTICAL]. The parameters of this cell differ from those of the analytical tier in main-text Fig. 3b. The horizontal dashed line in (b) marks the transport limit $\alpha_{\max} = 0.02\,\mathrm{Np/cell}$. The hatched region marks frequencies above the validity limit of the tee reference-plane model, where an ideal tee is used instead. **(c)** Pump-off transmission $|S_{21}|$ of the fabricated line from two VNA traces, Mem7 (light grey), which is the trace shown in main-text Fig. 3b, and a second trace, Trc2 (dark grey). The marker gives the value of Mem7 at $f_s = 3.1\,\mathrm{GHz}$. Since the boards on which both traces were recorded are not documented, an isolated overlap with a predicted stopband does not validate the model for that specimen.

Figure S4 shows two results quoted in the main text, the formal exceptional-point depth of the as-built cell and the dependence on length of the idealized design based on Varactor A. In panel (a), the calibration points of the on/off ratio cover only 3.0–3.1 GHz, so all other signal frequencies lie outside the calibrated range. In panel (b), the pump-on/off ratio varies with length as a Fabry–Pérot comb, damped by the vias. As the line becomes longer, harmonic balance eventually fails to converge. This onset of non-convergence is recorded as an unresolved point rather than as an

established instability, and it decreases by only $-5.4\,\mathrm{dB}$ per decade of length ($R^2 = 0.18$), much more slowly than the $-20\,\mathrm{dB}$ per decade expected from a $1/N$ law.

The operating points are sorted into three classes, following the length-sweep record of the code release. An operating point is classed as robust if its pump-on/off ratio drops by no more than 6 dB when the pump is lowered 3 dB below its operating level, so that the response builds up smoothly with pump power. It is classed as a SPIKE if the ratio drops by at least 15 dB under the same test, or if the point lies at the edge of non-convergence. All remaining points are classed as moderate. The peak at $N = 18$, with an absolute out/in of $+21.1\,\mathrm{dB}$, is one of at least two steady states that coexist at the same drive, and it collapses when the pump is lowered by a single 3 dB step. The best robust design with net amplification is therefore $N = 15$. With the two-port source, it reaches an absolute out/in of $+1.18\,\mathrm{dB}$, and with a single $50\,\Omega$ port it reaches a pump-on/off ratio of $+10.56\,\mathrm{dB}$ and an absolute out/in of $+6.59\,\mathrm{dB}$ (code release v1.0.3-paper).

S4 FDTD companion: engine, checks and results

S4.1 Engine

The engine advances Eq. (S2) on a staggered grid with $\Delta z = c_0/(15\,\mathrm{GHz} \times 30) = 0.667\,\mathrm{mm}$ and $\Delta t = \Delta z/(2c_0) = 1.11\,\mathrm{ps}$. At the line velocity of $1.84 \times 10^8\,\mathrm{m/s}$, this corresponds to a Courant number of 0.307 and to about 92 points per line wavelength at f_s and 46 at f_m. The generic line has $L_0 = 373\,\mathrm{nH/m}$, $C_0 = 79.1\,\mathrm{pF/m}$, $R = 1.14\,\Omega/\mathrm{m}$ and $G = 2.6 \times 10^{-4}\,\mathrm{S/m}$, and it is operated at $f_s = 3\,\mathrm{GHz}$ with $f_m = 2f_s$. Both ends are terminated with absorbing boundaries, and the line is excited either by a Gaussian pulse or by two tones.

The modulation is prescribed in terms of the elastance, $S(t) = S_0[1 + m\sin(2\pi f_m t + \phi_0)]$, with $C = 1/S$. The companion thus modulates the elastance, whereas the analytical core uses the capacitance-linear law, and the two descriptions agree to first order in the depth (Fig. S1). For elastance modulation, the smallest capacitance is $C_{\min} = C_0/(1 + m)$, at which the Courant number rises to 0.343 for $m = 0.25$.

Four modes are implemented. The first applies a *prescribed uniform* modulation to a slab of line. The second applies a *delivered traveling pump*, that is, a prescribed traveling profile whose local depth follows the share of a co-propagating pump that is actually delivered, $m_0 T e^{-\alpha_p z}$, where T is the transfer and α_p the pump attenuation. In this mode, the launch phase is scanned over a full period, from 0 to 345° in 15° steps, and the modulation is prescribed rather than solved for. The third mode is a *Bloch-loaded* slab, in which periodic loading creates a stopband that contains $3f_s$. For one period, the passive ABCD matrix gives $|\mathrm{Re}(A + D)/2| > 1$ from 8.069 to 9.346 GHz, with an attenuation of 0.205 Np per period at 9 GHz, and a continuous passive tone at $3f_s$ is transmitted 30.6 dB below its level on the smooth slab. The period (9.33 mm) and the loaded width (4.67 mm) are fixed in physical units, independently of the grid. The fourth mode is self-consistent, with a voltage-dependent capacitance, and it is not used for any result in this paper. Such a differential capacitance would require the charge $Q(V) = \int C_{\mathrm{diff}}\,dV$ (Section S1). The ABCD matrix of the third mode includes the line losses R and G, so its half-trace criterion indicates the gap edges only approximately. The lossless criterion ($R = G = 0$) shifts each edge by only 151–152 Hz, far less than the 1 MHz step of the frequency grid, and therefore gives the same sampled edges.

The band metric is the peak of the Hann-windowed FFT of the transmitted probe within $\pm 15\,\%$ of the band. It is evaluated on the same grid, path and duration with the pump on and off, and each structure is compared with its own pump-off run.

S4.2 Charge form and legacy form

To isolate the effect of the term $v\,\partial_t C$, the charge-conserving update was compared with a legacy update that keeps a constant capacitance in the denominator, with both updates treating the

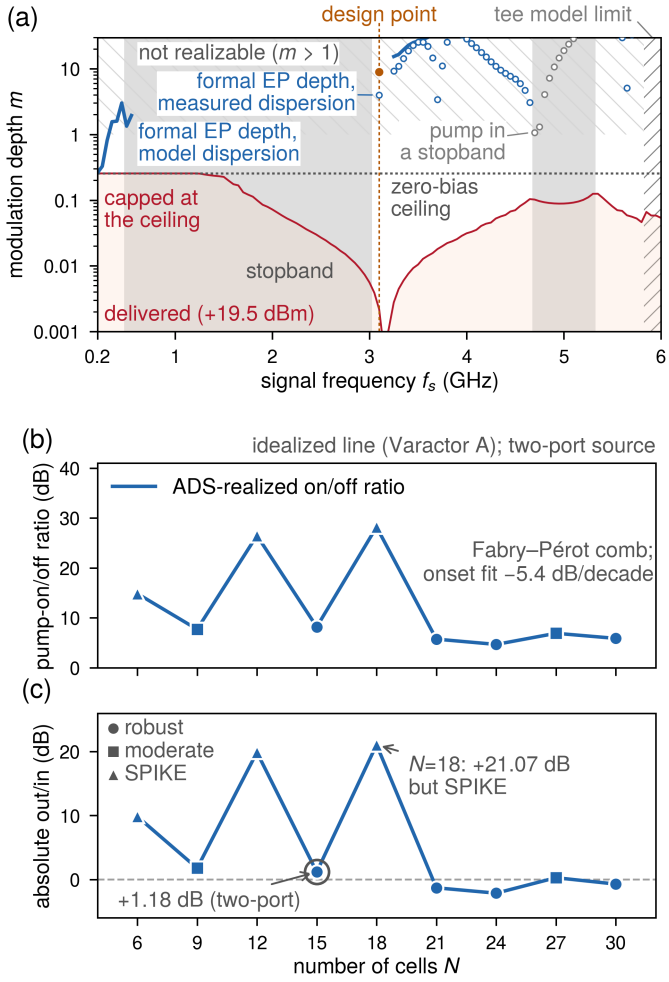


Figure S4: Formal exceptional-point depth and the effect of line length. **(a)** Modulation depth required to reach the signal–idler exceptional point of the as-built cell, as a function of signal frequency. The solid blue curve uses the model dispersion and model coupling, and the open circles combine the phase mismatch derived from the near-field ridge with the model coupling. This depth is a formal extrapolation that measures the phase mismatch and is not a threshold for finite amplification. The grey region above $m = 1$ cannot be realized. The red curve and the shaded area below it show the depth actually delivered at +19.5 dBm, which is capped at the zero-bias ceiling (dotted line). Grey vertical bands mark the stopbands, the orange dashed line marks the design point at 3.1 GHz, and the hatched region marks frequencies beyond the validity limit of the tee model. **(b, c)** Results for a separate idealized design based on Varactor A, with signal and pump at 2.0 and 4.0 GHz, as a function of the number of cells N. The matching is fixed at its 15-cell solution, and each length is simulated at its own pump level. **(b)** Pump-on/off ratio, which varies with length as a Fabry–Pérot comb. **(c)** Absolute out/in, with markers indicating the class of each operating point, robust (circles), moderate (squares) or SPIKE (triangles). The circled point at $N = 15$ is the best robust design, at +1.18 dB. The peak at $N = 18$ reaches +21.07 dB but is classed as a SPIKE, because it collapses when the pump is lowered by one step. The designs in (b) and (c) differ from the prototype model shown in (a).

losses R and G in the same way. For uniform modulation with $m = 0.5$, and with both results referred to the pump-off run of the charge form, the legacy and charge-conserving updates give pump-on/off ratios of 18.22 and 18.79 dB over two wavelengths, and of 62.28 and 62.78 dB over four wavelengths. In this example, the term $v\,\partial_t C$ alone therefore changes the ratio by 0.56 and 0.50 dB. A comparison in which the two updates also discretize the losses differently would mix both effects, so it is not used. Main-text Eq. (13) is written in charge form because this form is exact for a capacitance prescribed in time.

S4.3 Prescribed and delivered pump, and the Bloch escape

With a uniformly prescribed depth $m = 0.25$, the pump-on/off ratio grows steadily with length and reaches +15.0 dB at $4.0\lambda_0$. When the same slab is modulated instead by the delivered traveling profile in the passband, the ratio reaches at most +2.09 dB, at $2.5\lambda_0$ and a pump phase of 315°, even with full delivery and no pump decay ($T = 1$, $\alpha_p = 0$). Over the full period of launch phases, the upper envelope is 1.30, 1.82, 2.01, 2.08, 2.09, 1.92, 1.30 and 0.31 dB for lengths from 0.5 to $4.0\lambda_0$. At the worst phase, the ratio falls to −13.7 dB at $2.5\lambda_0$ and to −18.9 dB at $4.0\lambda_0$. With a delivery typical of a commercial varactor, $T = 0.05$ and $\alpha_p\lambda_0 = 1$, the upper envelope stays between +0.07 and +0.16 dB at all lengths. This limited growth applies to this modulation, depth and range of lengths, and it is not a general bound for smooth lines that carry their own pump.

Loading a slab of $3\lambda_0$ into a Bloch lattice whose stopband contains $3f_s$, with the pump matched to the loaded line, restores the signal band. Over the full pump-phase scan, the best ratio rises from +1.92 dB on the smooth slab (at 315°) to +7.48 dB on the loaded slab (at 300°). At the best phase, the product at $3f_s$ falls from +3.41 to −1.72 dB, relative to the pump-off level at f_s. A control with the same cells but a uniform mask, and hence no dispersion, reaches only +2.01 dB, with +1.41 dB at $3f_s$. The recovery is therefore caused by the dispersion and not by the geometry of the loading. Two further controls over the same phase scan separate the remaining ingredients. A slab in which the mean loading is spread uniformly, with the same modulation mask and pump pattern, reaches +1.84 dB (at 315°). The Bloch-loaded lattice with a pump pattern traveling at the smooth-line velocity v_{p0}, instead of the mean velocity of the lattice that defines the matched pump, reaches +1.35 dB (at 180°). The recovery thus requires both dispersion engineering and pump matching, and the stopband alone is not sufficient. With the physical period held fixed, the loaded result at 300° is 7.48, 7.43 and 7.44 dB for 30, 60 and 120 points per free-space wavelength at 15 GHz, which confirms convergence with respect to the grid.

In circuit-level records of the real Bloch designs, the $3f_s$ level is nearly constant from cell to cell, whereas the attenuation α would predict a decay of −0.07 to −1.86 Np/cell, and the pump-off levels lie far below [two-port source]. This behaviour is expected, because the product is generated in every cell rather than entering as a single wave. Its amplitude at the output is $A_3(L) = \int_0^L s_3(z)\, e^{-(\alpha_3 + j\Delta\beta)(L-z)}\, dz$, which a single input wave decaying as $e^{-\alpha_3 L}$ cannot describe. For a constant source s_3, this gives $A_3(L) = s_3\,[1 - e^{-(\alpha_3+j\Delta\beta)L}]/(\alpha_3 + j\Delta\beta)$, which approaches the flat level $s_3/(\alpha_3 + j\Delta\beta)$ once $\alpha_3 L \gg 1$. This level is set by the transport loss and the phase mismatch, and it reduces to s_3/α_3 only when $\Delta\beta = 0$. Here $\Delta\beta$ denotes the phase mismatch of the $3f_s$ product, not the pump mismatch of Eq. (S8). A flat profile therefore cannot distinguish generation from transport, and the suppression of harmonics on such lines is assessed with circuit-level simulations.

S5 Design campaigns: criteria, certification and records

S5.1 Screening criteria and buildability

Before any nonlinear simulation, candidate stub geometries are screened against the four criteria of main-text Eq. (11). Three of them are electrical, namely branch transfer (A), pump transport (B) and signal transport (D), all evaluated with the ADS-tabulated stub impedance and the thresholds

of Section S3. The fourth, buildability (C), is geometric and refers to the 150×30 mm design envelope. A design is buildable if the cascade fits along the envelope, with $N\times$ period ≤ 150 mm, and if the stub fits either straight within the half-clearance, (30 mm − main-line width)/2, or meandered within the per-cell slot, period × half-clearance. For the as-built cell, with a main-line width of 4.27 mm and a period of 6.11 mm, these limits are 12.865 mm and 78.6 mm^2. The designs that pass the screen are then certified as described in Methods. Certification requires a pump-on/off ratio paired with the absolute out/in, a single bounded lobe of a phase-sensitive amplifier, appropriate values of r_A and of the input-scaling slope, convergence at all pump phases, and an analytical-vs-realized parity $\varepsilon \leq 5$ dB (criterion D11).

S5.2 Redesign campaign records

Table S1 summarizes the campaign. Its SMV1245 and SMV2019 20 V rows come from the stage-0 redesign campaign at $f_s = 3.1$ GHz and $f_p = 6.2$ GHz, and were screened for buildability (C) against the design envelope. R17b, in contrast, is a separate design based on EM-modelled cells, and its 40-cell line, about 237 mm long, exceeds the length of the envelope. For the SMV1245, the scan yields no admissible candidate (Section S3). The $\omega_p C_0$ rule instead selects the SMV2019 at its 20 V limit, where its capacitance is 0.226 pF and $\omega_p C_0 = 8.8$ mS, compared with 287 mS for the SMV1245 at its as-built 0 V, both at $f_p = 6.2$ GHz. With this part, 781 of 3,083 evaluable geometries (25.3 %) pass at $T_{\mathrm{req}} = 0.174$, compared with none of 2,034 for the SMV1245. Both sets are drawn from the same 3,245 candidates, and both shares refer to the evaluable geometries of the scanned family.

Selecting designs by criteria A, B and C without signal transport (D), at the earlier screening threshold $T_{\mathrm{req}} = 0.1438$, where 25.9 % of the designs pass, yields the trap. Rerun with one 50 Ω source port, it gives a pump-on/off ratio of +15.6 dB, an absolute out/in of −25.5 dB and a parity error of 104 dB. The certified record comes from the separate campaign with EM-modelled cells (R17b, SMV2019 at 8.7 V, $N = 40$, $f_s = 1.26$ GHz). It reaches a pump-on/off ratio of +7.9 dB and an absolute out/in of +6.3 dB at the best sampled pump phase in the degenerate regime, with a parity of 0.09 dB. It is certified at a pump power of +20 dBm, which is therefore part of the record, and it passes the stability protocol (main-text Methods, deposited record `runs/r17b_004016/`) at the nominal design and at eight tolerance corners (Table S2). These are operational checks of finite duration. Convergence of harmonic balance does not prove stability, the transient runs at three of these points cover about 590 pump periods (234 ns) after settling, and the pole fits are inconclusive. Stability for unlimited time is therefore not claimed. The record also has its own limitations. The SMV2019-079LF has been discontinued, and the via enters the random tolerance draws only through analytical increments of its resistance and inductance, so the distributed tolerance of the via and the yield of physical vias are not validated.

The campaign also led to two refinements. First, the rule for choosing the varactor has two dimensions. The depth per volt of the branch capacitance, $k = |C'/C| = [M/(V_J + V)]\, C_j/(C_j + C_P)$, contains a junction factor that does not depend on C_{J0} and a package factor that decreases with C_P. For the SMV2019 at 8.713 V, $C_j = 0.391$ pF, $k_j = 0.1146\ \mathrm{V}^{-1}$ and $k_{\mathrm{tot}} = 0.0972\ \mathrm{V}^{-1}$, and the analytical generator of the certified record uses k_{tot}. With a capacitance limit of 1 pF, the eligible parts all reach nearly the same pump loading, about 38.96 mS (38.953–38.961 mS in the scan at the policy bias). The bias then trades admissibility against modulation depth. In a bias scan with 100 designs per family at $T_{\mathrm{req}} = 0.174$, the admissible share grows from 0.3 % at 0 V to 24.7 % at 20 V, compared with 25.3 % in the full screen of 3,083 designs above, while the junction-only depth per volt k_j falls from 0.400 to $0.060\ \mathrm{V}^{-1}$. The stage-1 figure of merit A_1 of the best admissible candidate peaks at 2 V on the scan grid, which is below the policy bias of 3.08 V, the lowest bias at which $C_0 \leq 1$ pF. Second, a guard against very small denominators in the pump-on/off ratio is essential. In one case, a ratio above 260 dB was recorded on a reference of only about 10^{-15} V and had to be rejected [two-port source].

Table S1: Redesign under the screening rule. The criteria and admissible shares come from the screening stage, and the pump-on/off ratio, absolute out/in and parity error ε come from harmonic balance. All three are given in dB. The SMV1245 and SMV2019 20 V rows belong to the stage-0 redesign campaign at $f_s = 3.1$ GHz, $f_p = 6.2$ GHz and $N_{\text{model}} = 15$, and were screened for buildability (C) against the 150×30 mm design envelope. For the SMV1245, no design passes. R17b is a separate design based on EM-modelled cells, at $f_s = 1.26$ GHz, and its 237.1 mm line exceeds the 150 mm length of the envelope. The parity error ε is the difference between the analytical and harmonic-balance results at the operating point, with a threshold of 5 dB (criterion D11). The trap was selected among the designs that passed the screen at the earlier threshold $T_{\text{req}} = 0.1438$, where 25.9 % of the SMV2019 designs pass, whereas the 25.3 % shown in the table results from re-screening at the adopted threshold of 0.174. The values of the trap (+15.6, −25.5 and 104 dB) come from a rerun of the campaign design at a pump power of +35 dBm, with one 50 Ω source port carrying both tones, as in the R17b record. The original campaign record of the same design, obtained with the two-port source, gives +21.5, −23.6 and 97.9 dB. The R17b record is certified at a pump power of +20 dBm, and its pump-off out/in is −1.56 dB.

part	criteria	admissible	on/off ratio	out/in	ε	verdict
SMV1245, 0 V	A∧B (+C)	0/2034; all N 9–21	—	—	—	infeasible
SMV2019, 20 V	A, B, C	25.3 %	+15.6	−25.5	104	trap: starved signal path
SMV2019, 8.7 V (R17b, EM cells, $N = 40$)	A, B, D	—	+7.86	**+6.30**	0.09	**certified**

S5.3 Buildability sweep (main-text Fig. 4c) and design maps (main-text Fig. 4d)

The stub-scale sweep multiplies the lengths of all stub sections by a factor $s \in [0.4, 8.0]$ and covers pump frequencies $f_p = 2$–16 GHz in steps of 10 MHz. For each value of s, the stub impedances are recomputed in ADS, the final criteria are applied at every pair (s, f_s), and the widest relative bandwidth that satisfies them is recorded. Recomputing the stub tables reproduces the earlier stub sweep exactly, with $\max|\Delta| = 0$ (deposited record `runs/r2_cascade_fign6/REPORT.md`, Section 3, SHA-256 prefix `e111f0154d7a9c2c`). At the adopted threshold $T_{\text{req}} = 0.174$, admissible windows exist only for $s = 3.95$–5.70, which corresponds to total stub lengths of 51.8–74.8 mm and signal frequencies $f_s = 1.0$–1.4 GHz. The buildable range, defined by the geometric screen of the 150×30 mm design envelope, ends at $s = 1.689$, that is, at a stub length of 22.2 mm meandered within the 78.6 mm^2 slot. Two further remarks apply. The window frequencies lie far below the operating band of the as-built line, and the lower edge of the window is set by the end of the scanned grid rather than by the physics of the line.

Figure S5 applies the same rule at zero bias over signal frequency and stub scale, using the stub model and frequency grid of the released tool, as in main-text Fig. 4d. Table S3 lists the individual criteria at the points marked in these maps. This map, main-text Fig. 4d and Fig. S14 are all computed with the released tool and its stub-model extension (code release v1.0.3-paper) for the nominal branch model with $N = 20$. In all three, frequencies $f_s \geq 5.85$ GHz are hatched because a reduced tee model is used there, and the buildable scales end at $s = 1.689$.

Figure S14 shows the as-built stub ($s = 1$) at 0 V, where the limiting criterion follows the bands of the line itself. Signal transport limits the design in the lower stopband (0.95–2.90 GHz), delivery limits it around the design point (2.95–4.60 GHz), and pump transport limits it at the second branch resonance (4.65–5.55 GHz). At 3.1 GHz, the map gives $T = 0.0016$, which is 40.7 dB below the threshold and cannot be recovered by any bias up to 5 V. The tabulated stub gives a slightly higher value of $T = 0.0026$. Overall, the rule is met only up to 0.25 GHz at 0 V and up to 0.65 GHz at 5 V.

In main-text Fig. 4d, buildable cells pass the rule up to 0.75 GHz. Above 2.8 GHz, they pass

Table S2: Certification of the R17b record. The checks are harmonic-balance results based on EM-modelled passive cells [ADS-REALIZED], whereas the rules and the derived quantities are analytical [ANALYTICAL]. The design uses the SMV2019-079LF varactor in $N = 40$ cells at $f_s = 1.26\,\mathrm{GHz}$ and 8.713 V reverse bias, with a +20 dBm pump at $2f_s$ and a −40 dBm signal, both supplied through one $50\,\Omega$ source port. The best pump phase is selected from eight fixed steps of 45°, and the absolute out/in is referred to the nominal available signal power. In check (iii), harmonic balance is repeated with each junction charge replaced by its tenth-order Taylor expansion around the bias point and compared with the full SPICE diode model. In check (vii), the signal is detuned to $f_s + 50\,\mathrm{MHz}$ while the pump remains at $2f_s$. Check (viii) compares the peak values, and away from the peak the analytical and harmonic-balance results differ by up to 2.39 dB. Check (x) covers all nominal runs, including the +23 dBm point of (ix) and the non-degenerate run of (vii). At the +20 dBm operating point alone, the junction voltages range from 5.19 to 13.50 V. The tolerance study comprises nine deterministic cases, namely the nominal design and variations of the etch by $\pm 25\,\mu\mathrm{m}$, of C_{J0} by ±10 %, of ε_r by ±2 % and of the pad and neck width by ±20 %, together with eight random draws with a per-cell spread of C_{J0} of $\sigma = 3\,\%$ and a spread of the via impedance of $\sigma = 30\,\%$. The pump-off row is obtained from the certified simulation with the pump amplitude set to zero. All checks are operational and of finite duration. Convergence of harmonic balance does not prove stability, the transient runs cover about 590 pump periods (234 ns), and the pole fits are inconclusive. The junction voltage limit is a margin in the model, not a measured margin to oscillation, and the tolerance cases form a sampled study rather than a yield estimate.

check	value	rule	verdict
(i) absolute out/in	+6.30 dB at 270°	best of eight phases ≥ 5 dB	pass
(ii) HB orders 5, 7, 9	maximum difference 0.042 dB	≤ 0.5 dB (matched phases and peaks)	pass
(iii) Taylor charge vs full diode	0.013 dB	≤ 1 dB (matched phases and peaks)	pass
(iv) pump-phase lobe	one maximum (270°), one minimum (90°)	one circular maximum and one minimum	pass
(v) signal off	$r_A = 0.001$ (−100 dBm seed)	$r_A < 0.5$	pass
(vi) input scaling	slope 1.00 dB/dB at −52, −46, −40 dBm	≥ 0.5 dB/dB	pass
(vii) signal/idler	2.55 dB (absolute out/in +1.12 dB at $f_s + 50\,\mathrm{MHz}$)	≤ 3 dB at $f_s + 50\,\mathrm{MHz}$	pass
(viii) analytical vs HB	peaks 6.39 vs 6.30 dB, difference 0.09 dB	≤ 5 dB (peaks)	pass
(ix) pump sweep	+11 to +23 dBm: 1.29 to 9.32 dB, all convergent	five points convergent and finite	pass
(x) junction voltage	3.85–15.97 V, every cell reverse biased	reverse biased and ≤ 22 V (rating)	pass
tolerance (17 cases)	+5.8 (C_{J0} −10 %) to +6.7 dB (C_{J0} +10 %), all at 270°	best phase ≥ 3 dB; junctions 0–20 V at all eight phases	pass
stability protocol (main-text Methods)	nominal and eight corners (pole fits at the operating pump inconclusive, reported, not a pass criterion): junction ceiling (first junction at its reverse rating) 7.5 dB above the operating pump at the nominal design and seven corners, 7.0 dB at the worst corner (pad/neck width −20 %); signal off (−100 dBm seed, eight phases, +20 dBm to 0.5 dB below the ceiling in steps of at most 1 dB): output ≤ −83.9 dBm, growth ≤ 2.6 dB per step (on the protocol's four listed pumps alone, adjacent-pump growth is 9.9–11.4 dB); transient runs at three cases (nominal, pad/neck width −20 %, etch +25 μm; +20 dBm at eight phases, +22 dBm at one phase each)	ceiling ≥ 2 dB; signal-off output ≤ −80 dBm and, above −90 dBm, growth ≤ 6 dB per step (1-dB-step rule: pumps added in 1-dB steps between the protocol's four listed pumps); no transient growth	pass
pump-off reference	−1.56 dB at every phase; pump-on/off ratio +7.86 dB	paired with (i)	reported

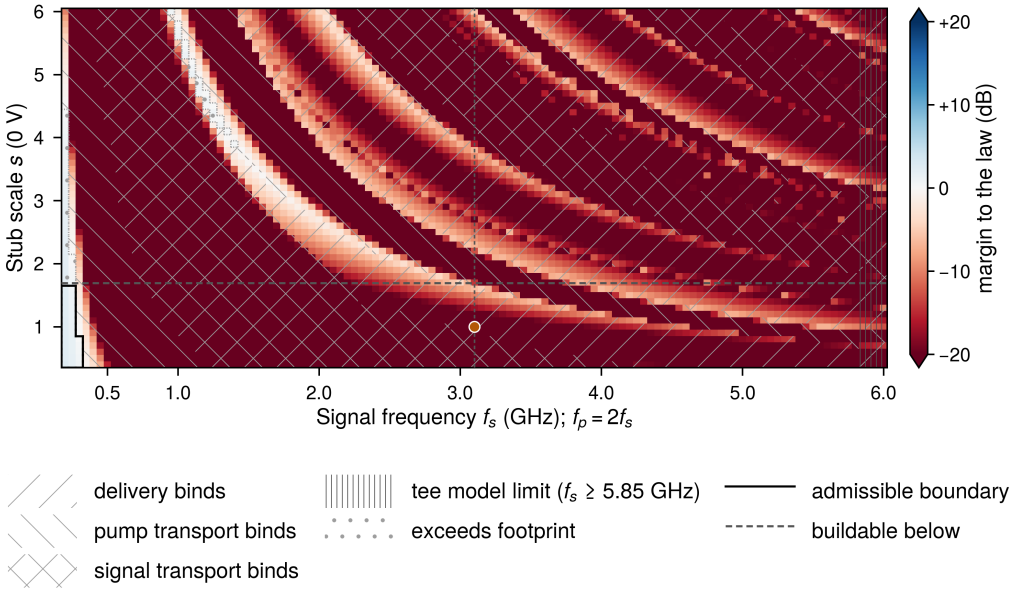


Figure S5: Limiting criteria over stub scale. Margin to the delivery rule for the nominal SMV1245 model of main-text Fig. 4d (20 cells) at 0 V, as a function of signal frequency and stub scale s. Blue indicates a positive margin, where the rule is met, and red a negative margin, where it is not. The hatching shows which criterion limits the design at each point, with single diagonal lines for delivery, the opposite diagonal for pump transport and cross-hatching for signal transport. The vertical hatching above 5.85 GHz marks the range where a reduced tee model is used, and the dotted pattern marks cells that exceed the footprint. The solid black contour marks zero margin, below the dashed line the cells are buildable, and the orange dot marks the as-built cell at 3.1 GHz. This map and main-text Fig. 4c use different stub models and frequency grids. The individual criteria at the marked points are listed in Supplementary Table S3.

only with a reverse bias of at least 3 V and with stubs that are 30–60 % longer than the as-built stub ($s = 1.3$–1.6). At 3.1 GHz in particular, they pass only for $s = 1.5$–1.6 and from 5 V upwards, and in all these cases the margins are at most 1.3 dB. The predicted absolute out/in at +10 dBm is positive only up to 0.75 GHz. Every admissible cell above 1 GHz is predicted to attenuate the signal, the best case giving −0.2 dB, and at +10 dBm and biases of 0–5 V these points lie outside the calibration points of the proxy.

In Fig. S5, apart from a window below 0.3 GHz that is shared by all buildable scales, admissible windows exist only for $s = 3.9$–6.0, at 0.95–1.40 GHz. The upper value of s is the end of the grid, and the whole range lies beyond the buildable scales. All these maps are predictions rather than records. Before any new specimen can be certified, it requires its own calibration, matching and bias network.

S6 Measurement reference and near-field characterization

Fixture reference. The VNA fixture comparison was carried out with the same setup, connectors and calibration as the two-tone experiment, as confirmed by the experimenter. Port 1 of the VNA was connected to the board through the directional coupler, and the opposite end of the board was connected to port 2. All signal and pump powers quoted in this work are generator settings. The measured coefficients of the coupler are listed in Table S4. To obtain the power at the board, they must be applied to the corresponding physical paths, but the assignment of the numbered ports and the corrections for cables and de-embedding were not provided. The trace of the bare board supplied with this comparison gives $S_{21} = -2.28$ dB at 3.1 GHz.

Calculations of the local power were also supplied. They show the power profiles along the

Table S3: The criteria at the marked points of main-text Fig. 4d and Fig. S5. Delivery A, pump transport B and signal transport D (the rule the maps use), the footprint C (checked apart from the rule) and the predicted absolute out/in (proxy at +10 dBm, extrapolated: every point lies outside the calibration anchors); ✓ passes, × fails. The fabricated columns are the sign test of the fabricated cell (tabulated stub), which records no $\alpha(f_s)$: D is not recorded there, and their rule verdicts follow from A or B alone. Generated from the Fig. 4d record.

criterion	fabricated 3.1 GHz, 0 V	fabricated 5.3 GHz, 0 V	stub $s = 1.7$ 2.9 GHz, 2 V	best buildable $s = 1.6$ 0.25 GHz, 0.5 V	stub $s = 1.6$ 3.05 GHz, 3 V
A: $T(2f_s) \geq 0.17378$	0.0026 ×	0.469 ✓	0.151 ×	1.42 ✓	0.174 ✓
B: $\alpha(2f_s) \leq 0.02$ Np/cell	0.0024 ✓	1.80 ×	0.019 ✓	0.009 ✓	0.018 ✓
D: $\alpha(f_s) \leq 0.02$ Np/cell	not recorded	not recorded	0.0030 ✓	0.0016 ✓	0.0031 ✓
the rule, A·B·D	fails (A)	fails (B)	fails (A)	passes	passes
C: footprint $\leq 78.6\,\text{mm}^2$	46.5 ✓	46.5 ✓	79.1 ×	74.5 ✓	74.5 ✓
predicted out/in (dB) (proxy, extrapolated; +10 dBm)	—	—	−0.58	+1.63	−0.56

A, B and D form the rule that the maps of Fig. 4d and Figs. S5 and S14 use; C, the footprint in the 150×30 mm design envelope, is checked apart. Fabricated columns: $T(2f_s)$ and $\alpha(2f_s)$ of the sign test (tabulated stub), which records no $\alpha(f_s)$; the map's own cell there gives $\alpha(f_s) = 0.032$ (3.1 GHz) and 0.0020 (5.3 GHz) Np/cell. Proxy calibrated at 3.0–3.1 GHz, +19.5 dBm, 0 V (≤ 0.12 dB); all points lie outside these anchors.

20 cells at 3.1 and 3.8 GHz and the power at the last varactor as a function of frequency up to 20 GHz. These calculations document that the power distribution was analysed. However, the input normalization and the comparison with the rating of the soldered component are not specified, so the calculations are not used to claim a power-handling margin for the pump drive of the experiment.

Figure S6 compares the measured reflection of the fabricated line, VNA memory trace Mem8, with the four model tiers of main-text Fig. 3b.

Table S4: Measured coupler coefficients supplied with the fixture comparison. The entries are S-parameter magnitudes in dB; S_{21} is the through coefficient and S_{41} the coupled coefficient. The coupling depends on frequency; the mapping of the numbered ports to the coupled port drawn in Supplementary Fig. S12 is not recorded, so these coefficients are not sufficient for the incident power at the board plane without the physical port assignment.

Frequency (GHz)	S_{21} (dB)	S_{11} (dB)	S_{41} (dB)
3.1	−0.348	−28.4	−35.4
6.2	−0.466	−42.6	−29.8

S6.1 Two-tone controls: signal-added test

Table S5 lists the 11 points of Supplementary Fig. S15, all measured with $f_p = 2f_s$. For each point, the table gives four quantities. The first is the pump reduction A predicted by the two pump-off traces, Mem7 and the trace of the bare board from the fixture comparison (August 2026). The second is the change Δ_{device} of the line at f_s when the signal is added. The third is the window of changes expected at any relative phase for a signal at its pump-off level. The fourth, given only at the five points where both traces give $A \geq 27$ dB, is the fraction p_k of relative phases at which a signal 10 dB stronger would have produced a change within the range observed over all 11 points. The products of p_k for independent phases are denoted Π_5, over these five points, and Π_3, over the three points where both traces give $A \geq 31$ dB ($f_s = 5.2$, 5.3 and 7.8 GHz).

Suppose that the signal were amplified by a constant +10 dB at every phase and that each phase were uniformly distributed. The joint fraction of phase space consistent with the observa-

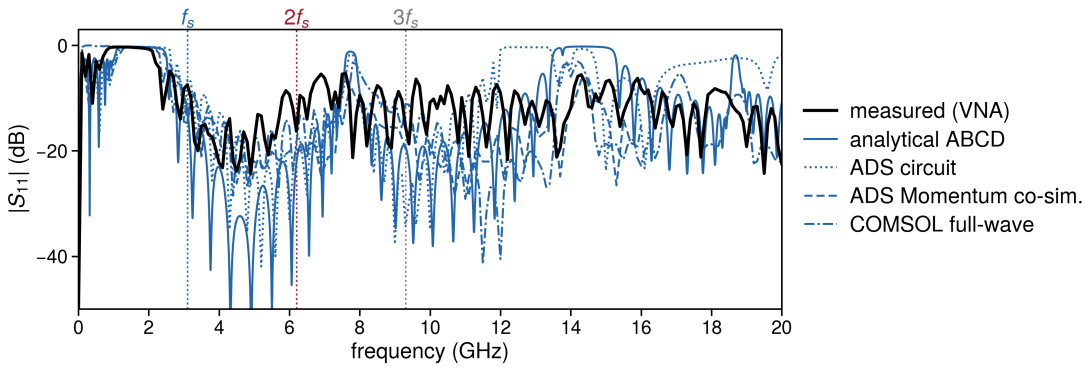


Figure S6: Pump-off reflection of the fabricated line. $|S_{11}|$ of the 20-cell line from the four model tiers of main-text Fig. 3b, analytical ABCD, ADS circuit, Momentum and COMSOL models, against measured VNA in Fig. 3b. Vertical lines mark f_s, $2f_s$ and $3f_s$.

tions is then at most $\min p_k = 0.151$, whatever the dependence between the phases, and equals $\Pi_5 = 2.6 \times 10^{-4}$ if the phases are independent. These are fractions of phase space under a stated hypothesis, not a model-free bound on amplification. In particular, they do not exclude amplification of a phase-sensitive amplifier at its best phase, because such an amplifier has a gain that varies with the same phase, and a constant gain at every phase is therefore not the best-phase hypothesis. Moreover, the phases were neither measured nor randomized, and the measurement model itself is not verified, as discussed below.

The 11 points lie in three bands, $f_s = 3.0$–3.2, 5.2–5.4 and 7.5–7.9 GHz, where the line produced by the pump alone is 8–17 dB above the signal. With the device in place, all 11 changes, between -2.3 and $+1.2$ dB, lie within their windows. The window assumes that the signal reading is not reduced when the pump is present. In some of the control runs, however, both with and without the device, the reading at f_s is lower with the pump than without it. This includes device runs for which the passive pump-off traces predict a pump reduction of up to 20 dB. That reduction is only a prediction of the passive traces, not an attenuation established under the strong pump, and it is not known where in the measurement chain, or by which mechanism, the reading is reduced. This is the reason for the passive selection criterion $A \geq 27$ dB. The five points that meet it, with $A = 27.7$–50 dB, are valid by this criterion only.

A passive pump-off attenuation does not establish that the receiver responds linearly under the pump. At the five selected points ($f_s = 3.2$, 5.2, 5.3, 7.7 and 7.8 GHz), the pump-only readings at f_p with the device differ from those with the through connection by $+0.002$, -0.455, -10.259, -0.028 and -0.212 dB, whereas the passive traces predict reductions of 27.7–50 dB. Since the analyzer is overdriven, these readings are not true powers, and the discrepancy remains unresolved. At $f_s = 3.1$ GHz, with the pump at 6.2 GHz (main-text Fig. 3d), the line at f_s has the same level with and without the signal (difference $+0.23$ dB), and the measurement chain alone, without the device, also produces a line at this frequency.

The pump-only runs as a function of pump frequency are shown in Fig. S7. The shaded bands mark the notches caused by branch resonances at 6.3–7.1, 10.3–10.9 and 15.0–15.9 GHz, and the low band at 1.0–2.6 GHz.

S6.2 Near-field scans

Near-field scans were recorded with a monopole probe on an XY stage, 1.5 mm above the board and on a 0.5 mm grid. Two types of scan were made. The first used the recorded near-field fixture, in both the unmodulated and the modulated regime. The second used a VNA as the source to record S-parameter field maps, with the far port terminated in 50 Ω. The dispersion derived from the ridge, shown in Supplementary Fig. S11 and main-text Fig. 3c, is extracted from these VNA scans. It is obtained as the ridge of the spatial FFT of the scan along the line, after smoothing,

Table S5: Signal-added test of Supplementary Fig. S15, generated from the deposited control tables. The phase fractions p_k assume a constant amplification of +10 dB at every phase and uniformly distributed phases (Section S6.1).

f_s (GHz)	f_p (GHz)	A_{Mem7} (dB)	A_{2026} (dB)	Δ_{device} (dB)	window (dB)	p_k	valid
3.0	6.0	7.3	7.2	+0.26	$[-2.39, +1.88]$	—	no
3.1	6.2	11.0	13.2	+0.23	$[-1.94, +1.59]$	—	no
3.2	6.4	27.8	31.7	−1.03	$[-1.34, +1.16]$	0.275	yes
5.2	10.4	31.1	34.8	−2.06	$[-2.63, +2.02]$	0.158	yes
5.3	10.6	50.0	47.5	−2.09	$[-2.24, +1.78]$	0.177	yes
5.4	10.8	27.7	26.9	−2.33	$[-3.37, +2.42]$	—	no
7.5	15.0	16.2	18.4	+1.15	$[-4.19, +2.82]$	—	no
7.6	15.2	20.3	25.3	+0.32	$[-2.32, +1.83]$	—	no
7.7	15.4	27.7	28.0	−1.67	$[-2.81, +2.12]$	0.151	yes
7.8	15.6	34.4	32.7	−1.34	$[-1.66, +1.39]$	0.226	yes
7.9	15.8	20.8	22.0	−1.99	$[-2.40, +1.88]$	—	no

$\Pi_5 = 2.6 \times 10^{-4}$
$\Pi_3 = 6.3 \times 10^{-3}$
min p_k (five) = 0.151
min p_k (three) = 0.158

p_k against the union observed window $[-2.33, +1.15]$ dB; signal taken 10 dB stronger at every phase.

with 12 wavenumber bins, so that each bin spans about 16.3° per cell.

The raw scan files are labelled S11, but it is not established whether this label refers to the actual reflection observable or was inherited from an earlier trace. The port mapping and the software that computed the FFT are not recorded either. The smoothed ridge is therefore shown exactly as it was supplied. Given the bin width, the agreement between its value of 24.8° per cell and the 26.1° per cell of the analytical cell at 3.1 GHz cannot be established at the level of 1.3°. The power maps at 3.1 GHz for both regimes, with the signal alone and with the signal and the pump at 6.2 GHz, are deposited as tables, and the raw scans at each point are provided as a linked dataset (Data availability).

S7 Component models and unit-cell geometry

S7.1 Two published models of one varactor

Table S6 lists two published models of the SMV1245 package. Their branch self-resonances differ by a factor of 1.5, and their local instability scales $2/Q$ by a factor of 2. All ADS records of the SMV1245 line use the SPICE model. At 0 V, where $C_0 = 7.37$ pF, the branch self-resonance $f_{\mathrm{SRF}} = 1/2\pi\sqrt{L_S C_0}$ of this model is 2.216 GHz with the 0.7 nH package inductance of the SMV1245-079LF in its SC-79 package (Fig. S8). With the 1.5 nH inductance of the SMV1245-011LF in its SOD-323 package, it falls to 1.514 GHz. Which of these package variants was actually soldered on the board remains to be confirmed, and reconciling the models with it is an open item of the fabrication record.

The SMV2019 rows of Table S1 use a different part, the SMV2019-079LF, through its SPICE model, with $C_{J0} = 2.25$ pF, $V_j = 3.5$ V and $M = 1.4$, and a package of 0.7 nH, 4.8 Ω and $C_P = 0.07$ pF. In the R17b record, this part is combined with Momentum EM unit cells (Methods), whose geometry is given in Table S8. All numbers related to the delivery rule in this work use SPICE parameters, namely the SMV1245 values from ADS for the fabricated line and the values of each screened part for the redesign.

Table S6: Varactor branch models in use, and their derived scales. The series RLC is the varactor model given in the folder note of the COMSOL data; the lumped element of the recorded COMSOL model file is in capacitor mode, which uses only $C = 7\,\mathrm{pF}$. f_{SRF} and $2/Q$ are local resonance scales of the isolated branch, not thresholds for finite gain of the line (Section S2).

model	C_0	R_s	L_{pkg}	$f_{\mathrm{SRF}} = 1/2\pi\sqrt{LC_0}$	$m_{\mathrm{inst}} = 2/Q$
ADS SPICE (SMV1245-079LF)	7.37 pF	2.0 Ω	0.7 nH	2.22 GHz	0.41
series RLC (COMSOL folder note)	7 pF	1.5 Ω	1.7 nH	1.46 GHz	0.19
idealized proxy (Varactor A)	0.3 pF	0.5 Ω	0.3 nH	16.8 GHz	0.03

Table S7: Fabricated unit-cell geometry (Rogers RO4003C, $\varepsilon_r = 3.38$, $\tan\delta \approx 0.0027$, thickness 1.54 mm; board 152.7×42.7 mm; 20-cell prototype). On both boards the ground connections are ground wires soldered through drilled holes. Varactor gap 1.50 mm; drilled holes of 0.8 mm diameter at the two stub steps.

parameter	p	w	w_2	l_2	w_3	l_3	w_4	l_4
value (mm)	6.11	4.27	1.93	6.83	1.41	5.25	1.98	1.04

S7.2 Varactor C–V characteristic

Figure S8 shows the package model of the SMV1245 and the capacitance–voltage characteristic used in the symbolically defined device (SDD) model.

S8 Degenerate phase-sensitive certification

Figure S9 compares the pump-phase sweep of the certified R17b design in harmonic balance with the two-term model $|H_s| \pm |H_i|$. The response forms a single bounded lobe, with a sampled peak of $+6.30\,\mathrm{dB}$ at 270° and a sampled trough of $-7.96\,\mathrm{dB}$ at 90°. The sampled maximum and minimum, at 270° and 90° on the 45° grid, are 180° apart, and so are the fitted extrema at 256° and 76°, as required by Eq. (S9). These extrema are values of the absolute, input-referenced quantity $G_{\mathrm{mod,in}}$, defined as the power of the output signal family divided by the available source power. They are therefore not pump-on/off ratios.

A complex least-squares fit of the eight output amplitudes, referred to the available source wave, gives $|H_s| = 1.197$, $|H_i| = 0.882$ and $|H_i|/|H_s| = 0.737$. It reproduces all eight points to within $9 \times 10^{-6}\,\mathrm{dB}$. On a continuous phase axis, the fitted lobe peaks at $+6.36\,\mathrm{dB}$ (at 256°) and falls to $-10.02\,\mathrm{dB}$ (at 76°). This fit only checks that the harmonic-balance samples follow the two-term form, and it is not an independent prediction. The independent analytical lobe is shown in main-text Fig. 5d. It gives 6.39 dB at the best sampled phase, which differs from harmonic balance by 0.09 dB at the peak and by up to 2.39 dB away from it.

The frequency response of the certified design is shown in Fig. S16. With a pump that tracks the signal frequency, the absolute out/in stays between 4.0 and 9.0 dB over ±10 % of f_s. With a fixed pump, it lies between −0.9 and +2.1 dB, so no bandwidth is claimed for operation with a fixed pump. The certification thresholds, a value of 0.5 for both r_A and the input-scaling slope, are the same for all designs and were not adjusted for any of them. The oscillator side of this classification is anchored by the MAVR example (Section S2).

S9 Evidence tiers, reproducibility and released records

Every circuit amplification record gives the pump-on/off ratio together with the absolute out/in. The two-tone controls are stated as differences at a fixed analyzer setting, and the linearity of the analyzer under the pump was not verified (main-text Fig. 3d). The drive levels form part of every record, and points at which harmonic balance does not converge are kept and marked as

Table S8: Unit-cell geometry of the certified simulated design (SMV2019-079LF, 40 cells; main-text Fig. 5a): electromagnetic verification geometry, not released for fabrication. Period $p = 5.93\,\mathrm{mm}$, main-line width $w = 0.767\,\mathrm{mm}$; stub sections in order from the varactor, positions along the stub axis measured from the main-line center line.

section	width (mm)	length (mm)	from–to (mm)
stub section 1	3.76	9.77	3.66–13.43
stub section 2	3.49	14.82	13.43–28.25
stub section 3	1.55	12.84	28.25–41.09
stub section 4	3.73	12.52	41.09–53.61

unresolved. The analytical core has two independent formulations, one based on an impedance matrix and one based on a tridiagonal system in the charge frame, and they agree to a residual of 1.8×10^{-16} in the release test.

The tables of the dense maps in the data deposit store the margins to the criteria at the earlier screening threshold $T_{\mathrm{req}} = 0.1438$. Main-text Fig. 4d, Fig. S14 and the text re-screen these margins at the adopted threshold $T_{\mathrm{req}} = 0.17378$, whereas Fig. S2 shows the stored margins directly. The margin of these maps, expressed in dB for criteria A, B and D, is

$$\min\{20\log_{10}[T(2f_s)/T_{\mathrm{req}}],\ 8.686\,N[\alpha_{\max} - \alpha(2f_s)],\ 8.686\,N[\alpha_{\max} - \alpha(f_s)]\},$$

so that a design meets the rule when its margin is at least $0\,\mathrm{dB}$.

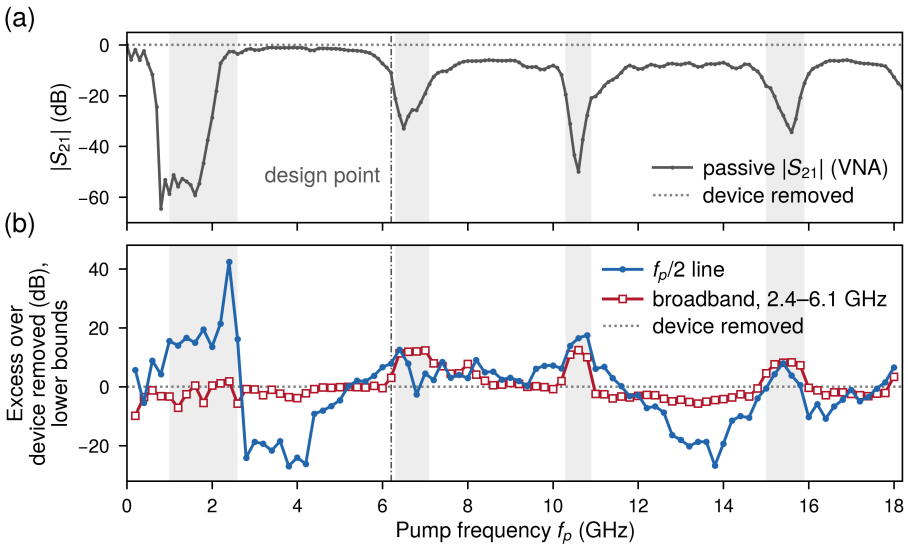


Figure S7: Spectral activity with the pump alone, as a function of pump frequency. All levels are uncalibrated analyzer readings. **(a)** Passive transmission $|S_{21}|$ from VNA memory trace Mem7, the trace shown in main-text Fig. 3b. The state of the board at the time of measurement is not documented. **(b)** Level of the line at $f_p/2$ above that of the run without the device. The shaded regions mark the notches caused by branch resonances and the low band. Relating this activity to the passive notches is possible only if the VNA trace was recorded on the same physical specimen.

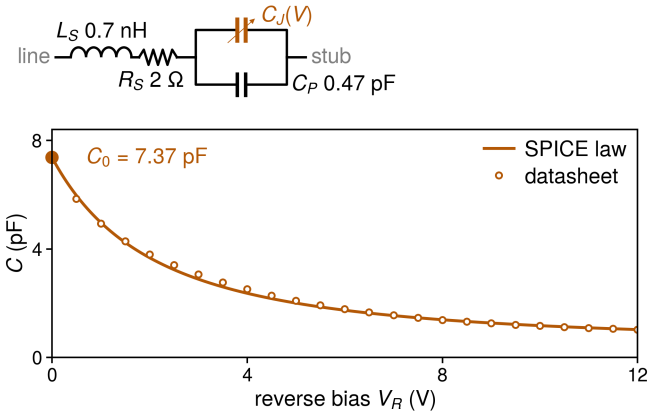


Figure S8: SMV1245 package model and C–V law. SMV1245-079LF package and junction parameters from the datasheet, as drawn; curve, the SPICE $C(V)$ law; circles, the datasheet table. The line capacitance is $C_0 = C_J + C_P$ (junction plus package). The fabricated line operates at 0 V without a bias network. Model parameters: Supplementary Table S6.

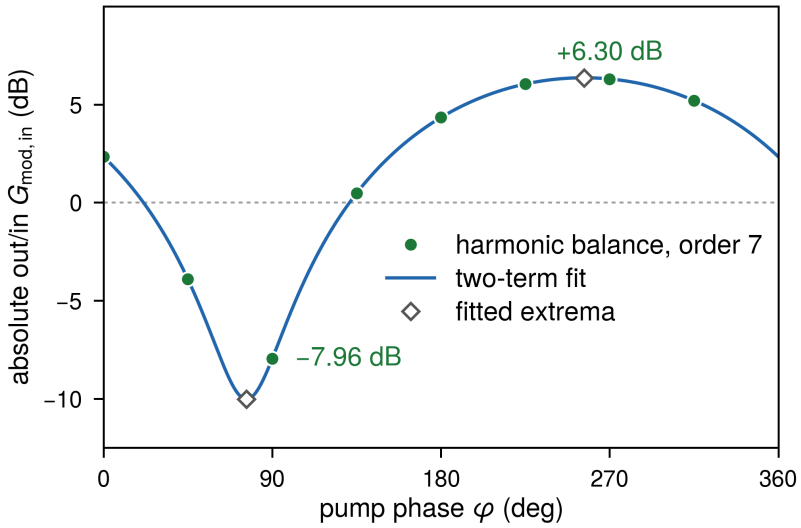


Figure S9: Degenerate phase sensitivity of the certified simulated design. Harmonic-balance samples on the 45° pump-phase grid against a two-term fit of the harmonic-balance lobe (curve), a consistency check, not an independent prediction (the independent analytical lobe is main-text Fig. 5d), at the operating conditions of Supplementary Table S2. The ordinate is the absolute output relative to available input power, not an on/off ratio. Printed labels give the extrema of the sampled phases; open diamonds, the fitted extrema between the samples. Certification rests on the tests of Table S2.

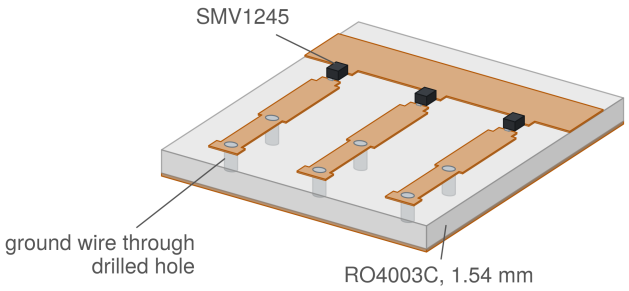


Figure S10: Three cells of the fabricated line in 3D. Varactors, ground wires through drilled holes and substrate (copper thickness exaggerated). Complete dimensions: Supplementary Table S7.

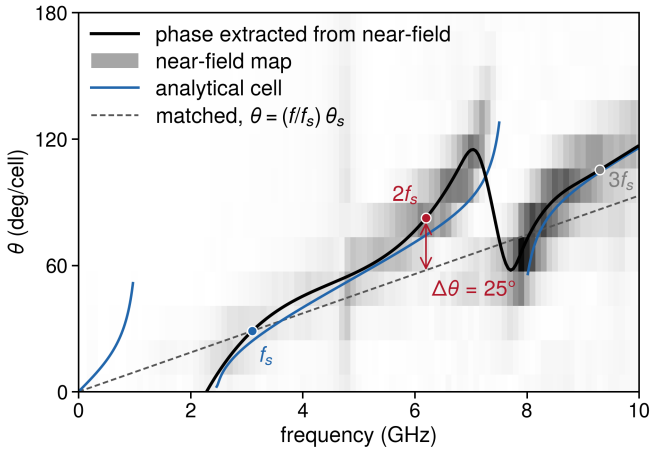


Figure S11: Near-field spatial spectrum of the fabricated line. Spatial spectrum of the supplied near-field scan along the line, together with the supplied smoothed ridge, both shown as recorded. The spectrum uses 12 wavenumber bins, each spanning about 16.3° per cell. The analytical Bloch phase and the phase-matching line are shown for comparison. The ridge can be read as a Bloch phase only with caution, because the measured quantity has not been identified. The raw scan files are labelled S_{11}, but it is not established whether this label refers to the reflection itself or was inherited from an earlier trace. The port mapping and the software that computed the FFT are not recorded, and the board used for the scan is not documented.

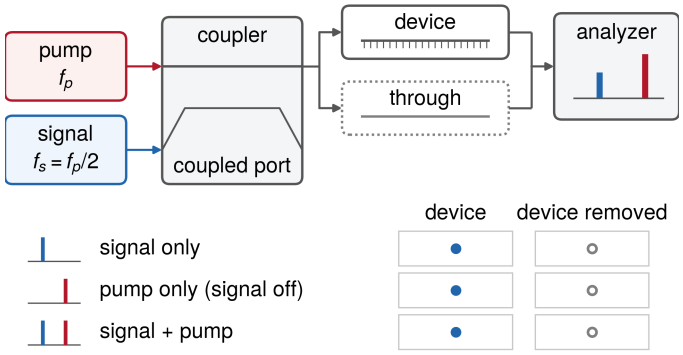


Figure S12: Two-tone control protocol. The line is measured with the signal alone, the pump alone and both together, and each measurement is repeated with a through connection in place of the device. The signal enters through the coupled port of the directional coupler, whose coupling depends on frequency (Supplementary Table S4).

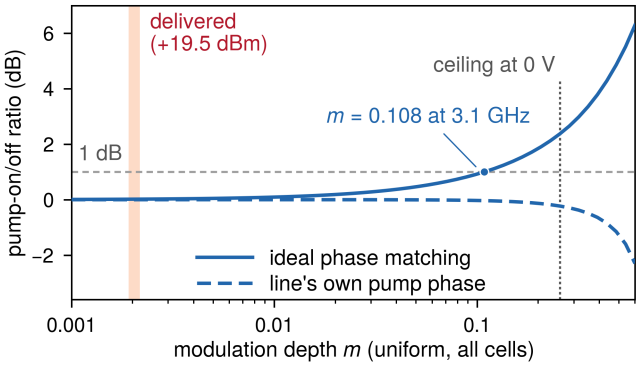


Figure S13: Modulation depth required by the fabricated line. Pump-on/off ratio as a function of a uniform modulation depth m, for the 20-cell SMV1245 model at 3.1 GHz. The solid curve assumes ideal phase matching and reaches 1 dB at $m = 0.108$. The dashed curve uses the pump phase of the line itself and does not reach 1 dB even at $m = 0.6$. The shaded region marks the depth actually delivered at a pump power of +19.5 dBm, and the dotted line marks the assumed depth limit at zero bias. Both curves are analytical calculations with this model and represent distinct cases.

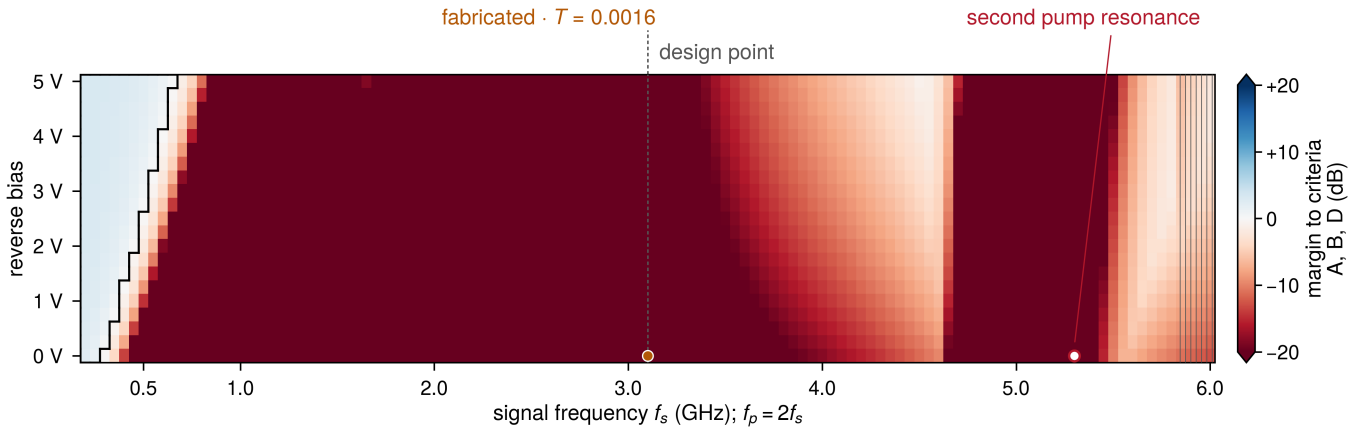


Figure S14: Predicted margin map of the as-built line. Margin to the delivery rule for the SMV1245 line with the nominal branch model, 20 cells and the as-built stub ($s = 1$), computed with the released tool (v1.0.3-paper) and its stub-model extension. The colour gives the smallest of the three margins in dB, for delivery (A), pump transport (B) and signal transport (D). The black contour marks zero margin, the hatching marks the range where the reduced tee model is used, and the open marker indicates the second pump resonance. At the design point of 3.1 GHz, the map gives $T = 0.0016$ with the stub model of the tool, which is 40.7 dB below the threshold and cannot be recovered by any bias up to 5 V. The rule is met only up to 0.25 GHz at 0 V and up to 0.65 GHz at 5 V. Passing this electrical screen does not imply a positive absolute transmission. All values are analytical predictions.

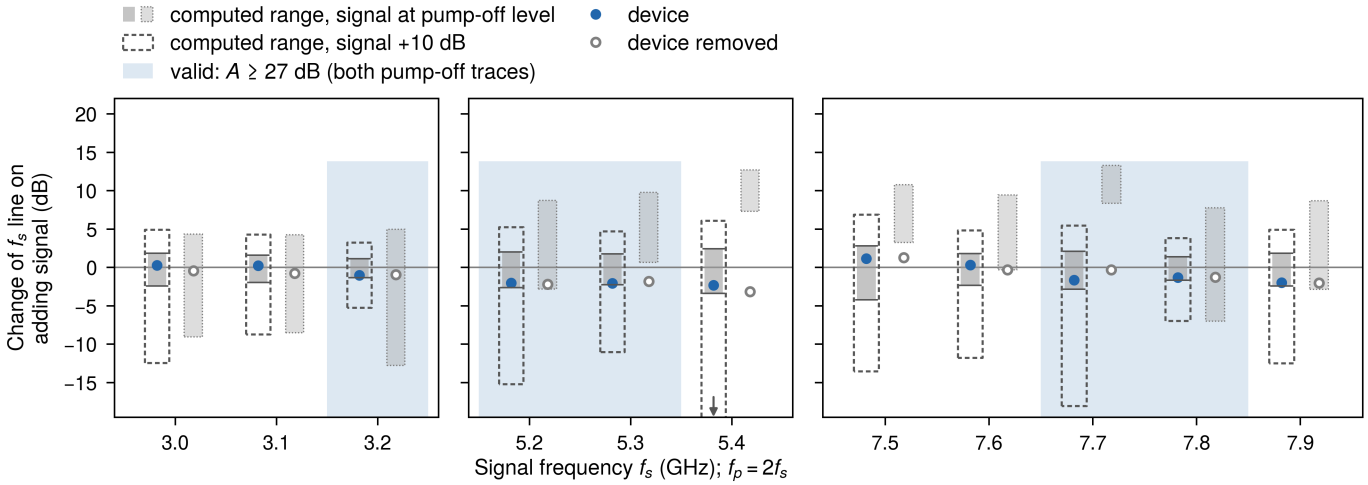


Figure S15: Signal-added test on the fabricated line [MEASURED]. Change of the line at f_s when the signal is added, measured with the device (filled symbols) and with the device removed (open symbols). The bars show the range of changes expected at any relative phase for a signal at its pump-off level, which is not a measurement uncertainty, and the dashed bars show the same range for a signal 10 dB stronger. The shaded points are the five at which both pump-off traces give a pump reduction $A \geq 27$ dB, assuming that these traces represent the measured board. The point at 3.1 GHz is not among them. The analyzer levels are uncalibrated, and the pump overdrives the analyzer. This analysis is not a statistical proof that no amplification occurs. The gain of a phase-sensitive amplifier varies with the same phase, so a hypothesis of constant gain at every phase does not represent its best-phase performance. Further details are given in Supplementary Section S6.

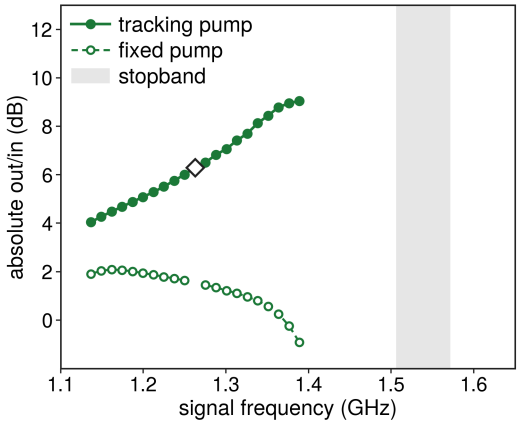


Figure S16: Frequency response of the certified design. Absolute out/in of the signal, computed by harmonic balance, for two pump configurations. With a pump that tracks the signal, $f_p = 2f_s$, the out/in stays between 4.0 and 9.0 dB over 21 points spanning ±10 % of f_s. With the pump fixed at 2×1.26 GHz, it lies between −0.9 and +2.1 dB over 20 detuned points. In this case the signal and the idler are resolved separately, and the two branches are not joined through the centre. The diamond marks the certified point of main-text Fig. 5c. The shaded region marks the first stopband of main-text Fig. 5a, at 1.506–1.572 GHz, and its edge does not indicate a stability margin. No bandwidth is claimed for operation with a fixed pump. All results are simulations, not measurements.

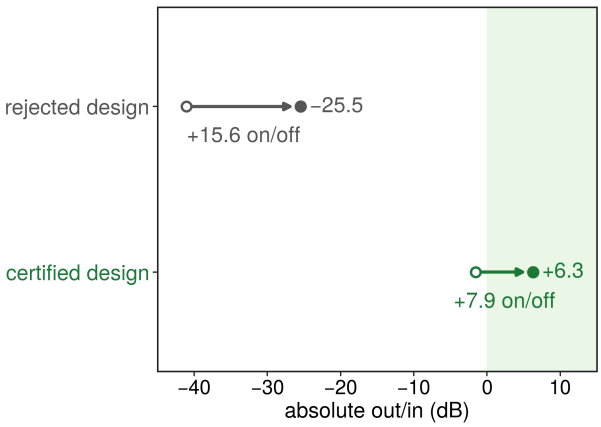


Figure S17: Pump-on/off ratio versus absolute transmission for a rejected and the certified design. Both quantities are computed by harmonic balance with one 50 Ω source port. For each design, the open circle marks the pump-off level and the filled circle the pump-on level, so that the segment between them represents the pump-on/off ratio. The green field marks positive absolute out/in. The rejected design (grey) uses the SMV2019 at 20 V in 15 cells, at 3.1 GHz and a +35 dBm pump, and is an audit run rather than a certified record. It shows a pump-on/off ratio of +15.6 dB but an absolute out/in of only −25.5 dB, so its pump-off level, obtained as the absolute value minus the ratio, is −41.05 dB. The certified design (green) uses 40 cells at 1.26 GHz and a +20 dBm pump, and at the best sampled pump phase it shows a ratio of +7.9 dB together with an absolute out/in of +6.3 dB. A pump-on/off ratio alone therefore does not establish amplification. All results are simulations, not measurements.