The Fable Computer: A Room-Temperature Terahertz Half Adder on a Regenerative Graphene-Plasmon Logic Fabric

Concept, Architecture, and Reduced-Order Feasibility

Ryoji Furui · www.ryoji.info · June 2026

*The machine described here is named the **Fable Computer**: a half adder — two input bits in, Sum and Carry out — realized on a clocked, regenerative graphene-plasmon logic fabric.*

> **Authorship and Status Disclaimer.** This manuscript is derived from the author's original paper [1]. Generation of content involved the use of an AI, Anthropic's Fable 5, at the author's direction. The author wishes to declare that they lack formal expertise in graphene-plasmonic device fabrication or full-wave electromagnetic simulation, and that the document has not received independent expert review. Consequently, the architectural and quantitative claims presented herein must be treated solely as a starting point for discussion and not as peer-reviewed engineering guidance. Formal expert review is required prior to any practical implementation.

Abstract

We present the Fable Computer, a machine model for room-temperature terahertz logic in which the bit is a gate-screened (acoustic) graphene plasmon, every logic gate is the same physical object — a resonant Dyakonov–Shur (DS) gain cell biased below its self-oscillation threshold — and the system clock is a phase-stable pulse train launched at the fabric edge by a microcomb-photomixing chain. Logic levels are restored at every gate by saturating regenerative gain; inversion is a bias choice, not a separate mechanism. The design obeys two locked constraints: no cryogenics, and chip temperature ≤ 80 °C in standard use. The machine presented is a **half adder**: five regenerative cells realize Sum and Carry at one addition per 4-ps clock slot (2.5×10^{11} additions per second, wave-pipelined), at ~0.6 fJ per addition on the fabric side, in a block of roughly 12 × 16 μm — small enough to be analyzed honestly, complete enough to exercise every architectural mechanism: gain, restoration, inversion, fan-out, clocking, and read-out. A five-pass reduced-order feasibility chain (transparent models, all parameters exposed) converges on a single self-consistent operating point: carrier density $n = 10^{12}$ cm^{-2} ($E_F \approx 117$ meV), cell length $L \approx 575$ nm tuned to 1 THz, drift bias $M = v_0/s \approx 0.7 \cdot M_{th}$. In-model, one cell then delivers +10 dB net gain per 3-cycle pulse at a 0.25 THz symbol rate with 0.38 dB pulse-to-pulse spread; two cascaded cells preserve +8.4–9.0 dB per cell through junctions as lossy as −6 dB; the noise figure sits at the 3 dB thermal floor; cell yield against measured literature disorder is 98–100 %; thermal duty reaches ~95 % at one-third area fill; kinetic corrections to the hydrodynamic tier are bounded at the ~5 % scale; and parasitic coupled-cell oscillation through the charge channel is excluded at every coupling strength. What remains open is

exactly two items: a Boltzmann–Maxwell solve owning radiative/electromagnetic coupling and absolute calibration, and a bench experiment, for which a five-gate pass/fail protocol is specified.

---

1. Introduction: the problem of the picosecond bit

A graphene plasmon is an attractive carrier for logic. It is programmable — a gate voltage sets, point by point, whether a region of the sheet supports the wave at all — and at terahertz frequencies the gate-screened (acoustic) mode confines the field to deep-sub-wavelength scales on a chip that needs no cryostat. The operating frequency is chosen by one number: the quality factor $Q = \omega\tau$ decides whether the bit is a wave or an overdamped relaxation. With hBN-encapsulated graphene ($\tau \approx 1$ ps at 300 K), Q reaches ≈ 6 at 1 THz, and the optical-phonon channel that would collapse τ at higher frequencies sets a natural room-temperature ceiling at 0.5–1 THz. The Fable Computer operates at that knee: carrier ≈ 1 THz, where the plasmon oscillates several times within its lifetime.

The same number states the central obstacle. A lifetime of a picosecond means the wave loses ~4.3 dB of power per plasmon wavelength at 300 K (5.1 dB at the 80 °C cap): launched passively, a bit decays below any usable read threshold within a handful of gates. A machine — as opposed to a one-shot demonstration — therefore needs four things a bare plasmonic circuit does not have: **distributed gain** that at least cancels the loss at every gate; **level restoration**, so that decisions made early in a cascade do not erode; a **clock** that can be distributed across millimetres with femtosecond-class phase; and **memory**, since nothing plasmonic survives between operations. Each requirement is answered here by assigning it to hardware that can physically deliver it.

Gain is electrical: a DC drift current through each cell drives a Dyakonov–Shur plasmon instability — the only mechanism with experimentally demonstrated room-temperature, continuous-wave plasmonic gain [2]; optically pumped gain cannot be sustained beyond a sub-picosecond transient [3] — and biasing the cell *below* its self-oscillation threshold turns it into a regenerative amplifier whose saturating response doubles as the logic nonlinearity (Section 4). Timing is optical: the bias plane's RC corner (~1 MHz at chip scale, ~10 GHz even per 10-μm segment) is two to six orders of magnitude below 1 THz, so no electrically modulated "clock bias" is physically available; instead a microcomb-photomixing chain at the fabric edge launches the clock as a pulse train, and optics holds femtosecond phase (Section 5.2). Restoration is intrinsic: every gate is the same gain cell, and its rails re-clamp the levels (Sections 4, 6.4). Memory is electronic: bits persist as gate charge, not as plasmons (Section 5.4).

The machine built from these answers in this manuscript is deliberately the smallest complete one: a **half adder** — five regenerative cells computing Sum and Carry — specified to the level of cell map, timing, gain budget, and energy in Section 5.6, and feasibility-checked end to end by the five-pass model chain of Section 6. The closing sections give the energy accounting honestly (Section 7) and reduce everything still unproven to five bench gates and one remaining solve (Section 8).

2. Physical platform and drive chain

2.1 Substrate, stack, and the programmable gate lattice

The chip is built on sapphire ($\varepsilon_r \approx 9.4$, $\tan\delta < 10^{-4}$) carrying a segmented Au gate lattice, hBN-encapsulated monolayer graphene, and a ZnO top cap — a stack demonstrated for on-chip terahertz plasmon transport by Yoshioka et al. [4]; the cell cross-section is shown in Figure 3a. The graphene channel is dual-gated: a global back gate sets the working density ($n = 10^{12}$ cm^{-2}, $E_F \approx 117$ meV), while an electrode-mapping controller writes a per-segment voltage map onto the top-gate lattice. That density map *is* the program: it defines which cells are doped (active, amplifying) and which sit at the charge-neutrality point (dark, absorbing guard cells). The lattice pitch follows the gain cell of Section 4 — DS quarter-wave elements of $L \approx 575$ nm — comfortably within deep-UV lithography.

2.2 The photonic clock chain: microcomb → DFB pair → photomixer

The clock generator is a microcomb-photomixing chain of the kind demonstrated by Tokizane et al. [5] as a chip-scale, fiber-bonded, phase-coherent terahertz source (soliton operation sustained > 24 h at 1 W pump; 112 Gbps wireless transmission at 560 GHz). Two adjacent comb lines injection-lock a pair of DFB lasers (one IQ-modulated for data), which are recombined and photomixed in a uni-traveling-carrier photodiode to produce the carrier (Figure 1). The comb's free spectral range is engineered to the **symbol rate, 0.2–0.5 THz**, so the chain directly emits the few-cycle 1-THz-carrier pulse train that constitutes the system clock (Section 5.2). Phase is held by the separate-arm locking method of Hirota et al. [6], whose demonstrated ±9 mrad residual corresponds to **1.4 fs at 1 THz against the fabric's 42 fs (±15°) budget** — thirty-fold margin.

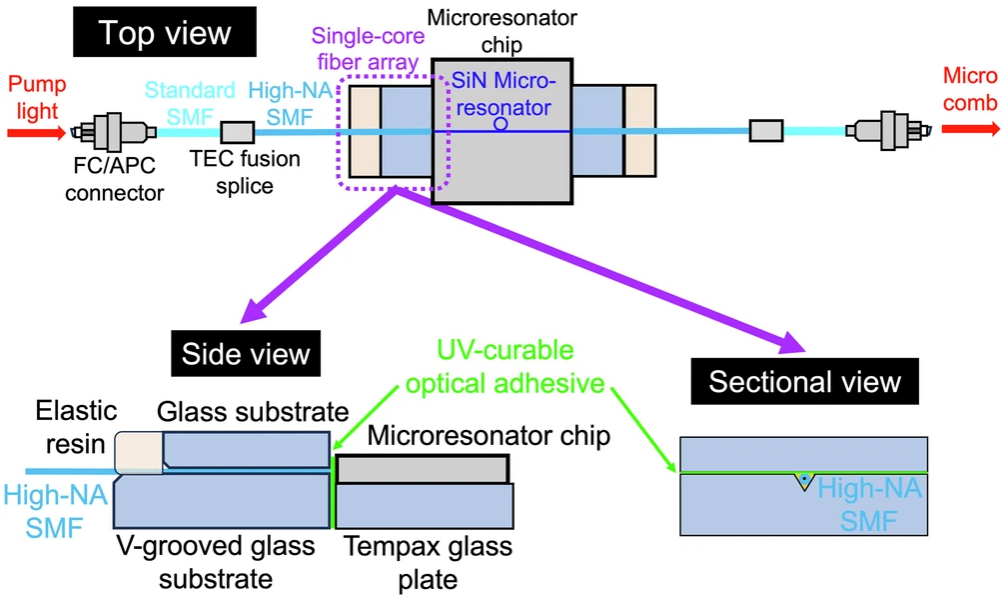

Figure 1: The drive/clock chain, reproduced from Tokizane et al. [5] (CC BY 4.0): a fiber-coupled SiN soliton microcomb feeds bandpass filters selecting two comb lines; these

injection-lock two DFB lasers (one IQ-modulated); the recombined beams photomix in a UTC-PD to emit the carrier. Here the comb FSR is engineered to the 0.2–0.5 THz symbol rate.

2.3 Launch, read-out, and in-situ reference

On-chip terahertz launch and detection follow the monolithic optoelectronic circuit of Kusyak et al. [7]: a generator photoconductive switch at the center of a coplanar stripline with two symmetric detector switches, providing odd-mode propagation, galvanic isolation between generation and detection, and — critically for a fabric that must be trimmed against fabrication spread — **simultaneous in-situ measurement of signal and reference under identical conditions** (Figure 2). The Fable Computer uses this one-launcher / two-detector / one-reference topology at the fabric edges for launch and Sum/Carry read-out, and additionally **reuses the detector/reference elements mid-fabric as the optical taps of the re-sync stations** (Section 5.2). A graphene patch antenna of the kind designed by Haque et al. [8] remains the free-space port option, dimensionally scaled to the port frequency. One honest flag: the Kusyak circuit's demonstrated −3 dB bandwidth is 440 GHz; its extension toward 1-THz carriers is an engineering assumption that the bench protocol tests directly.

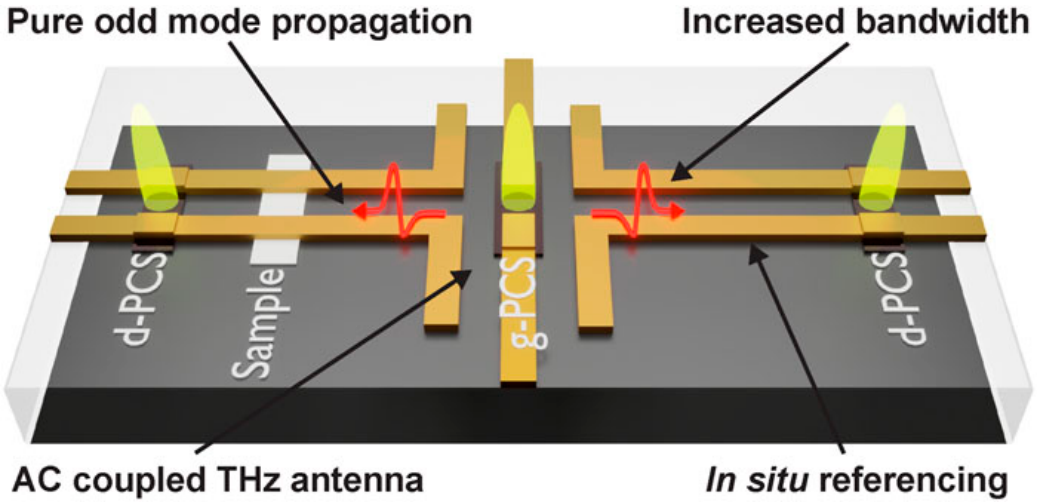

Figure 2: Monolithic launcher/detector/reference circuit, reproduced from Kusyak et al. [7], APL Photonics 10, 076117 (2025), CC BY-NC 4.0: generator switch (g-PCS) centered on a coplanar stripline with two symmetric detector switches (d-PCS) — galvanic isolation plus in-situ referencing. Used at the fabric edges for launch/read-out and reused mid-fabric as re-sync clock taps.

3. Operating point

Parameter	Value	Note
Carrier frequency	~1 THz ($\hbar\omega \approx 4.14$ meV)	room-temperature Q-knee of hBN-encapsulated graphene
Symbol/clock rate	**0.2–0.5 THz** (0.25 THz baseline)	few-cycle pulses; set by comb FSR
Plasmon mode	acoustic / gate-screened	required for confinement and gate

		programmability
Quality factor Q = ωτ	6.3 (300 K) / 5.3 (80 °C)	τ ≈ 1.0 ps → 0.85 ps at the cap
Carrier density n	**10^{12} cm^{-2} (E_F ≈ 117 meV)**	selected by disorder from below, heat from above (§3 text)
Plasmon speed s	2.33×10^{6} m/s	gated-plasmon dispersion at d = 10 nm
Plasmon wavelength λ_p	**2.33 µm**	s/f_0
Logic cell	**DS quarter-wave resonator, L ≈ 575 nm**	tuned to 1 THz at the operating bias
Drift bias	**$M = v_0/s \approx 0.12 \approx 0.7\cdot M_{th}$** ($v_0 \approx 2.8\times10^{5}$ m/s)	calibrated in situ by approach to self-oscillation
Lithography class	**~500 nm**	deep-UV
Chip temperature	**≤ 80 °C, no cryostat**	locked requirement; 55 K budget over 25 °C ambient

Two of these rows carry the design. The **frequency** row is the Q = ωτ argument of Section 1: at 1 THz the bit is a genuine wave, and passive loss (~4.3 dB per λ_p at 300 K, 5.1 dB at 80 °C) is steep enough to force distributed gain but shallow enough for a regenerative cell to beat.

The **density** row is the most constrained choice in the design, squeezed from both sides. From above: at n = 10^{13} cm^{-2} the drift-gain route fails twice over — the saturation velocity caps the drift Mach number below the gain threshold, and the drive heat would allow < 0.32 % duty under the 80 °C cap (Sections 6.1–6.2). From below: at n = 10^{11} cm^{-2} the logic swing becomes comparable to the charge-puddle inhomogeneity of even the best material, and cell yield collapses to zero (Section 6.4). n = 10^{12} cm^{-2} survives all constraints simultaneously, with the regenerative bias (Section 6.3) making that density thermally viable.

4. One cell, all gates: the regenerative threshold logic family

Every logic gate in the Fable Computer is the same object: a resonant DS cell — a quarter-wave gated-graphene cavity with asymmetric boundary conditions (source side density-clamped, drain side current-clamped) carrying a DC drift $v_0 = M\cdot s$, biased below its self-oscillation threshold M_th (Figure 3a). Its compact transfer model is a sigmoid with saturating rails:

A_out = A_pump · g(n_local), with g(n) = g_max · ½[1 + tanh((n − n_th)/Δn)] − ℓ, and n_local = n_bias − β·A_in.

Logic follows from bias: park n_bias above n_th and a HIGH input depletes the cell below threshold (gain off → output LOW) — inversion as a bias choice. AND and OR are the same cell with the threshold applied to the sum of two input fields: a high threshold fires only on

both inputs (AND), a low threshold on either (OR). Three quantitative properties make the family viable, all established in the model chain of Section 6:

- **Regeneration condition.** The cell restores levels only if the threshold sharpness k = β·A_swing/Δn exceeds ~5; the design rule is **k ≥ 8 with on/off extinction ≥ 10 dB**, giving a static noise margin ≈ 0.2 of the logic swing (0.26–0.35 at k = 16, 13–20 dB extinction).
- **Saturating rails.** The nonlinear solve shows 1-dB gain compression at $\delta\rho/\rho_0 \approx 1$ % and ~10 dB compression at 10 %: the logic swing belongs at the knee, and the rail is what re-clamps levels gate after gate (observed directly in the cascade solve, Section 6.4).
- **Bias ceiling.** Above ~0.8·M_th the resonator rings between clock pulses (5 dB pulse-to-pulse spread at 0.94·M_th): the usable bias window is roughly 0.5–0.8·M_th, with 0.7·M_th the working point.

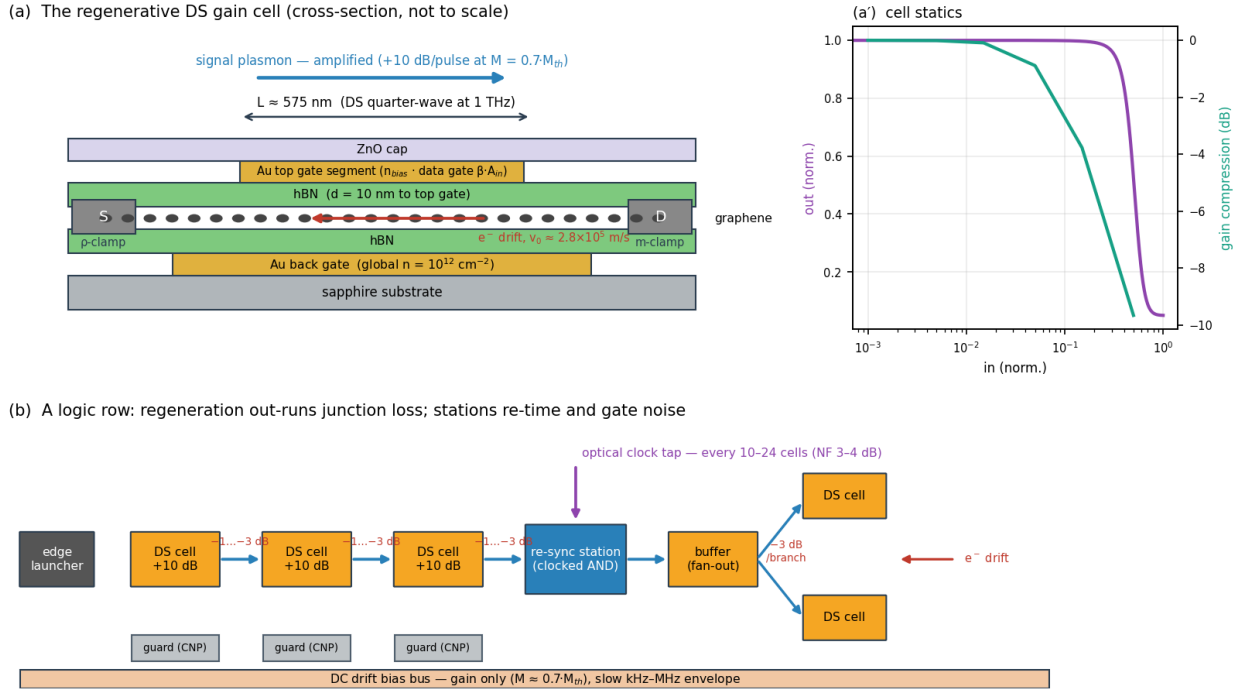

Figure 3: The regenerative DS gain cell and its use in a logic row. (a) Cross-section on the Section 2.1 stack: the quarter-wave cavity (L ≈ 575 nm) between a density-clamped source contact and a current-clamped drain, with the signal amplified against the electron drift. (b) A logic row: per-cell regeneration out-runs junction loss, a clocked re-sync station re-times pulses and blocks ASE every 10–24 cells, and fan-out branches from buffer cells. Insets: the inverting tanh transfer (k ≥ 8, ≥ 10 dB extinction) and the saturating rail.

5. Architecture

Figure 4 assembles the whole machine: the photonic chain of Section 2.2 on the left generates clock and data pulses; the fabric of DS cells amplifies and routes them under the

gate-map program; the DC bias bus below supplies gain only; the Kusyak elements read out on the right; and an electronic layer wraps the fabric for storage and trim.

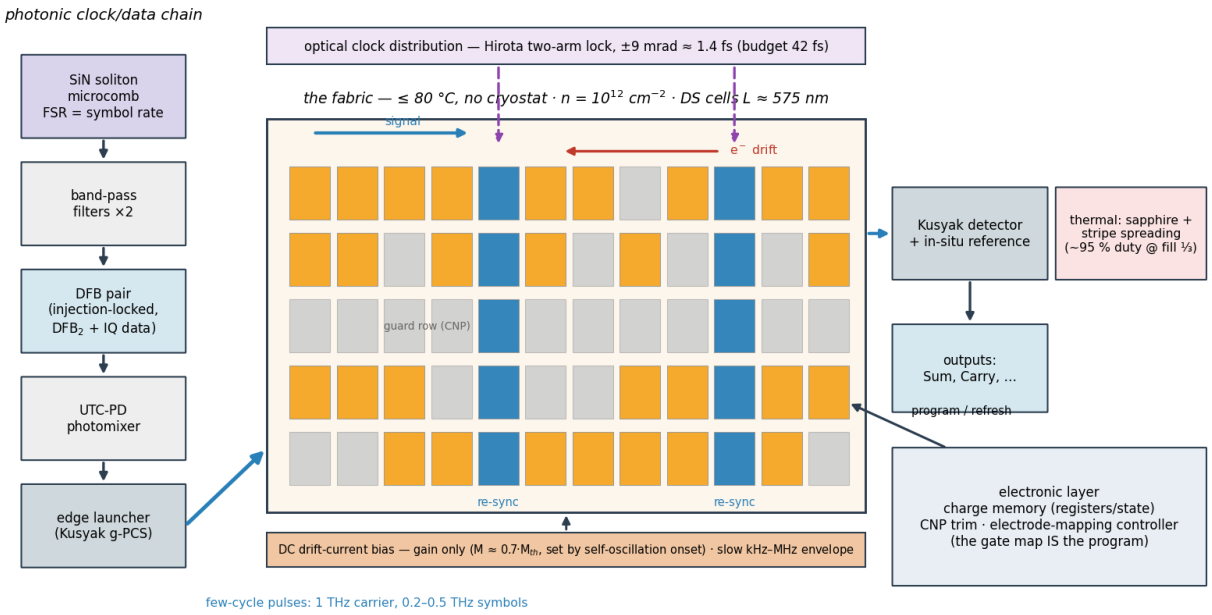

Figure 4: System view of the Fable Computer. Left: microcomb → injection-locked DFB pair (IQ data) → UTC-PD photomixer → edge launcher; few-cycle pulses at 1 THz carrier, 0.2–0.5 THz symbol rate. Center: the ≤ 80 °C fabric — rows of regenerative DS cells (one cell = one gate) with re-sync stations fed by optical clock taps, CNP guard rows, signal propagating against the electron drift. Bottom: the DC drift bias (gain only) with slow envelope gating. Right: monolithic read-out and the electronic charge-memory/trim layer.

5.1 Gain: a DC drift bias, directional and self-isolating

A uniform DC drift bias supplies the gain; the gate map programs which cells couple to it. Cells parked at the charge-neutrality point are dark and absorbing — the gate pattern defines both the signal paths and the gain, suppressing cross-talk. The amplifying wave propagates against the electron drift, so each region has a preferred amplification axis; the reverse direction is extra-damped by the same Doppler factor, $\gamma\cdot(1 + v_d/v_p) \approx -3.3$ to -3.8 dB per gate at $v_d/v_p = 0.3–0.5$. The directional gain therefore doubles as a built-in isolator: any reflective loop loses ~5 dB per cell pair even at forward transparency, structurally suppressing the classic self-oscillation failure of cascaded CW amplifiers.

5.2 Clock: the launcher pulse train

Timing is generated at the fabric's left-edge pulse generator — the Section 2.2 chain — emitting few-cycle 1-THz pulses at a 0.2–0.5 THz repetition rate. The DC bias carries no timing; it is gated only by a slow envelope (kHz–MHz, comfortably inside the bias-plane RC corner) for power management and noise flushing. The phase budget closes with thirty-fold margin (Section 2.2).

Three riders attach. (i) **Re-sync stations:** a launcher-only clock is a source-synchronous wave pipeline, so jitter accumulates; clocked-AND re-timing cells, fed optically by the same comb through mid-fabric Kusyak taps, are inserted every 10–24 gates — the spacing set by the noise budget below. (ii) **ASE gating:** the stations pass energy only during pulse slots, flushing amplifier noise at the same period. (iii) **Gain-recovery ISI:** with recovery $\sim\tau \approx$ 0.85 ps against 2–5 ps slots, pattern-dependent gain is benign in-model (< 1 dB spread up to 1 THz at ≤ 30 % per-pulse depletion; 0.38 dB measured directly in the pulse-train solve at 0.25 THz).

5.3 Fan-out

A Y-split costs 3 dB per branch plus junction loss. The regenerative window (Section 6.3) reaches +9–13 dB per cell at M = 0.12–0.15, so fan-out of 2 with restoration headroom is in-model feasible at the operating bias. Higher fan-out uses dedicated buffer tiles (Figure 3b).

5.4 Memory: compute in plasmons, store as charge

Stated plainly: there is no native plasmonic storage at ~1 ps lifetimes. The bit's persistent home is the charge on a cell's gate (DRAM-like, slow, refresh-limited); pumped bistable cells and recirculating loops serve only small fast registers. The Fable Computer is a plasmonic datapath wrapped in an electronic memory layer — a hybrid by design, not by concession.

5.5 Subsystem precedents

Every subsystem rests on a demonstrated technology; the fabric is the integration.

Subsystem	Demonstrated basis
Clock/data generation	microcomb-photomixing chain, chip-scale, phase-coherent [5]
Clock-distribution phase control	separate-arm optical locking, ±9 mrad [6]
Launch, read-out, in-situ reference	monolithic generator/detector stripline circuit [7]
Substrate and plasmon transport	sapphire / segmented-gate / hBN-graphene stack [4]
Free-space port (option)	graphene terahertz patch antenna [8]
Per-cell gain mechanism	current-driven plasmon amplification at 300 K [2]

5.6 The machine: a half adder

The Fable Computer computes the textbook decomposition S = (A OR B) AND NOT(A AND B), C = A AND B, mapped onto **five active cells** — OR (low threshold), AND (high threshold, doubling as the carry source), NOT (inverting bias), one delay buffer, and the output AND — plus two optional buffers that align the carry to the sum's slot (Figure 5). Both fan-outs in the netlist (the inputs feeding two gates each; the carry AND feeding both C and the NOT) are fan-out-2, exactly the budget Section 5.3 provides. The two inputs of the output AND

arrive through equal depth (OR→DLY vs AND→NOT), so the design needs no asymmetric timing trim; total logic depth is three cells.

A	B	OR	AND (= C)	NOT	**S**
0	0	0	0	1	**0**
1	0	1	0	1	**1**
0	1	1	0	1	**1**
1	1	1	1	0	**0**

The budgets close with room. **Gain:** the worst path (launch → −3 dB split → cell → −1 to −3 dB junction → cell → junction → cell → read-out) alternates +10 dB regeneration with single-digit losses, and the saturating rails re-clamp the level at every stage — the two-cell cascade solve (Section 6.4) is precisely this path's physics. **Timing:** at the 0.25 THz baseline symbol rate each stage adds well under one slot of latency, and the cascade solve preserved slot cadence through two cells; end-to-end latency is about one 4-ps slot, with one addition completing per slot thereafter (2.5×10^{11} additions/s per block, wave-pipelined). **Noise:** depth three sits far inside the 10–24-cell re-sync spacing, so the block needs no internal station — one station at the read-out boundary suffices. **Footprint:** ~11 × 7 lattice sites at λ_p-scale row pitch ≈ **12 × 16 µm**, including guard structure, clock tap, and read-out boundary. **Drive energy:** seven biased cells at ~1.7 kW/cm^2 over $\sim10^{-8}$ cm^2 each dissipate ~0.14 mW; at one addition per 4 ps that is **~0.6 fJ per addition** on the fabric side — with Section 7's insistence that the *system* energy (photonics, not fabric) dominates.

The block's defining property is that its outputs are launch-grade: S and C leave at full rail amplitude, re-timed at the boundary station, and can drive the inputs of a following block — two half adders and an OR compose a full adder; full adders compose a ripple-carry word. Nothing about the five cells is special to addition: they are the generic gate vocabulary (threshold-AND/OR, bias-inversion, buffering) of the fabric at large. The half adder is simply the smallest machine that exercises all of it — which is why it, and not a larger demonstrator, is what the feasibility chain of Section 6 and the bench protocol of Section 8 are sized against.

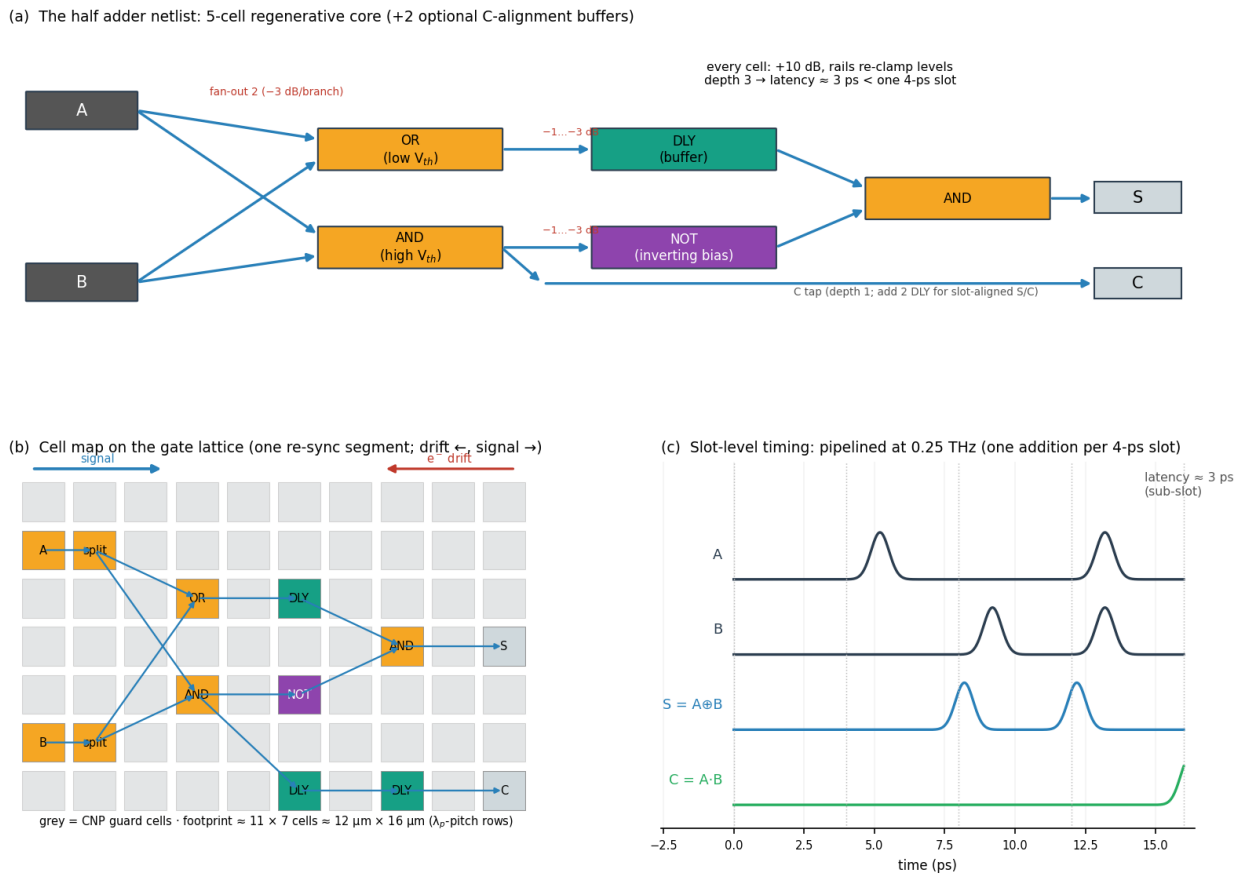

Figure 5: The Fable Computer half adder. (a) The netlist: five regenerative cells (OR, carry-AND, inverting-bias NOT, delay buffer, output AND) with fan-out-2 splits and per-edge junction losses annotated; two optional buffers slot-align the carry. (b) The cell map on the gate lattice — one re-sync segment, signal propagating against the electron drift, gray CNP guard cells. (c) Slot-level timing at 0.25 THz: the four input combinations of one pipelined add stream, with $S = A \oplus B$ and $C = A \cdot B$ emerging at sub-slot latency.

6. Feasibility: the five-pass chain

Each pass is a deliberately transparent reduced-order model with every parameter exposed. The chain is ordered so that each pass either kills the design or sharpens the next pass's question.

6.1 Pass 1 — bulk gain and brute thermal limits (the double no-go)

A drift-Doppler cascade model established two independent failures at the highest candidate density, $n = 10^{13}$ cm^{-2} — the natural first choice for a robust, well-confined plasmon. Bulk convective gain is impossible: graphene's saturation velocity sits below the acoustic-plasmon phase velocity ($v_{sat}/v_p \approx 0.3–0.6$) [9, 10], so uniform drift can at best halve the per-gate loss (net ≈ −1.3 dB/gate at $v_d = v_{sat}$) — transparency is unreachable, and the demonstrated 9 % room-temperature amplification [2] must be understood as a resonant boundary (Dyakonov–Shur) effect, not bulk gain. Independently, the drive heat at

that density (~174 kW/cm² at v_sat) caps the duty cycle at 0.32 % under the 80 °C budget. Both failures point the same way: lower density, resonant cells.

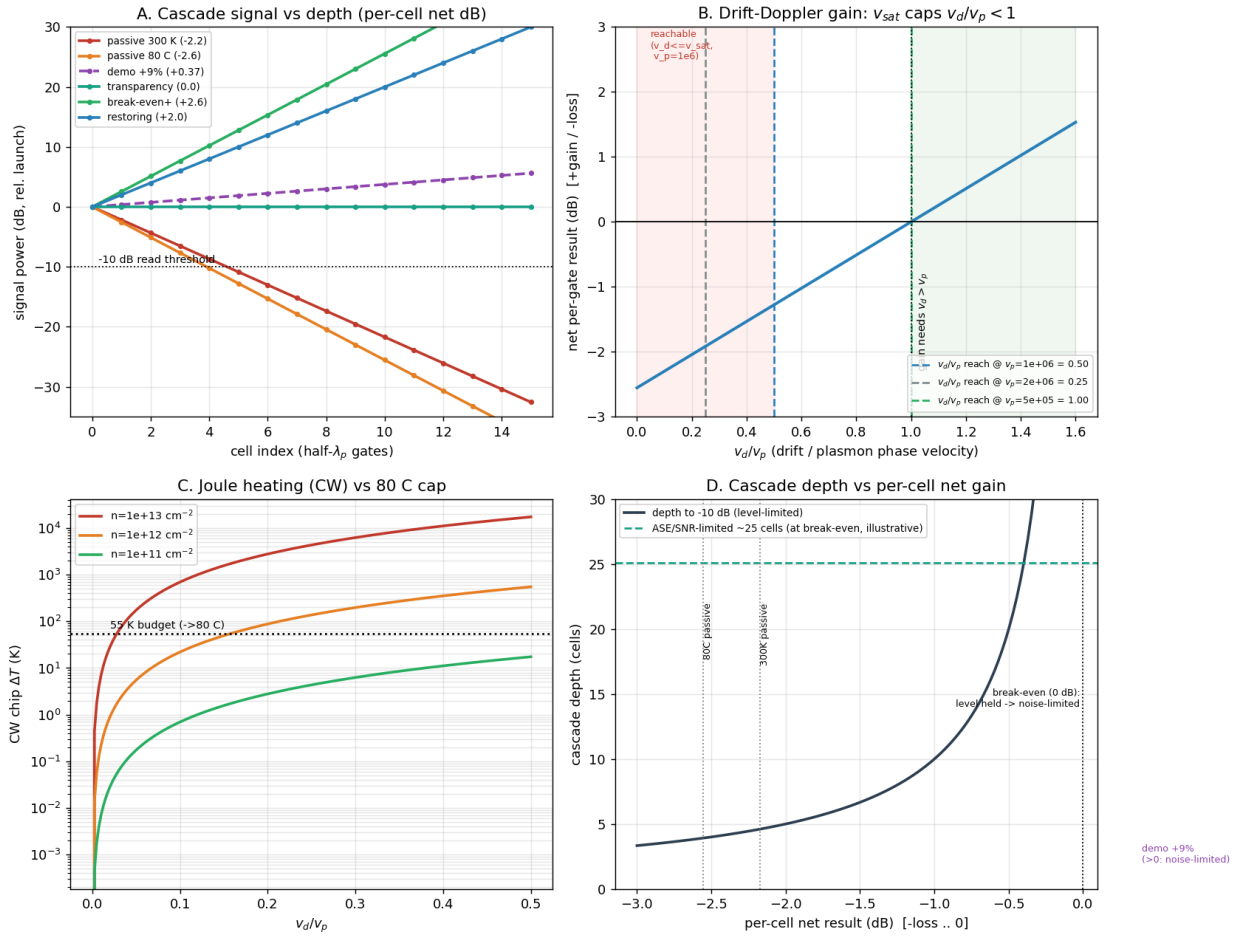

Figure 6: Pass-1 cascade model — passive decay, the v_sat < v_p ceiling, Joule heating against the 55 K budget, and depth-versus-gain.

6.2 Pass 2 — analytic resonant cells: break-even only at low density

Applying the DS instability increment [11] to a 1-THz-tuned quarter-wave cell, single-transit net gain at 80 °C reaches +2.97 dB at 10^{11} cm^{-2} and +2.55 dB at 3×10^{11} — but −0.15 dB at 10^{13}, where v_sat caps the Mach number at 0.13. The low-density requirement is thus a gain requirement, not only a thermal one. The same pass set the logic-margin design rule ($k \geq 8$, extinction ≥ 10 dB), bounded gain-recovery ISI as benign, and derived the re-sync spacing law: with launch SNR 20 dB and 6 dB margin, station spacing ≤ 24/8/3 cells for per-cell noise figures of 3/6/10 dB — making **NF ≲ 4 dB** the sharpest requirement on the cell. Single-transit fan-out gain (9–12 dB) was unreachable everywhere — a conclusion the next pass overturned.

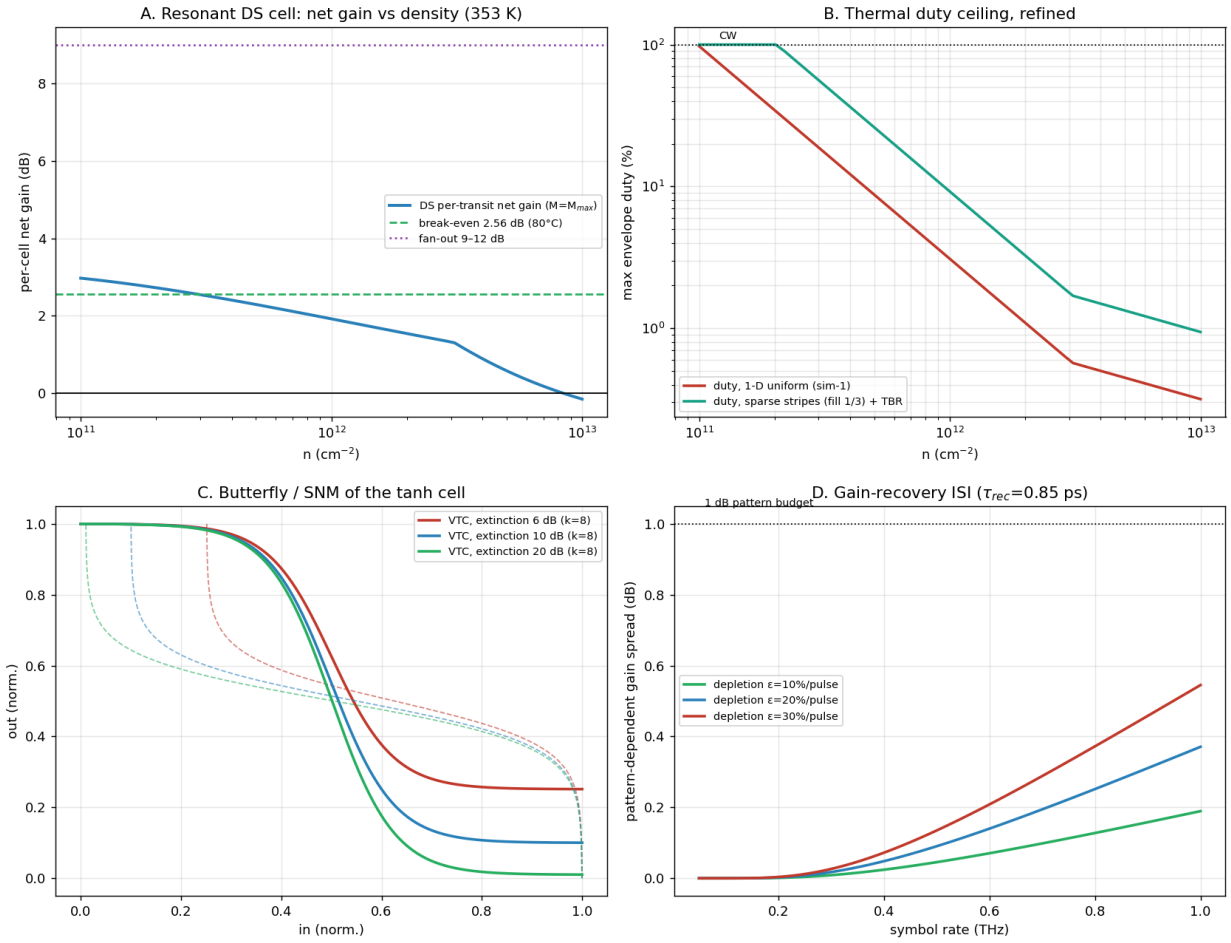

Figure 7: Pass-2 panels — DS transit gain vs density, refined duty ceilings, butterfly noise margins, ISI vs symbol rate.

6.3 Pass 3 — the regenerative operating point (nonlinear time-domain solve)

A nonlinear shallow-water solver with true DS boundary conditions revealed what the transit picture misses: below its oscillation threshold the cell is a multi-pass regenerative amplifier. Net gain referenced to the identical passive cell rises from break-even at M ≈ 0.05 through **+9.4 dB (CW) at M = 0.12** toward formal divergence at threshold (numerical M_th ≈ 0.176 vs analytic 0.146, the offset consistent with measured numerical dissipation). Driven with the actual clock waveform — 3-cycle pulses at 0.25 THz — the cell delivers **+10.0 dB per pulse with 0.38 dB spread at M = 0.12**, while near-threshold bias rings (5 dB spread at 0.94·M_th): gain-bandwidth resolves at moderate bias. Because the operating drift is far below v_sat, Joule heat falls ~25-fold against the at-v_sat assumption; the exact periodic-slab thermal solution gives **~95 % duty at n = 10^{12} with one-third fill**. Heat stops selecting the density — disorder takes over that role, and the Monte-Carlo selected **n = 10^{12} cm^{-2}** (91–100 % cell yield at realistic puddle levels; 0 % at 10^{11}, where the logic swing equals the puddle scale).

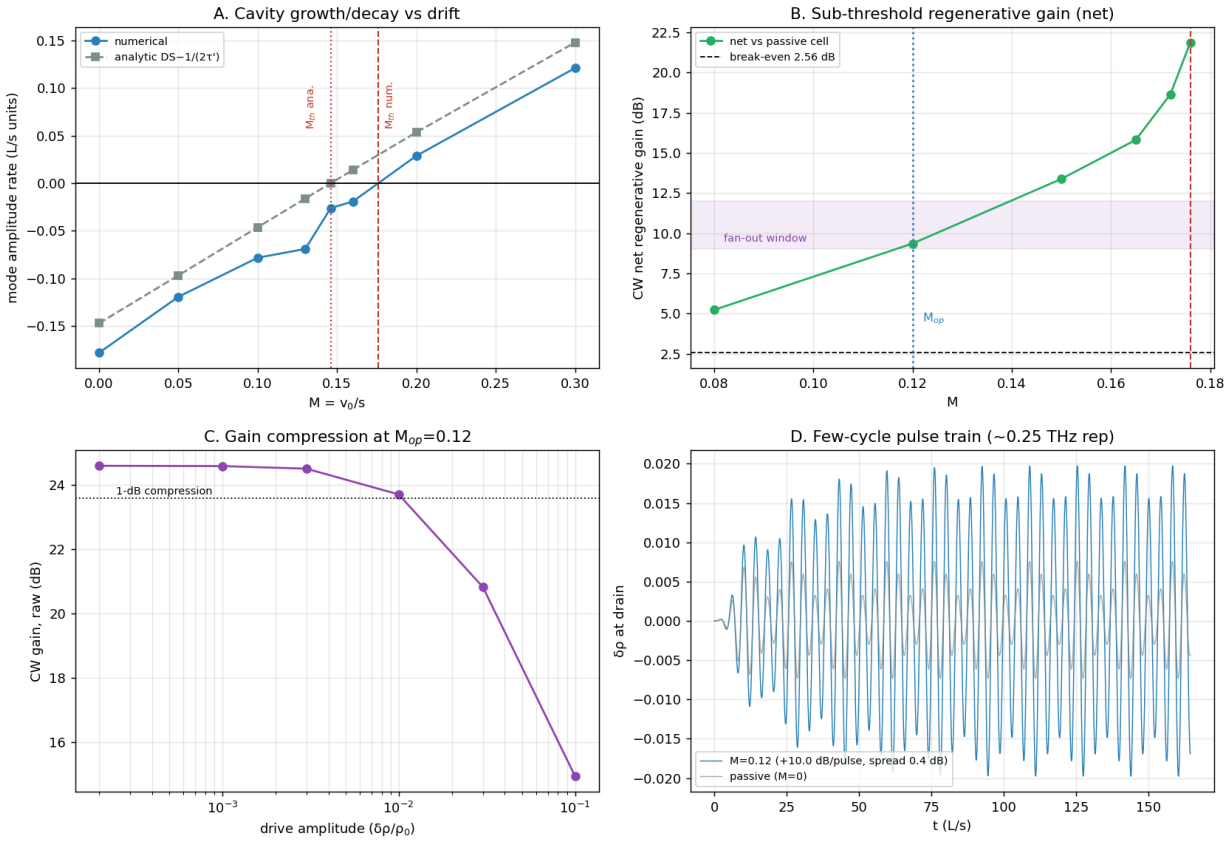

Figure 8: Pass-3 panels — cavity growth vs drift with both thresholds marked, sub-threshold regenerative gain through the fan-out window, gain compression (the logic rail), and the clean pulse-train eye at M = 0.12.

6.4 Pass 4 — cascade, noise, and measured disorder

The architecture's load-bearing claim — cascadability — was then tested directly: two cells chained through junction transmissions of 0/−1/−3/−6 dB sustain **+8.4 to +9.0 dB net per cell** (two-cell total +17–18 dB over the passive chain), with cell 2 performing identically on cell 1's distorted output as on a clean train, and per-cell gain rising slightly as junction loss moves the input below the saturation knee — the restoring rail in action. Distortion accumulates sub-linearly (0.38 → 0.66 dB over two cells). A Langevin run found the cell amplifies channel noise by +6.8 ± 1.9 dB against +9.3 dB signal gain — **no excess above the 3.0 dB thermal floor** (F = 2 − 1/G at matched temperature), so the NF ≤ 4 dB requirement is attainable in-model. Finally, literature disorder values ($\sigma \approx 2.5\times10^{9}$ cm^{-2} for graphene on ultra-flat hBN [12]; $\sim4\times10^{10}$ cm^{-2} typical encapsulated [13, 14]) place cell yield at n = 10^{12} at **98–100 % across the entire realistic range**; 3×10^{11} is a premium-material option (graphite-gate class).

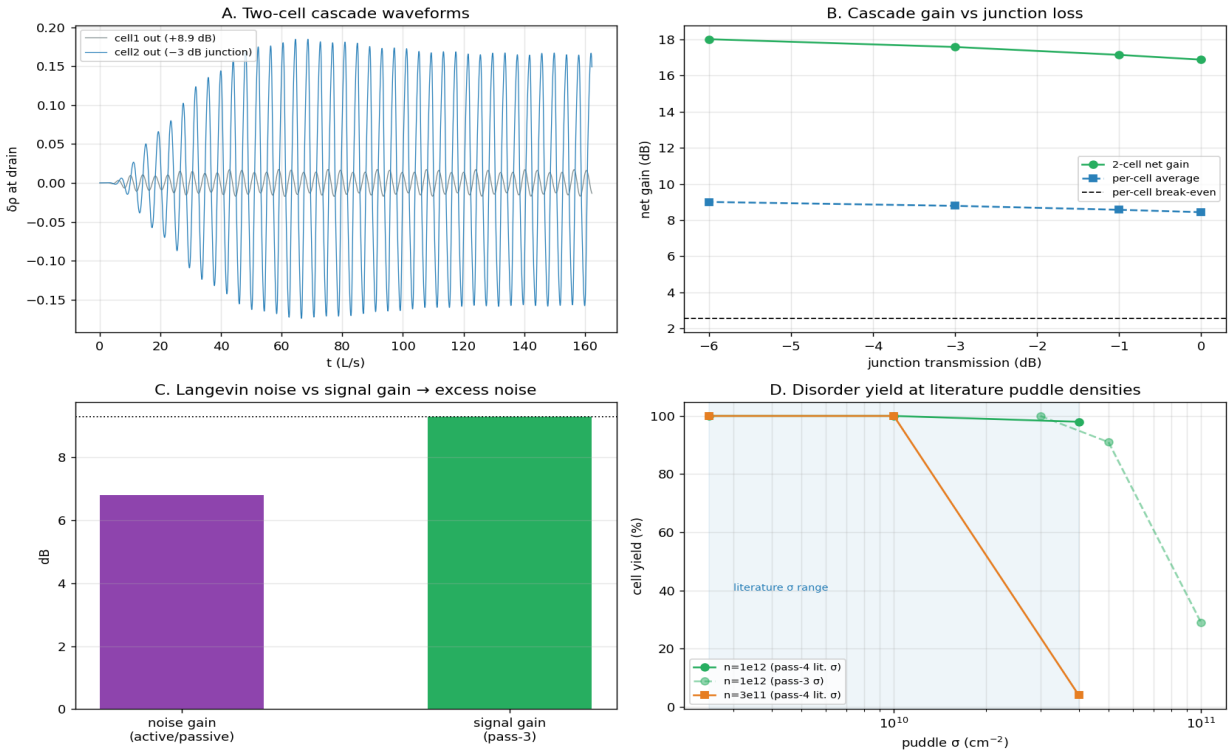

Figure 9: Pass-4 panels — two-cell cascade waveforms, net gain vs junction loss, noise-vs-signal gain, and disorder yield against the literature σ band.

6.5 Pass 5 — validity bounds and the last parasitic channel

Two checks close the reduced-order chain. Kinetic validity: at the operating point $\omega\tau_{ee} = 0.053$, electron-viscous damping is 1.3 % of $1/\tau$, and $L/\ell_{mfp} = 0.68$ — the hydrodynamic description is the correct tier at 1 THz [15], its gain numbers carrying ~5 %-scale kinetic uncertainty, not factors. Coupled-cell stability: two DS-active segments on one continuous channel, with a midpoint density clamp swept from hard contact to fully free, are **stable at every coupling strength**, even at $0.9 \cdot M_{th}$ — the DS increment's $1/L$ scaling protects long parasitic cavities. Charge-channel parasitics are therefore excluded in-model; the only remaining coupling channel is radiative/electromagnetic, which is precisely scoped to the Maxwell tier.

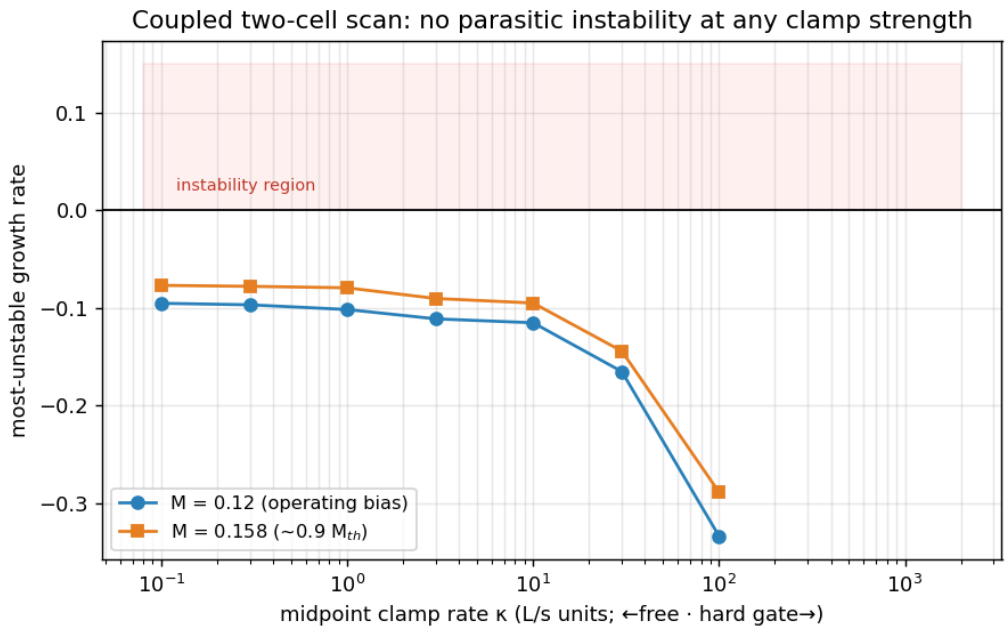

Figure 10: Pass-5 coupled two-cell scan — most-unstable growth rate vs clamp strength at both bias points; no instability region is entered.

6.6 Scoreboard

Question	Status after five passes
Net gain ≥ loss at 1 THz, ≤ 80 °C	in-model yes: regenerative cell, +10 dB/pulse at 0.7·M_th
Cascadability	in-model yes: 2 cells, +8.4–9.0 dB/cell through −6 dB junctions
Noise	NF at the 3 dB thermal floor; re-sync every 10–24 gates
Disorder	98–100 % yield at n = 10^{12} for literature σ; measured on-die in the protocol
Thermal	~95 % duty at fill ⅓; not binding at the regenerative bias
Kinetic-tier risk	bounded at ~5 %
Coupled-cell parasitics	charge channel stable at all couplings; EM channel open
Remaining	Boltzmann–Maxwell solve (EM coupling, absolute calibration) and the bench

7. Energy and thermal honesty

The intrinsic field energy is ~10 aJ per bit, and with 10^4–10^5 active cells a 1 W-class drive corresponds to ~10–100 aJ/op — but only at full utilization, and the explicit 0.2–0.5 THz symbol rate already places the realistic floor 2–5× higher. The system energy — microcomb wall-plug, DFB lasers, photomixers, detectors — is far larger; any “orders below CMOS” claim survives only for the intrinsic field energy, never for the system [16,

17]. The half adder's own accounting (Section 5.6) makes the same point concretely: ~0.6 fJ per addition on the fabric side is CMOS-class, and the photonic chain above it costs more. Thermally, the ≤ 80 °C cap bounds drive power, envelope duty and active-area density; at the operating point those bounds are met with margin (ΔT = 57.6 K at worst-case fill ⅓ against a 55 K budget at ~95 % duty), and the positive-feedback loop (T → τ → Q → loss → more drive) is defanged by the 25-fold heat reduction of threshold-proximate bias.

The cap also has a cold side. Re-running the pass-3 cell with the $\tau(T) = \tau(300\ K)\cdot(300/T)$ loss model and the pass-5 kinetic bounds shows the operating point improving monotonically under moderate cooling — per-gate loss falls from 2.17 dB at 300 K to 1.45 dB at 200 K and 1.08 dB at 150 K, the thermal noise floor drops with T, and the servo-biased cell ($0.7\cdot M_{th}(T)$, re-calibrated by approach to self-oscillation exactly as the Section 8 protocol prescribes) remains regeneratively stable at every temperature swept, pulse-to-pulse spread staying far inside the protocol's 1 dB gate down to 115 K. That stability is conditional on bias tracking: M_{th} scales roughly as T, so a fixed bias calibrated at 353 K crosses the $\sim 0.8\cdot M_{th}$ ringing onset after only ~45 K of cooling — for an untracked bias, cooling, not heating, is the unstable direction. The floor is set by model validity rather than by any instability: $\omega\tau_{ee}$ grows as $1/T^2$, reaching 0.5 near 114 K and 1.0 near 80 K, where electron-viscous damping simultaneously overtakes phonon loss and the viscous-corrected threshold, minimal near 105 K, turns upward — the hydrodynamic tier that certifies every gain number here expires at roughly 80–115 K. In-model, then, stable operation extends from the 353 K cap down to ≈ 115 K under tracked bias: Peltier and dry-ice stages sit comfortably inside the envelope, while 77 K lies just below the certifiable floor and belongs to the Boltzmann–Maxwell solve of Section 8. (Below ~150 K the 1/T phonon scaling is itself optimistic — impurity scattering saturates τ — but the validity floor, set by τ_{ee} alone, does not move.)

8. Falsifiability: the experiment and the solve

The bench protocol — a lab-ready measurement plan that fixes the device, the drive and read-out chain, the calibration sequence and the pass/fail gates in advance, each threshold parameterized by the model predictions so that any large deviation is itself a result that recalibrates the model chain rather than a tuning failure — specifies the decisive test: an hBN-encapsulated dual-gated DS cell on the Section 2.1 stack, L ≈ 575 nm, d = 10 nm, n = 10^{12} cm^{-2}, with M_{th} calibrated in situ by approach to self-oscillation and bias set at $0.7\cdot M_{th}$. Five gates: G1 — single-cell per-pulse net gain ≥ +2.6 dB at ≤ 80 °C (target +10 dB); G2 — pulse-to-pulse spread < 1 dB at 0.2–0.25 THz; G3 — NF ≤ 6 dB (target ~3 dB); G4 — two cascaded cells ≥ 1.7× the single-cell decibel gain (target ~2×); G5 — thermal compliance at operating duty. The die plan includes passive twins for differential measurement, junction variants, a d = 5 nm variant ($s \propto \sqrt{d}$ is a +2.6 dB-class gain lever), and a graphite-gated disorder de-risk die; the calibration sequence measures on-die every quantity this manuscript has had to assume, beginning with the disorder σ itself.

The remaining modeling tier — a Boltzmann–Maxwell grating-gate solve, the method demonstrated for drift-driven terahertz gain in [18] — owns four questions with fully specified inputs: the regenerative G(M) curve under kinetic and electromagnetic

corrections; real junction loss; radiative coupling between cells (the one surviving parasitic channel); and absolute noise calibration against the 3 dB floor.

9. Honest bottom line

The design's feasibility rests on a single stated physics question — can room-temperature graphene deliver net plasmon gain at ~1 THz? Five passes of reduced-order modeling have answered it as well as reduced-order physics can: yes in-model, by operating a resonant DS cell as a sub-threshold regenerative amplifier, at a density chosen by disorder rather than heat, clocked optically because electrical clocking was never physically available, cascaded without gain degradation, at the thermal-noise floor. None of these statements is yet a measurement. The Fable Computer should be judged by the protocol gates of Section 8, and the most informative outcome — pass or fail — is G4: two cells, one junction, the first concatenation of room-temperature plasmonic gain ever attempted. It is the same property that lets the half adder of Section 5.6 drive its successor — the property that makes five cells a computer rather than a demonstration. The model chain is exhausted; the design is now exactly as falsifiable as it was always meant to become.

References

1. R. Furui, "A 100-GHz Room-Temperature Half Adder Using Graphene Plasmons as Bits, Driven by a Microcomb-Photomixing Carrier: Architecture, Quantitative Roadmap, and Machine-Learning-Assisted Design Optimisation," https://github.com/r-coin/basic/blob/master/hardware%20solusions/THz_Half_Adder_100GHz_RT.pdf
2. S. Boubanga-Tombet, W. Knap, D. Yadav, A. Satou, D. B. But, V. V. Popov, I. V. Gorbenko, V. Kachorovskii, T. Otsuji, "Room-Temperature Amplification of Terahertz Radiation by Grating-Gate Graphene Structures," Phys. Rev. X 10, 031004 (2020).
3. V. Ryzhii, M. Ryzhii, T. Otsuji, "Negative dynamic conductivity of graphene with optical pumping," J. Appl. Phys. 101, 083114 (2007).
4. K. Yoshioka, G. Bernard, T. Wakamura, et al., "On-chip transfer of ultrashort graphene plasmon wavepackets using terahertz electronics," arXiv:2311.02821 (2023).
5. Y. Tokizane, H. Kishikawa, T. Kikuhara, M. Talara, Y. Makimoto, K. Yamaji, Y. Okamura, K. Nishimoto, E. Hase, I. Morohashi, A. Kanno, S. Hisatake, N. Kuse, T. Nagatsuma, T. Yasui, "Beyond 350 GHz: Single-channel 112 Gbps photonic wireless transmission at 560 GHz using soliton microcombs," Commun. Eng. 5, 77 (2026), doi:10.1038/s44172-026-00659-8.
6. K. Hirota, T. Kashiwazaki, G. Ha, T. Yamashima, P. Jaturaphagorn, T. Suzuki, K. Takahashi, A. Kawasaki, A. Inoue, W. Asavanant, M. Endo, T. Umeki, A. Furusawa, "Generation of 10-dB squeezed light from a broadband waveguide optical parametric amplifier with improved phase locking method," Opt. Express 34, 7958–7972 (2026), doi:10.1364/OE.585323.
7. K. Kusyak, B. Schulte, T. Matsuyama, G. Kipp, H. M. Bretscher, M. W. Day, G. Meier, A. M. Potts, J. W. McIver, "Monolithic optoelectronic circuit design for on-chip terahertz applications," APL Photonics 10, 076117 (2025), doi:10.1063/5.0278183.

8. Md A. Haque, Md S. Akhter, I. Das, T. U. Zaman, J.-J. Tiang, N. S. S. Singh, A. D. Algarni, A. A. Ateya, “Graphene based terahertz MIMO antenna with machine learning regression for 6G communications,” Sci. Rep. 16, 3900 (2026), doi:10.1038/s41598-025-32487-9.
9. V. E. Dorgan, M.-H. Bae, E. Pop, “Mobility and saturation velocity in graphene,” Appl. Phys. Lett. 97, 082112 (2010).
10. “Transition from acoustic plasmon to electronic sound in graphene,” Sci. Adv. 9, eadi0415 (2023).
11. M. Dyakonov and M. Shur, “Shallow water analogy for a ballistic field effect transistor: New mechanism of plasma wave generation by dc current,” Phys. Rev. Lett. 71, 2465 (1993).
12. J. Xue et al., “Scanning tunnelling microscopy and spectroscopy of ultra-flat graphene on hexagonal boron nitride,” arXiv:1102.2642 (charge fluctuations $\sigma \approx 2.5\times10^{9}$ cm^{-2}).
13. “Estimation of residual carrier density near the Dirac point in graphene through quantum capacitance measurement,” arXiv:1304.3957.
14. “Ultra-long carrier lifetime in neutral graphene-hBN van der Waals heterostructures under mid-infrared illumination,” Nat. Commun. 11, 906 (2020).
15. A. Lucas and K. C. Fong, “Hydrodynamics of electrons in graphene,” J. Phys.: Condens. Matter 30, 053001 (2018).
16. D. A. B. Miller, “Attojoule optoelectronics for low-energy information processing and communications,” J. Lightwave Technol. 35, 346–396 (2017).
17. International Roadmap for Devices and Systems (IRDS), 2023 ed., IEEE.
18. “Terahertz Amplification Induced by Electron–Phonon Interactions in Gated Graphene Plasmonic Systems,” Research (2025), DOI 10.34133/research.1023.