COMMENT

# Beyond the Pockels coefficient: bringing electro-optic materials from lab to fab

**Electro-optic modulators are becoming a bottleneck for AI interconnects and quantum hardware, and newer materials promise much stronger responses than today's. Whether any of them reaches the market will depend less on its Pockels coefficient than on where its wafers come from, a question that research, funding and investment rarely address directly.**

---

In 2018, integrated thin-film lithium niobate (TFLN) modulators were shown to run at CMOS-compatible voltages and carry 210 Gb/s [1]; by March 2026 it had a qualified high-volume line and a foundry partnership for 8-inch production [2]. Barium titanate ($BaTiO_3$, BTO) integrated on silicon was reported within months, with a largest Pockels coefficient more than 25 times that of lithium niobate [3]. Its path to volume since then has run largely through one quantum-computing company and, from 2026, a 300-mm process developed by an equipment maker with imec [4,5], and designers outside those partnerships have had little access to an open foundry process.

Both materials address a growing need. Transceivers at 1.6 Tb/s are entering AI data centres, with 3.2 Tb/s modules and co-packaged optics in preparation [6]; photonic quantum computers need fast, low-loss switches to route single photons [7]; and LiDAR and free-space links add further demand. In all of these the modulator accounts for much of the power and bandwidth budget.

Why, then, did the material with the weaker effect reach volume first? Lithium niobate had real advantages, among them decades of device experience and a stable single-domain crystal. Our view, though, is that above a certain performance threshold adoption depends mainly on something papers seldom measure: whether a material can obtain wafers, process tools and qualification data at a price and pace set by a much larger industry.

In investment diligence on early-stage photonics companies, one of us has met founders who could describe their device performance in detail but not where their wafers would come from at volume, and investors who raised the concentration of TFLN wafer supply as a concern of their own.

## Every platform borrows a supply chain

Silicon photonics shows the pattern most clearly. It runs on silicon-on-insulator wafers, more than three million of which are made each year for microelectronics [8]; photonics inherited that wafer supply, together with the tools and process control, without paying for it. Wang and colleagues make the same point about lithium tantalate, whose insulator-bonded wafers are already produced for 5G radio-frequency filters (pro-

jected capacity, 750,000 wafers a year), and attribute TFLN's higher wafer cost to the absence of a comparable host market [8].

We see three ways in which emerging electro-optic materials reach scale. The first is to inherit a supply chain outright, as silicon and lithium tantalate have done. The second is to enter as an add-on to a host platform: electro-optic polymers and plasmonic slots are processed in the back end of existing silicon photonics lines, so the wafer, the foundry and much of the qualification come with the host [9,10]. The third is to find an anchor customer prepared to build the supply chain. PsiQuantum needed low-loss switches that operate at cryogenic temperature, built a 300-mm BTO epitaxy line for the purpose and integrates those wafers with a commercial foundry flow [5,7].

Inheritance ties a material to its host's wafer specifications, and an add-on must work within the host's limits; a polymer, for example, has to survive the packaging that silicon photonics already uses. An anchor customer may build a supply base tuned to its own requirements and of limited use to others. Because silicon and indium phosphide modulators continue to improve, a newcomer will also be compared with the incumbents as they stand when it reaches the market.

## What thin-film lithium niobate teaches

TFLN does not fit neatly into any of these routes, which is why its rise is worth examining. Its wafers are bespoke, made by ion slicing and bonding at a handful of specialist vendors; what it borrowed was fab capacity and people. HyperLight, a Harvard spin-out, qualified a high-volume line at Wavetek, a compound-semiconductor foundry, and in March 2026 added the 8-inch capacity of Wavetek's parent company, UMC [2]. CCraft, spun out of CSEM, is building a dedicated TFLN foundry with a target of up to 20,000 200-mm wafers a year [11]. The process step that opened the field, low-loss dry etching of the film [1], came from academic laboratories, but it was the spin-outs that made manufacturing their main problem, and demand from AI interconnects happened to arrive as their lines came up.

TFLN devices still show d.c. bias drift and photorefractive effects, which complicate bias control in deployed links, and concern about lithium contamination keeps the material out of many mainstream silicon fabs. Wafer supply is also concentrated. In February 2026 the US Air Force Research Laboratory contracted Raytheon to establish domestic TFLN wafer production, citing consolidation among suppliers [12].

## The next contenders

All three platforms now have some route to wafers, but none has yet built up the reliability data and open design kits that TFLN is starting to accumulate (Table 1). Of the three, the add-on platforms are closest to products, mainly because they borrow so much from silicon photonics. BTO has by far the strongest effect, but its 300-mm supply still rests largely on one buyer's investment, and merchant wafers together with an open process would allow other companies to design with it.

**Table 1 | Emerging electro-optic platforms by route to scale, late 2026**

| Platform | Headline performance | Route and status | Gate that remains |
|---|---|---|---|
| BTO on silicon | $r_{42} \approx 923$ pm $V^{-1}$ on Si, 50 Gb/s [3]; operation down to 4 K [13] | Anchor customer: 300-mm MBE line backed by a $100M CHIPS award [5]. Also a 300-mm process from Veeco and imec [4] and an 8-inch merchant wafer facility with initial production planned for 2026 [11] | Open foundry processing and design kits [11]; ferroelectric domain and hysteresis control; long-term stability data |
| Electro-optic polymers | Foundry reference designs targeting 110 GHz and beyond [9] | Add-on: designs entering Tower Semiconductor's PH18 design kit, engineering tape-outs in 2026 [9] | Degradation above about 80 °C and need for hermetic sealing [11]; de-poling over lifetime |
| Plasmonic–organic hybrids | Low-loss plasmon-assisted modulation [14] | Add-on: transmit chips on a 200-mm silicon photonics line sampled to three transceiver makers [10]; developer acquired by Marvell in 2026 [6] | Optical loss budget, stability of the organic layer and, above all, volume reliability data |

## Free space raises the stakes

Most of the discussion above concerns waveguide modulators, in which a weak electro-optic effect can be offset by a longer device. Free-space modulators have no such option, because light crosses a film only a fraction of a micrometre thick. A single pass through 500 nm of lithium niobate at 1,550 nm, driven by a CMOS backplane, gives a phase shift of about 0.002 rad [15]. Resonant metasurfaces and guided-mode resonances can make up the difference, at some cost in optical bandwidth and fabrication tolerance.

Commercial spatial light modulators have megapixel apertures but switch at kilohertz rates [15]. Gigahertz free-space modulation has been shown with organic electro-optic layers on silicon metasurfaces [16], and TFLN has been bonded to a megapixel CMOS backplane with a modulation bandwidth of 1.6 GHz [15], but megapixel scale and gigahertz speed have yet to be combined in one device. Such a device would find uses in free-space optical links, atom- and ion-based quantum computers, imaging through scattering media and optical computing.

Here a large coefficient does help, and BTO is one of several candidates. The supply question still applies, however. The most likely host is the CMOS backplane already used in liquid-crystal-on-silicon displays [15], and whether an electro-optic film can be grown or bonded onto such a backplane with uniform properties over a centimetre-scale aperture is a question for materials scientists.

## What each community can do

**Researchers.** We suggest reporting a minimum set of industry-relevant metrics alongside the headline coefficient (Box 1), with statistics over many devices as well as the best one. Bias drift over hours, damp heat and optical power handling are what cause failures in the field and deserve measurement. Authors could also state which route to scale a result assumes and note when it relies on a process step that no foundry offers, and journals could request such information whenever a paper claims technological relevance.

**Funders.** Where no host industry exists, supply has to be paid for by someone, and public money is often the only realistic source at the start. Shared pilot lines and multi-project wafer runs for BTO and TFLN, with open process design kits, would let academic groups and start-ups test manufacturability without building fabs of their own. Reliability testing also deserves support as research, since it rarely wins grants on novelty. The 2026 public contracts for TFLN wafers [12] and 300-mm BTO [5] show that governments will pay when the strategic case is clear; a useful next step would be to open such capacity to more than one user.

**Investors and industry.** Diligence on a materials company should establish which route it is on and who pays for its wafers, as well as how its devices perform. Joint benchmarking among material suppliers, foundries, transceiver makers and end users would produce reliability data years earlier than any of them could alone. Strategic acquirers are already buying device-level innovation [6], so exit paths exist, but timing is a problem. TFLN took about eight years to go from a landmark demonstration to a foundry partnership for 8-inch production, whereas a venture fund investing in its third year typically has about seven years left. Patient capital, milestone-based financing and co-investment with strategic partners suit such companies better than standard early-stage terms.

**Box 1 | A minimum lab-to-fab reporting panel for electro-optic materials**

| Category | Report |
|---|---|
| Wafer | Diameter, source, film thickness and its spread across the wafer |
| Loss | Propagation loss (dB $cm^{-1}$) with method and waveguide length; insertion loss |
| Efficiency | $V_\pi$ at stated device length, $V_\pi L$, electro-optic bandwidth |
| Statistics | Number of devices measured, mean and spread, yield |
| Stability | Bias drift under sustained d.c. voltage, with duration and temperature; optical power handling |
| Environment | Operating temperature range; damp-heat result (for example 85 °C, 85% relative humidity) if tested |
| Process | Maximum thermal budget; non-standard steps; compatibility with CMOS lines |

Thin-film lithium niobate reached high-volume foundry production within a decade of its landmark demonstration. For the materials that follow, we hope the question of where the wafers will come from is asked as early as the question of how large the coefficient is.

## References


1. Wang, C. et al. Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages. *Nature* **562**, 101–104 (2018).
2. HyperLight, UMC, and Wavetek announce strategic partnership for high-volume foundry production of TFLN Chiplet platform. *Business Wire* (11 March 2026).
3. Abel, S. et al. Large Pockels effect in micro- and nanostructured barium titanate integrated on silicon. *Nat. Mater.* **18**, 42–47 (2019).

4. Veeco and imec develop 300mm-compatible process to enable integration of barium titanate on silicon photonics. imec press release (27 January 2026).
5. PsiQuantum executes definitive $100 million CHIPS Act award with U.S. Department of Commerce. *Quantum Computing Report* (8 September 2026).
6. Marvell acquires Polariton Technologies in bid to enhance its silicon photonics capabilities. *Lightwave* (April 2026).
7. PsiQuantum Team. A manufacturable platform for photonic quantum computing. *Nature* https://doi.org/10.1038/s41586-025-08820-7 (2025).
8. Wang, C. et al. Lithium tantalate photonic integrated circuits for volume manufacturing. *Nature* **629**, 784–790 (2024).
9. Lightwave Logic and Tower Semiconductor announce development agreement. *Newswire* (March 2026).
10. The 400G-per-lane reckoning: Polariton ships plasmonic modulator samples. *Fibre Systems* (February 2026).
11. Nikitski, I. Materials beyond silicon: the industrial push for advanced integrated photonics. *Laser Focus World* (8 December 2025).
12. RTX's Raytheon to accelerate domestic supply of critical material for commercial and defense applications. RTX press release (17 February 2026).
13. Eltes, F. et al. An integrated optical modulator operating at cryogenic temperatures. *Nat. Mater.* **19**, 1164–1168 (2020).
14. Haffner, C. et al. Low-loss plasmon-assisted electro-optic modulator. *Nature* **556**, 483–486 (2018).
15. Trajtenberg-Mills, S. et al. LNoS: lithium niobate on silicon spatial light modulator. Preprint at arXiv:2402.14608 (2024).
16. Benea-Chelmus, I.-C. et al. Gigahertz free-space electro-optic modulators based on Mie resonances. *Nat. Commun.* **13**, 3170 (2022).