

Group-Delay Tolerances for Coherent Combining in Tiled Apertures Driven by Photonic Integrated Circuits for Optical Communication

Deniz Onural
NovaLink Space Inc.
Cambridge, MA, USA
deniz@novalink.space

Paul Serra
NovaLink Space Inc.
Cambridge, MA, USA
paul@novalink.space

Idan Ezer
NovaLink Space Ltd.
Holon, Israel
idan@novalink.space

Andrey Nazarov
NovaLink Space Ltd.
Holon, Israel
andrey@novalink.space

Itay Tamir
NovaLink Space Ltd.
Holon, Israel
itay@novalink.space

Yekutieli Uliel
NovaLink Space Ltd.
Holon, Israel
kuti@novalink.space

Elimelech Keller
NovaLink Space Ltd.
Holon, Israel
elimelech@novalink.space

Eyal Wohlgemuth
NovaLink Space Inc.
Potomac, MD, USA
eyal@novalink.space

Abstract—Photonic integrated circuits (PICs) enable coherent signal distribution and per-element phase control across tiled aperture arrays for high-throughput free-space optical communication. Physical path-length variation across the PIC waveguide routing, polarization and filtering stages, fiber fan-out, and mechanical assembly produces inter-element differential group delays on the order of picoseconds. On-chip phase shifters provide only femtoseconds of group-delay equivalence, insufficient to bridge the gap through such phase control alone. The resulting delay spread degrades combining efficiency in proportion to signal bandwidth. As channel bandwidths grow while element counts increase to close link budgets, the bandwidth-delay interaction becomes a scaling constraint specific to the tiled architecture, absent in single-aperture terminals.

We derive a closed-form expression for the bandwidth-dependent combining efficiency loss from inter-element delay variance. We decompose the per-element group delay into PIC, fiber, amplifier, and free-space contributions. The fiber interconnect emerges as the dominant random source. We quantify achievable length-matching tolerances for standard and precision assembly methods. For aggregate bandwidths on the order of hundreds of gigahertz, sub-millimeter path mismatches produce multi-decibel combining penalties. The model returns the maximum tolerable delay spread for a given signal bandwidth and target efficiency, with implications for transceiver channel bandwidth, modulation order, and equalization requirements. We compare the analysis with delay-variance measurements from a seven-element coherent combining terminal, confirming the predicted delay decomposition across the full optical path.

Where the measured delay variance exceeds the tolerable bound for target link performance, compensation is required. We evaluate approaches at three levels of the system: fiber-level path trimming, on-chip switched delay lines, and transceiver-domain equalization. For each, we quantify the trade-off between delay-variance reduction and the penalties introduced: insertion loss, switch extinction, and added amplitude spread for the photonic approaches; bandwidth and complexity overhead for electronic equalization. The resulting framework provides a unified delay budget that spans from fiber assembly tolerances through on-chip photonic correction to transceiver equalization, applicable to any PIC or fiber driven coherently combined aperture architecture.

The delay budget translates into fiber assembly specifications, PIC delay allocation, and transceiver equalization depth for a given target data rate.

Index Terms—free-space optical communications, photonic integrated circuits, coherent combining, group delay, tiled aperture, aperture scaling, optical phased array

I. INTRODUCTION

Optical inter-satellite links are moving from tens of gigabits toward terabit class. Of the few terms available to raise link gain, aperture is the most effective, since received power scales as D_{eff}^4 for a symmetric link between identical terminals [1], with D_{eff} [m] the effective aperture diameter. A tiled aperture supplies that diameter without a large monolithic telescope, using an array of small sub-apertures fed by a photonic integrated circuit (PIC) through a fiber-coupled antenna and trading optical fabrication for photonic integration and repeated optical channels [1]–[3]. The efficiency of such an aperture decomposes as

$$\eta_{\text{ap}} = \eta_{\text{pattern}} \eta_{\text{cbc}} \eta_{\text{point}} \eta_{\text{delay}}, \quad (1)$$

in which η_{pattern} captures array-factor spill and scan loss, η_{cbc} coherent combining loss from amplitude nonuniformity and residual carrier phase, η_{point} pointing and fiber-coupling loss, and η_{delay} the bandwidth-dependent loss from differential group delay across elements. All four are dimensionless and bounded by unity. The alignment terms η_{pattern} and η_{point} are developed for this architecture in the companion work [3] and η_{cbc} in [1], both of which scope η_{delay} out. This paper develops η_{delay} .

The reason the bandwidth term separated in the first place is that combining analysis is monochromatic while the signal is not. Each element of a tiled aperture terminates a distinct physical path, and the paths differ in length by whatever the amplifiers, the fiber interconnects and the mechanics leave




Fig. 1. The link. Each terminal carries a transmit chain, in which a coherent line is amplified, split K -to- N , phase shifted per element and routed through the fiber bundle to the tiled antenna, and a receive chain of the same form in reverse on the other wavelength of the pair; the two directions are drawn in separate colors. This paper analyzes the transmit chain only.

behind. At the carrier those differences are absorbed by per-element phase control. Across the modulation band they are not, because a phase shifter applies a constant ϕ_0 [rad] while a path applies $2\pi f\tau$ [rad], with f [Hz] the frequency offset from the carrier and τ [s] the group delay of that path. The scales involved are far apart. The integrated phase shifters in the terminal described here reach just over 3π at their drive limit, which is 8.3 fs of group-delay equivalence, while a fiber interconnect built from standard connectorized parts contributes hundreds. The consequence has already been observed on hardware: the same seven-element prototype that recovers 94.2% of its coherent far-field envelope with a continuous-wave source recovers only 66.5% with a 100 Gbit s⁻¹ coherent transceiver, and the difference is attributed primarily to path-length mismatch between the seven channels [3]. As channel bandwidth grows to carry the throughput and element count grows to close the budget, the bandwidth-delay product becomes a scaling constraint specific to the tiled architecture. A single-aperture terminal has one path and no delay spread at all.

The perturbative treatment of combining efficiency for mismatched sources is established for filled-aperture high-power arrays [4], [5], which write the temporal dephasing term without assuming a lineshape and then specialize to a Gaussian one, but do not extend it to tiled arrays carrying a communication waveform. Coherent combining of a data waveform has been demonstrated on a 32-element transmit array over 10 km [6], and that work states the requirement directly: the 32-branch splitting tree must hold the end-to-end optical path difference within fractions of the symbol unit interval, met by feedback splicing to a 1 mm standard deviation. The requirement is stated qualitatively and the tolerance is not derived. The highest-rate PIC array demonstrations do not treat group delay at all [7], [8], and where a tolerance is stated it is for a receiver in optical-bandwidth terms at a modulation bandwidth far below a transceiver line rate [9]; that work inserts variable optical delay lines between the fiber section and the PIC, an acknowledgement that a mixed fiber-and-PIC feed does not path-match itself. On receive the problem is

solved digitally: signals from separate apertures are digitized, temporally aligned and superposed, so per-aperture delay costs an equalizer tap and nothing else [10]. On transmit the delays are physical, and the array sums in the optical field before any detector.

We derive η_{delay} in closed form and show that it takes the same algebraic form as η_{cbc} , with a delay coherence factor in place of the phase coherence factor and the same amplitude concentration ratio setting the incoherent floor. That coherence factor is the band average of a frequency-resolved one, which separates the penalty a carrier suffers across its own occupied width from the penalty it suffers from its offset in a multi-carrier comb, a distinction that decides the specification once a terminal carries terabit-class bandwidth. We build an estimated delay budget for the optical line system of a real terminal and show that the driver PIC removes itself from it, leaving the amplifier stage and the fiber interconnects. We apply the resulting tolerance to the seven-element prototypes and to the 61 and 91 element terminals that follow them, bound the delay spread of the present prototype by inverting its measured continuous-wave and modulated combining efficiencies [3], and specify a switched delay stage that absorbs a standard fiber assembly without precision trimming.

II. OPTICAL LINE SYSTEM AND DELAY BUDGET

A. Terminal architecture

The delay budget is a property of the whole optical line system and not of the PIC alone. The line system comprises four active subsystems with a fiber interconnect between each pair: a coherent transceiver generating the modulated line signal; an amplifier stage raising it to the level the array requires; a bidirectional driver PIC that splits the transmit signal into N nominally equal parts, applies per-element phase control, and performs the reciprocal combination on receive; and a tiled antenna of N fiber-fed elements that launches and collects the free-space beams [3]. Only the segments that differ between elements enter the budget. Everything upstream of the split is common and produces a global time shift, which changes nothing: an array delayed as a whole is an array with

a different absolute time reference. The contributing segments are therefore the amplifier stage, the three fiber interconnects, the per-element paths inside the PIC, and the free-space path inside each antenna element.

Combining quality is measured across the link rather than within a terminal, Fig. 1. A terminal cannot observe the far-field quality of its own transmit array, because its tap and photodiode sit after its receive combining and so measure the beam arriving from the far terminal. That error belongs to the far terminal's transmit array, and it is returned over the link on the other wavelength of the pair to close that array's phase lock.

Two seven-element terminals have been built [3]. The first is an all-fiber implementation in which a 1-to-8 fiber splitter feeds seven chains of an inline piezoelectric fiber phase shifter, a variable optical attenuator and a motorized delay line, with connectorized interconnects throughout. The second replaces the splitter and the phase shifters with a driver PIC and a seven-fiber output array, retaining the per-branch attenuator and delay line. Both feed the same active antenna, seven populated positions of a 91-element head on a hexagonal center-plus-six pattern, each 20 mm-class element carrying a fiber collimator, a beam reducer and a microelectromechanical fast-steering mirror. The full-scale terminal is 91 elements forming a 300 mm-class aperture at an aggregate 1.6 Tbit s^{-1} , with 61 elements as the intermediate step. Throughout the line system the fiber is standard PANDA polarization-maintaining fiber for the C band with a mode field diameter of $10.5 \mu\text{m}$, selected there for its fabrication tolerance at the antenna.

B. Driver PIC

The per-element phase shifters on the driver PIC are the array's operational actuators: they co-phase the elements for coherent combining, steer the beam and carry the dither that closes the lock. They do not set path length, which on chip is fixed at mask level. Nothing in this paper is specific to one material system: the analysis is written in delay and the circuits in Sec. VI in splitting ratio, so any platform serves that closes the link budget, meaning enough optical power handled per element, low enough loss end to end through the waveguides, the phase shifters and the fiber-to-chip coupling, and no nonlinearity strong enough to distort the modulated line at the power the budget demands.

The present driver PIC is built on thin-film lithium tantalate (TFLT), selected for its phase-shifting mechanism. Phase control is electro-optic, so a state costs no static hold power and settles without thermal lag, and the mechanism drives stably from kilohertz to tens of megahertz, which covers both the steering loop and the dither. Relative to thin-film lithium niobate the platform holds a static bias without feedback and tolerates higher optical power [11], both of which matter on transmit, where per-element power is high and the loop must not spend its range chasing drift. The shifter is low loss and so is the propagation, about 9 dB m^{-1} [12].

One property of the platform enters the delay arithmetic directly. The waveguide is a shallow rib, 100 nm of slab

carrying a 100 nm rib under an oxide cladding, so the mode is weakly confined and most of the field sits in the oxide: solving it for the fundamental transverse-electric mode at 1550 nm gives $n_{\text{eff}} = 1.62$ and $n_g = 1.95$, against a bulk ordinary index of 2.118 and a material group index of 2.17. On-chip length therefore converts to delay at 6.50 ps mm^{-1} and not the 7.24 ps mm^{-1} the bulk figure implies, an 11 % error if the material index is used by mistake.

Every arm of the splitting tree and every phase-shifter path is drawn on the same mask, so the optical path from the input coupler to each output, through the tree, the shifters, polarization management and the wavelength-selective transmit and receive routing to the output fiber array, can be equalized at layout to better than 100 nm of physical length. That is 0.75 fs of group delay, less than one optical period at the carrier, so the integrated part of the line system removes itself from the delay budget more completely than any assembly process could. An unbalanced tree, whose outputs traverse different numbers of stages, would otherwise impose a deterministic offset, but the compensating length is added at layout at no cost in area, loss or runtime control.

The topology does enter the delay problem, but through amplitude rather than length. Because the hexagonal element counts are never powers of two, the tree is either uniform with unused ports or non-uniform with a non-flat splitting ratio [1], and the resulting amplitude distribution sets the concentration ratio ρ of (8), which fixes the incoherent floor of both the carrier-phase and the delay terms.

C. Per-element delay decomposition

Each element $i \in \{1, \dots, N\}$ terminates a distinct physical path and accumulates a total group delay τ_i [s]. Only the spread about the mean is observable, so we define the array mean delay $\bar{\tau}$ [s] and the realized inter-element delay spread s_τ [s] as

$$\bar{\tau} = \frac{1}{N} \sum_{i=1}^N \tau_i, \quad s_\tau^2 = \frac{1}{N} \sum_{i=1}^N (\tau_i - \bar{\tau})^2. \quad (2)$$

Throughout, s_τ denotes the realized spread of one built array and σ_τ [s] the standard deviation of the population from which the τ_i are drawn; the two differ by a finite-sample factor given below. For independent segment contributions the variances add, so the decomposition is a budget rather than a list.

Delay converts to length through the group index of the medium carrying it. One picosecond is $299.8 \mu\text{m}$ of optical path in vacuum, $204.2 \mu\text{m}$ of fiber, or $153.7 \mu\text{m}$ of TFLT waveguide; equivalently, 1 mm of fiber is 4.90 ps. The fiber figure uses $n_g = 1.4682$ at 1550 nm [13], the standard single-mode value, taken as a proxy because no PANDA datasheet specifies a group index. Delay is the physical quantity, but length is what is cut, spliced and laid out.

Table I is the resulting budget. Variances add, so the two terminal rows are the root sum of the segments above them, counting the fiber link three times. The fiber link is quoted as two grades, coarse and precision, each as a range, because its contribution depends on how the bundle is built and terminated




Fig. 2. The seven-element build as measured. (a.1) the terminal with only the transmit chain active, drawn in color, the unused receive chain in outline; length matching is marked where it enters the chain. (a.2) the bench that replaces the far terminal: the combined free-space beam is reduced and split, one arm collected into single-mode fiber for the transceiver test set and one imaged onto a far-field profiler, with a tap to the photodiode that closes the combining lock locally rather than through the far terminal. (b.1) the populated aperture, seven elements on a 20 mm pitch across a 60 mm head. (b.2) the optical path inside one element, from fiber collimator through the steering mirror to the output relay.

TABLE I
PER-ELEMENT DELAY BUDGET

Segment	Medium	σ_L [mm]	σ_τ [ps]
Transceiver, common feed	fiber	0	0
Driver PIC	TFLT	10^{-4}	0.0008
Antenna, fiber to aperture	air	0.05	0.17
Amplifier, as built	fiber	20	98
Amplifier, trimmed	fiber	0.2	0.98
Fiber link, coarse	fiber	5 to 58	25 to 283
Fiber link, precision	fiber	0.2	0.98
Terminal, coarse		22 to 102	107 to 499
Terminal, precision		0.40	1.97

rather than on any single component. The coarse grade runs from fibers cabled and cut as a group to the standard patch-cord length tolerance of ± 100 mm [14]; the precision grade is set by what a comparable terminated bundle has achieved, as Sec. II-D argues.

The PIC and the antenna are negligible, the first by layout and the second because its fiber-to-aperture path is a machined free-space assembly [3]. The amplifier stage is the largest single active term and the least well characterized, since the internal fiber length of a module is not a published parameter and no measured distribution across nominally identical modules appears to exist. Its estimate rests on the design statement that amplifier paths, typically more than 10 m long, are difficult to match to better than a few centimeters [15], and its trimmed row on the practice of holding each channel to better than 1 mm peak to peak with an adjustable delay line ahead of the amplifier [16].

D. Length matching in a terminated fiber bundle

The precision fiber grade is an order of magnitude looser than the precision with which a single fiber can be cut, and the limit is structural rather than metrological. A motorized stage under reflectometric feedback holds a single fiber to better than $50 \mu\text{m}$ [17], so neither cutting nor measurement binds. What binds is termination. Routing to the antenna is in

fiber, so an N -element terminal carries $2N$ fibers between the driver PIC and the antenna, a transmit and a receive path per element, 14 now and 182 at full scale; one end of that bundle forms a fiber array for the PIC facet and the other a multi-fiber connector to the antenna, so every fiber is terminated twice and the assembly carries $4N$ terminations that must all be low loss. Length and loss are then coupled objectives on the same piece of fiber, and loss is the binding one: a splice whose measured loss is unacceptable is cleaved back and remade, which shortens that fiber by of order 10 mm [18], or 49 ps, more than ten times the tolerance at 31.5 GHz and more than a hundred times it at 400 GHz.

Length matching therefore cannot be completed before termination. It has to survive a rework loop whose iteration count is a per-fiber random variable, and polarization-maintaining fiber lengthens that loop, each splice also requiring rotational alignment of the stress rods and carrying roughly twice the loss of a standard splice [18]. A six-channel astronomical beam combiner of the same construction, terminated in a v-groove array at one end and connectors at the other and length matched by a specialist cable house, reached a largest inter-fiber difference of 2.38 mm on its first bundle and under 0.7 mm on a second [19], while a spliced 32-branch tree without connectors reached 1 mm standard deviation [6]. Sub-millimeter is reachable on a terminated bundle and tens of micrometers is not, which is what the precision grade of Table I records.

E. Per-element actuator range

The per-element actuators of Sec. II-A are phase shifters: the integrated TFLT shifters on the driver PIC of the second build, and the inline fiber piezoelectric shifters of the all-fiber build. Both apply a frequency-independent ϕ_0 , whereas the physical paths of Table I apply $2\pi f\tau_i$, linear in frequency, so the two agree at exactly one frequency. Equating them at the carrier gives the group-delay equivalence of a shifter at full range,

$$\tau_{ps} = \frac{\Delta\phi_{\max}}{2\pi f_c}, \quad (3)$$

TABLE II
GROUP-DELAY RANGE OF THE PER-ELEMENT ACTUATORS

Actuator	Range	Delay	Native length
carrier period, reference	2π	5.17 fs	0.7 μm TFLT
TFLT shifter, 40 V	3.2π	8.3 fs	1.1 μm TFLT
fiber PZT, 40 V	14.7π	37.9 fs	7.8 μm fiber
fiber PZT, 150 V	55π	142 fs	29 μm fiber
4-stage switched bank	16 states	3.38 ps	0.52 mm TFLT
motorized delay line	150 mm	500 ps	150 mm air

with $\Delta\phi_{\max}$ [rad] the available phase range and f_c [Hz] the optical carrier frequency. At $f_c = 193.4$ THz one full 2π of carrier phase is 5.17 fs.

Table II evaluates this for the devices in the two builds. The TFLT shifter has a half-wave voltage near 12.5 V and is driven to 40 V, giving just over 3π and 8.3 fs. The inline fiber piezoelectric shifter is a genuine path-length actuator and the most capable phase device in either terminal: at the same 40 V drive limit it reaches 14.7π , and at its full 150 V rating 55π , which is 142 fs [20].

Read against Table I, every phase shifter falls short by three to five orders of magnitude: the most capable of them supplies 142 fs against 1.97 ps for a precision assembly and hundreds of picoseconds for a coarse one. The gap is set by the ratio of the optical period to the mismatch rather than by device engineering, so no phase-shifter design closes it, and the correction has to come from a device that changes path length rather than phase. The two such devices in Table II are the ones treated here. The motorized line is what the present builds use and it has range to spare, but it does not scale. Each element carries its own stage with 150 mm of travel, and that volume is not practical in a terminal of 61 or 91 elements. The setting is also mechanical rather than deterministic, since returning to a previously used position does not reproduce the same path length exactly, so a state has to be re-established against a measurement rather than recalled. The switched bank of Sec. VI is the proposed replacement, reaching the same range in millimeters of waveguide on a chip the terminal already carries. Both remove the spread contributed by the fiber and connector run between the driver PIC and the antenna, which Fig. 2(a.1) marks as the point of length matching in the transmit chain.

III. COMBINING EFFICIENCY UNDER INTER-ELEMENT DELAY SPREAD

A. Array transfer function and frequency-resolved efficiency

In the carrier-domain treatment each element contributes a field $E_i = a_i e^{j\Delta\phi_i}$, with a_i [$\sqrt{\text{W}}$] the field amplitude of element i and $\Delta\phi_i$ [rad] its residual phase error after control [1], [3]. For a modulated signal each element additionally applies its own group delay τ_i , so on the aggregate signal the array acts as the linear channel

$$h(t) = \sum_{i=1}^N a_i \delta(t - \tau_i), \quad H(f) = \sum_{i=1}^N a_i e^{-j2\pi f \tau_i}, \quad (4)$$




Fig. 3. The two seven-element builds behind the delay budget of Table I. (a) the all-fiber implementation, seven inline piezoelectric fiber phase shifters connectorized at both ends. (b) the driver-PIC implementation, the packaged thin-film lithium tantalate circuit at right, its fiber array fanned out across a routing tray to the bulkhead connectors at left.

with $h(t)$ the impulse response of the aperture in the far field on boresight and $H(f)$ its transfer function, f measured from the optical carrier.

Normalizing to ideal coherent summation of the same amplitudes and averaging over the signal power spectral density $S(f)$ [W Hz⁻¹] referred to the carrier,

$$\eta_{\text{delay}} = \frac{1}{P_s} \int S(f) \frac{|H(f)|^2}{(\sum_i a_i)^2} df, \quad P_s = \int S(f) df, \quad (5)$$

with P_s [W] the total signal power. A flat spectrum of full occupied width B [Hz] is assumed throughout, reducing this to a band average over $\pm B/2$. For a single carrier of symbol rate R [Bd] and transmit roll-off α , $B = (1 + \alpha)R$. The roll-off is carried as a variable, because the implementation agreement specifies spectral masks and not a transmit roll-off [21] and because every tolerance below scales exactly as $1/B$, so a change in α is one multiplication. Values of $\alpha = 0.05$ to 0.2 are in common use, the larger roll-off giving the wider band and the tighter requirement. Three properties keep the decomposition of (1) clean: with all delays equal, $\eta_{\text{delay}} = 1$ for any amplitudes, so amplitude nonuniformity enters η_{cbc} and not η_{delay} ; the quantity is defined at the aperture, before any receiver processing; and it is invariant to a delay common to all elements.

Evaluated at a single frequency, with the τ_i independent and normally distributed with standard deviation σ_τ , the expectation of $|H(f)|^2$ gives a frequency-resolved efficiency

$$\eta_{\text{delay}}(f) = \gamma(f) + (1 - \gamma(f))\rho, \quad \gamma(f) = e^{-(2\pi f \sigma_\tau)^2}, \quad (6)$$

in which $\gamma(f)$ is a dimensionless coherence factor and ρ is defined in (8). Delay is frequency-dependent phase error, exactly: (6) is the carrier-phase result of [1] with an effective residual phase standard deviation $\sigma_\phi(f) = 2\pi f \sigma_\tau$ [rad] growing linearly with frequency offset. The coherence spectrum is Gaussian, with a half-power half-width of 132.5 GHz at $\sigma_\tau = 1$ ps.

The penalty this represents is a deficit in received signal power, and the reason the receiver cannot recover it decides what a measurement has to look at. The channel $H(f)$ acts

identically on both polarizations, so it lies in the identity component of the Jones channel, whereas a coherent receiver estimates differential group delay between the two principal states of polarization [21]. A co-polarized delay is invisible to that estimator and appears only in bit error ratio. A linear equalizer can invert $H(f)$, but at the cost of noise enhancement that diverges wherever the response vanishes, so η_{delay} is the correct pre-processing penalty; and tap span bounds what is reachable at all, since the impulse response of (4) must fit an equalizer that for a ZR-class receiver spans near two symbol periods [21]. Per-element equalization would remove the restriction but requires N coherent front ends, which is the receive-side architecture [10] and not this one.

B. Band-averaged efficiency in closed form

Averaging (6) over a flat band of width B gives the band-averaged result

$$\mathbb{E}[\eta_{\text{delay}}] = \Gamma_\tau + (1 - \Gamma_\tau) \rho, \quad (7)$$

in which

$$\rho \triangleq \frac{\sum_i a_i^2}{(\sum_i a_i)^2} = \frac{1}{N_{\text{eff}}} \quad (8)$$

is the dimensionless amplitude concentration ratio of [1], equal to $1/N$ for uniform amplitudes and larger for any taper, and

$$\Gamma_\tau \triangleq \frac{1}{B} \int_{-B/2}^{B/2} \gamma(f) df = \frac{\sqrt{\pi}}{2} \frac{\text{erf}(\beta)}{\beta}, \quad \beta \triangleq \pi \sigma_\tau B \quad (9)$$

is the band average of the coherence factor of (6), with dimensionless argument β , the bandwidth-delay product in natural units. That Γ_τ is an average separates the two ways a wide aggregate bandwidth can be occupied. A single carrier of width B suffers the average; a comb of narrow carriers spanning the same width does not, since each suffers $\gamma(f_0)$ at its own offset f_0 from the phase-lock frequency and the outermost ones set the requirement. For a 1.6 THz comb at $N = 91$ and $\sigma_\tau = 0.1$ ps the band average is 0.35 dB while the edge carrier loses 1.08 dB, so sizing the assembly from the average understates the requirement on the carriers that fail first.

Expanding (5) to second order in the delay deviations, without any distributional assumption, gives a result that holds for any signal spectrum,

$$1 - \eta_{\text{delay}} = \langle \Delta\omega^2 \rangle s_\tau^2 + \mathcal{O}(s_\tau^4), \quad \langle \Delta\omega^2 \rangle = (2\pi)^2 \sigma_f^2, \quad (10)$$

where σ_f^2 [Hz²] is the second central moment of $S(f)$. The penalty depends on the signal spectrum only through its RMS bandwidth, so the tolerance is independent of pulse shape. A flat spectrum of width B has $\sigma_f^2 = B^2/12$, giving $1 - \eta_{\text{delay}} = (\pi^2/3) B^2 s_\tau^2$; a Gaussian of full width at half maximum Δf has $\sigma_f^2 = \Delta f^2/(8 \ln 2)$, recovering the tolerance of [4]. The realized and population spreads are related by $\mathbb{E}[s_\tau^2] = (1 - 1/N) \sigma_\tau^2$, a factor that matters only at small N : it is 0.926 in amplitude at $N = 7$ and within three percent of unity above $N = 19$.

For a measured delay set the ensemble assumption can be dropped altogether. For uniform amplitudes and a flat band, (5) evaluates to

$$\eta_{\text{delay}} = \frac{1}{N^2} \left[N + 2 \sum_{i < j} \text{sinc}(B(\tau_i - \tau_j)) \right], \quad (11)$$

with $\text{sinc}(x) = \sin(\pi x)/(\pi x)$, which converts a measured delay map into a predicted efficiency. Its two-element case is exact:

$$\eta_{\text{delay}} = \frac{1}{2} [1 + \text{sinc}(B\Delta\tau)], \quad s_\tau = \Delta\tau/2, \quad (12)$$

with $\Delta\tau = \tau_2 - \tau_1$ the branch delay difference. The response has a null at $f = 1/(2\Delta\tau)$, which reaches the band edge when $\Delta\tau = 1/B$, and there $\eta_{\text{delay}} = 1/2$ exactly. The 3 dB point of a two-branch combiner therefore sits at one inverse bandwidth of delay difference, with no approximation and no fitted constant.

C. Relation to carrier-phase loss

Delay and residual carrier phase enter the array the same way and differ only in their coherence factor. In the notation used here the combining efficiency of [1] reads $\mathbb{E}[\eta_{\text{cbg}}] = \Gamma_\phi + (1 - \Gamma_\phi) \rho$, with a phase coherence factor $\Gamma_\phi = e^{-\sigma_\phi^2}$ set by the residual phase standard deviation σ_ϕ [rad] the control loop holds, which is (7) with Γ_ϕ in place of Γ_τ . Both factors are unity for a perfect array and both drive the efficiency to the same incoherent floor ρ in the fully decohered limit, since $\Gamma_\tau \rightarrow (\sqrt{\pi}/2)/\beta$ for large β . The amplitude distribution set by the splitting tree therefore bounds them identically, which is the coupling anticipated in Sec. II-B.

Carrying both through the same expectation gives $\mathbb{E}[\eta_{\phi\tau}] = \Gamma_\phi \Gamma_\tau + (1 - \Gamma_\phi \Gamma_\tau) \rho$, so the coherence factors multiply inside the same structure rather than the two efficiencies multiplying each other. The multiplicative decomposition of (1) is therefore a lower bound, short by $\rho(1 - \rho)(1 - \Gamma_\phi)(1 - \Gamma_\tau)$, which is second order in the individual impairments and first order in ρ . At $N = 7$, where ρ is largest, an aggressive $\sigma_\phi = 0.2$ rad with $s_\tau = 1$ ps at $B = 400$ GHz leaves a gap of 1.7×10^{-3} , falling to 1.5×10^{-4} at $N = 91$, so treating the terms of (1) as independent budget lines is accurate to better than 0.2% and errs conservatively.

D. Tolerance inversion and scaling

Both impairments invert through the same dimensionless target

$$\tilde{\eta} \triangleq \frac{\eta_0 - \rho}{1 - \rho}, \quad \eta_0 > \rho, \quad (13)$$

with η_0 the target efficiency, the residual-phase bound of [1] is $\sigma_\phi^2 \leq -\ln \tilde{\eta}$ and the delay bound is its exact counterpart,

$$\beta \leq \Gamma_\tau^{-1}(\tilde{\eta}), \quad \sigma_\tau^{\text{max}} = \frac{\Gamma_\tau^{-1}(\tilde{\eta})}{\pi B}. \quad (14)$$

Written against the line rate rather than the occupied band, this reads $\sigma_\tau^{\text{max}} = \Gamma_\tau^{-1}(\tilde{\eta})/[\pi(1 + \alpha)R]$, so the tolerance falls as the reciprocal of both the symbol rate and one plus the

TABLE III
DELAY TOLERANCE, $\eta_0 = 0.95$, UNIFORM AMPLITUDES

B [GHz]	$\sigma_\tau^{\max}$ [ps]		path, $N = 91$	
	$N = 7$	$N = 91$	fiber [μm]	TFLT [μm]
31.5	4.344	4.029	823	619
62.8	2.179	2.021	413	311
100	1.368	1.269	259	195
200	0.684	0.635	130	98
400	0.342	0.317	65	49
800	0.171	0.159	32	24
1600	0.086	0.079	16	12
3200	0.043	0.040	8	6

roll-off. At $\eta_0 = 0.95$ and $N = 7$ the numerator is 0.430, and 0.399 at $N = 91$, which makes the rule arithmetic: moving from $\alpha = 0.05$ to 0.2 tightens every tolerance by 13 %. For η_0 near unity, $\Gamma_\tau^{-1}(\tilde{\eta}) \simeq \sqrt{3(1 - \tilde{\eta})}$ to 2 %. The amplitude distribution enters through $\tilde{\eta}$ alone, so a modeled or measured ρ and a target η_0 fix the delay spread the assembly must achieve, exactly as the phase bound fixes the residual variance the control loop must maintain.

Only the product $B\sigma_\tau$ appears in β , so bandwidth and delay spread are exactly interchangeable and a doubling of line rate halves the length tolerance. The dependence on group index is linear through $\tau = n_g L/c$, so a high-index platform buys compact delay elements and pays with a tighter on-chip length tolerance for the same delay. The dependence on element count is weaker than it appears. A tiled aperture on this lattice takes the centered hexagonal counts $N(n) = 3n^2 - 3n + 1$ with effective diameter $D_{\text{eff}} = d(2n - 1)$ for a sub-aperture pitch d [m] [3], and element count enters (7) through ρ alone. Since ρ shifts the incoherent floor rather than the initial slope, a large array saturates far worse than a small one, the floor moving by 11.1 dB from $N = 7$ to $N = 91$, while over that same series the tolerance tightens by a factor of only 1.08 at $\eta_0 = 0.95$, at every bandwidth.

IV. DELAY TOLERANCE AND ACHIEVABLE BANDWIDTH

The tolerance $\sigma_\tau^{\max}$ depends on N , B and η_0 only, so a specification set on either seven-element build carries to the scaled terminals nearly unchanged. Table III inverts (14) across the bandwidths of interest, from a single current-generation carrier to the 1.6 Tbit s⁻¹ aggregate of the full-scale terminal [3]; B is the total occupied optical bandwidth in each case, so the table is independent of the roll-off; a line rate enters only through $B = (1 + \alpha)R$. The length columns are quoted at $N = 91$, the tighter of the two.

At the aggregate bandwidth the tolerance falls below the mode field diameter of the fiber carrying the signal, 8 μm against 10.5 μm [3], and no cleaving or splicing process reaches that, so at terabit rates the requirement cannot be met as a length tolerance at all. If the aggregate is carried on a comb, it restates per carrier through $\gamma(f)$ of (6) and the binding quantity becomes the offset of the outermost carrier: at $N = 91$ over 1.6 THz, holding the edge carrier to 1 dB needs σ_τ below about 0.1 ps, or 20 μm of fiber. Otherwise

TABLE IV
ACHIEVED SPREAD AND THE BANDWIDTH IT SUPPORTS, $N = 7$,
 $\eta_0 = 0.95$

Assembly	σ_τ [ps]	B [GHz]
coarse	107 to 499	1.3 to 0.27
precision	1.97	70
precision, plus motorized line	0.017	8200
precision, plus 4-stage bank	0.065	2107

the residual must be removed actively. Either way the delay budget, rather than the modulator or the amplifier, sets the maximum bandwidth a tiled terminal can carry.

Read in the other direction, the same inversion applied to the terminal totals of Table I gives the bandwidth each assembly grade supports, which Table IV states for $N = 7$.

A coarsely assembled build cannot carry a modulated line at any rate of interest, which is why the all-fiber prototype carries a motorized delay line in every channel [3]: those lines are load-bearing rather than instrumentation. Once the driver PIC replaces the splitter and the shifters it contributes nothing to the budget, so the ceiling is set entirely by the amplifier stage and the fiber interconnects. And a precision bundle on its own reaches only about 70 GHz, short of a single current-generation carrier and far short of the aggregate, so fiber assembly alone does not close the problem at any bandwidth the terminal is intended to carry.

The last two rows are the two ways to supply it, and both clear the terminal's aggregate bandwidth with room to spare. The motorized line reaches further, its 5 μm repeatability [22] being a finer floor than the quantization of a switched bank, so delay is not what separates them and the choice is made on volume, stability and where in the chain each can sit.

V. DELAY COMPENSATION

Two layers reduce the spread optically, applied in order: length matching at assembly, and a per-element adjustable delay set after assembly. They differ in cost, in permanence and in how far they reach, and the second absorbs what the first leaves.

The first layer is matching the fiber interconnects and trimming the amplifier modules, which is true time delay: no bandwidth limit, no leakage, no added loss beyond the splice. It moves the budget from coarse to precision, a factor of fifty and the largest single improvement available. It is also one-time, not adjustable, and its cost is labor that scales with the $4N$ terminations of Sec. II-D rather than with N . Its floor is set by the rework loop described there, which is why the precision grade stops at a few picoseconds rather than the few hundred femtoseconds a single trimmed fiber would suggest. Both the labor scaling and that floor argue for moving the fine correction onto the chip.

The second layer is an adjustable delay in each element, set after assembly and converting a millimeter-scale mismatch into a residual of tens of micrometers. Both present builds




Fig. 4. Delay tolerance. (a) $\mathbb{E}[\eta_{\text{delay}}]$ of (7) against σ_τ for a family of occupied bandwidths, solid at $N = 7$ and dotted at $N = 91$. Each curve falls from unity to the incoherent floor ρ , $1/7$ for the solid family and $1/91$ for the dotted one, and the marker at $\eta_0 = 0.95$ is the allocation used throughout; the two element counts differ by little over the useful range. (b) the coherence factor $\gamma(f)$ of (6) for several σ_τ , against offset from the phase-lock frequency, with the edge of a 1.6 THz comb at 800 GHz, where a carrier is penalized by its own offset and not by the total span.

implement it as a motorized free-space delay line per channel. The proposal here replaces that with a switch network selecting among fixed waveguide legs. It is a quasi-static trim rather than a tunable delay, its states set once at calibration and revisited only for temperature and aging, so no switching-speed requirement arises. It can be co-integrated into the driver PIC between the splitting tree and the phase shifters, which adds no fiber interfaces, or built as a separate fiber-coupled chip, which costs two coupling interfaces per element but lets an existing driver PIC be retrofitted. In either case it sits upstream of the per-element amplifier so that its insertion loss is paid in gain rather than in transmitted power, the same placement used for the same reason in high-power combining arrays that trim with bulk delay lines ahead of each amplifier [16].

Transceiver-domain equalization takes what those two leave. It needs channel memory spanning the composite spread and only a quasi-static update rate, but by Sec. III-A it cannot restore power lost where the response vanishes, so it bounds the residual rather than replacing the optical correction.

Range sets the division between the two optical layers. Fiber assembly must bring the spread within what the switched stage can span, and the stage sets the residual. For the terminals considered here that is a precision bundle plus a four-stage bank, which together reach 2.1 THz against about 70 GHz for the bundle alone, a thirtyfold gain in carried bandwidth for half a millimeter of waveguide per element.

VI. DESIGN OF THE ON-CHIP SWITCHED DELAY STAGE

The stage proposed here is a switched delay line, one per element, integrated on the driver PIC. The delay is not tuned. Each Mach-Zehnder cell, Fig. 5(c), carries a phase shifter, but it is driven between two states to route the field, and the delay itself comes from waveguide legs of fixed length, so

what a state delivers is set by lithography rather than by the accuracy of the drive or by a position that has to be held. Its two specifications, the finest delay step and the total span, are set independently, which is what lets one design be scaled to a terminal rather than redesigned for it.

The topology is a two-rail binary cascade, Fig. 5(b). A reference rail and a delay rail run the length of the block; stage k places a leg of excess delay $2^k \Delta\tau_0$ on the delay rail, and a Mach-Zehnder cell before and after each leg, Fig. 5(c), selects which rail carries the field. K stages therefore need $K + 1$ cells and give

$$R = (2^K - 1) \Delta\tau_0, \quad s_{\tau,q} = \frac{\Delta\tau_0}{\sqrt{12}}, \quad (15)$$

where R [s] is the total range, $\Delta\tau_0$ [s] the unit delay increment carried by the first stage, and $s_{\tau,q}$ the residual spread left by quantization. The delay-bearing waveguide summed over the stages is R regardless of how many states the cascade provides, and $\Delta\tau_0$ fixes the residual while K fixes the span, neither affecting the other. The two rules below therefore take different inputs and can be applied in either order.

The increment follows from the bandwidth. When the incoming spread lies inside R , quantization is the only residual the stage leaves, so substituting $s_{\tau,q}$ for σ_τ^{max} in (14) makes the stage count cancel and leaves

$$\Delta\tau_0 \simeq \frac{0.36}{B} \quad (16)$$

for a 0.15 dB allocation, with only the weak ρ dependence of Sec. III-D left in it, or $\Delta\tau_0 \simeq 0.36 / [(1 + \alpha)R]$ against a symbol rate. Since $1/B$ is the two-branch 3 dB delay, the rule reads: make the unit increment about one third of the delay that halves a two-element array. At 1600 GHz that is 0.225 ps, the value drawn in Fig. 5(a).




Fig. 5. Binary length switch. (a) the four-stage design point, sized by (16) and (17) for $B = 1600$ GHz and $500 \mu\text{m}$ of peak-to-peak fiber mismatch. (b) the general K -stage cascade, stage k carrying $2^k \Delta\tau_0$ on the delay rail and needing $K + 1$ cells. (c) the switch cell, two nominally 50% couplers and a push-pull arm phase. Legs are annotated in delay; the drawn length follows from (18).

The stage count follows from what the assembly leaves. The cascade adds delay to the shorter paths only, so R must span the full peak-to-peak spread across the array rather than half of it. Writing that spread as a physical mismatch ΔL_{pp} in the medium the assembly is specified in, and ΔL_0 for the unit increment expressed in the same medium,

$$K = \left\lceil \log_2 \left(\frac{\Delta L_{pp}}{\Delta L_0} + 1 \right) \right\rceil. \quad (17)$$

Table V evaluates (17) for interconnect fiber. Each doubling of the mismatch costs one stage and each doubling of the bandwidth costs one stage, and the ceiling supplies between one and two times margin, so no separate allowance is needed. A tolerance quoted as $\pm X$ enters as $\Delta L_{pp} = 2X$; one quoted as an RMS σ_τ enters as $\Delta L_{pp} \simeq 4\sigma_\tau$ at $N = 7$ or $6\sigma_\tau$ at $N = 61$ to 91, being the expected range of that many normal samples plus two standard deviations of it.

Both rules are stated in delay, so they hold on any integration platform; only the drawn length depends on the waveguide, through

$$\Delta L_0 = \frac{c \Delta\tau_0}{n_g}. \quad (18)$$

The index in (18) is that of the guided mode and not of the substrate, and Table VI gives both for the two candidate platforms and for the two media an assembly tolerance is quoted in. Silicon halves the drawn length for the same delay, which begins to matter once K reaches seven or eight, but its switch cells hold state thermally, whereas the electro-optic

TABLE V
STAGES K REQUIRED, FOR MISMATCH IN INTERCONNECT FIBER

ΔL_{pp}	B [GHz]			
	400	800	1600	3200
	$\Delta\tau_0$ [ps]			
	0.900	0.450	0.225	0.113
50 μm	1	1	2	2
100 μm	1	2	2	3
250 μm	2	2	3	4
500 μm	2	3	4	5
1 mm	3	4	5	6
2 mm	4	5	6	7
5 mm	5	6	7	8

TABLE VI
BULK AGAINST GUIDED-MODE INDEX, 1550 nm, AND THE INCREMENT EACH IMPLIES

Medium	bulk	guided mode		ΔL_0 , 1600 GHz
	n	n_{eff}	n_g	
air	1.000	1.000	1.000	67.4 μm
PANDA PM fiber	1.444	1.447	1.468	45.9 μm
TFLT rib, 100/100 nm	2.118	1.62	1.95	34.6 μm
Si, 220 nm strip	3.476	2.44	4.20	16.1 μm

platform holds state at no static power, which across $N(K+1)$ cells is the term that decides.

The two rules fix the design point for the terminals of Sec. IV. The aggregate bandwidth is 1600 GHz, so (16) sets the increment at 0.225 ps, which is 45.9 μm of interconnect fiber by Table VI. The bundle is specified to hold the array within 500 μm peak to peak, which (17) and the corresponding row of Table V turn into four stages. That gives 16 states and 3.38 ps of span, enough for 689 μm of fiber mismatch and so carrying the specified bundle with margin, drawn as 0.52 mm of delay waveguide on the tantalate platform. The residual is the quantization floor of 0.065 ps, which supports 2107 GHz at $N = 7$ and 1954 GHz at $N = 91$ against about 70 GHz for the bundle alone. Figure 6 shows that floor against the spread the assembly leaves.

Stage count is the parameter to revise once assembly data exist. A bundle delivering instead the 1.61 mm peak to peak at $N = 7$ that the precision row of Table I implies would need six stages on the same increment, four times the waveguide and no change in the bandwidth carried. A coarse assembly is out of reach either way, needing 0.3 to 1.4 ns of span, or 46 to 215 mm of waveguide per element, so the fiber layer has to deliver a precision bundle before the stage is worth building.

Both penalties the cascade adds are common mode, which keeps them out of the combining budget. Every state traverses the same $K + 1$ cells, so cell loss does not enter the amplitude spread of [1], and the path-dependent remainder across the leg set at 9 dB m^{-1} is under 0.01 dB. Finite extinction behaves the same way, because only one rail carries the field at a time: a fraction that leaks at one cell crosses to the opposite rail and stays there, and the final cell rejects it through its bar operation,




Fig. 6. Simulated per-element group delay across a 61-element array, before and after each element's switched stage is set to its nearest state, on a common scale. Only the spread is observable, so the origin is the array mean and a delay common to every element would not appear. The scatter the assembly leaves collapses onto the quantization floor $\Delta\tau_0/\sqrt{12}$, uniform because every element sees the same increment. That floor is set by $\Delta\tau_0$ alone, so adding states does not beat it and adding elements does not widen it. The stage sized for the seven-element prototype therefore serves the 61 and 91 element terminals unchanged.

which nulls for any splitting ratio common to that cell's two couplers. Leakage is therefore state-dependent loss and not delay error, 0.010 dB across states at the 34.9 dB extinction demonstrated on a thin-film lithium niobate switch [23], and simulating six stages at $N = 91$ and 1600 GHz gives 0.149 dB of delay penalty against 0.152 dB for ideal cells. What this does leave is a layout requirement, that the two couplers of a cell, Fig. 5(c), be identical and adjacent. Erring in opposite directions they leave the bar operation imperfect, and an independent $\kappa = 0.55$ against 0.45 costs 0.086 dB; neither has to be accurate. The bank is fixed at fabrication, so that requirement and the two rules above have to be right before the mask is drawn.

VII. DELAY SPREAD ON BUILT HARDWARE

The companion measurement on the seven-element antenna shows the effect already limiting hardware [3]. With a continuous-wave source the coherently combined far field reaches 94.2% of the array envelope calculated from the measured individual channels. Substituting a 100 Gbit s⁻¹ coherent transceiver for that source, a ZR pluggable running its oFEC DP-QPSK mode at 30.07 GBd near 1547 nm and boosted before the split, through the same optics and the same phase-control loop, the combined beam reaches 66.5% of the same envelope. The difference is attributed primarily to optical path-length mismatch between the seven channels, with residual polarization crosstalk secondary. It appears as a difference between the two runs rather than as a power loss in either, which is the signature of η_{delay} . Coherent combining co-phases the carrier, so combined power does not fall from group delay alone and the penalty acts only on the modulated signal. The control loop is a stochastic parallel gradient descent lock driving the piezoelectric shifters alone, which by Table II have no authority over group delay.

Because the envelope is constructed from the measured individual channel profiles, amplitude nonuniformity and element pointing are already removed from both figures, and the residual phase error of the control loop is common to the

two runs. The ratio therefore isolates the bandwidth-dependent term, giving $\eta_{\text{delay}} \geq 0.665/0.942 = 0.706$, an inequality rather than an equality because polarization crosstalk takes some of the same loss. That crosstalk is a polarization quantity rather than an inter-element one. PANDA fiber is birefringent by construction, carrying roughly 1 ps per meter between its slow and fast axes, and imperfect axis alignment converts part of that into a per-element amplitude and phase error, which the co-polarized channel of (4) does not otherwise see.

Inverting (7) at $N = 7$ bounds the realized spread of that build at $\sigma_\tau \leq 12.6$ ps, which is 2.6 mm of fiber, taking $B = 31.5$ GHz for that line at $\alpha = 0.05$; at $\alpha = 0.2$ the same measurement gives 11.0 ps, so the conclusion is insensitive to the choice. The 12.6 ps falls between the two assembly grades of Table I, an order of magnitude below the coarse one and a factor of a few above the precision one, which is where a build with delay lines set by inspection rather than against a length reference belongs. Being read through the model rather than measured, it fixes the scale of the problem on built hardware at a few millimeters of fiber without testing the model itself.

A direct test on the same bench splits the transceiver line into two branches, imposes a known co-polarized delay on one and recombines. At 0 ps, 5 ps and 8 ps of imposed delay the line-side pre-FEC error ratio reads 1.1×10^{-3} , 3.8×10^{-3} and 4.3×10^{-3} , but the optical signal-to-noise ratio was 31.9 dB on the matched reference against 26.3 dB and 26.0 dB on the swept points, so (12) has to be carried into the error ratio at each recorded level, where it reproduces all three. Delay does little of the work over this span. At $B = 31.5$ GHz the two-branch penalty reaches 1 dB only at 16.8 ps and 3 dB at 31.7 ps, so 8 ps costs a predicted 0.23 dB against the 5.6 dB of noise loading separating the points. The sweep locates the knee rather than resolving the delay term, which would need matched signal-to-noise points carried past 25 ps.

In a coherently combined pair of fiber amplifiers, normalized fringe visibility was measured against optical path difference for a phase-modulated source of 4.5 GHz Gaussian full width at half maximum, falling to $1/\sqrt{2}$ at 20.8 mm [24].

That measurement tests the two-beam limit of (5) to an accuracy this bench does not reach. Evaluating it at $N = 2$ with the Gaussian spectrum in place of a flat one puts the same point at 20.79 mm, and agreement to better than a tenth of a percent confirms both the transfer function of (4) and the claim of Sec. III-B that spectral shape enters only through the spectrum.

VIII. CONCLUSION

The bandwidth-dependent efficiency η_{delay} completes the aperture decomposition with the alignment terms of [3] and the carrier-phase term of [1], sharing the latter's algebraic form with a delay coherence factor in place of the phase one. Because that factor is a band-averaged, frequency-resolved coherence spectrum, each carrier in a terabit comb is penalized by its own offset from the phase-lock frequency and not by the total span, and the tolerance scales as $1/B$ until, at a full-scale terminal's aggregate bandwidth, it falls below the mode field diameter of the fiber itself. Within that budget the driver PIC removes itself, its arms matched in layout to under a femtosecond, leaving the amplifier stage and the fiber interconnects. Phase shifters reach neither, the shortfall being fixed by the ratio of the optical period to the mismatch rather than by device design, and the gap is already visible on hardware, where one prototype recovers 94.2% of its coherent envelope under continuous-wave illumination against 66.5% when modulated, bounding that build at 12.6 ps.

Two devices address a mismatch that exceeds the phase-control range. The motorized delay line of the present builds has step resolution to spare, but one translation stage per element does not scale in volume to 61 or 91 elements, and its setting is mechanical rather than reproducible, so a state has to be found against a measurement rather than recalled. We therefore propose a binary switched delay network on the driver PIC, its increment set by the occupied bandwidth and its stage count by the peak-to-peak path-length variation across the array. Four stages and 0.52 mm of delay waveguide per element follow, and with a precision bundle they carry 2.1 THz against about 70 GHz for the bundle alone, leaving only a residual for transceiver-domain equalization. The correction then depends on a layout tolerance rather than an assembly one, and scales by lithographic replication as the element count grows.

ACKNOWLEDGMENT

This work was supported in part by the Israel Innovation Authority under Grant 90038. The authors thank Sammy Weber and Jonah Klein for developing the test equipment control software used in the measurements reported here.

REFERENCES

- [1] D. Onural, P. Serra, and E. Wohlgenuth, "Scalable tiled apertures with photonic backends for coherently combined free-space lasercom," in *IEEE Summer Topicals Meeting Series (SUM)*, 2026.
- [2] P. Serra, Y. Uliel, and E. Wohlgenuth, "Tiled-aperture optical phased array terminals for inter-satellite lasercom," in *International Conference on Space Optical Systems and Applications (ICSOS)*, 2025.
- [3] P. Serra, R. Kalfon, A. Nazarov, Y. Uliel, D. Onural, I. Tamir, E. Keller, I. Ezer, G. Alin, and E. Wohlgenuth, "Prototyping of a 35 cm flat phased array aperture for 1.6 Tbps laser communication," in *Proc. International Conference on Space Optics (ICSO)*, 2026.
- [4] G. D. Goodno, C.-C. Shih, and J. E. Rothenberg, "Perturbative analysis of coherent combining efficiency with mismatched lasers," *Optics Express*, vol. 18, no. 24, pp. 25 403–25 414, 2010.
- [5] G. D. Goodno, C.-C. Shih, and J. E. Rothenberg, "Perturbative analysis of coherent combining efficiency with mismatched lasers: errata," *Optics Express*, vol. 20, no. 21, p. 23587, 2012.
- [6] O. Balasiano, E. Wohlgenuth, I. Attia, A. Roizman, T. Falk, R. Vered, and D. Sadot, "Demonstration of coherent beam combining for atmospheric free space optical communication over 10 km," *Journal of Lightwave Technology*, vol. 42, no. 20, pp. 7085–7094, 2024.
- [7] C. V. Poulton, M. J. Byrd, P. Russo, E. Timurdogan, M. Khandaker, D. Vermeulen, and M. R. Watts, "Long-range LiDAR and free-space data communication with high-performance optical phased arrays," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 25, no. 5, p. 7700108, 2019.
- [8] X. Ma, M. Yuan, J. Li, H. Zhang, B. He, P. Zhang, Y. Jiang, H. Xiao, G. Ren, A. Mitchell, Y. Su, and Y. Tian, "Multi-target and ultra-high-speed optical wireless communication using a thin-film lithium niobate optical phased array," *Nature Communications*, vol. 17, p. 969, 2026.
- [9] V. Billault, J. Bourderionnet, L. Leviandier, P. Feneyrou, A. Maho, M. Sotom, and A. Brignon, "Evaluation of a multimode receiver with a photonic integrated combiner for satellite to ground optical communications," 2022, arXiv:2208.08869.
- [10] D. J. Geisler, T. M. Yarnall, M. L. Stevens, C. M. Schieler, B. S. Robinson, and S. A. Hamilton, "Multi-aperture digital coherent combining for free-space optical communication receivers," *Optics Express*, vol. 24, no. 12, pp. 12 661–12 671, 2016.
- [11] C. J. Axline, S. Gamper, P. M. Tengdin, M. Businger, N. Brusadin, R. Giroud, G. Alma, M. A. Roquet, and L. G. Villanueva, "Stable, bidirectional electro-optic transduction in thin film lithium tantalate," 2026, arXiv:2606.12726.
- [12] K. Powell, X. Li, D. Assumpcao, L. S. Magalhães, N. Sinclair, and M. Lončar, "DC-stable electro-optic modulators using thin-film lithium tantalate," *Optics Express*, vol. 32, no. 25, pp. 44 115–44 122, 2024.
- [13] Corning Incorporated, "Corning SMF-28 ultra optical fiber, product information PI-1424-AEN," 2019, effective group index of refraction 1.4682 at 1550 nm.
- [14] OZ Optics Ltd., "Polarization maintaining patchcords and connectors, DTS0071," 2024.
- [15] G. D. Goodno, "System and method for coherent beam combination," 2011, U.S. Patent 7,884,997 B2, Northrop Grumman Systems Corp.
- [16] C. X. Yu, S. J. Augst, S. M. Redmond, K. C. Goldizen, D. V. Murphy, A. Sanchez, and T. Y. Fan, "Coherent combining of a 4 kW, eight-element fiber amplifier array," *Opt. Lett.*, vol. 36, no. 14, pp. 2686–2688, 2011.
- [17] R. Liu, S. Zhao, Y. Gu, R. Yu, P. Deng, and M. Zhao, "Fabrication method for the high-accuracy optical fiber delay line with specified length," *Photonics*, vol. 10, no. 2, p. 176, 2023.
- [18] Fujikura Ltd., "FSM-100P+ specialty optical fiber fusion splicer, product specification," 2025.
- [19] N. Anugu, J.-B. Le Bouquin, J. D. Monnier, S. Kraus, B. R. Setterholm, A. Labdon, C. L. Davies, C. Lanthermann, T. Gardner, J. Ennis, K. M. Johnson, B. Kloppenborg, M. J. Creech-Eakman, C. Bennett, and Y. Kobayashi, "MIRC-X: A highly sensitive six-telescope interferometric imager at the CHARA array," *Astron. J.*, vol. 160, no. 4, p. 158, 2020.
- [20] Thorlabs, "FVP155LP low-loss piezo-based PM fiber phase shifter, 1440 to 1625 nm," 2026, datasheet DOC-103601 Rev B.
- [21] Optical Internetworking Forum, "Implementation agreement 400ZR," OIF, Tech. Rep. OIF-400ZR-03.0, 2024. [Online]. Available: <https://www.oiforum.com/wp-content/uploads/OIF-400ZR-03.0.pdf>
- [22] Advanced Fiber Resources, "Optical delay line with driver (ODL series), rev 11b," 2026.
- [23] W. Ke, Y. Lin, M. He, M. Xu, J. Zhang, Z. Lin, S. Yu, and X. Cai, "Digitally tunable optical delay line based on thin-film lithium niobate featuring high switching speed and low optical loss," *Photonics Research*, vol. 10, no. 11, pp. 2575–2583, 2022.
- [24] B. Anderson, A. Flores, R. Holten, and I. Dajani, "Comparison of phase modulation schemes for coherently combined fiber amplifiers," *Optics Express*, vol. 23, no. 21, pp. 27 046–27 060, 2015.