EQUALISATION IN HIGH-SPEED SERIAL LINKS

From a measured channel to a closed link budget — characterisation, transmit and receive equalisation, and the road to 224 Gb/s

September 2026 • Prepared for Brendan Lynskey

What this report is about, in plain terms. A serial link is a filter you did not design and cannot change, followed by a filter you did design and can. This report takes one concrete channel — 28.8 inches of differential stripline across two line cards and a backplane, carrying 28 GBd — and works it all the way through: what the S-parameters say, what the pulse response says, why the eye is completely shut before any equalisation, what each equaliser block can and cannot fix, and what the noise budget looks like once they are all in place. The answer, for this channel, is that the link lands 1.36 dB short of a 10^{-12} error rate, so the report also works through which of the available remedies actually recovers that margin — and which, despite being the obvious thing to reach for, does not. It is a companion to 'Matrix Methods in Network Parameters' and 'Matrix Concepts in Digital Filter Design': the channel is described by the mixed-mode S-parameters of the first, and every equaliser in it is an instance of the optimal-tap problem of the second.

The channel, in one line. 28 GBd NRZ, so a unit interval of 35.7 ps and a Nyquist frequency of 14 GHz. 28.8 inches of 100 Ω differential stripline on Dk 3.7 laminate at 163 ps per inch, split across a transmit package, a 6-inch line card, an 18-inch backplane, a 4-inch line card and a receive package, with two backplane connectors and the vias that go with them.

1 What closes an eye

A transmitter launches a clean rectangular symbol and a receiver, twenty-nine inches later, sees something that barely resembles it. Four mechanisms are responsible, and it is worth separating them because each is fixed by a different block.

The first and largest is **frequency-dependent loss**. Copper loss rises as the square root of frequency because current crowds into an ever thinner skin, and dielectric loss rises roughly linearly because the laminate's molecules absorb more energy the faster you ask them to reorient. For the laminate used here the two together come to 0.83 dB per inch at 14 GHz. A square pulse is a wide band of frequencies, so a channel that attenuates the high ones and passes the low ones does not merely shrink the pulse — it smears it across many symbol periods. That smearing is **intersymbol interference**, and it is the dominant impairment in every channel in this report.

The second is **reflection**. Every impedance discontinuity — a via, a connector footprint, a package ball — sends part of the wave back, and part of that returns again from the next discontinuity. The result is energy arriving late, which is ISI with a longer memory and a less predictable shape than loss-induced smearing. The third is **crosstalk** from neighbouring lanes, which unlike the first two is not correlated with your own data and therefore behaves as noise rather than as ISI. The fourth is **jitter**, which moves the sampling instant rather than the level, and which converts into amplitude error in proportion to the slope of the signal at the moment you sample it.

Why not just send less far, or use better laminate? Both help, and §8 prices them. But the economics of a backplane system are set by the chassis, the connector and the laminate cost, all of which are decided long before the SerDes is chosen, and none of which scale with each doubling of data rate. Equalisation is the part of the problem that lives on silicon, which is the only part that gets cheaper every generation. That is why every rate increase for twenty years has been paid for with more equalisation rather than with better copper.

2 A short history

The first serial links worth the name ran at a few gigabits per second over channels that were, by later standards, almost transparent. XAUI in 2002 ran 3.125 GBd over a channel with perhaps 6 dB of loss at Nyquist, and a fixed two-tap transmit de-emphasis — implemented as a switchable current in the output driver — was enough. There was nothing adaptive about it; a board designer chose a setting from a table.

The step that made modern links possible was 10GBASE-KR in 2007, which standardised three things at once: a three-tap transmit feed-forward equaliser, a receive decision-feedback equaliser, and — most importantly — a **training protocol** by which the receiver tells the transmitter which way to move its taps. Until then equalisation had been open-loop and set by hand. After it, a link could negotiate its own settings at start-up, and channels with 20 dB of loss became routine.

CEI-28G, from around 2011, added the continuous-time linear equaliser as a standard receive block and pushed decision feedback from one or two taps to a dozen or more. This is the generation the worked example in this report belongs to, and it is the last one in which a designer can reasonably follow every block by hand. What came next changed the character of the problem: 400 Gigabit Ethernet and CEI-56G, standardised from 2017, moved from two-level signalling to four-level PAM4 and made forward error correction mandatory rather than optional. The 9.5 dB that PAM4 gives away (§9) is simply not recoverable by equalisation, so the industry stopped trying to reach 10^{-12} raw and started budgeting to a pre-correction error rate of a few times 10^{-4} instead.

IEEE 802.3ck and CEI-112G, from 2022, completed the transition by replacing the analogue slicer with an analogue-to-digital converter and doing the equalisation in the digital domain. Once the waveform is in a register file, the number of taps stops being a question of silicon area per tap and becomes a question of digital power, and algorithms that were impractical in analogue — long feed-forward filters, maximum-likelihood sequence estimation — become available. The 224 Gb/s generation now being standardised as 802.3dj continues in that direction.

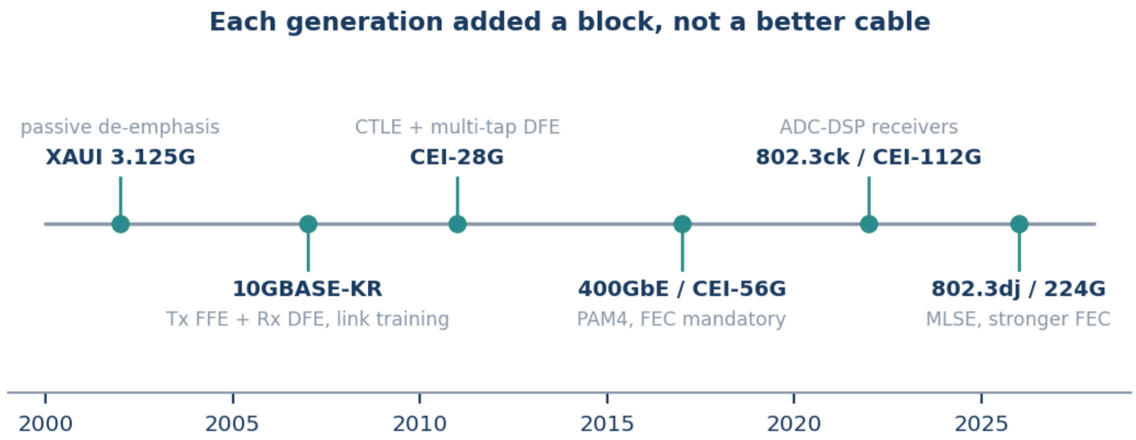

Figure 1 — Twenty-four years of serial links. Each generation bought its rate increase with more equalisation on silicon, not with better copper.

3 The channel: layout, and what the S-parameters say

3.1 The physical channel

The channel is a conventional two-card backplane topology. A transmitter drives out through its package escape, across six inches of line-card stripline, down through a via into a backplane connector, along eighteen inches of backplane, up through a second connector, across four inches on the receiving line card, and into the receiver's package. Twenty-eight point eight inches of copper in total, on a laminate with a dielectric constant of 3.7, which gives 163 picoseconds of delay per inch and a total flight time of 4.7 nanoseconds — a hundred and thirty-one unit intervals. There are always more than a hundred symbols in flight at once, which is worth remembering when reasoning about what a reflection does.

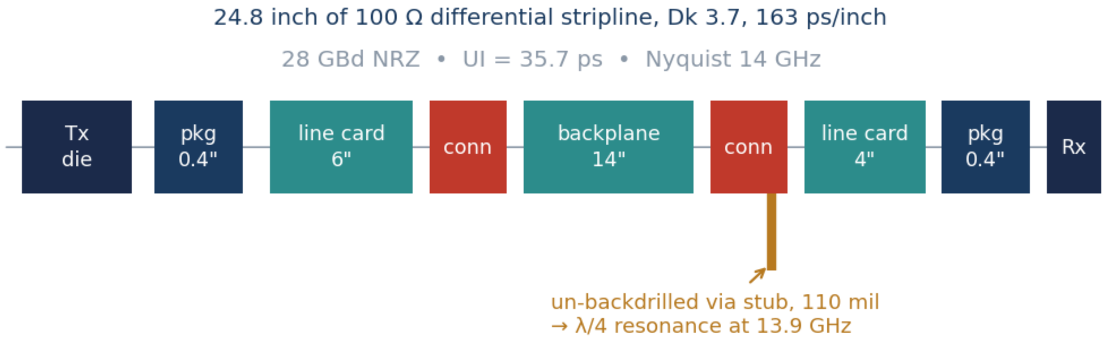

Figure 2 — The channel. Two connectors, four via transitions and one of them left un-backdrilled.

3.2 Insertion loss, and a via stub in exactly the wrong place

The differential insertion loss S_{dd21} is the single number a link budget starts from. Here it is -4.0 dB at 1 GHz and -33.3 dB at the 14 GHz Nyquist frequency — a hard channel, at the limit of what 28 GBd NRZ is normally asked to cross, but not an unreasonable one.

That figure assumes the vias have been back-drilled. If one of them is left with the unused barrel intact, the result is a 110 mil open-circuited stub hanging off the through path, and an open stub is a quarter-wave short at the frequency where its length equals a quarter wavelength. At 163 ps per inch that lands at 13.9 GHz, which is within a percent of Nyquist. The stub adds 25 dB of loss at exactly the frequency the link can least afford it, taking the channel from -33.3 dB to -58.2 dB. No amount of equalisation recovers that. Back-drilling is not an optimisation; on this channel it is the difference between a link and a brick.

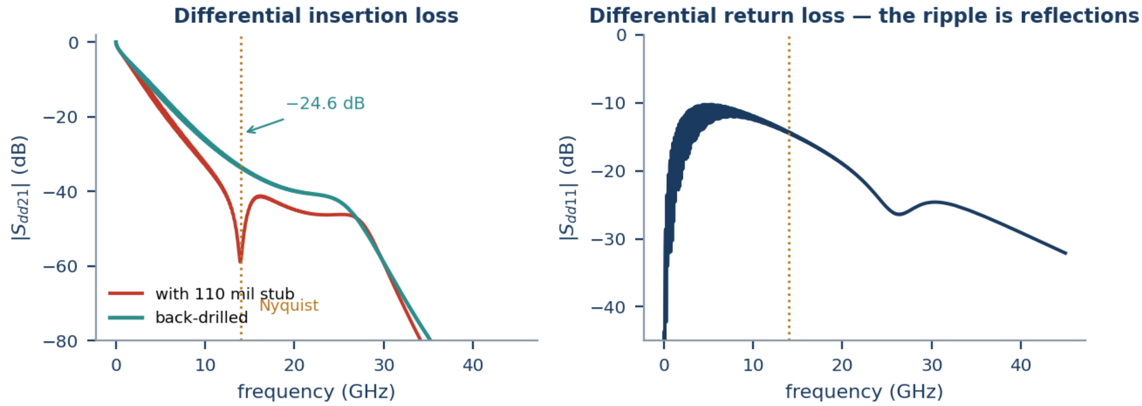

Figure 3 — Differential insertion loss with and without the via stub, and the differential return loss. The stub notch sits within a percent of Nyquist.

3.3 Return loss, crosstalk and mode conversion

Return loss matters less directly than insertion loss but is not ignorable. Here S_{dd11} peaks at about -15 dB near Nyquist, and the ripple visible on the insertion-loss curve is the signature of energy bouncing between the two connector transitions. Each bounce is a delayed copy of the signal, and because the round trip between connectors is many unit intervals long, those copies land as isolated ISI terms far from the cursor — exactly where a decision-feedback equaliser with a handful of taps cannot reach them.

Crosstalk does not appear in a two-port measurement at all, which is one reason characterisation is done on a multiport fixture. Near-end crosstalk from lanes driving in the opposite direction and far-end crosstalk from lanes driving alongside are aggregated into a single figure of merit, the integrated crosstalk noise, which is what enters the budget in §7 as 3.0 mV rms. And mode conversion, the S_{cd} block of the companion report, matters here for a reason specific to equalisation: converted energy is energy the receiver never sees, so it reduces the cursor without appearing anywhere in the differential insertion loss.

4 From S-parameters to a pulse response

Insertion loss tells you how much of a sine wave survives. It does not directly tell you what a receiver sees, because a receiver does not sample sine waves — it samples one symbol at a time, in the presence of all the others. The object that answers that question is the **pulse response**: the voltage waveform produced by sending a single one-unit-interval symbol and nothing else. Take the inverse Fourier transform of S_{dd21} to get the impulse response, convolve with a one-UI rectangle, and you have it.

Sample that waveform once per unit interval, at the phase that maximises the peak, and the resulting sequence of numbers is the whole story. The largest sample is the **cursor**: the part of the symbol that arrives when it is supposed to. Samples before it are **pre-cursors**, energy that has arrived early; samples after it are **post-cursors**, energy still arriving one, two, ten symbols later. Every one of them lands on top of a neighbouring symbol.

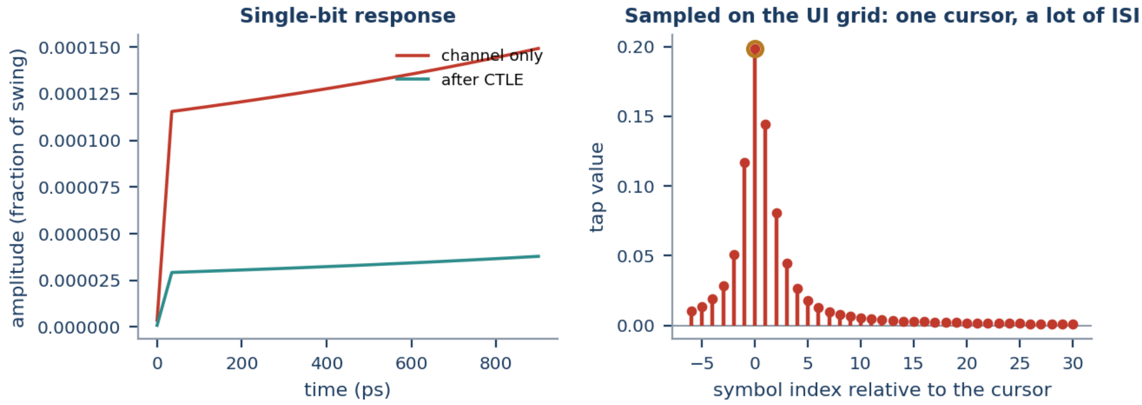

Figure 4 — The single-bit response before and after the CTLE, and the same waveform sampled on the unit-interval grid. The circled sample is the cursor; everything else is interference.

The closed eye, in numbers. For this channel the cursor is 0.199 of the transmitted amplitude, and the sum of the absolute values of every other tap is 0.636 — nearly three times as large. The worst-case eye opening is twice the difference between the two, which comes out at -0.87. A negative eye opening is not a small eye; it means that for some patterns of neighbouring data the receiver sees the wrong sign entirely. The first post-cursor alone is 0.144 and the first pre-cursor 0.117, each comparable with the cursor itself. Nothing about this link works without equalisation.

5 The equalisers, one block at a time

5.1 The continuous-time linear equaliser

The CTLE is an analogue filter in the receiver front end, usually a degenerated differential pair whose source degeneration is bypassed by a capacitor so that gain rises with frequency. It is the cheapest equalisation there is: a handful of devices, no clock, no adaptation loop beyond selecting one of a dozen settings. What it does is approximately invert the channel's slope over the band of interest. What it cannot do is add gain — a CTLE with a boost of 8.3 dB achieves it by attenuating low frequencies, not by amplifying high ones, so the cursor shrinks along with the ISI. That is less damaging than it sounds, because the front-end noise and the crosstalk arriving alongside the signal are attenuated by the same filter: over the receiver's noise bandwidth this setting has a gain of -5.6 dB, so what a CTLE really trades is signal against ISI rather than signal against noise.

That is exactly what happens here. The CTLE brings the ISI sum down from 0.636 to 0.139, a real improvement, but the cursor falls from 0.199 to 0.070 at the same time. The eye opening improves from -0.87 to -0.14 — still negative, still shut. The CTLE has done useful work and has not, on its own, produced a link. The setting used here is the one a receiver's adaptation converges on — the best of the sixteen or so a real part offers, found by sweeping them and keeping whichever gives the highest signal-to-noise ratio at the slicer.

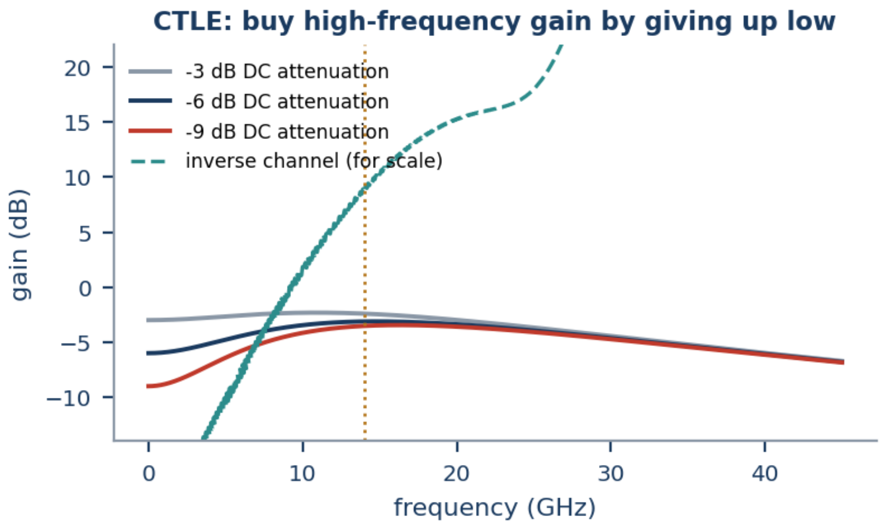

Figure 5 — The CTLE family. More boost is bought by giving away more low-frequency gain. The dashed curve is the inverse channel, for scale: a single pole-zero pair cannot track it far.

5.2 The transmit feed-forward equaliser

The transmitter can pre-distort. A feed-forward equaliser is a short FIR filter running at the symbol rate, and because it operates before the channel it shapes the signal without amplifying anything the channel or the neighbours add afterwards. That is its unique advantage, and it is the reason transmit equalisation survives even though receivers have become far more capable.

Its unique disadvantage is the peak-power constraint. A transmitter has a maximum output swing, so the tap magnitudes must sum to no more than one; every unit of correction is subtracted from the main cursor. The three-tap filter designed here, with two pre-taps set to zero-force the first two pre-cursors, comes out as [0.040, -0.340, 0.620], and the price is visible immediately: the cursor falls again, from 0.070 to 0.032. The eye opening finally turns positive, at -0.007, but only just.

Why the transmitter is asked to handle pre-cursors specifically. A decision-feedback equaliser works by subtracting the known contribution of symbols it has already decided. It therefore cannot touch pre-cursor ISI, which is caused by symbols it has not decided yet. A receive feed-forward filter can, at the cost of noise enhancement. A transmit FFE can, at the cost of swing. Splitting the work so that the transmitter handles the pre-cursors and the feedback path handles the post-cursors is not arbitrary — it puts each impairment where the cheapest remedy is.

5.3 The receive feed-forward equaliser, and what it costs

The receiver can filter too, and a symbol-rate FIR there is the most flexible linear block available: it sees the signal after the channel has done its worst, it can have many taps, and in a digital receiver it adapts continuously. It is the block that distinguishes a modern ADC-DSP receiver from the analogue one this worked example is built around, and §8 prices exactly what adding it would be worth here.

Its cost is noise enhancement. A receive filter amplifies whatever arrives with the signal — thermal noise from the front end, crosstalk from the neighbours — in proportion to the Euclidean norm of its tap vector, while amplifying the signal only in proportion to its effect on the cursor. The meaningful figure is the ratio of the two, and for a nine-tap minimum-mean-squared-error design on this channel it is 0.75 dB. That is modest, but it is not free, and it is the single reason a decision-feedback tap is preferred to a feed-forward tap wherever a decision-feedback tap will do.

5.4 The decision-feedback equaliser

The decision-feedback equaliser is the one block that removes interference without amplifying noise at all. Having decided a symbol, the receiver knows exactly what that symbol will contribute to the next few, and subtracts it. Because the subtraction is driven by a clean two-level decision rather than by the noisy waveform, no noise rides along with it.

The five taps used here are [0.541, 0.133, 0.038, 0.027, 0.026], and they reduce the residual ISI to 0.349 of the cursor. Two caveats come with them. The first is **error propagation**: a wrong decision is subtracted with the wrong sign, which makes the next decision more likely to be wrong. With a tap as large as 0.54 that is not a remote possibility, and it is one reason forward error correction and decision feedback are usually specified together. The second is that the feedback path must settle within one unit interval — 35.7 ps here — which puts the first tap on the critical timing path of the whole receiver and is why it is often unrolled into a speculative, parallel form.

Block	Removes	Cannot remove	Costs
CTLE	Broadband slope, cheaply	Anything with structure — reflections, isolated echoes	Cursor amplitude; it attenuates rather than amplifies
Transmit FFE	Pre-cursor and near post-cursor ISI, before crosstalk is added	Anything it cannot see; it has no knowledge of the receiver's noise	Transmit swing, through the peak-power constraint
Receive FFE	Pre-cursor ISI and long tails, with many taps available	Nothing in principle, but each tap costs noise	Noise enhancement — here 0.75 dB
DFE	Post-cursor ISI exactly, with no noise penalty	Pre-cursor ISI, ever; it acts only on decided symbols	Error propagation, and a hard timing path on the first tap

6 The worked design, stage by stage

Putting the four blocks in series and tracking the pulse response through them gives the table below. The cursor falls monotonically, which is the part that surprises people new to this: equalisation is not about making the signal bigger but about making the interference smaller faster than the signal shrinks.

Stage	Cursor	Σ|ISI|	Worst-case eye	Verdict
Channel only	0.199	0.636	-0.87	Shut, and inverted for some patterns
+ CTLE	0.070	0.139	-0.14	Still shut
+ transmit FFE	0.032	0.036	-0.007	Barely open
+ DFE, five taps	1.000 (nor malised)	0.349	1.30	Open, but not by much

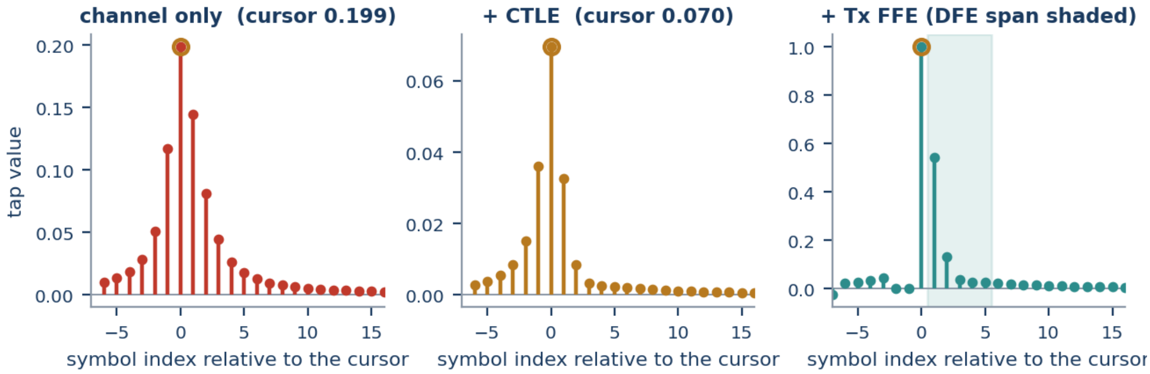

Figure 6 — The sampled pulse response at each stage. Note the vertical scales: the CTLE halves the cursor while it halves the ISI, and the final panel is normalised to its own cursor with the five feedback positions shaded.

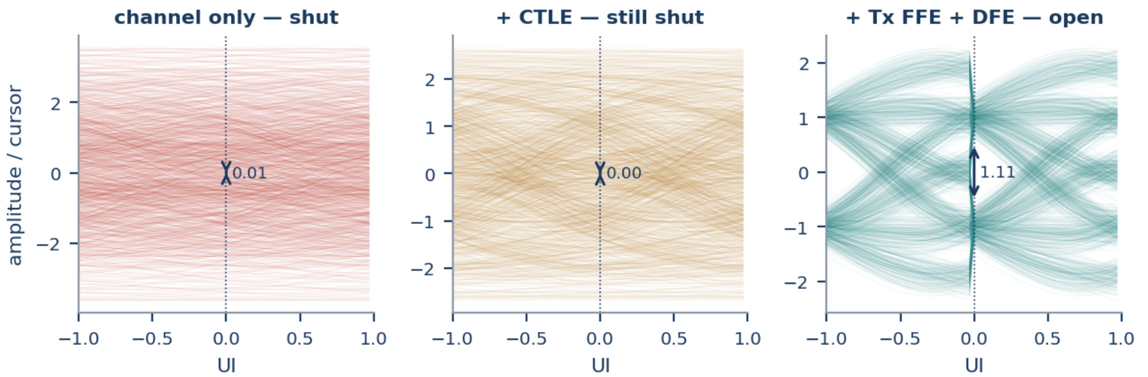

Figure 7 — The same three stages as eye diagrams, built by driving four thousand random symbols through the modelled chain. The arrow marks the opening at the sampling instant, as a fraction of the cursor.

7 The noise and jitter budget

An open eye is necessary but not sufficient. What decides the error rate is the ratio of the eye opening to everything that moves the received level around, and those contributions have to be referred to a common point. Referring them to the receiver input, where the cursor is 12.9 mV out of a 800 mV transmitted swing, gives the budget below.

Contribution	Value	Where it comes from
Receiver noise	0.95 mV rms	1.8 mV of front-end noise, shaped by the CTLE, whose gain over the noise bandwidth is -5.6 dB
Crosstalk	1.58 mV rms	3.0 mV of integrated crosstalk noise from neighbouring lanes, shaped by the same filter
Jitter	0.10 mV rms	350 fs rms of random jitter, converted to amplitude through the slope of the waveform at the sampling instant
Residual ISI	1.12 mV rms	What the equalisers left behind, treated statistically rather than worst-case
Total	**2.15 mV rms**	Root-sum-square of the above

The cursor is 12.9 mV and the total noise 2.15 mV rms, so the quality factor is Q = 6.01. That corresponds to a bit error rate of about 10^{-9}. The conventional target for an uncorrected link is 10^{-12}, which needs Q = 7.03, so this link is **1.36 dB short**.

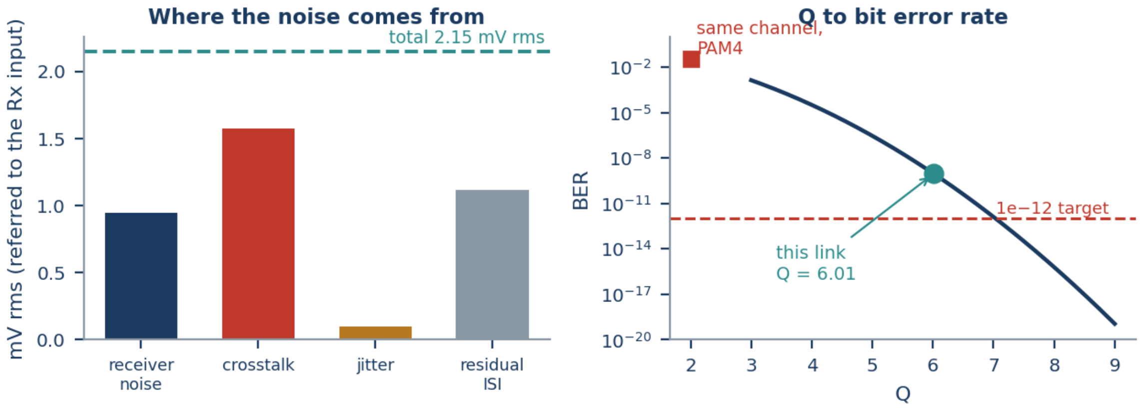

Figure 8 — Where the noise comes from, and what the resulting Q buys. Crosstalk dominates; residual ISI is close behind it, and receiver noise third.

Read the budget before choosing a fix. The largest single contribution is crosstalk at 1.58 mV, with residual ISI second at 1.12 mV and receiver noise third. That ordering matters, because it says the link is short of margin mostly because of its neighbours and only partly because of what the equalisers left behind. Removing every last scrap of residual ISI would take the total from 2.15 mV down to about 1.8 mV, worth roughly four tenths of a decibel — nothing like the shortfall. §8 checks that arithmetic against the alternatives.

8 It does not close: what would actually fix it

Seven candidate remedies, each re-run through the whole chain from the channel model outwards, so that the margin figures below are computed rather than estimated.

What you change	Loss at Nyquist	Q	Margin to 10^{-12}
As built	-33.3 dB	6.01	-1.36 dB
Halve the crosstalk (connector + pair pitch)	-33.3 dB	7.77	+0.87 dB
Low-loss laminate (Df 0.005 → 0.003)	-28.0 dB	8.98	+2.12 dB
Shorten the backplane 18" → 13"	-29.3 dB	8.52	+1.67 dB
Ten DFE taps instead of five	-33.3 dB	6.20	-1.10 dB
Add a 9-tap receive FFE (the ADC-DSP answer)	-33.3 dB	6.34	-0.90 dB
Low-loss laminate + halved crosstalk	-28.0 dB	11.38	+4.18 dB

The table divides cleanly, and the division is the point. Everything that changes the board is worth between one and a half and two and a half decibels. A lower-loss laminate lifts the cursor by reducing the slope the equalisers have to undo; five inches off the backplane does much the same and is often cheaper if the mechanical design still has any freedom in it; halving the crosstalk attacks the largest single term in the budget directly. Any one of the three closes the link on its own.

Everything that changes the silicon is worth a fraction of that. Doubling the decision-feedback taps from five to ten buys 0.26 dB. Adding a nine-tap receive feed-forward equaliser — the block that defines a modern ADC-DSP receiver, and the obvious thing to reach for — buys 0.46 dB, and still leaves the link short of target. Neither is useless, and on a channel dominated by intersymbol interference rather than by crosstalk both would look very different. But on this channel the budget said the problem was the neighbours, and the budget was right.

The general lesson. The block that fixed the last problem is rarely the block that fixes the next one. A link short of margin is short for a reason, the budget names the reason, and the instinct to reach for more equalisation is the instinct to solve the problem you solved last time. Build the budget first: it costs an afternoon, and it tells you which of seven candidate fixes are worth the meeting.

9 The same channel at 56 Gb/s: PAM4 is not free

The obvious way to double throughput without doubling the Nyquist frequency is to send two bits per symbol. PAM4 does exactly that, with four amplitude levels instead of two, at the same 28 GBd and therefore across the same 14 GHz of channel. The channel does not get any worse. The signalling does.

Four levels spread across the same peak-to-peak swing leave three eyes, each one third the height of the single NRZ eye. That is a loss of 20 log 3 = 9.5 dB of signal-to-noise ratio before anything else is considered, and it is not recoverable by equalisation because it is a property of the constellation rather than of the channel. On this channel the Q of 6.01 becomes 2.00, and the raw error rate rises from 10^{-9} to about $10^{-1.5}$.

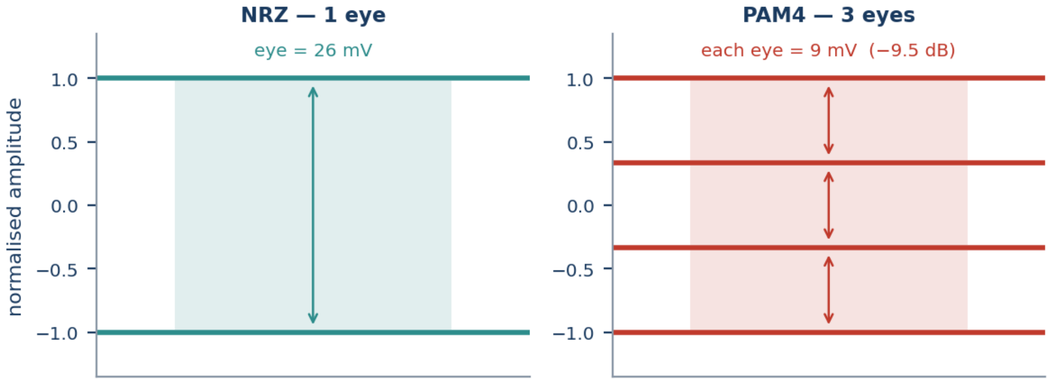

Figure 9 — The same swing, divided two ways. Three eyes of one third the height cost 9.5 dB before the channel is even considered.

An error rate of one in thirty is not a link, which is why PAM4 generations made forward error correction mandatory rather than optional. The Reed–Solomon KP4 code, RS(544,514) over GF(2^{10}), takes a pre-correction error rate of about 2.4e-04 down below 10^{-15}, and PAM4 links are budgeted to that threshold instead of to 10^{-12}. Even so, this channel does not reach it: the required Q is 3.60 and the available Q is 2.00, a shortfall of 5.1 dB.

Channel change	Loss at Nyquist	Q	Margin to the KP4 threshold
PAM4, as built	-33.3 dB	2.00	-5.08 dB
PAM4 + low-loss laminate	-28.0 dB	2.99	-1.60 dB
PAM4 + low-loss + 13" backplane	-25.1 dB	3.97	+0.84 dB
PAM4 + all of the above, plus a receive FFE	-25.1 dB	6.30	+4.87 dB

The lesson is the one the industry learned between 2015 and 2020. PAM4 does not let you run twice the data over the channel you already have; it lets you run twice the data over a channel roughly 8 dB better than the one you already have, in exchange for a coding layer, a larger latency budget and a considerably more complicated receiver.

10 Adaptation and training

Nothing in §5 can be set from a datasheet. Manufacturing spread, temperature, the particular board and the particular connector all move the channel, so every tap in a modern link is adapted. The mechanism is the least-mean-squares algorithm of the companion report on digital filters, usually in its sign-sign form: take the sign of the error between the sampled level and the decided level, take the sign of the data in the relevant tap position, multiply, and accumulate. That requires no multiplier and no knowledge of the channel, and it converges to the same minimum of the same quadratic bowl that the Wiener–Hopf equations describe.

Receive-side taps can be adapted continuously because the receiver has the error signal in hand. Transmit-side taps cannot, because the transmitter does not know what the receiver saw, so the standards define a **link training** phase: before data flows, the two ends exchange a control channel over which the receiver requests increment, decrement or hold on each transmit tap, and the transmitter reports what it did. 10GBASE-KR introduced this and every backplane standard since has kept it. Training typically takes a few milliseconds and is re-run on every link-up.

Adaptation interacts with the clock. The clock and data recovery loop chooses the sampling phase, and the equaliser adapts to whatever phase it is given, while the phase detector sees a waveform shaped by whatever the equaliser is currently doing. The two loops can and do fight, and a link that trains to a stable but poor setting — the wrong lock point, an eye centred half a unit interval away from the best one — is a common and deeply unpleasant failure. Bandwidth separation between the loops is the usual defence.

11 Where this is going

The 224 Gb/s per lane generation now in standardisation keeps PAM4 and doubles the symbol rate again, to around 112 GBd, which puts Nyquist at 56 GHz. Copper at that frequency is unforgiving, and the practical consequence is that the reach of an electrical channel keeps shrinking: a backplane of the kind modelled here is already out of the question, and even the few inches from a switch die to a front-panel cage have become a design problem rather than a piece of routing.

Three responses are visible. The first is more sophisticated receivers: maximum-likelihood sequence estimation, which instead of equalising the channel and then slicing, searches for the most likely transmitted sequence given the known channel response, and which is worth a few dB over a decision-feedback equaliser at the cost of considerable digital power. The second is stronger coding, with concatenated and soft-decision schemes replacing the hard-decision Reed–Solomon codes that have served since 2017. The third, and the most consequential, is to stop sending electrical signals so far: co-packaged optics move the electrical-to-optical conversion into the switch package, and linear-drive pluggable optics remove the retimer from the module so that the host SerDes drives the optical engine directly. Both change the problem from equalising twenty inches of copper to equalising two.

For all of that, the structure of the problem has not changed since 10GBASE-KR. There is a channel you cannot alter, a pulse response that describes it completely, a set of filters whose taps are the solution of a quadratic minimisation, and a budget that decides whether the result works. Everything in this report transfers; only the numbers move.

12 Cheat sheet

Quantity	This channel	What it tells you
Insertion loss at Nyquist	-33.3 dB	The headline difficulty. Above about −30 dB, NRZ needs everything in §5
Cursor, unequalised	0.199	Fraction of the transmitted amplitude arriving on time
Worst-case eye, unequalised	-0.87	Negative means some data patterns invert the decision
Noise enhancement of the receive FFE	-5.59 dB	The price of each linear tap; compare with a DFE tap, which is free
Cursor at the slicer	12.9 mV	Out of 800 mV launched
Total noise	2.15 mV rms	Dominated by crosstalk, not by residual ISI
Q and margin	6.01, -1.36 dB	Short of 10^{-12}; see §8 for what recovers it
PAM4 penalty	9.5 dB	Structural, not recoverable by equalisation

Working order for a new channel. First, get the S-parameters and check them: passive, causal, reciprocal, correctly de-embedded, and with the mixed-mode port map the right way round. Second, look for structural damage — a via stub notch near Nyquist, a connector resonance, an unterminated branch — because no equaliser fixes a notch. Third, compute the pulse response and read the cursor and the first few pre- and post-cursors; that tells you which blocks you need before you design any of them. Fourth, build the budget and find out what actually dominates. Only then choose the equalisation, and re-check the budget afterwards, because the equaliser you added changed the noise as well as the signal.

Glossary

COM. Channel operating margin: the statistical figure of merit IEEE 802.3 uses to decide whether a channel is compliant, computed from the S-parameters and a reference receiver.

CTLE. Continuous-time linear equaliser. An analogue filter with rising gain, which flattens the channel by attenuating low frequencies.

Cursor. The sample of the pulse response at the chosen sampling phase — the part of a symbol that arrives when it should. Everything else is ISI.

DFE. Decision-feedback equaliser. Subtracts the known contribution of already-decided symbols, so it removes post-cursor ISI without amplifying noise.

FFE. Feed-forward equaliser. A symbol-rate FIR filter, in the transmitter (limited by peak power) or the receiver (limited by noise enhancement).

ICN. Integrated crosstalk noise: near- and far-end crosstalk aggregated into a single rms voltage for the budget.

ISI. Intersymbol interference. Energy from one symbol landing on another, caused by frequency-dependent loss and by reflections.

KP4. The RS(544,514) Reed–Solomon code used by 100G-per-lane Ethernet, which takes a pre-correction error rate of roughly 2×10^{-4} below 10^{-15}.

MLSE. Maximum-likelihood sequence estimation. Searches for the most likely transmitted sequence rather than slicing symbol by symbol.

PAM4. Four-level pulse amplitude modulation: two bits per symbol, three eyes, and a 9.5 dB signal-to-noise penalty relative to NRZ at the same swing.

Pulse response. The waveform produced by a single symbol. Sampled on the UI grid it gives the cursor and the ISI taps, which is everything an equaliser design needs.

Q. Ratio of half the eye opening to the rms noise. Q = 7.03 corresponds to a bit error rate of 10^{-12} for a Gaussian budget.

Via stub. The unused remainder of a plated through-hole. A quarter-wave resonator that puts a deep notch in the channel, removed by back-drilling.

References

[1] Bogatin, Signal and Power Integrity — Simplified, 3rd ed., Prentice Hall, 2018.

[2] Hall and Heck, Advanced Signal Integrity for High-Speed Digital Designs, Wiley, 2009.

[3] Proakis and Salehi, Digital Communications, 5th ed., McGraw-Hill, 2007 — the MMSE decision-feedback equaliser and its analysis.

[4] IEEE Std 802.3-2022, Clause 93 (100GBASE-KR4) and Clause 162 (100G per lane) — link training and the channel operating margin method.

[5] OIF CEI-56G and CEI-112G Implementation Agreements.

[6] Gustavsen and Semlyen, 'Rational approximation of frequency domain responses by vector fitting', IEEE Trans. Power Delivery, 1999 — turning measured S-parameters into a simulatable model.

[7] Grivet-Talocia and Gustavsen, Passive Macromodeling, Wiley, 2016.

[8] Haykin, Adaptive Filter Theory, 5th ed., Pearson, 2013 — LMS and its sign-sign variants.