AI data centre network trends: forward error correction

A deep-dive into forward error correction (FEC) in data centre networks and how new photonic architectures can reduce the compute cost of scale-out networks.

Optalysys

Optalysys

IN THIS REPORT

Errors in communication lines →

Forward error correction (FEC) →

FEC: the hidden hero of the AI data centre →

From link reliability to fleet performance →

The cost of correction →

Rethinking where error correction happens →

Building the next generation of data center infra →

Optalysys

Errors in communication lines

The ability to send a message — a sequence of bits — over a network and have that same sequence arrive correctly at its destination is one of the unsung heroes of the information age. Without it, everything from communication within a rack to the internet itself would be impossible.

Digital systems work on bits. When it comes to computing, we tend to take it for granted that the physical systems involved behave reliably: a bit that should be 0 will remain 0, and a 1 will remain 1. Within a silicon chip, the properties of semiconductors mean that we can be reasonably — albeit not completely — certain that this will be the case.

Networking is different. When we send a representation of a bit over a distance, we are manipulating physics to produce a signal. That is as true of a tin-can telephone as it is of the latest optical transceiver. The signal might take the form of an electrical voltage, an optical pulse, a radio wave or even a vibrating length of string.

Fig.1 | An analogue communications channel

Fig.2 | Also an analogue communications channel

Whatever the medium, these are fundamentally analogue phenomena and are susceptible to their environment. Noise and distortion can enter a signal in many different ways. If we did not account for them, the messages emerging at the receiver would rapidly become garbled.

One solution is to make the signal “louder,” increasing its strength so that it is less likely to be drowned out by noise. Unfortunately, this significantly increases the energy required for transmission and does not address every source of error, such as distortion, crosstalk or dispersion.

Another approach is to change the physical channel, using a medium that is intrinsically less prone to interference.

This approach is one of the driving factors in the adoption of optical communications in data centres.

But even optics doesn't eliminate the problem: at the leading edge of data transmission, where an individual optical lane may carry 200 billion bits per second, we are pushing the physical channel extremely hard.

Errors still occur. We therefore need something more effective.

Suppose we send a sequence of bits over a physical channel. We convert those bits into a signal, allow it to propagate through the channel and then reconstruct the bits at the receiver.

During transmission, there's a good chance that the signal may be distorted.

How can the receiver tell whether the message has changed if it does not already possess the original? And if it detects an error, how can it reconstruct the correct message?

This is where forward error correction comes in...

Optalysys

Forward error correction (FEC)

The technology that allows us to detect and correct these errors is known as forward error correction, or FEC.

FEC is everywhere in the digital world: in mobile phones, laptops, Wi-Fi routers, satellites and telecommunications networks.

Within the data centre, FEC is foundational to the mass movement of data that underpins contemporary AI.

Although there are many forms of FEC, the general principle is to add carefully structured redundant information to a message.

This structure allows the receiver to:

- Detect that the message contains an error
- Determine where the error has occurred
- Correct the error before passing the data on

Without this structure, a corrupted sequence of bits simply looks like another potentially valid message. FEC separates valid messages within a wider mathematical space. If noise moves a received message away from its proper location, the decoder can identify the closest valid result and effectively “snap” it back into place.

There are, naturally, limits. If the signal is too noisy or too many elements of the message are corrupted, the distance between valid messages is no longer sufficient to recover the original reliably.

Optalysys

For a channel with a given signal-to-noise ratio and available bandwidth, there is a theoretical maximum rate at which information can be communicated reliably. This limit is described by the **Shannon-Hartley theorem**.

The related noisy-channel coding theorem tells us that, if the information rate remains below the capacity of the channel, there exists a coding method capable of reducing the probability of an uncorrected error to an arbitrarily low level. The objective of practical coding schemes is to come as close as possible to this limit while minimising the redundancy, energy, silicon and latency required to do so.

This is important wherever bandwidth is scarce or expensive. It is particularly important when we need to move as much data as possible, as quickly as possible and using as little energy as possible.

Modern AI infrastructure is exactly such an application.

FEC: the hidden hero of the AI data centre

Training and serving large AI models requires enormous amounts of data movement.

Accelerators exchange model parameters, gradients, activations, tensors and intermediate states.

During inference, infrastructure may also need to move data such as requests, tokens and KV-cache information among different machines or groups of machines.

The resulting networks are expected to provide three things simultaneously:

- Very high bandwidth
- Very low and predictable latency
- Low energy consumption per useful bit

These requirements appear at several different levels of the infrastructure:

“Scale-up” networks connect accelerators and memory within a tightly coupled compute domain, allowing a group of devices to behave more like one large machine.

“Scale-out” networks connect those domains across racks and across the wider data centre.

They enable much larger workloads, but they are also responsible for pooling resources, routing traffic, controlling congestion, working around failures and sharing the accelerator fleet among many simultaneous jobs and users.

“Scale-across” networks extend connectivity even further, between data centres or campuses.

Optical communications already dominate long-distance data transmission. They are also are widely used for rack-to-rack Ethernet links through 400G and 800G pluggable transceivers, with 1.6T technology now emerging.

As optical engines move closer to switches and computing packages, optics is also beginning to penetrate further into accelerator-to-accelerator and accelerator-to-memory connections.

But data must not only arrive quickly; it must also arrive correctly.

For short (1-3m), tightly controlled optical scale-up links, a high-quality physical channel may make it possible to use a lightweight FEC mode or avoid an additional optical FEC stage. This is attractive because a tightly coupled accelerator or memory fabric is extremely sensitive to even small amounts of fixed latency.

Scale-out presents a different balance of requirements. Links must work across a broader range of fibre routes, connectors, temperatures and component characteristics. They must also interoperate across equipment and remain reliable as components age.

At the same time, increasing the number of bits carried by each optical lane makes the signal progressively harder to distinguish.

Moving from two-level NRZ signalling to four-level PAM4, for example, allows more data to be transmitted per symbol but reduces the separation between signal levels. Faster lanes leave less margin for noise, distortion and receiver uncertainty.

For many current high-speed Ethernet interfaces, Reed–Solomon FEC — particularly RS(544,514) denoting a scheme that adds 30 error correction symbols to every 514 data symbols — is the workhorse that makes these physical links practical.

For emerging 200-Gb/s-per-lane interfaces, more demanding channels may employ a concatenated architecture in which an “outer” Reed-Solomon code is supplemented by a lightweight “inner” Hamming code.

The scale of the correction being performed can be surprising. At a representative pre-FEC bit-error rate of 1 in 10,000, a 400-Gb/s physical stream presents the decoder with approximately 40 million raw bit errors per second. FEC corrects these errors before they become corrupt Ethernet frames, allowing the higher layers of the network to experience an effectively error-free link.

This makes FEC one of the technologies that enables modern optical Ethernet.

Without it, operators would need to accept more packet loss or compensate with more powerful transmitters, more sensitive receivers, tighter manufacturing tolerances and more restricted physical channels.

From link reliability to fleet performance

The importance of FEC is clear in terms of what the scale-out network is doing in a real deployment. During training, scale-out carries collective operations such as all-reduce, reduce-scatter, all-gather and all-to-all between racks.

These operations are often synchronised: progress depends on the slowest participating transfer.

During inference, the same network may be supporting many users simultaneously. It connects pools of compute, moves requests to available model instances and carries intermediate data between disaggregated stages of the service.

Its performance helps determine user-facing measures such as time to first token, inter-token latency and tail latency.

A physical error that is corrected by FEC is invisible to this system. An error that escapes correction may instead cause a frame to be discarded, a packet to be retransmitted, a collective operation to wait or traffic to be rerouted. The consequences can extend beyond the original flow when multiple users share the same links and queues.

Effective FEC therefore contributes to:

- Stable collective performance and job-completion time
- More predictable tail latency
- Higher useful network throughput
- Improved accelerator utilisation
- Greater operating margin and link reach
- Fewer link failures and maintenance events

It also provides a margin of tolerance against the imperfections of a physical deployment, including laser degradation, connector contamination, thermal variation, crosstalk and reflections.

At the scale of a hyperscale data centre, “rare” events do not remain rare for long

A very low probability of failure on one link may translate into a routine event when multiplied across tens or hundreds of thousands of links. FEC helps prevent these physical imperfections from continually becoming network-level disruptions.

Optalysys

The cost of correction

FEC is both essential and usually much cheaper with respect to power consumption than the physical alternatives. But it isn't free.

A receiver must accumulate the entirety of a FEC codeword before it can start to decode it. It must then determine whether errors are present, locate them and calculate the corrections before releasing validated data introducing a fixed latency at each physical link termination. This latency is the most significant cost of most FEC implementations.

When data traverses several switch hops, it encounters several independently encoded and decoded links.

The latency of FEC can therefore accumulate alongside serialisation delay, switch processing, NIC processing and, most significantly under load, queuing.

Encoding is relatively inexpensive, but decoding at hundreds of gigabits per second requires considerable parallelism.

The corresponding logic consumes silicon area and power in every relevant port. It also produces heat that must ultimately be removed by the data centre’s cooling infrastructure.

FEC also has an impact on usable bandwidth, because it adds redundant symbols to the transmitted stream. Ethernet generally preserves the advertised data rate by increasing the physical signalling rate, but those additional symbols still require electrical and optical bandwidth and still consume energy.

The result is a fundamental trade-off. Insufficient FEC produces packet loss, retransmissions and unpredictable disruption.

Excessively powerful FEC imposes more latency, power and complexity than the physical channel requires.

Optalysys

The ideal implementation would provide the necessary coding gain while adding only a small amount of latency, consuming minimal energy and silicon area, and supporting a lower-latency mode when the physical channel is sufficiently clean.

Rethinking where error correction happens

Optical communications are evolving rapidly.

Traditional pluggable transceivers place substantial signal-processing capability within the module.

Linear pluggable optics reduce that processing and rely more heavily on the host system. Near-packaged and co-packaged optics move the optical engine closer still to the switch ASIC.

These changes do not remove the need for reliable data. They change where the work required to produce reliable data can take place.

As the optical engine moves closer to the switch package, we have an opportunity to reconsider not only the location of physical-layer computation, but also how that computation is performed.

At Optalysys, we are changing both.

Optalysys

Reed-Solomon encoding and decoding are normally implemented using specialised digital pipelines. However, Reed-Solomon decoding can also be formulated using transform-domain methods, which expose structured finite-field transforms and linear operations, offering a highly parallel and streaming-friendly computational structure.

This creates an opportunity to shift part of the computational and power load away from conventional digital processing and into the optical data path.

Optalysys is bringing these ideas together through our photonic Compute-in-Transit technology.

Our work focuses on a hybrid photonic–analogue architecture capable of executing selected parts of the FEC pipeline while data is still moving through the system.

Photonics and analogue electronics can be applied to structured transforms and fixed linear operations. Digital logic retains responsibility for control, validation and the exact operations that are better suited to conventional silicon. The result is a division of labour that places each part of the algorithm on the hardware best suited to execute it.

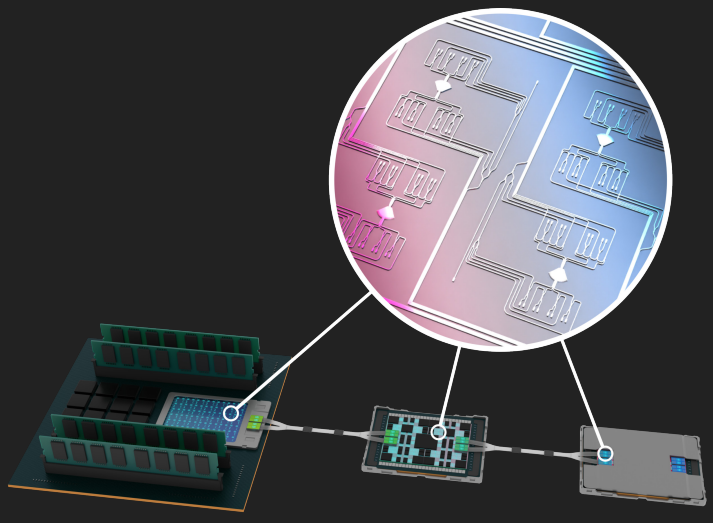

Fig.3 | Optalysys Waveguide Receivers

This is particularly relevant to RS(544,514), which is widely used in high-speed data centre networking. Transform-domain implementations open the possibility of performing important elements of decoding with greater parallelism and fewer sequential processing constraints than conventional DSP-style electronic architectures.

The objective is not to introduce a proprietary wire format or relax Ethernet's reliability requirements. The interface presented to the rest of the system remains bit-exact and standards-compatible. Clean codewords pass through unchanged, correctable errors are repaired, and uncorrectable inputs are identified without silent corruption.

The opportunity lies inside that interface: reducing the latency, energy, silicon area or throughput cost required to provide the same reliable result.

Building the next generation of data center infra

We are working with hyperscale operators, switch and optical-DSP developers, optical-interconnect companies and specialists in FEC to integrate and validate this approach within representative network architectures.

For hyperscale operators, the opportunity is to translate greater physical-layer efficiency into fleet-level value: lower port power, higher density, more predictable network performance, improved reliability and better accelerator utilisation.

Optalysys

For silicon and DSP developers, it is an opportunity to reduce the cost of FEC while retaining a standards-compatible interface and a practical integration boundary with existing switch and networking architectures.

For optical-interconnect companies, Compute-in-Transit adds new capability within linear, NPO and CPO systems: not only transporting data, but performing useful physical-layer computation as that data moves.

These partnerships are essential because FEC exists within a larger system. Its value is ultimately measured not by the performance of an isolated operation, but by the performance, efficiency and reliability of the complete link and the network built from it.

The next generation of AI networks will not be defined only by how many bits can be placed on a fibre, but by how efficiently those bits can be made reliable, routed through the system and delivered as useful work.

FEC is one of the least visible parts of that process. As lane rates, network sizes and infrastructure power continue to grow, it is also becoming one of the most consequential places to rethink how the computation is done.

by Joseph Wilson
Principal Field Applications Engineer at Optalysys

The future runs on light.

Build it with us →