



Whitepaper

**HDMI 2.1 & Beyond:
FRL and Link Training**

1. INTRODUCTION

As display interfaces push beyond 18 Gbps, the limitations of legacy HDMI TMDS signaling have become structural rather than incremental. HDMI 2.2's Fixed Rate Link (FRL) and mandatory link training are not optional enhancements, they are essential architectural changes that enable higher bandwidth, improved robustness, and scalable future evolution. This paper explains the technical fundamentals and why the industry's long-term trajectory makes FRL and link training unavoidable.

Table of Contents

1.	Introduction	1
	Table of Contents	1
2.	HDMI Before FRL: TMDS and Its Limits.....	2
	2.1 Format-determined Transmission	2
	2.2 Dedicated Clock Lane.....	3
	2.3 No Link Training.....	3
	Real-life use case: TMDS gaming gone wrong	5
3.	FRL: An Architectural Shift	6
	3.1 Lane Reconfiguration.....	6
	3.2 Standardized link rates	6
	3.4 Mandatory link training.....	7
	3.5 Forward Error Correction (FEC)	8
4.	Why FRL and Link Training Are Inevitable	9
	4.1 Limits of the fixed equalizer.....	9
	4.2 Adaptive Link Is Required for Future Scaling	9
	4.4 Industry Convergence.....	9
	4.5 Real-World Channels Are Not Ideal.....	11
	Real life use case: FRL gaming done right	12
5.	New Challenges for vendors	13
	5.1 Implementation Complexity.....	13
	5.2 From TMDS to FRL: Overcoming Legacy Design Assumptions.....	13
	5.3 Interoperability Challenges.....	13
6.	Lessons from DisplayPort Link Training	15
	5.1 DP Has Always Assumed an Adaptive Channel	15
	5.2 DP Link Training in Practice	15
	5.3 What FRL Borrows - and What It Does Not.....	15
7.	Testing FRL	16
	Link Training on UCD Console Software	16
	HDMI Link Analyzer	16

2. HDMI BEFORE FRL: TMDS AND ITS LIMITS

HDMI up to version 2.0 relies on **Transition Minimized Differential Signaling (TMDS)**, using three data lanes plus a dedicated clock lane to transmit video. This design is simple, predictable, and broadly compatible, enabling plug-and-play operation at moderate resolutions and refresh rates.

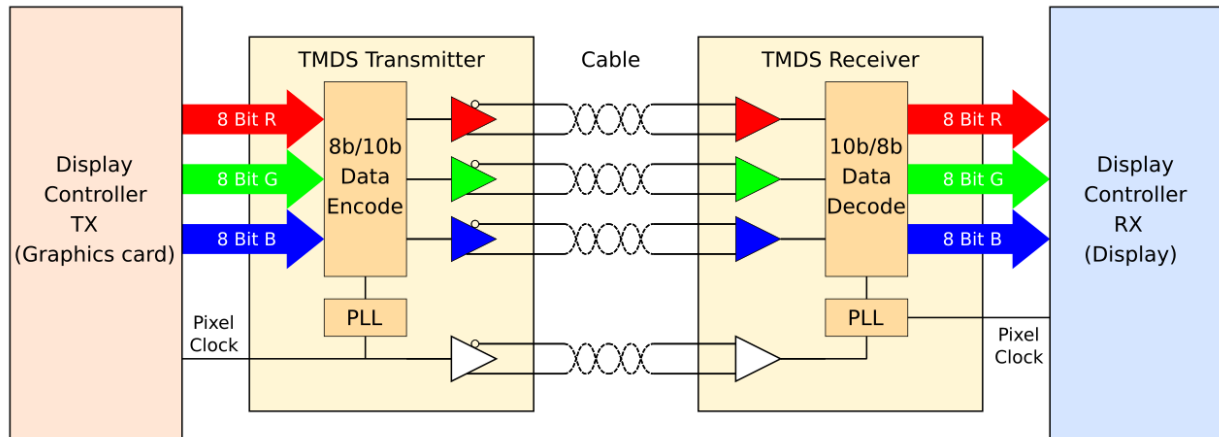


Figure1: "Schematic TMDS link"

Figure1 by wdwd is licensed under CC BY-SA 4.0. Source: https://commons.wikimedia.org/wiki/File:Schematic_TMDS_link.svg

Key design characteristics of TMDS include:

- Pixel clocks scaled by the TMDS multiplier set the link speed based on color depth and chroma subsampling.
- Dedicated Clock Lane. Simplifies timing recovery but reduces available bandwidth and tightens timing margins at higher data rates.
- No Link Training. The source assumes ideal conditions, leaving the sink to compensate for cable or connector limitations.

While TMDS works well for moderate speeds, its design limits bandwidth and flexibility, leaving little room for higher resolutions, refresh rates, or HDR content. In modern networking, if the signal is weak, sink and source negotiate a slower speed. TMDS is not capable of this. Once the video format is chosen (e.g., 4K @ 60Hz), the source simply sends data down the cable at the required speed. There is no consideration for poor cable quality for example. It always assumes a perfect connection.

As a result, TMDS features several limitations:

- No dynamic rate adaptation. The link cannot downshift to a lower rate if errors occur.
- Because TMDS permanently occupies three data lanes and one clock lane, it does not support the power-saving mechanisms that have become a key feature in today's industry
- No channel-aware equalization negotiation. Transmitter parameters are largely static and predefined.
- Designs must assume poor cables and sinks to guarantee interoperability.

The upside is that TMDS is easy to implement, predictable, and broadly compatible. The downside is limited bandwidth, inflexibility and potentially high costs for the end customer if things go wrong.

2.1 Format-determined Transmission

In TMDS, the link operates at a single constant frequency determined by the *pixel clock* and *TMDS multiplier*. This format-determined link architecture simplifies design and ensures predictable operation, but it also imposes strict constraints on timing, bandwidth, and signal integrity. Because TMDS lacks dynamic link training, the transmitter applies a *fixed equalization* to the signal, which cannot adapt to variations in cable quality or the receiver. This static equalizer is tightly coupled to the non-adaptive link architecture, helping maintain interoperability while limiting flexibility at higher data rates or over imperfect cables.

2.1.1 Pixel clock

A key aspect of TMDS is that the transmission rate is fully determined by the video format. The pixel clock defines the base rate at which pixels are transmitted, essentially setting the tempo for the video signal. This rate dictates how quickly pixel data must flow through the cable to display a complete frame on the screen at the correct refresh rate. Higher resolutions, faster frame rates, or increased color depth require a higher pixel clock, because more pixel data must be sent in the same amount of time. Unlike modern adaptive links, TMDS does not adjust the transmission rate dynamically, the pixel clock is fixed for a given format, and the source sends pixels at that constant rate regardless of cable quality or signal conditions. This makes the pixel clock the fundamental driver of timing, bandwidth requirements, and overall link performance in TMDS.

2.1.2 TMDS Multiplier

The TMDS multiplier is the factor that scales base pixel clock to determine the actual data rate on the link. It depends on how much color data is carried per pixel. Color depth and chroma subsampling define the TMDS multiplier, because they determine how many bits must be transmitted for each pixel. A higher data-per-pixel requirement results in a higher multiplier applied to the base clock.

Here are some examples:

- 8-bit color is the baseline (1.0x multiplier)
- 10-bit color requires a 1.25x speed multiplier
- Chroma Subsampling (4:2:0) is a compression technique that removes some color information effectively cutting the bandwidth requirement in half, acting as a 0.5x multiplier.

Combined with timing details of the video stream:

- For example, the bandwidth needed for 4K @ 60Hz, 8-bit, 4:4:4 is **17.82 Gbps**. As 8-bit color is the baseline (1.0x multiplier) the 17.82 Gbps just fits within TMDS's maximum bandwidth of 18Gbps
- For 4K @ 60 Hz, **10-bit**, 4:4:4, you need a 1.25x multiplier. $17.82 \text{ Gbps} \times 1.25 = \mathbf{22.28 \text{ Gbps}}$. This is outside TMDS's scope.

2.1.3 Transmitter-side equalization

Once set, the link operates at a single fixed frequency. This is the most critical constraint of TMDS. Because TMDS lacks any form of link training or feedback, transmitter-side equalization is also fixed. The source applies a predefined amount of signal boost that does not adapt to the electrical characteristics of the cable or the receiver. This static equalizer approach is tightly coupled to TMDS's architecture and is central to its simplicity and interoperability.

2.2 Dedicated Clock Lane

In TMDS, the clock lane exists because the receiver has no reliable way to recover timing from the data stream itself at lower data rates. Each data lane is decoded independently, and the clock provides a shared timing reference so that the sink knows exactly when to sample incoming bits.

This design works well at moderate speeds, but one lanes that could be used for data is permanently consumed by the clock. Further, managing timing becomes increasingly difficult as data rates increase. The bit period shrinks, which reduces the available timing margin for correct sampling. Jitter, skew, and clock inaccuracies that were negligible at lower speeds become a significant portion of the unit interval. In other words, higher data rates tighten the eye opening and make the link far more sensitive to timing errors.

2.3 No Link Training

At the time TMDS was created, simplicity and plug-and-play interoperability were the primary goals. Early HDMI resolution targets, first Full HD and later 4K at moderate refresh rates were achievable within the signal integrity margins of TMDS. Data rates up to 6 Gbps per lane enabled 4K @60Hz with limited color space options. It also meant that:

- In TMDS, the sink has the responsibility of cleaning up a defected signal without help from the source. Sinks use receiver chips to automatically boost high frequencies and sharpen the "blurred" bits caused by e.g. long cables.
- Transmitter settings could be fixed or selected from a small set of presets. Because there is no feedback loop between the source and the sink, the source device cannot tune its output to match the cable's quality. Instead of dynamic adjustments, it transmits using a "one-size-fits-all" signal strength, assuming the receiver will handle any signal degradation.

- Clock recovery was trivial due to the dedicated clock lane. This architecture simplified early electronics because timing was provided on a separate wire. The receiver didn't need complex logic to guess the data rhythm; it simply multiplied the clean signal from the dedicated clock lane to instantly synchronize with the video stream.

Up until HDMI 2.0, adding link training would have increased complexity without delivering clear benefits.

However, the lack of link training comes with significant tradeoffs:

- Marginal links fail completely rather than partially as there is no mechanism to reduce rate, lane count, or encoding overhead when errors appear.
- Variation in cable quality is handled outside the protocol. Certification, labeling, and user guidance substitute for technical adaptation.
- No feedback loop between sink and source. The transmitter has no visibility into real channel conditions or receiver margins.



Real-life use case: TMDS gaming gone wrong

You have just bought a PS5 and connect it to your 4K TV using an old HDMI cable purchased in 2012. The TV does not support HDMI 2.1, so the link between the PS5 and the TV uses TMDS signalling.

You plug it in and the system menu works without issue: it runs at 4K @ 60 Hz in SDR (8-bit color) and typically uses 4:2:0 chroma subsampling. This requires roughly 9 Gbps, just within the cable's ~10 Gbps rating.

However, when a modern HDR game is launched, the video output switches to 10-bit color with 4:2:2 chroma. The required bandwidth increases to approximately 18 Gbps. As TMDS lacks any mechanism to adapt to higher data rates, the PS5 attempts to transmit the full signal regardless of cable limitations. As a result, the bits arrive corrupted, checksums fail, and the TV drops the signal entirely, producing a black screen.

This gaming-gone-wrong scenario perfectly shows TMDS's limitations in a real-world situation. The black screen occurs because TMDS operates at a fixed rate determined by the selected video format, lacks dynamic rate adaptation or link training, and relies on static transmitter equalization, leaving no mechanism to negotiate a lower speed or compensate for the cable's limited bandwidth.

3. FRL: AN ARCHITECTURAL SHIFT

Introduced with HDMI 2.1, Fixed Rate Link (FRL) and Extended in HDMI 2.2, FRL represents a fundamental architectural shift in the HDMI specification. FRL introduces several architectural differences compared to TMDS. This paper focuses on four key aspects: **Lane reconfiguration**, **Standardized link rates**, **Link training** and **Forward Error Correction**.

3.1 Lane Reconfiguration

Under TMDS architecture, the physical interface used three data lanes and one dedicated clock lane. FRL allows the HDMI interface to repurpose the physical TMDS Clock channel as a fourth high-speed data lane. Since the dedicated clock lane is repurposed for data, similar to DisplayPort, the Sink must now recover the clock embedded within the data stream itself. If the FRL link training fails or the device does not support FRL, the system can revert back to the standard 3-lane + Clock TMDS mode.

Using four lanes provides several benefits:

- Higher total bandwidth. 18Gbps in TMDS vs. 96Gbps in FRL (HDMI 2.2)
- Better scaling across different cable qualities
- Adaptive fallback to lower standardized link rates or reduced lane counts when channel conditions require it

3.2 Standardized link rates

Unlike TMDS, where the clock frequency varies continuously according to the pixel clock of the selected video format (e.g., exactly 148.5 MHz for 1080p), FRL operates at one of six standardized link rates. These discrete operating modes simplify the design of SerDes blocks and equalization circuitry, since the hardware only needs to be optimized for a defined set of frequency targets rather than an open-ended range.

The use of standardized rate steps also makes link training more predictable. Instead of attempting to operate at arbitrary frequencies, the source and sink evaluate a small number of known FRL modes and select the highest one that satisfies signal integrity requirements.

Additionally, operating at defined rates improves EMI control and overall signal integrity. PCB trace geometry, termination networks, and equalization settings can be tuned for each supported mode, reducing variability and improving robustness over marginal cables

Mode	Lane Configuration	Data Rate per Lane	Aggregate Raw Bandwidth
TMDS (HDMI 1.4)	3 lanes + 1 clock	3.4 Gbps	10.2 Gbps
TMDS (HDMI 2.0)	3 lanes + 1 clock	6 Gbps	18 Gbps
FRL1	3 lanes	3 Gbps	9 Gbps
FRL2	3 lanes	6 Gbps	18 Gbps
FRL3	4 lanes	6 Gbps	24 Gbps
FRL4	4 lanes	8 Gbps	32 Gbps
FRL5	4 lanes	10 Gbps	40 Gbps
FRL6	4 lanes	12 Gbps	48 Gbps
FRL7	4 lanes	16 Gbps	64 Gbps
FRL8	4 lanes	20 Gbps	80 Gbps
FRL9	4 lanes	24 Gbps	96 Gbps

Figure2. TMDS and FRL link rates

3.4 Mandatory link training

FRL introduces a mandatory link initialization sequence before video transmission begins. Unlike TMDS, which immediately transmits pixel data once configured, FRL requires the source and sink to establish a validated operating mode through structured link training. This ensures that the selected data rate and lane configuration match the electrical characteristics of the cable and receiver before high-bandwidth video traffic is sent.

Before video transmission, the source and sink perform a negotiation and channel evaluation procedure to determine the highest reliable FRL mode. Here's a high-level training sequence overview:

1. **Capability Exchange**
The sink advertises its supported FRL rates and lane configurations through EDID and control channel signaling.
2. **FRL Mode Selection Attempt**
The source begins with the highest mutually supported FRL rate (e.g., FRL9 at 24 Gbps per lane over four lanes).
3. **Transmission of Link Training Patterns**
Instead of video data, the source transmits predefined training symbols. These patterns are carefully constructed to stress specific aspects of the channel.

3.4.1 What Link Training Patterns Test

The training patterns are designed to evaluate:

Parameter	Test Purpose
Clock recovery stability	Patterns with frequent transitions allow the receiver's CDR (Clock Data Recovery) circuit to lock accurately to the incoming data stream.
High-frequency channel loss	Long runs of alternating bits (e.g., 101010...) emphasize high-frequency spectral components, exposing attenuation and inter-symbol interference.
Equalizer performance	The sink adjusts its adaptive equalizer while receiving known patterns and evaluates eye opening, jitter margin, and bit error rate.
Lane alignment and skew	Multi-lane correlation patterns ensure that all active lanes are phase-aligned within tolerance.

3.4.2 Step-Down Mechanism

During training, the sink monitors error metrics and reports status back to the source. If the channel cannot sustain the attempted FRL mode, the source steps down to the next lower standardized rate shown in Figure 2. The process repeats until:

- Clock recovery is stable
- Error rate falls within specification
- Lane alignment is verified


Only after successful validation does the link transition to active video transmission.

3.5 Forward Error Correction (FEC)

Even after link training, FRL operates at multi-gigabit per second data rates where random bit errors can occur due to thermal noise, crosstalk, jitter, and cable attenuation. Rather than relying solely on margin, FRL incorporates Forward Error Correction to improve reliability.

FRL uses block-based FEC coding operating continuously and transparently. At 24 Gbps per lane across four lanes, the aggregate data rate reaches 96 Gbps. Even a very low raw bit error rate can result in visible artifacts without correction. FEC allows FRL to tolerate occasional bit errors, maintain video continuity without retransmission, reduce visible glitches under marginal channel conditions. Here are the responsibilities on both sink and source side:

Step	Source Action	Sink Action
1	Groups data into fixed-length blocks	Incoming blocks are buffered
2	Computes parity symbols using a structured mathematical code (Reed-Solomon-type coding)	Syndrome calculations detect whether errors occurred
3	Appends redundancy to each block before transmission	If the number of corrupted symbols is within correction capability, the original data is reconstructed without retransmission



4. WHY FRL AND LINK TRAINING ARE INEVITABLE

By the time HDMI 2.0 reached its maximum aggregate bandwidth of 18 Gbps, the system was already operating close to the edge of what TMDS could reliably support. Small degradations in cable quality, connector wear, or board layout could push a link from fully functional to completely unusable. These limitations were somewhat manageable at 18 Gbps but they left no room for increased resolutions or framerates. As video resolutions, refresh rates, and color depths increase, legacy TMDS signaling reaches its physical limits. Static equalizers and non-adaptive link can no longer guarantee reliable transmission in real-world situations. FRL with link training solves this by measuring the channel, adjusting signal parameters, and ensuring stable high-bandwidth video. This chapter explores why FRL is an inevitable change, covering:

1. **Limits of fixed equalization:** High-speed links are sensitive to cable and connector variations.
2. **Adaptive links for future scaling:** FRL enables higher resolutions, refresh rates, and robust error correction.
3. **Industry convergence:** HDMI now aligns with other high-speed standards.
4. **Real-world channel challenges:** Cable and connector imperfections make adaptive training mandatory.

4.1 Limits of the fixed equalizer

In the era TMDS, system engineers could often rely on fixed-parameter solutions. A receiver could simply apply a standard boost to the signal and assume it would work for almost any compliant cable.

With FRL pushing speeds to 24 Gbps per lane, this approach no longer works. We are no longer limited by chip logic, but by the fundamental physics of copper cable transmission. At these frequencies, the channel behaves less like a wire and more like a complex, variable filter. At 24 Gbps, the margin for error is microscopic. A fixed equalizer setting fails because the variance between good and bad channels is too wide to cover with a single static value.

- **Under-Equalization:** If you apply a standard, moderate boost (e.g., +5 dB) to a long or cheap cable, the high-frequency signal is lost in the copper resulting in total signal loss.
- **Over-Equalization:** If you apply a massive, aggressive boost (e.g., +15 dB) to ensure long cables work, you will break short cables. On a short, high-quality cable, this boost amplifies Crosstalk and noise, causing the receiver to misinterpret noise as data.

Therefore, the system cannot prefer a setting, it must mathematically derive the setting that matches the physics of the attached cable. Link Training exists precisely to solve this problem: to measure the channel, compute optimal parameters, and adapt the link.

4.2 Adaptive Link Is Required for Future Scaling

By moving to an adaptive link, HDMI gains the flexibility needed to push beyond physical bottlenecks. While the four-lane architecture already provides a strong foundation for higher bandwidth and scalability, it is not sufficient on its own. An adaptive link is the key that unlocks these three critical capabilities:

- **Higher Refresh Rates:** To reach refresh rates up to and beyond 120Hz the timing margins become vanishingly small. An adaptive link can dynamically tune out jitter, allowing the interface to maintain sync at speeds where a static link would drift and lose the picture.
- **Resolutions up to 10K:** As resolution increases, higher signalling rates create significant electromagnetic interference (EMI) challenges on the PCB. FRL addresses this through Link Training, which selects a signal amplitude strong enough to maintain link integrity while minimizing EMI.
- **Robust Error Handling:** At aggregate data rates above 40 Gbps, bit errors are unavoidable. FRL mitigates this by combining adaptive link optimization with Forward Error Correction (FEC), maintaining reliable video transmission despite physical-layer errors.

4.4 Industry Convergence

For many years, HDMI was an outlier. While the rest of the world moved toward packet-based, clock-embedded, adaptive protocols, HDMI stuck to its legacy TMDS architecture.

With the adoption of FRL, HDMI is aligning with the industry standard. This convergence is driven by the reality that at multi-gigabit speeds, physics is universal. Whether the data is video (HDMI), file transfers (USB), or raw compute instructions (PCIe), the method of moving bits reliably at is mathematically identical.

4.4.1 The Universal Physical Layer (PHY)

Modern System-on-Chip architects design chips for laptops, consoles, and GPUs. In the past, they had to design a specialized, unique physical block just for HDMI TMDS. Now, because HDMI FRL uses the same fundamental signaling principles as other standards, engineers can use Unified SerDes PHYs.

A single high-speed PHY block on a silicon die can be configured via firmware to act as a DisplayPort lane, a USB4 lane, or an HDMI FRL lane. This lowers manufacturing costs and accelerates time-to-market.

4.4.2 The Convergence Peers

HDMI 2.2 borrows the best practices that have been battle-tested in *DisplayPort* and *PCIe*.

DisplayPort has used a Main Link with embedded clock and micropacket architecture since its inception. HDMI FRL effectively adopts this topology. Both standards now use Link Training to negotiate lane counts and speeds.

HDMI's FRL Link Training State Machine (LTSM) mimics the robustness of PCIe. Just as a GPU negotiates with a motherboard to find the fastest stable speed (e.g., PCIe 4.0 x16), an Xbox now negotiates with a TV to find the fastest stable FRL rate. They both use Linear Feedback Shift Registers (LFSR) for scrambling to reduce EMI.



4.5 Real-World Channels Are Not Ideal

In a textbook or simulation, an HDMI link is often modeled as a perfect differential pair. In reality, the channel is a challenging environment for high-speed signals. Without Link Training, the Source is essentially blind, unable to compensate for the unique properties of the attached channel.

4.5.1 Cable Length and Construction

The cable is the dominant variable in the link quality.

- **The Physics:** All copper cables suffer from Insertion Loss, the weakening of the signal as it travels. This loss is frequency-dependent. High-frequency transition are attenuated much faster than low-frequency patterns.
- **The Construction Factor:** Thicker wires (24 AWG) have less resistance but make the cable stiff and heavy. Thinner wires (30 AWG) are flexible but suffer massive signal loss.

4.5.2 Connector Quality

The connector is often the weakest link in the signal chain.

- **Impedance Discontinuity:** The geometry of the HDMI connector pins is physically different from the twisted pair wires inside the cable. This transition point creates a "bump" in impedance.
- **Return Loss:** Every time the signal hits an impedance mismatch (e.g., going from wire to a connector), part of the signal energy bounces back toward the Source. These reflections collide with incoming data, creating noise known as Inter-Symbol Interference.
- **Tolerances:** A "loose" port or a budget connector with poor gold plating can introduce contact resistance, further degrading the signal.

These real-world impairments make a static, one-size-fits-all signaling scheme inefficient and unreliable at very high data rates. FRL addresses these challenges by introducing mandatory link training and adaptive equalization. During link training, the Sink evaluates the incoming signal quality and instructs the Source to adjust transmit parameters such as voltage swing and pre-emphasis to compensate for insertion loss and high-frequency attenuation. This directly mitigates cable-length and construction variability.

FEC in turn allows the link to tolerate a controlled level of bit errors caused by reflections, return loss, and connector imperfections. Instead of requiring a perfectly clean channel, FRL is designed to operate reliably in the presence of realistic channel noise.

Real life use case: FRL gaming done right

Now consider the same PS5 and cable, as in the example [shown in chapter 2](#), but connected to a TV that supports HDMI 2.1. The PS5 can now switch from TMDS to FRL mode.

When the game requests high-bandwidth output, the PS5 and TV perform link training. The PS5 transmits a test pattern at the requested rate, while the TV measures signal quality. The cable fails at the highest speed, but instead of dropping the signal, the PS5 automatically retries at progressively lower FRL rates until a stable link is established.

The user sees a working image. The PS5 may adjust settings, such as limiting the refresh rate to 4K @ 60 Hz, to match the cable's capabilities. The key difference is that FRL with link training adapts to real-world limitations, providing a stable experience even when the physical medium is marginal.

5. NEW CHALLENGES FOR VENDORS

HDMI 2.1, followed by HDMI 2.2, introduces a major shift, moving from the legacy TMDS signaling to the more flexible FRL architecture. This change enables support for modern high-resolution, high-frame-rate, and HDR content, but also introduces new design challenges for hardware vendors.

In this chapter, we examine three key challenges:

1. **Implementation Complexity:** How link training, state machines, and forward error correction increase the design and validation effort.
2. **Legacy Design Inertia:** How the shift from TMDS to FRL impacts PHY, controller, and firmware architectures.
3. **Interoperability Challenges:** Real-world issues arising from cable limitations, timing-sensitive training, and early device variations.

5.1 Implementation Complexity

The transition from TMDS to FRL in HDMI 2.1 with 2.2 extending that further, introduced significant implementation complexity. Unlike TMDS, where the data rate was tied directly to the pixel clock and resolution, FRL uses fixed lane rates that are negotiated during an initial handshake. This requires both the source and sink to perform link training, testing signal integrity before video transmission begins.

As described in chapter 3.4, link training involves multiple coordinated steps and the link is considered stable only when the link training succeeds. Because FRL operates at very high bandwidths, the standard also mandates FED, allowing the receiver to detect and correct bit errors in real time without retransmission.

To support these mechanisms, vendors must implement complex state machines that manage multiple training phases, retries, fallbacks, and strict timing requirements. The logic must handle lane alignment, equalization, rate negotiation, and power transitions reliably. High-speed signaling increases sensitivity to jitter, crosstalk, process-voltage-temperature (PVT) variations, and cable quality, all of which must be accounted for during validation.

5.2 From TMDS to FRL: Overcoming Legacy Design Assumptions

In TMDS, signaling rates were closely tied to video resolution and pixel clock: moving from 1080p to 4K required proportionally increasing the clock frequency. FRL instead uses fixed lane rates, independent of resolution or frame format. This shift enables higher bandwidths but requires a packetized architecture, insertion of null packets for lower-resolution content, and more sophisticated timing and buffering logic.

The change impacts every layer of HDMI implementation:

- **Physical Layer:** Lanes are AC-coupled, the clock is embedded in the data stream, and clock-data recovery circuitry must handle self-clocking signals.
- **Controller Layer:** Video data is encoded using 16b/18b schemes, packetized, framed, and protected with FEC, making the controller more like a data transport engine than a simple pixel pipeline.
- **Software/Firmware Layer:** Dynamic link management handles link training, retries, fallbacks, and readiness signaling. HDCP and other content-protection layers must synchronize with the packetized data flow.

These architectural shifts required vendors to rethink long-established design assumptions, resulting in higher development effort and silicon complexity.

5.3 Interoperability Challenges

The architectural leap from TMDS to FRL led to early real-world interoperability issues. Because HDMI 2.2 allows for optional feature subsets, early devices often implemented only parts of the FRL specification. Some chipsets supported only 32 Gbps or 40 Gbps instead of the 48 Gbps, limiting support for uncompressed 12-bit 4K @ 120 Hz video in the early days of HDMI 2.1 introduction. Differences in FEC handling, state machine timing, and training behavior caused inconsistencies across devices.

FRL link training is highly time-sensitive. If the source and sink fail to complete training within the specified limits, video output may never start. Marginal cables or environmental variations could push training beyond timeout thresholds, causing repeated retraining attempts, flickering, or black screens. Vendors often introduced proprietary workarounds to stabilize the link within their ecosystem, but these could create failures with third-party devices.

These challenges underscore why HDMI 2.2's FRL, while solving the black-screen problem inherent in TMDS for high-bandwidth content, also requires careful implementation, testing, and validation to ensure reliable operation across diverse real-world setups.



6. LESSONS FROM DISPLAYPORT LINK TRAINING

5.1 DP Has Always Assumed an Adaptive Channel

From its first revision, DisplayPort assumed that the physical channel would vary in quality and require active management. DisplayPort link training was built on the expectation that signal conditions could change due to cable quality, board layout, or operating environment. Rather than relying on fixed signaling assumptions, the protocol supports scalable lane counts and data rates that can be adjusted during link bring-up. This adaptive mindset shaped a more resilient approach to high-speed link operation, one that HDMI later adopted with FRL.

5.2 DP Link Training in Practice

DisplayPort link training is divided into well-defined phases that progressively stabilize the link before video transmission begins. The process starts with a **clock recovery (CR)** phase, where the sink verifies that it can reliably extract timing information from the incoming data stream. Once a stable clock is established, the link moves into the **channel equalization (EQ)** phase, where the sink evaluates signal quality and requests adjustments to improve eye opening and margin.

A key strength of DisplayPort is its support for per-lane tuning, allowing voltage swing and pre-emphasis to be adjusted independently on each lane. DP's structured, feedback-driven training model has allowed DisplayPort to scale from early data rates of 1.62 Gbps to more than 20 Gbps per lane while maintaining interoperability and robustness across a wide range of platforms and cable conditions.

5.3 What FRL Borrows - and What It Does Not

HDMI FRL borrows many proven concepts from DisplayPort link training while adapting them to the constraints of consumer electronics. Although the two approaches share the same fundamental goals, their implementations reflect different design priorities.

5.3.1 Similarities

Aspect	Description
Sink-driven feedback	Like DisplayPort, FRL relies on the sink to evaluate signal quality during link training. The display reports errors or margin issues, guiding the source toward stable link settings.
Fixed link rates	Both standards define a fixed set of supported lane counts and data rates. The link is established at standardized speeds rather than arbitrary frequencies, improving interoperability.
Explicit training states	FRL introduces a formal training sequence with defined states and transitions, ensuring both source and sink follow a common procedure before video transmission begins.

5.3.2 Differences

Aspect	Description
Appliance-oriented design	FRL is designed for consumer devices such as TVs, consoles, and AV receivers, where cost, power consumption, and simplicity are critical. DisplayPort targets PC-class systems with fewer such constraints.
Limited tuning flexibility	FRL exposes fewer tuning parameters than DisplayPort. While this reduces implementation complexity, it also limits how precisely marginal links can be optimized.
Backward compatibility requirements	HDMI must maintain seamless support for legacy TMDS devices and cables. FRL training and fallback behavior are therefore tightly integrated with backward compatibility logic, adding constraints that DisplayPort does not face.

7. TESTING FRL

Link Training on UCD Console Software

Unigraf's UCD-422 is a FRL capable video generator and analyzer. With UCD Console Software it allows FRL link training, link configuration and status monitoring. Users can configure FRL link parameters and initiate Link Training. The result of the link training is shown alongside FRL mode and other parameters.

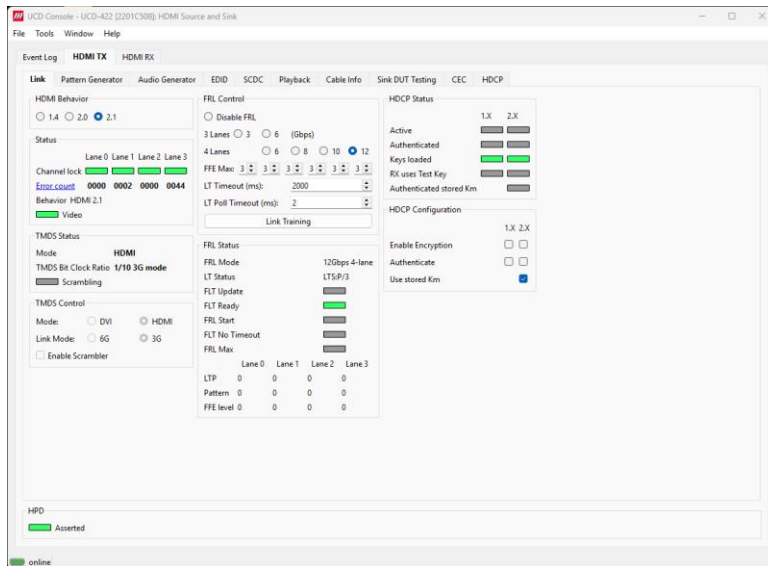


Figure 3. HDMI 2.1 Link Status and Configuration in UCD Console GUI on HDMI TX role.

HDMI Link Analyzer

HDMI Link Analyzer is a tool for capturing Main-link Data Events and I2C Transactions. For each bulk capture, the data is parsed in a timeline view presenting the captured frames. Within each frame, users can have a deep view of the events and metadata that occurs in each line. Users can measure distances in nanoseconds and symbol size to identify the length of problems in between events. Link analyzer enables observing descrambled and scrambled symbols in hexadecimal format in each of the active lanes.

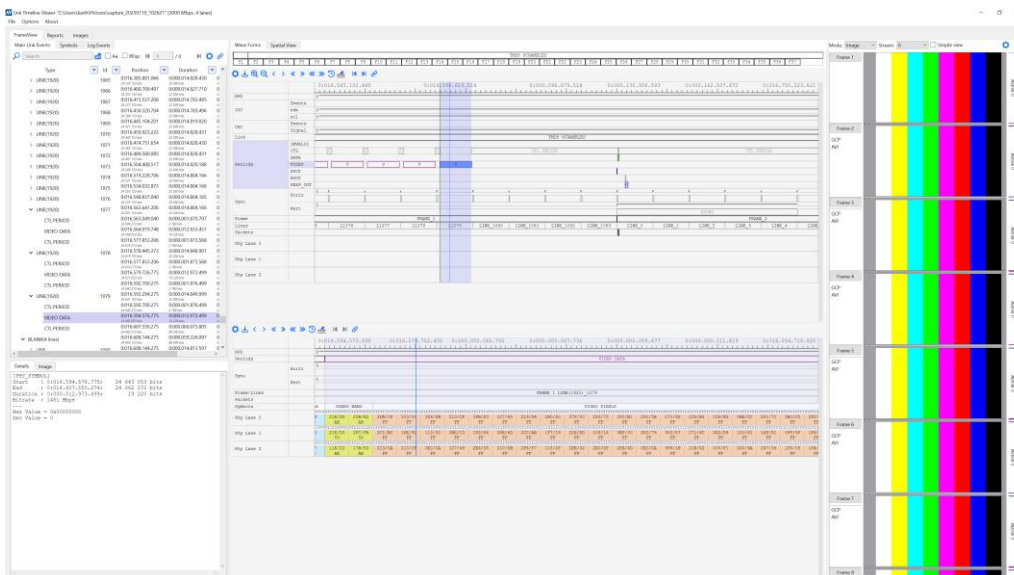


Figure 4. HDMI Link Timeline viewer UI



www.unigraf.fi

