

DEMYSTIFYING PAM SIGNALING

Products:

- R&S®RTP
- R&S®RTO6
- R&S®RTP-K135
- R&S®RTO6-K135

Dhanshree Lade | 1SL415 | Version 0e | 10.2025

<https://www.rohde-schwarz.com/appnote/1SL415>



Contents

1	Introduction: The Rise of PAM Signaling	3
2	Understanding PAM Signaling	3
2.1	PAM-N Fundamentals.....	4
2.2	PAM in Communication Standards.....	6
2.3	Practical Considerations: PAM vs. NRZ in Test and Measurement	6
3	Implementation Nuances for PAM Signaling	9
3.1	AC vs. DC Coupling	9
3.2	Unidirectional vs. Bidirectional Communication in PAM	9
3.3	Single-Lane vs. Multi-Lane Interfaces	10
3.4	Line Coding Schemes (8B/10B, 64B/66B, 128B/130B).....	10
3.5	Scrambling and its Measurement Implications	11
3.6	Forward Error Correction (FEC)	11
3.7	Optimizing Testing Strategies for PAM Signaling Across Diverse Standards....	12
4	PAM-N Signal Integrity Analysis	13
4.1	Signal Integrity Foundations: From NRZ to PAM-N.....	13
4.2	Advanced Eye, Jitter Analysis and Equalization with Rohde & Schwarz Oscilloscopes	15
4.2.1	PAM Eye Diagram and Eye Isolation	16
4.2.2	Transition-Level Jitter	17
4.2.3	Equalization Techniques for PAM-N.....	19
5	Summary and Outlook	21
6	Literature	22
7	Appendix	23
A	Appendix 1: Comprehensive Overview of PAM-Based Standards	23

1 Introduction: The Rise of PAM Signaling

As serial communication systems evolve to meet growing data demands, traditional binary signaling formats such as Non-Return-to-Zero (NRZ), also referred to as PAM-2, have historically been used in serial standards such as PCI Express® Gen1–Gen5 [1], USB 3.x SuperSpeed interfaces [2], etc. because of their simplicity and large signal margins. However, as serial data rates continued to increase, NRZ began to face fundamental limitations in bandwidth scalability and signal integrity. These limitations have prompted the adoption of Pulse Amplitude Modulation (PAM), a multilevel signaling approach that transmits multiple bits per symbol using discrete voltage levels.

PAM-N refers to a generalized form of PAM where N denotes the number of amplitude levels. For instance, PAM-4 uses four levels to encode two bits per symbol, while PAM-8 and PAM-16 can encode three and four bits per symbol, respectively. Higher-order formats improve spectral efficiency by encoding more bits per symbol at the same symbol rate, but they also introduce new challenges in terms of design complexity, noise susceptibility, and increased sensitivity requirements for test and measurement equipment.

Today, PAM-4 is widely adopted across multiple high-speed serial communication standards, including various generations of Ethernet (such as 100G, 200G, and 400G) [3] [4] [5], PCI Express (PCIe) Gen 6 [1], and electrical interface specifications developed by the Optical Internetworking Forum (OIF) [6]. These standards use PAM-4 to double the bit rate per lane without doubling channel bandwidth. The gain in efficiency comes with reduced eye margins, greater noise and jitter sensitivity, and stricter equalization needs, making PAM-4 a compromise between higher throughput and signal integrity limits.

This educational note provides a technical overview of PAM-N signaling from a test and measurement perspective. It introduces the foundational principles of multilevel modulation and the relationship between bit rate and symbol rate. Furthermore, PAM is compared with NRZ in terms of signal structure, eye diagram behavior, clock recovery challenges, coding schemes, and equalization techniques. Finally, it explains how R&S® test solutions enable meaningful analysis of PAM-N signals, providing insight into signal integrity, transition-level jitter, and equalization techniques.

2 Understanding PAM Signaling

Pulse amplitude modulation improves bandwidth efficiency by coding multiple bits per symbol with discrete voltage levels. In earlier generations of high-speed serial communication, the adoption of PAM was limited by technology constraints such as transmitter linearity, Digital-to-Analog Converter (DAC) resolution, and available equalization techniques.

During this time, NRZ (PAM-2) remained the dominant signaling format. Its simplicity, large signal margins, and compatibility with existing infrastructure enabled it to scale into the multi-gigabit range across many serial interfaces. However, as system requirements pushed data rates upward, NRZ signaling encountered increasing signal-integrity constraints, particularly from high-frequency attenuation and inter-symbol interference (ISI). For example, at data rates approaching 25 Gb/s, NRZ signals over copper channels experience significant ISI-induced eye closure, limiting channel reach and degrading timing margin.

These limitations motivated the shift toward PAM-N signaling, where information is conveyed using multiple discrete amplitude levels. While this improves spectral efficiency, it also reduces the vertical and horizontal eye margins, increasing noise sensitivity and measurement complexity.

The following subsections introduce key concepts such as PAM bit encoding, symbol versus bit rate, and standard implementations. These foundations support later discussions on signal behavior, equalization, and test strategies.

2.1 PAM-N Fundamentals

In Pulse Amplitude Modulation (PAM-N), information is conveyed by mapping each transmitted symbol to one of N discrete voltage levels. This allows multiple bits to be encoded per symbol. The number of bits per symbol follows:

$$\text{Bits per symbol} = \log_2(N)$$

For example, PAM-2 encodes 1 bit per symbol using two voltage levels, PAM-4 uses four levels for 2 bits, and PAM-16 uses sixteen levels for 4 bits. The relationship between PAM levels and bits per symbol is summarized in Table 1:

Modulation	Levels (N)	Bits/Symbol	Number of Unique Bit Combinations	Example Bit Combinations
PAM-2 (NRZ)	2	1	2	0, 1
PAM-4	4	2	4	00, 01, 10, 11
PAM-8	8	3	8	000 to 111
PAM-16	16	4	16	0000 to 1111

Table 1: Relationship Between PAM Levels and Bits per Symbol

This encoding property allows PAM signaling to increase data throughput without increasing the symbol rate. Understanding the distinction between symbol rate and bit rate is essential when configuring test equipment, performing compliance checks, and interpreting signal behavior in PAM-based systems:

- Bit rate (Gbps): The number of data bits transmitted per second.
- Symbol rate (GBd or baud): The number of symbols transmitted per second, where each symbol may represent more than one bit.

The relationship between bit rate and symbol rate for PAM-N systems is defined as:

$$\text{Bit Rate} = \text{Symbol Rate} \times \log_2(N)$$

Higher-order PAM schemes achieve higher bit rates at the same symbol rate, which can be beneficial in bandwidth-constrained environments. Note that this is the theoretical bit rate under ideal conditions. In practice, the effective data rate is the net throughput available for user data, which is lower because of overhead from line coding (Section 3.4), scrambling (Section 3.5), or error correction such as forward error correction (FEC, Section 3.6).

Table 2 shows how increasing the PAM order can achieve higher bit rates without increasing the symbol rate, which is especially useful in environments with limited bandwidth. A constant symbol rate of 28 GBd (a common reference in high-speed communications standards) has been considered for better illustration.

Modulation	Levels (N)	Bits/Symbol	Symbol Rate (GBd)	Resulting Bit Rate (Gbps)
PAM-2 (NRZ)	2	1	28	28
PAM-3	3	~1.58	28	~44.2
PAM-4	4	2	28	56
PAM-5	5	~2.32	28	~65
PAM-8	8	3	28	84
PAM-16	16	4	28	112

Table 2: Example Calculations of Bit Rate vs Symbol Rate for Various PAM Formats

(Note: PAM-3 and PAM-5 bit rates are approximate because they encode non-integer bits per symbol, requiring block encoding (e.g., ternary-to-binary mapping or redundancy). Their effective throughput depends on the coding scheme used.)

In practice, the symbol rate influences timing-related characteristics in digital communication such as unit interval UI (time duration allocated for one symbol), eye width (Section 4.1), jitter sensitivity, latency, synchronization time and required channel bandwidth. These parameters are crucial for the performance and reliability of communication systems and must be considered when planning and implementing networks. The bit rate reflects actual data throughput and affects bit-level functions like scrambling, line encoding, and FEC.

To transmit these bits efficiently, PAM systems map bit combinations to amplitude levels using schemes such as Gray coding. Gray coding minimizes error propagation, improving decoding accuracy in the presence of noise. In Gray-coded PAM, adjacent voltage levels differ by only one bit, helping reduce the probability of multi-bit errors caused by small voltage variations. Table 3 shows a typical, Gray-coded mapping for PAM-4, and Figure 1 illustrates how PAM-4 compares with NRZ in the resulting signal waveform and eye diagram.

Bits	Amplitude Level
00	-3
01	-1
11	1
10	3

Table 3: Bit-to-Level Mapping Example for PAM-4

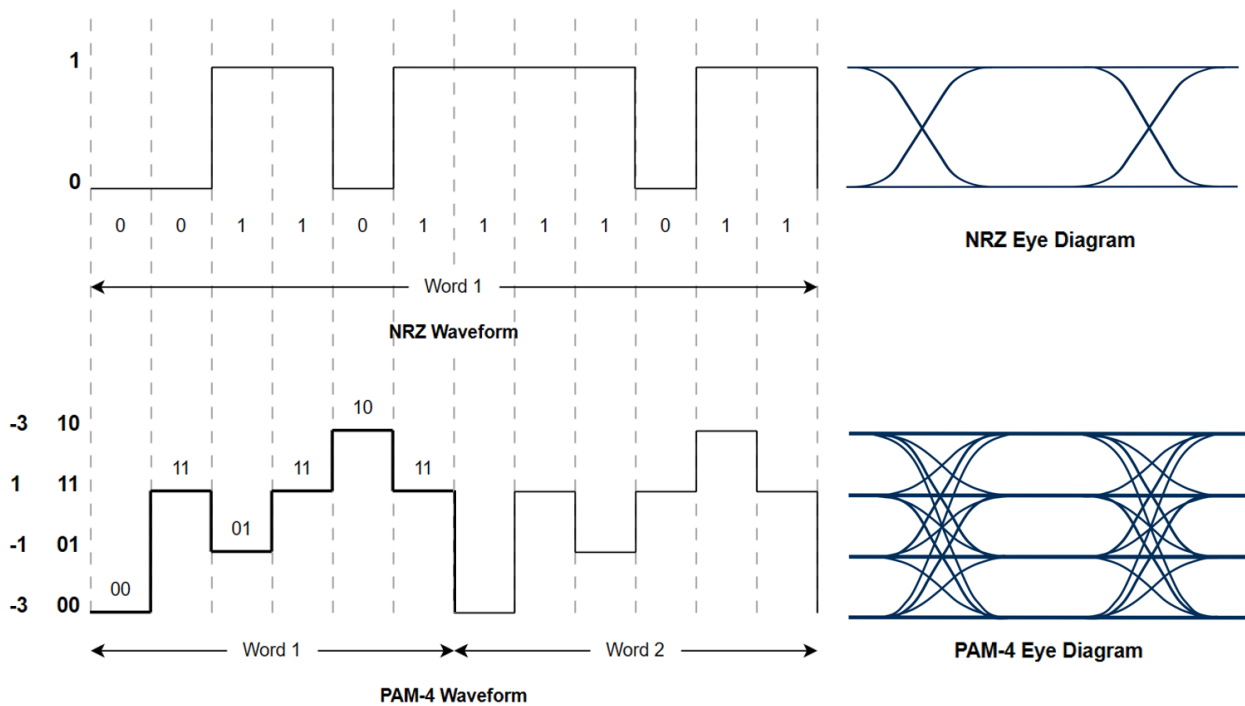


Figure 1: NRZ vs. PAM-4 Signal and Eye Representation (Source: own illustration)

As the number of voltage levels increases, both the vertical (amplitude) and horizontal (timing) margins shrink. Narrower amplitude spacing reduces noise immunity, while multiple transition magnitudes create varying slew rates and edge timing, increasing jitter sensitivity and clock-recovery complexity.

Figure 1 illustrates how increasing PAM order affects signal behavior. NRZ produces a single eye opening between two levels, whereas PAM-4 forms three stacked eyes between four levels. Each eye must be analyzed individually in terms of timing and amplitude margin. In Figure 1, “word1” and “word2” represent two consecutive groups of bits mapped to amplitude levels, highlighting how multibit PAM-4 symbols differ from single-bit NRZ words.

These signal characteristics directly influence jitter decomposition, equalization strategy, and compliance testing. For pulsed amplitude modulated signals with increasing order, measurement tools like the Rohde & Schwarz Oscilloscope R&S® RTP® or R&S® RTO6 equipped with the PAM-N Analysis software option (R&S® RTP-K135/RTO-K135) can reliably capture and interpret stacked eye openings, transition-dependent dynamics, and the reduced noise margins of multilevel signals. All these aspects are explored further in later chapters.

2.2 PAM in Communication Standards

PAM signaling has been adopted across a range of communication standards, spanning industrial and automotive Ethernet, high-speed serial links, and consumer protocols. The basic idea of sending multiple bits in one symbol using different amplitude levels remains the same, but the selection of PAM order, symbol rate, and bandwidth can change a lot based on the specific application. Representative examples of PAM-based standards are given in Table 4. A more extensive list of PAM-based standards is provided in the Appendix 1.

Standard	PAM Order	Data Rate	Symbol Rate	Channel Bandwidth	# Pairs / Lanes	Application
1000BASE-T [7]	PAM-5	1 Gbps	125 MBd	~62.5 MHz	4 twisted pairs	Standard Ethernet
1000BASE-T1 [8]	PAM-3	1 Gbps	750 MBd	~375 MHz	1 twisted pair	Automotive single-pair Ethernet
PCIe 6.0 [1]	PAM-4	64 GT/s	32 GBd	~16 GHz	1 lane	High-speed computer interconnect
USB4 v2 [9]	PAM-3	80 Gbps	~25.6 GBd	~12.8 GHz	2 lanes	USB over Type-C
100GBASE-DR [3]	PAM-4	100 Gbps	53.125 GBd per lane	~25 GHz (optical)	1 optical lane	Data center serial link

Table 4: PAM-Based Communication Standards

This comparison illustrates how system design decisions drive the choice of PAM order. For instance, both 1000BASE-T and 1000BASE-T1 achieve 1 Gbps data throughput but use different approaches. 1000BASE-T operates over four twisted pairs using PAM-5, resulting in a relatively low symbol rate (125 MBd per pair) and modest channel bandwidth (~62.5 MHz). In contrast, 1000BASE-T1 transmits over a single pair using PAM-3, requiring a much higher symbol rate (~750 MBd) and wider bandwidth (~375 MHz). This choice favors reduced cabling in automotive environments, where EMI robustness and physical simplicity take priority. While PAM-3 and PAM-5 have been adopted effectively in such contexts, they are less common in newer high-speed interfaces. One reason is that the gain in data transfer per symbol is lower compared to PAM-4, and the use of an odd number of levels introduces additional coding and decoding complexity, such as requiring specialized non-binary line coding (e.g., ternary-to-binary mapping or trellis-coded modulation), increasing digital complexity and overhead [8] [7]. As a result, standards such as PCIe 6.0 and various generations of Ethernet use PAM-4 to enable data rates beyond 25 Gbps per lane.

2.3 Practical Considerations: PAM vs. NRZ in Test and Measurement

Multilevel PAM formats, such as PAM-4 and PAM-8, enable higher data throughput without proportionally increasing the symbol rate. While PAM formats improve throughput, they also introduce new considerations for system design, signal integrity, and test methodology. This section outlines key differences between PAM-N and NRZ (PAM-2) signaling from a test and measurement perspective, focusing on signal complexity, bandwidth efficiency, and characterization strategies.

- **Signal Margins and Timing Behavior**

In NRZ signaling, the voltage swing is divided between two levels, resulting in a single eye opening with relatively large vertical separation and strong noise resilience. PAM-N formats divide the same voltage range into multiple levels, reducing the vertical margin between adjacent symbol states. For example, PAM-4 produces three eye openings, each typically one-third the height of an NRZ eye (assuming equal level spacing), increasing susceptibility to voltage noise, power-supply variation, and crosstalk.




Figure 2: Comparison of NRZ and PAM-4 eye diagrams

Figure 2 compares NRZ and PAM-4 eye diagrams, where the PAM-4 signal shows three smaller stacked eyes that each require separate evaluation for signal quality. Horizontal (timing) margins within the constant unit interval (UI, the time allocated for one symbol) are also affected. Although the UI is defined by the symbol rate and remains unchanged, PAM's varying transition magnitudes (e.g., small 2→3 versus large 0→3 transitions) introduce different slew rates and edge characteristics, reducing effective eye width and increasing jitter sensitivity.

These effects cause additional signal variability such as transition-dependent jitter and inter-symbol interference (ISI, distortion caused by one symbol interfering with subsequent ones). Test equipment must account for these factors by performing symbol-aware, level- and edge-dependent analysis of jitter and eye diagrams. These measurement considerations are examined in greater detail in Chapter 4, which covers multi-level eye analysis, transition-specific jitter decomposition, and equalization techniques for PAM-N signals.

- **Clock Recovery and Pattern Awareness**

Clock recovery methods differ significantly between NRZ and PAM-N systems. In NRZ, clock and data recovery (CDR, a circuit that extracts timing information from the incoming data stream) circuits typically rely on regular edge transitions to maintain timing lock. PAM-N systems, however, often contain flat regions or small-amplitude transitions that make traditional edge-based CDR less reliable. This increases reliance on statistical, symbol-aware, or pattern-based CDR techniques that analyse symbol history or use encoding patterns to infer timing. Such methods are particularly important when working with scrambled or line-encoded PAM signals, where direct observation of transitions may not yield sufficient information for stable clock recovery.

- **Measurement Strategy**

PAM-N signals require a more granular approach to signal analysis. While NRZ systems typically use a single eye diagram and one decision threshold, PAM-N requires multiple vertically stacked eye diagrams and corresponding decision thresholds for each symbol level. Each eye needs to be evaluated independently in terms of vertical closure, timing margin, and jitter.

Test equipment must support additional metrics specific to PAM systems, such as vertical eye closure (VEC, a measure of the vertical eye margin reduction due to impairments), ratio of level mismatch (RLM, indicating amplitude imbalance among PAM levels), and per-level bit error rates. Eye mask testing becomes more detailed, with distinct templates applied to each eye. Jitter decomposition may also require transition-aware separation, especially when assessing performance across varying edge profiles.

- **Selection of suitable Measurement Equipment**

PAM-N's stacked eyes, reduced vertical margins, and transition-dependent jitter demand higher precision in both resolution and analysis. Oscilloscopes and bit error rate testers (BERTs, instruments that measure bit error rates by generating and analysing test patterns) must therefore provide high vertical resolution, symbol-aware CDR, and per-level eye and jitter analysis. Although PAM-N signals may appear more degraded due to reduced margins, equalization and forward error correction (FEC – Section 3.6) help maintain reliability. The required analysis depends on the test objective: transmitter tests emphasize pre-FEC eye and jitter performance, while receiver/system tests require post-FEC BER or FEC emulation.

Table 5 summarizes key differences between NRZ and PAM-N signaling from a test and measurement perspective.

Aspect	NRZ (PAM-2)	PAM-N (N > 2)
Voltage Levels	2 levels	Multiple discrete levels (e.g., 4 for PAM-4)
Vertical Margin	Affected by dispersion and Inter-symbol interference (ISI), but large voltage margin improves tolerance to eye closure.	Same as NRZ plus reduced margin per level due to voltage swing division
Horizontal Margin	Slew rate is affected, consistent transitions	Narrower per-eye timing, transition-dependent variability
Transitions	2 symmetric transitions	Multiple transitions of varying amplitude and direction
Eye Diagram	Single eye opening	Stacked eye openings requiring independent analysis
Jitter Analysis	Global jitter metrics	Transition-dependent jitter decomposition required
CDR Requirements	Edge-triggered, pattern/symbol lock not necessarily needed	Symbol-aware or pattern-based; flat regions complicate recovery
Measurement Metrics	Eye opening in height and width, jitter, eye mask	Same as NRZ plus VEC, RLM, per-level BER, per-eye masks
Eye Diagram Generation	Single decision threshold and basic resolution	Multiple decision thresholds (level-based separation) with high vertical resolution and SNR
Evaluation Techniques	Visual inspection sufficient; single transition makes ISI easy to spot. Equalization may be applied if needed.	Multiple overlapping transitions reduce clarity; FEC and equalization are often needed. Per-level eye views help apply familiar NRZ metrics.

Table 5: Summary of test and measurement differences between NRZ and PAM-N

Understanding these signal-level challenges, such as reduced vertical and horizontal margins, stacked eyes, transition-dependent jitter, and the need for symbol-aware clock recovery, is only part of the picture. To effectively analyze and address them, one must also understand the underlying system, including its encoding schemes, symbol structure, and channel behavior. These signal-level challenges set the stage for understanding how architectural elements such as line coding, scrambling, and error correction influence performance, which are introduced in Chapter 3.

3 Implementation Nuances for PAM Signaling

PAM-N signaling is influenced not only by its multilevel encoding but also by surrounding system architecture, which includes coupling, line coding, error correction, signal coupling, clock recovery, and link directionality. These aspects affect signal integrity, receiver complexity, test requirements, and standard compliance. This chapter explores these implementation details.

3.1 AC vs. DC Coupling

PAM signals can be transmitted using either AC or DC coupling, depending on the system architecture and physical interface requirements. The choice between these coupling methods has implications for both signal integrity and test methodology.

DC coupling allows the full voltage swing of the signal to pass through without alteration, preserving absolute voltage levels. This is typically used in short-reach, controlled environments such as chip-to-chip or backplane links, where DC balance is maintained, and baseline wander is minimal.

AC coupling blocks the DC component of a signal using series capacitors, allowing only the time-varying portion to pass. It is commonly used in differential high-speed links over long distances and backplanes where voltage offsets may exist between transmitter and receiver. While AC coupling improves interoperability and protects against ground loops, it can introduce baseline wander, which is a slow drift in the signal's zero level caused by long sequences of symbols with similar voltage levels. To minimize this effect, PAM systems typically employ line coding and scrambling to maintain transition density and suppress low-frequency content. However, some residual drift may persist, potentially affecting the accuracy of eye diagram measurements. In such cases, test equipment must apply baseline correction or DC restoration to recover the intended waveform shape.

Understanding the coupling method in use is essential for selecting the appropriate measurement strategy: with DC coupling, absolute voltage levels can be measured directly, while with AC coupling, the blocked DC component can cause baseline wander, so baseline correction or DC restoration is required to obtain accurate eye margins and to place decision thresholds correctly across multiple levels.

3.2 Unidirectional vs. Bidirectional Communication in PAM

PAM-based systems may support either unidirectional or bidirectional communication, depending on the protocol design (Figure 3).

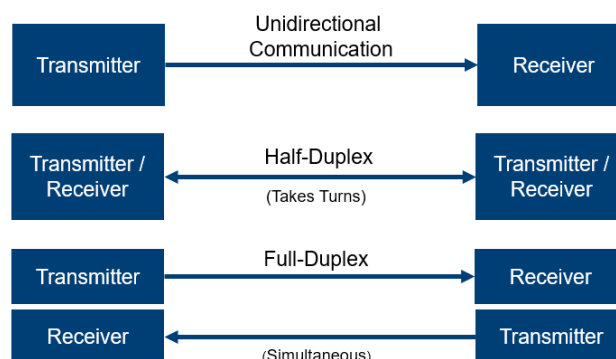


Figure 3: Directionality in PAM-based Communication System

In unidirectional communication, data flows in one direction only, from a transmitter to a receiver. This is typical for links where the signal path does not require a return channel, such as display interfaces or point-to-point data delivery.

In bidirectional communication, the physical layer (PHY) supports both transmission and reception of data. Bidirectional communication can operate in two modes. This can be realized by using one of the two operation modes described below:

- **Full duplex:** Transmission and reception occur simultaneously. This requires additional PHY complexity, such as echo cancellation, to handle overlapping signals on the same link.
- **Half duplex:** Transmission and reception alternate in time. Only one side communicates at a time, and timing coordination is essential to avoid collisions. This is often implemented using time-division multiplexing (TDM).

It's important to note that bidirectionality refers to the capability of the PHY, not necessarily the number of physical lanes or wires used. Definitions related to single-lane and multi-lane configurations are discussed separately in Section 3.3.

3.3 Single-Lane vs. Multi-Lane Interfaces

A PAM-based interface can be categorized as single-lane or multi-lane, based on the number of independent channels used to carry data.

Single-lane interfaces transmit all signal data over one physical lane. Depending on system design, single-lane configurations may operate in unidirectional or bidirectional modes (see also Section 3.2).

Multi-lane interfaces use several parallel lanes to increase aggregate throughput without increasing the symbol rate per lane. Lanes may operate independently or be tightly synchronized to ensure timing alignment and data reconstruction.

Understanding whether an interface is single- or multi-lane is essential for interpreting per-lane test results, characterizing skew, and aligning compliance measurements to standard-specific lane configurations. For example, in 1000BASE-T, the four twisted pairs act as four parallel lanes (as introduced in Section 2.2). Because the data is distributed across multiple lanes, each lane operates at a much lower symbol rate (125 MBd per pair) than in single-pair systems such as 1000BASE-T1, which must carry the full data stream over one lane at 750 MBd. [10]

3.4 Line Coding Schemes (8B/10B, 64B/66B, 128B/130B)

Line coding is used in high-speed serial links to maintain synchronization, support clock recovery, and ensure balanced transmission characteristics. In PAM-N systems, the selected line coding scheme affects not only data throughput but also transition density, error detection capability, and signal analysis complexity.

- **8B/10B Encoding**

This scheme maps 8-bit data blocks to 10-bit code words. It ensures frequent transitions and DC balance, which helps with clock recovery and baseline wander suppression. However, the 25% overhead makes it less efficient and is thus majorly used in NRZ.

- **64B/66B Encoding**

By mapping 64 bits of data into a 66-bit block, this format reduces overhead to approximately 3%. It maintains sufficient transition density, especially when combined with scrambling to further randomize data patterns.

- **128B/130B Encoding**

Overhead is just 1.5%. It includes embedded synchronization markers and is typically combined with scrambling and forward error correction to ensure signal robustness.

The choice of line coding affects the structure of the transmitted signal and the nature of its transitions. Efficient coding schemes reduce bandwidth overhead, and those with embedded framing or synchronization patterns can simplify symbol-aligned acquisition during testing by helping instruments lock onto block or symbol boundaries.

3.5 Scrambling and its Measurement Implications

Scrambling is a common technique used in high-speed serial links to randomize the bitstream before modulation. Unlike line coding, which structures data into defined blocks, scrambling introduces controlled randomness to improve spectral balance, reduce EMI, and maintain baseline stability. In multilevel PAM systems, it helps avoid long runs of identical symbols that could cause DC baseline drift and degrade clock recovery.

Scrambling is typically implemented using linear feedback shift registers (LFSRs) that apply a deterministic polynomial-based sequence to the data stream. For example, 64B/66B Ethernet systems use a self-synchronizing scrambler based on the polynomial $x^{58}+x^{39}+1$ [10], producing a long pseudo-random pattern that helps maintain consistent signal transitions and prevents repetitive patterns that could affect timing recovery.

While scrambling improves signal robustness, it also reduces pattern predictability. Measurement tools must therefore rely on statistical accumulation or symbol-aware acquisition, since conventional edge-based triggering is often unreliable. To overcome this, modern standards embed periodic synchronization symbols which enable pattern or block lock even in scrambled streams. These synchronization symbols allow test equipment to perform symbol-aware acquisition, facilitating stable measurements. Metrics such as vertical eye closure per level and jitter decomposition are applied to PAM-N signal analysis and become especially important when scrambling is used.

3.6 Forward Error Correction (FEC)

As the number of levels in PAM signaling increases, the vertical and horizontal margins between symbols become smaller, making the system more sensitive to noise, crosstalk, and jitter during transmission. To maintain reliable transmission in such conditions, modern PAM-N systems incorporate forward error correction schemes. While retransmission-based error correction can be effective in some systems, it introduces additional latency and relies on bidirectional communication. In high-speed serial links, where low latency and continuous data flow are critical, FEC is typically preferred. When FEC is used in a transmission system, the transmitter encodes the data to be transmitted in a redundant manner so that the receiver can detect and correct transmission errors without consulting the transmitter.

Common FEC schemes include block codes such as Reed-Solomon (RS) widely used in PAM-4-based Ethernet standards [3] [4] [5] and low-density parity-check (LDPC) codes, which are prevalent in many high-speed serial interfaces [1]. They offer robust correction capabilities with manageable latency and overhead. With these techniques PAM-N systems achieve the low bit error rates required by high-speed communication standards.

While FEC improves system robustness, it also increases receiver complexity and places additional demands on clock recovery and synchronization. These architectural and measurement considerations must be addressed during system design, performance evaluation and finally compliance testing.

From a test and measurement perspective, the presence of FEC affects how signal quality is evaluated. A signal with a degraded eye diagram may still pass system-level requirements if post-FEC error rates are within acceptable limits. As a result, test strategies must account for both pre-FEC and post-FEC performance. Measuring equipment should include FEC emulation or post-processing tools that estimate post-FEC bit error rate (BER) and signal margin.

3.7 Optimizing Testing Strategies for PAM Signaling Across Diverse Standards.

As a summary, the transmission of PAM signaling is implemented in various ways depending on the standard and its intended environment. Differences in modulation order, symbol rate, lane configuration, coding schemes, and error correction all influence how testing must be approached.

These implementation choices also translate into specific measurement constraints. Factors such as communication direction (unidirectional or bidirectional), lane configuration (single or multi-lane), use of FEC, line coding, scrambling, and coupling method impact on how signal quality can and should be evaluated. For instance, systems with bidirectional single-lane communication may require directional couplers or loopback configurations to isolate the transmitter during testing. FEC-equipped systems may mask impairments unless post-FEC BER analysis is available. Scrambling and block coding introduce acquisition challenges that necessitate symbol-aware tools capable of recognizing synchronization markers.

Understanding the specific characteristics and limitations of the system under test is essential. Test strategies must be tailored to the architecture and constraints of each standard to ensure accurate and meaningful results.

Table 6 highlights examples of common high-speed digital standards to illustrate how these implementation choices translate into practical test constraints. These details are critical for selecting appropriate test methodologies.

Standard	PAM Format	Symbol Rate	No. of Pairs / Lanes	Error Coding	Special Features
1000BASE-T	PAM-5	125 MBd/pair	4 twisted pairs	Trellis-coded modulation (TCM)	Full duplex, adaptive echo cancellation, scrambling, no FEC [7]
1000BASE-T1	PAM-3	750 MBd	1 twisted pair	Reed–Solomon	Automotive, lightweight cabling, EMI hardened, scrambling, block coding [8]
400GBASE-FR4	PAM-4	~53 GBd/lane	4 optical lanes	Reed–Solomon (RS(544,514))	Gray coding, symbol-specific testing, advanced FEC for long optical reach [3]
PCIe 6.0	PAM-4	64 GT/s	1 TX + 1 RX per lane	FEC + CRC (cyclic redundancy check, an error-detection code) + FLIT (flow control unit, a fixed-size data packet)-based retransmission	Low-latency design, high reliability, packetized transfer (FLITs) [1]
OIF CEI-112G	PAM-4	56 GBd	1 SerDes lane	Configurable (FEC optional)	Supports short to long reach, flexible equalization options [6]

Table 6: Implementation Variations Across PAM-Based Communication Standards

These implementation choices strongly influence the way signal quality is evaluated. Chapter 4 examines how multilevel signaling affects signal integrity, covering eye diagrams, jitter characteristics, and the use of equalization techniques

4 PAM-N Signal Integrity Analysis

As signaling schemes evolve from binary NRZ to multilevel Pulse Amplitude Modulation (PAM-N), the requirements for signal analysis and validation also become more demanding. Eye diagrams, jitter decomposition breakdown, crosstalk analysis, etc. are still the same signal integrity analysis tools used to validate PAM signal transmission and reception. However multilevel formats introduce overlapping eye patterns, symbol-dependent transitions, and narrower voltage and timing margins, which can affect measurement accuracy. Traditional techniques for jitter decomposition, clock recovery, and mask testing must be adapted to address these effects.

This chapter explores both the challenges, and the instrumentation solutions required for analyzing PAM-N signals. It begins by reviewing how foundational methods such as eye diagram analysis and jitter evaluation are extended to accommodate multilevel modulation formats.

4.1 Signal Integrity Foundations: From NRZ to PAM-N

Before exploring the signal analysis methods applicable to PAM-N systems, it is helpful to revisit the foundational signal integrity concepts associated with NRZ (PAM-2) signaling. The eye diagram remains one of the most widely used tools in serial link analysis. It is created by folding many consecutive unit intervals (UIs) of a digital signal on top of each other, producing a composite view of timing and amplitude behavior (Figure 4). Figure 5 shows an NRZ eye diagram with its key parameters and a histogram at the eye crossing point.




Figure 4: Formation of Eye Diagram (Source: Own Illustration)




Figure 5: Labeled Eye Diagram for NRZ Signal Analysis (Source: Own Illustration)

Some Key Eye Diagram Parameters:

- **Eye Height:** The vertical opening of the eye diagram, representing voltage margin. A taller eye indicates better noise immunity and a higher signal-to-noise ratio (SNR).
- **Eye Width:** The horizontal opening of the eye at the decision threshold, reflecting available timing margin. It narrows in the presence of jitter or ISI, which compromises bit decision reliability.
- **Zero Crossing Point:** The ideal crossing occurs at 50% of the UI. Deviations may suggest duty-cycle distortion or unequal rise/fall times.
- **Bit Period:** The bit period is the time taken to transmit one bit of data. In NRZ, one bit period corresponds to one unit interval (UI).
- **Jitter Histogram:** A histogram centered at the eye crossing point shows the distribution of transition timing. The spread of this histogram provides insight into the total jitter, including deterministic (DJ) and random components (RJ).

The measurement principles described here form the basis for PAM-N signal analysis. However, higher-order modulation formats introduce additional eye openings, symbol-dependent transitions, and asymmetric margins, requiring more advanced measurement techniques. These include symbol-aware slicing (i.e., evaluating eye and BER per level using level-specific thresholds), transition-specific jitter classification, and adaptive clock data recovery (CDR).

Clock recovery in PAM-N signals is particularly challenging because not all symbol changes result in distinct voltage transitions, which complicates timing recovery in multi-level signaling. Unlike NRZ systems, PAM symbols may repeat at the same voltage level, producing few or no transitions and making conventional edge-based CDR less reliable. To address this, test equipment must use symbol-aware CDR techniques that monitor level transitions across multiple unit intervals rather than relying solely on zero-crossings.

Rohde & Schwarz oscilloscopes (R&S® RTP® or R&S® RTO6 with K135 software option) implement such a symbol-aware CDR approach, which remains effective even when transitions are sparse or irregular. When combined with adaptive equalization, this method helps maintain timing lock and symbol alignment across a wide range of PAM formats and symbol rates, enabling robust signal reconstruction and accurate measurement in complex channel conditions.

Rohde & Schwarz oscilloscopes support both hardware-based and software-based clock recovery modes when generating eye diagrams. In hardware CDR, a dedicated real-time circuit recovers the clock directly from the incoming signal. This enables construction of a Live Eye, a continuously updating eye diagram that reflects receiver-like behavior during acquisition. In contrast, software CDR uses post-processed waveform data and configurable algorithms to reconstruct timing, offering flexibility in loop bandwidth and pattern tracking. While software CDR allows deeper analysis, hardware CDR ensures accurate representation of jitter and timing effects, especially during compliance or stress testing. The ability to switch between these modes gives flexibility when evaluating PAM-N signals.

4.2 Advanced Eye, Jitter Analysis and Equalization with Rohde & Schwarz Oscilloscopes

Section 4.1 introduced the signal integrity principles underlying multilevel signaling and the unique requirements of PAM-N clock recovery. Building on that foundation, this section focuses on how Rohde & Schwarz instruments implement practical tools for eye diagram generation, margin analysis, and jitter decomposition, specifically addressing the validation and verification challenges of PAM-N signals.

For the following measurements, a MultiGBASE-T1-capable DUT was used. The DUT is an automotive Ethernet PHY test board supporting MGBASE-T1 operation across 2.5GBASE-T1, 5GBASE-T1, and 10GBASE-T1 modes. For the measurements presented here, the board is configured for 5GBASE-T1, transmitting a PAM-4 signal at 2.8125 GBaud over a single twisted pair. The DUT provides the native single-ended outputs of the PAM-N signal, which were connected to Channel 1 and Channel 3 of the oscilloscope to obtain the differential waveform for analysis. The Signal Configuration panel allows users to specify the technology, PAM order, and baud rate, with automatic symbol rate estimation. This view also outlines the signal flow for embedding/equalization (see section 4.2.3), eye diagram generation, and jitter/noise decomposition (see Figure 6).

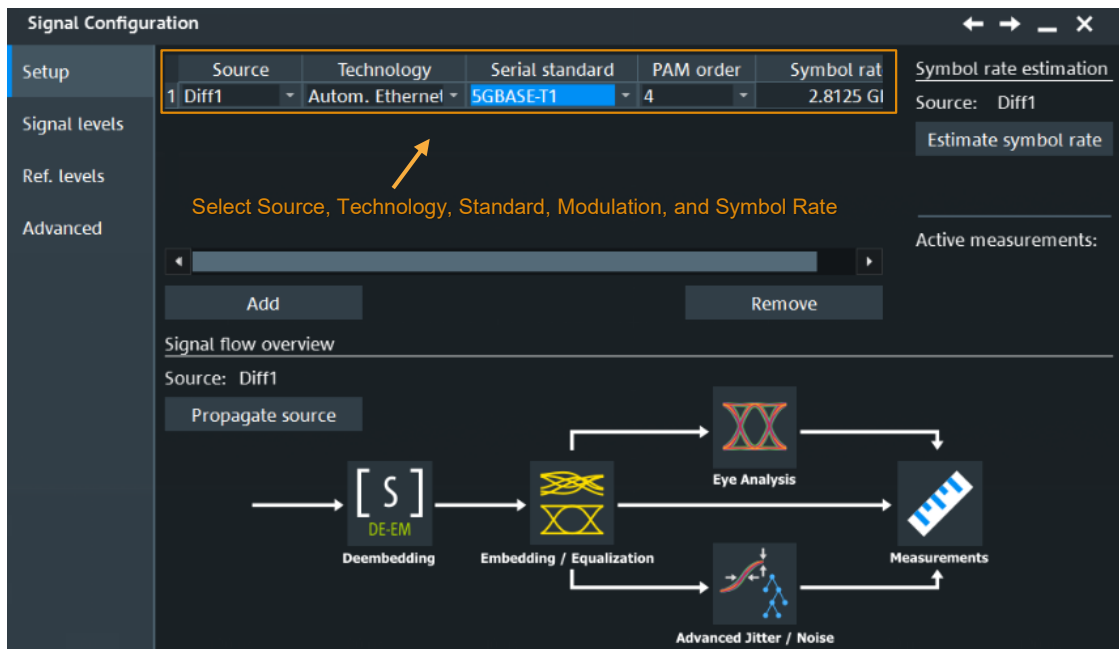
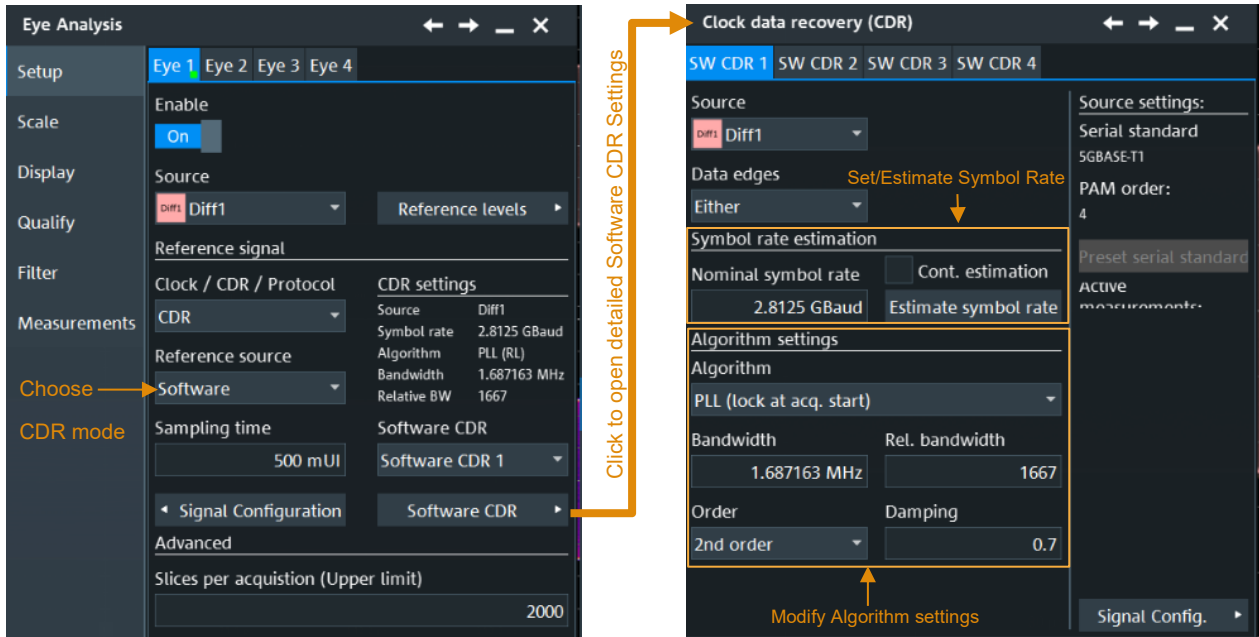


Figure 6: Signal Configuration in the R&S®RTP oscilloscope

4.2.1 PAM Eye Diagram and Eye Isolation

To enable accurate multilevel eye diagram capture, the R&S®RTP oscilloscope with the PAM-N analysis option offers configurable Clock Data Recovery (CDR) and eye diagram setup. Figure 7 shows the Eye Analysis and CDR configuration interfaces. In the *Eye Analysis* setup (Figure 7a), users define the CDR source and choose between hardware or software modes. Clicking the Software CDR button opens the detailed Software CDR settings panel (Figure 7b), where users can set or estimate the symbol rate and modify the algorithm settings, including bandwidth, damping, and order. These parameters directly affect the stability and lock behavior of the recovered clock, which is critical for PAM signals.



(a) Eye Analysis Setup

(b) Software CDR Settings

Figure 7: Eye Analysis and CDR Configuration on the R&S®RTP Oscilloscope

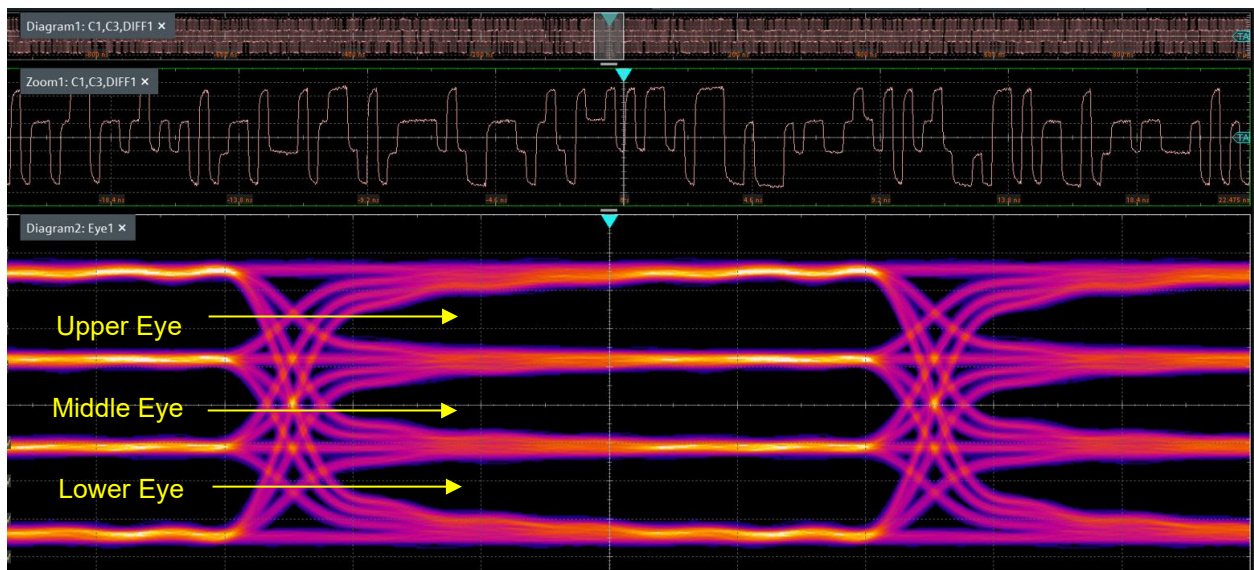


Figure 8: PAM-4 Eye Diagram

Figure 8 illustrates a PAM-4 eye diagram measured using the R&S®RTP oscilloscope with symbol-aware acquisition. This acquisition method identifies individual PAM levels and symbol transitions in real time rather than relying solely on edge timing, which improves the accuracy of multilevel eye diagram capture.

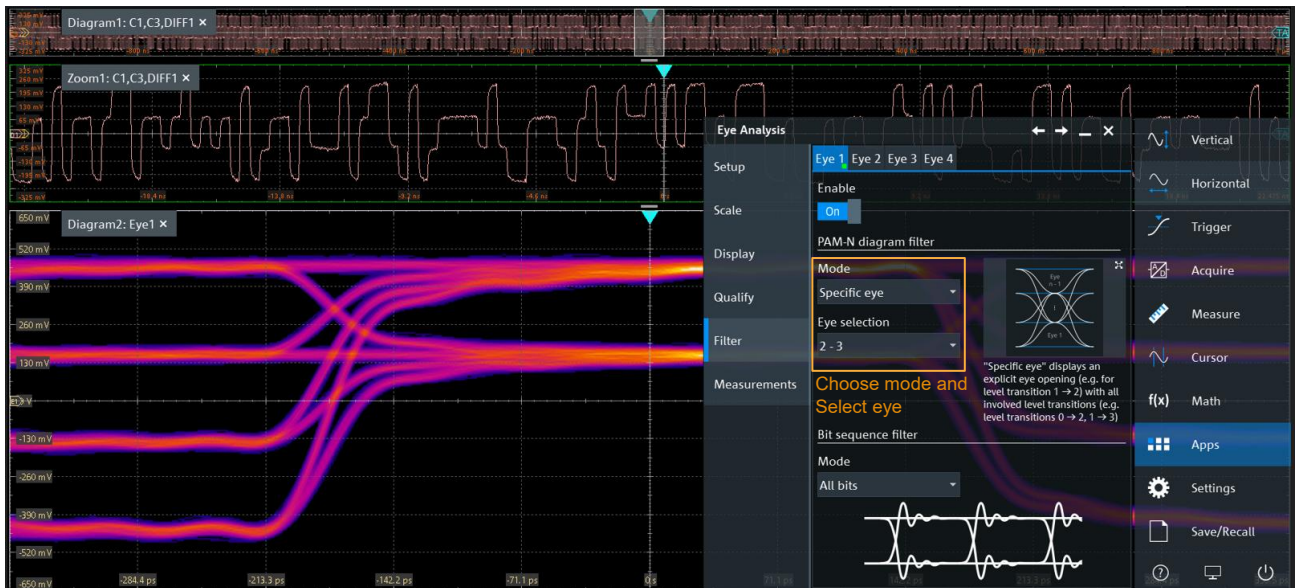


Figure 9: Eye Isolation in PAM-4

Furthermore, PAM transitions are not symmetric. For example, a 0→1 transition may behave very differently from a 1→2 or a 2→3 transition. R&S oscilloscopes allow users to selectively isolate individual eyes in a PAM-4 signal (see Figure 9). In the Eye Analysis menu, navigate to Filter and select Specific eye mode in the Mode field, then choose the desired eye in Eye selection. Each eye corresponds to a distinct set of symbol transitions. For example, the lower eye represents transitions from 0 to 1, the middle eye 1 to 2, and the upper eye 2 to 3. This selection enables targeted analysis of level-specific impairments. By isolating a single eye, users can examine issues such as asymmetric inter-symbol interference, transition-dependent jitter, or unequal rise and fall times without interference from adjacent levels. This level-aware isolation helps identify root causes that may be masked in a composite eye view.

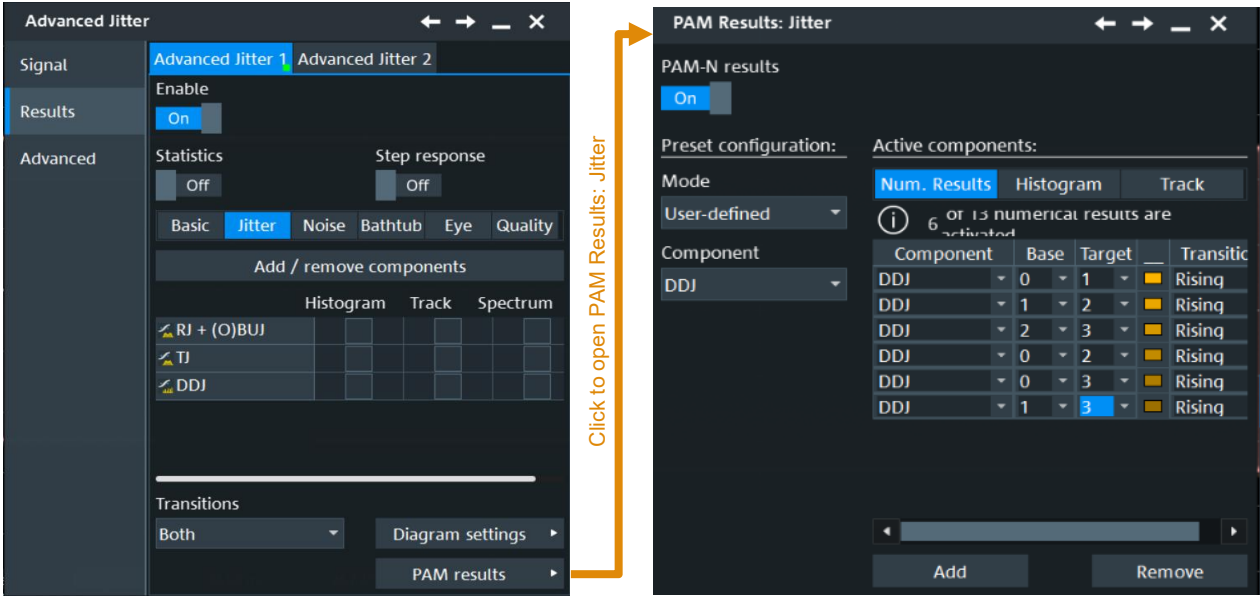
4.2.2 Transition-Level Jitter

Because PAM transitions differ in amplitude and slope, jitter also varies by transition. A global jitter metric, as used in NRZ, may obscure these transition-specific variations in timing uncertainty. Rohde & Schwarz oscilloscopes (R&S® RTP® or R&S® RTO6 with K135 software option) address this with per-transition jitter decomposition, breaking total jitter into deterministic (DJ), data-dependent (DDJ), and periodic jitter (PJ) components for each transition class.

Figure 10a shows the *Advanced Jitter* menu, where users can go to the *Results* tab to perform the usual jitter decomposition. For PAM signals, they can go further by selecting PAM results to access the transition-dependent jitter view shown in Figure 10b. In this mode, the oscilloscope is configured to measure DDJ for selected rising-edge transitions in a PAM-4 signal. The base and target symbol levels are explicitly chosen (e.g., 0→1, 1→2, 2→3) so the oscilloscope quantifies jitter for each transition type individually. This avoids averaging effects across transitions and enables targeted analysis of how specific symbol changes contribute to timing uncertainty.

The results of such a configuration are presented in Figure 11. Each row in the table corresponds to a specific transition, with the standard deviation (σ) and peak-to-peak (abs/pp) jitter values shown in milli UI (mUI). For the 5GBASE-T1 PAM-4 signal, the transition-specific DDJ values reveal a clear pattern: the adjacent-level transitions 0→1, 1→2, and 2→3 exhibit the highest jitter, with peak-to-peak values approaching roughly 20 mUI. In contrast, transitions involving larger amplitude steps such as 0→2, 1→3 or the full-swing 0→3 transition show narrower histograms and lower DDJ due to greater level separation. This behaviour reflects the expected PAM-4 characteristic that transitions with smaller vertical spacing are more

sensitive to noise and pattern-dependent effects, and it highlights which transitions contribute most to horizontal eye reduction and should be prioritized for transmitter optimization.



(a) Advanced Jitter Setup

(b) Transition-Dependent Jitter Configuration

Figure 10: PAM Jitter Analysis Options



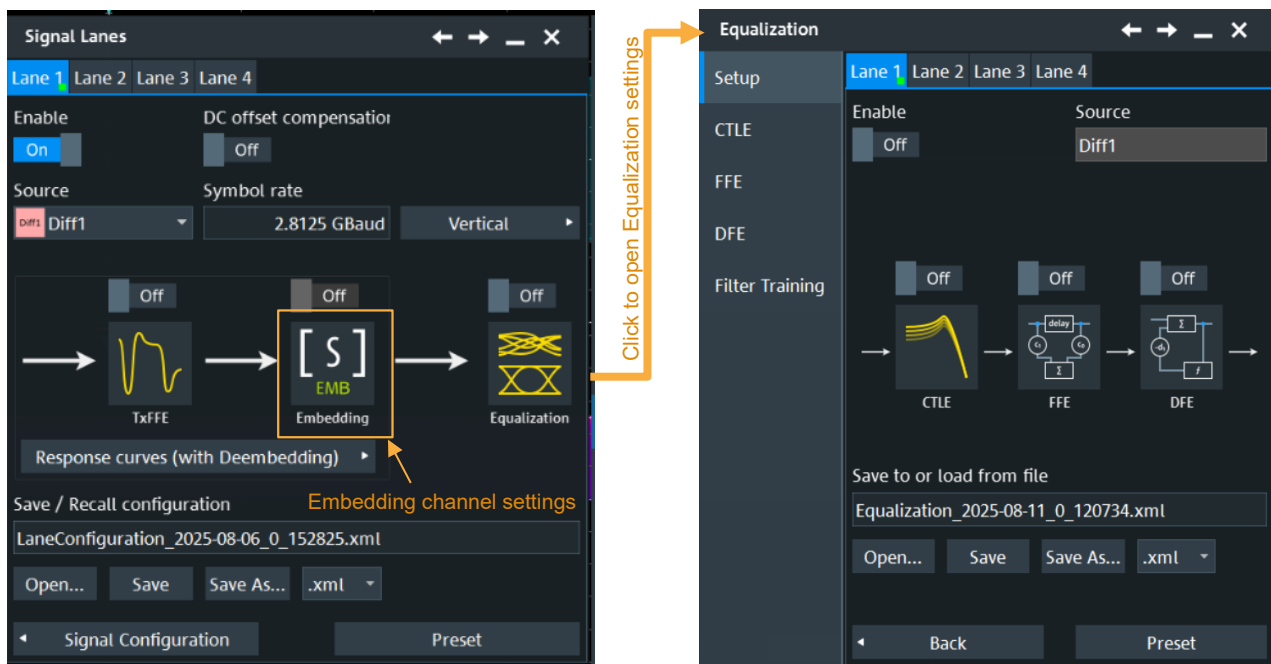
Figure 11: Transition-Level Jitter Decomposition – PAM4

By combining transition-aware jitter classification with selective eye isolation (section 4.2.1), the measurement system enables targeted analysis of signal degradation sources, particularly in environments with asymmetric channel conditions or level-specific impairments.

4.2.3 Equalization Techniques for PAM-N

As outlined in Sections 2.3 and 4.1, PAM-N signaling tends to exhibit reduced eye margins and increased sensitivity to ISI. Techniques such as equalization are used either on the receiver or the transmitter to mitigate these impairments and restore timing and margin requirements.

In compliance testing, particularly for transmitter validation, the oscilloscope acts as a golden receiver, emulating how a real-world receiver would interpret a degraded signal after transmission through a worst-case channel. To achieve this, the measurement setup applies a standard-defined equalizer following a representative channel model, creating a stressed link environment. This allows verification of whether the transmitted signal remains recoverable and meets compliance requirements under post-equalization conditions.



(a) Signal Lane Configuration

(b) Equalization Setup Window

Figure 12: Embedding and Equalization Workflow

Figure 12a shows the *Signal Lanes* dialog on the R&S®RTP, used to embed a worst-case channel model into the acquired signal. Clicking the *Embedding* block opens the channel embedding settings. From the same view, clicking *Equalization* opens the panel shown in Figure 12b, where users can configure CTLE, FFE, and DFE stages. Embedding a standard-defined worst-case channel and applying the chosen equalizer emulates a stressed receiver environment, enabling realistic margin and compliance testing

For high-speed links, there are several common equalizers employed. Rohde & Schwarz oscilloscopes (R&S®RTP® or R&S® RTO6 with K135 software option) support such types of equalization:

- **CTLE: Continuous-Time Linear Equalization**

CTLE is used to compensate for high-frequency loss by amplifying attenuated signal components. In PAM-N systems, where vertical eye openings are narrower and transition edges are shallower, finely controlled CTLE improves signal rise and fall times, reopening the eye vertically and reducing baseline wander.

- **FFE: Feed-Forward Equalization**

FFE is typically applied to cancel ISI through a series of weighted taps that act on past and future symbols. In PAM signaling, FFE works by pre-distorting the transmitted signal to counteract the channel's effects, effectively opening the signal "eyes" (visualization of signal levels) and improving the receiver's ability to accurately decode the transmitted data.

- **DFE: Decision Feedback Equalization**

DFE operates by subtracting ISI contributions from previous symbol decisions. While DFE is commonly implemented at PAM-4 receivers due to its strong capability to reduce ISI, it is less frequently employed in transmitter-focused compliance testing scenarios, as its nonlinear behavior complicates symbol-level eye diagram measurements and can lead to error propagation in test conditions.

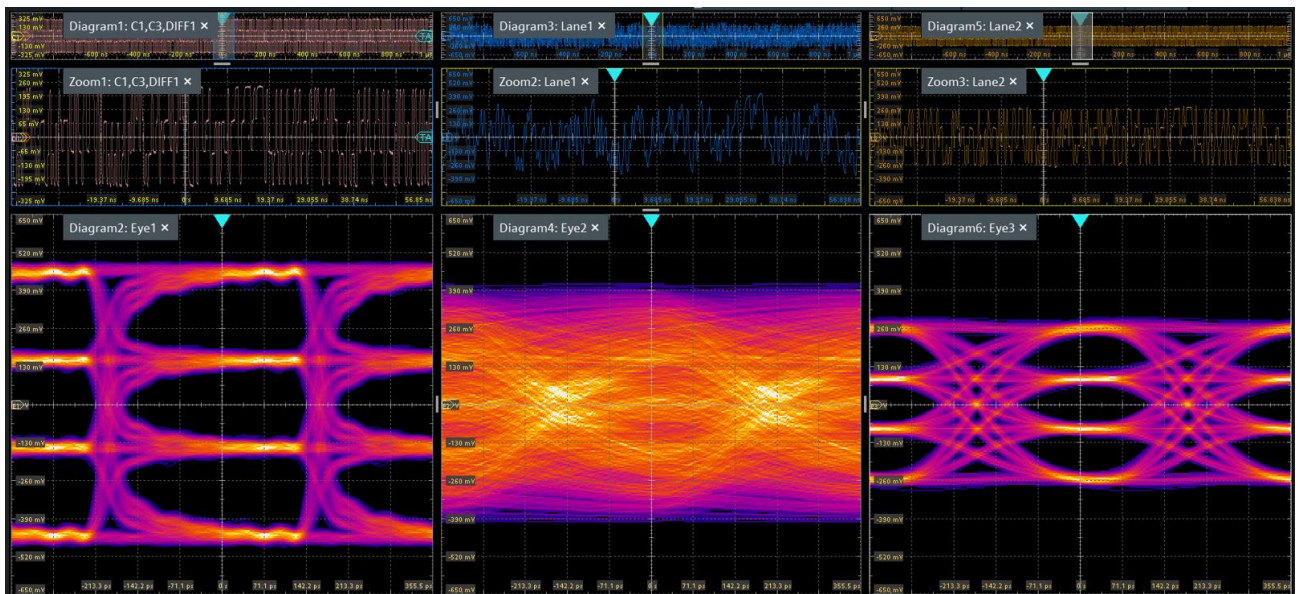


Figure 13: Equalization on PAM-4 signal

Figure 13 illustrates the effect of channel degradation and equalization on a PAM-4 signal using internal software processing. The left image shows the clean, original waveform without any added channel model, representing the transmitter output with clear, well-separated eyes. The middle image applies a worst-case channel model, introducing insertion loss, inter-symbol interference, and visible vertical and horizontal eye closure, resulting in significantly reduced eye height and narrowed eye width. These impairments emulate a stressed link condition in which the receiver would struggle to reliably distinguish between PAM-4 levels.

In the right image, equalization techniques such as CTLE or FFE are applied algorithmically within the oscilloscope software. While the equalized result does not fully restore the ideal eye height or width, the eye openings partially reappear, vertical level separation increases, transition slopes become steeper, and the horizontal openings widen relative to the stressed case. This demonstrates the practical role of equalization. It does not recreate the original waveform but recovers sufficient vertical and horizontal margin for reliable symbol decisions under degraded channel conditions. Quantitative eye height values are not reported for the equalized case, as the measured amplitude depends directly on the CTLE gain setting applied. Adjusting the gain rescales the eye vertically without reflecting a change in the underlying signal quality.

These features allow users to evaluate PAM-N performance in realistic receiver conditions without requiring external pattern generators or physical stress fixtures. Equalization settings can be applied manually or adaptively within the measurement software, giving visibility into each eye opening and per-transition behavior. The system supports real-time visualization of equalized eye diagrams for all symbol levels, enabling users to assess the impact of each equalizer stage. With fine-grained control over tap weights,

delay settings, and equalizer order, R&S instruments offer flexibility in emulating receiver behavior and optimizing signal recovery in bandwidth-constrained PAM-N systems.

5 Summary and Outlook

This educational note demonstrates how Pulse Amplitude Modulation has become essential to next-generation high-speed systems by overcoming traditional limitations of NRZ.

- Chapter 1 introduced the shift from NRZ to PAM formats, driven by increasing data rates and physical bandwidth constraints.
- Chapter 2 introduced the basics of PAM-N signaling, including symbol-to-level mapping and how multi-level amplitude encoding differs from traditional NRZ. It also outlined how these differences affect signal interpretation and waveform complexity, laying the groundwork for later discussions on measurement challenges.
- Chapter 3 detailed practical implementation nuances, emphasizing critical choices in coding schemes, scrambling, forward error correction (FEC), and multi-lane configurations.
- Chapter 4 illustrated advanced measurement methods required to tackle PAM's unique signal integrity challenges, including multi-level eye analysis, transition-specific jitter decomposition, and the critical role of adaptive equalization techniques. These techniques, available in R&S oscilloscopes and analysis software, provide powerful capabilities for both precise performance measurement and effective root-cause analysis in PAM systems.

As PAM-4 and higher-order PAM-N formats gain prominence in Ethernet, PCIe, automotive, and optical interfaces, engineers encounter tighter eye openings, increased jitter sensitivity, and more stringent compliance criteria. Upcoming standards such as Ethernet beyond 800G and PCI Express 7.0 will intensify these challenges further. Consequently, mastering both fundamental and advanced PAM measurement techniques will become indispensable.

Rohde & Schwarz directly addresses these evolving measurement needs with the R&S®RTP-K135 PAM-N Analysis Software for the R&S RTP and RTO6 oscilloscopes, providing engineers with specialized tools essential for reliable PAM-N analysis:

- Comprehensive support for PAM orders from 2 to 8, with automatic level detection ensuring rapid and accurate measurements.
- Advanced eye diagram analysis, including selectable “whole eye,” “specific eye,” and “selected eye” modes, complemented by built-in compliance masks tailored for PAM signals.
- Sophisticated jitter and noise decomposition at each PAM transition level, enabled by advanced PAM-aware software clock data recovery (CDR).
- Real-time hardware-based CDR, allowing “live eye” captures at symbol rates up to 3.8 GBaud, vital for real-world testing conditions.

6 Literature

- [1] PCI-SIG, "PCI Express Base Specification Revision 6.0," PCI-SIG, 2022.
- [2] USB Implementers Forum, "Universal Serial Bus 3.2 Specification," USB-IF, 2017.
- [3] IEEE, "IEEE Standard for 200 Gb/s and 400 Gb/s Ethernet," in *IEEE 802.3bs-2017*, IEEE, 2017.
- [4] IEEE, "IEEE Standard for 50 Gb/s, 100 Gb/s, and 200 Gb/s Ethernet," IEEE, 2018.
- [5] IEEE, "Physical Layer Specifications and Management Parameters for 100 Gb/s, 200 Gb/s, and 400 Gb/s Operation over Backplanes and Copper Cables," IEEE, 2022.
- [6] OIF, "CEI-112G Common Electrical I/O Implementation Agreement," Optical Internetworking Forum, 2021.
- [7] IEEE, IEEE Standard for 1000BASE-T: Gigabit Ethernet over Copper, IEEE, 1999.
- [8] IEEE, IEEE Standard for Ethernet Amendment 4: Physical Layer Specifications and Management Parameters for 1 Gb/s Operation over a Single Twisted Pair Cable, IEEE, 2016.
- [9] U. I. Forum, "USB4 Specification Version 2.0," USB-IF, 2022.
- [10] IEEE , "IEEE Standard for Ethernet, IEEE 802.3-2022, Clause 49.2.6," 2022.
- [11] IEEE, "IEEE Standard for Ethernet Amendment 1: Physical Layer Specifications and Management Parameters for 100 Mb/s Operation over a Single Balanced Twisted Pair Cable (100BASE-T1), IEEE 802.3bw-2015," IEEE, 2015.
- [12] IEEE, "IEEE Standard for Local and Metropolitan Area Networks: Supplement – Media Access Control Parameters, Physical Layer, and Management Parameters for 100 Mb/s Operation, Type 100BASE-T2, IEEE 802.3y-1997," IEEE, 1997.
- [13] IEEE, "IEEE Standard for Information Technology – Physical Layer and Management Parameters for 10 Gb/s Operation, Type 10GBASE-T, IEEE 802.3an-2006," IEEE, 2006.

7 Appendix

A Appendix 1: Comprehensive Overview of PAM-Based Standards

(Bandwidth values are indicative (Nyquist-based), not the formal mask limits from the respective standards.)

Standard	PAM Order	Data Rate	Symbol Rate	Channel Bandwidth	# Pairs / Lanes	Application
100BASE-T2	PAM-5	100 Mbps	25 MBd	~12.5 MHz	2 twisted pairs	Fast Ethernet over Cat 3
100BASE-T1	PAM-3	100 Mbps	66.7 MBd	~33 MHz	1 twisted pair	Automotive single-pair Ethernet
1000BASE-T	PAM-5	1 Gbps	125 MBd per pair	~62.5 MHz per pair	4 twisted pairs	LAN over Cat-5e
1000BASE-T1	PAM-3	1 Gbps	750 MBd	~375 MHz	1 twisted pair	Automotive single-pair Ethernet
10GBASE-T	PAM-16	10 Gbps	800 MBd	~400 MHz	4 twisted pairs	LAN over Cat-6A
PCIe 6.0	PAM-4	64 GT/s	32 Gbd	~16 GHz	1 lane	High-speed computer interconnect
USB4 v2 Gen4	PAM-3	80 Gbps (dual-lane)	~25.6 GBd per lane	~12.8 GHz	2 lanes	USB4 over Type-C
100GBASE-DR	PAM-4	100 Gbps	53.125 GBd per lane	~25 GHz (optical)	1 optical lane	Data center serial link
OIF CEI-56G	PAM-4	56 Gbps	~28 GBd	~14 GHz	1 SerDes lane	Chip-to-chip SerDes
OIF CEI-112G	PAM-4	112 Gbps	~56 GBd	~28 GHz	1 SerDes lane	High-speed SerDes

Rohde & Schwarz

The Rohde & Schwarz electronics group offers innovative solutions in the following business fields: test and measurement, broadcast and media, secure communications, cybersecurity, monitoring and network testing. Founded more than 80 years ago, the independent company which is headquartered in Munich, Germany, has an extensive sales and service network with locations in more than 70 countries.

www.rohde-schwarz.com

Certified Quality Management

ISO 9001

Rohde & Schwarz training

www.rohde-schwarz.com/training



Rohde & Schwarz customer support

www.rohde-schwarz.com/support

