White Paper

PAM4 Demands Accurate S-parameters

1. Introduction

As high speed serial data rates advance from 15 Gb/s to beyond 50 Gb/s, we face a fundamental shift. Conventional logic-emulating NRZ (non-return to-zero) signaling is being replaced by PAM4, a 4-level pulse amplitude modulation scheme that takes half the bandwidth to transmit the same payload as the equivalent NRZ signal.

In this paper we'll see how PAM4 challenges signal integrity, test, and design engineers who are responsible for SERDES (serializer/deserializer) components, interconnects, backplanes, cables, connectors, circuits, and complete systems. The problems solved by PAM4 outweigh the problems it introduces, but PAM4's increased complexity means that we must address a host of new issues.

The bottom line is that PAM4 requires more accurate measurements. Designs have to be verified before they're produced. Simulations and measurements have to agree on a point-by-point basis. Since SNR (signal to noise ratio) is about 10 dB worse for PAM4 than NRZ, qualitative agreement is no longer good enough. Simulation accuracy and channel evaluation begin and end with accurate measurements of S-parameters.

This paper will help you understand how to use S-parameters (scattering parameters), how to think of channel response and crosstalk in terms of S-parameters, and where to look for problems that might be caused by inaccurate measurements of S-parameters. We'll introduce the ins and outs of PAM4 in the next section. In Section 3 we'll discuss S-parameters in terms of inter-symbol interference and crosstalk. In Section 4 we give a short introduction to IBIS-AMI simulation, the keys to creating models that you can trust, and the use of S-parameters to evaluate equalization schemes, de-embed text fixtures from signal integrity analyses, and measure channel operating margin. We conclude in Section 5 with guidance on how to achieve VNA-quality S-parameter measurements on a tight a budget.

2. NRZ and PAM4 Bandwidth Demands at High Data Rates

Conventional, logic-emulating NRZ signaling has hit a bandwidth brick wall caused by high insertion loss, 35 dB and higher on a typical trace at next generation data rates.

The limits of NRZ signaling were first approached in 100 GbE (Gigabit Ethernet) with its 25-28 Gb/s signals on 4 channels that combine to 100 Gb/s data payloads. Since each bit period covers half a cycle, the fundamental frequency of a 26.6 Gb/s NRZ signal is 13.3 GHz. The insertion loss of conventional copper traces on standard FR-4 PCB (flame retardant type-4 printed circuit board) typically falls at about 10 dB per decade of frequency, depending on its length. Adding the effects of connectors, vias, BGAs (ball grid arrays) and all the other impedance mismatches in a real circuit, and the result is a frequency response with resonant peaks and absorption valleys as well, Figure 1.

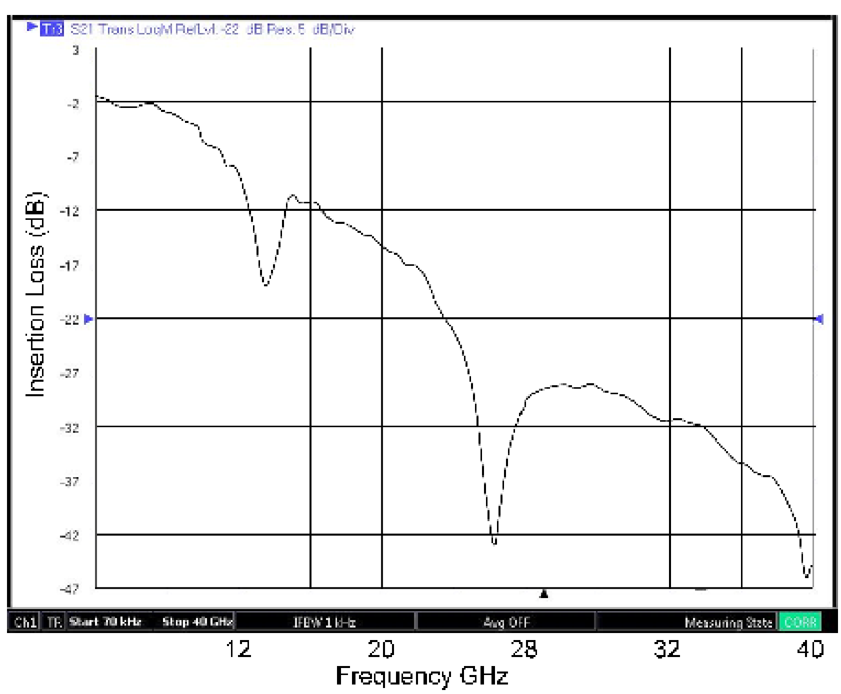

Figure 1: Frequency response of a typical 100 GbE channel, $S_{dd21}(f)$ or IL(f).

With 20 dB of loss at the 13.3 GHz fundamental harmonic, the channel depicted in Figure 1 can function with NRZ signaling at BER (Bit Error Rate) < 1E-12, the required 100 GbE bit error rate, but only if it has optimally tuned equalization at both the transmitter and receiver.

At 53 Gb/s, the channel rate required by 400 GbE, Figure 2, the fundamental frequency is 26.5 GHz. At 26.5 GHz this channel has over 40 dB of loss. In the presence of crosstalk, no amount of equalization wizardry can pull NRZ signals out of that noise.

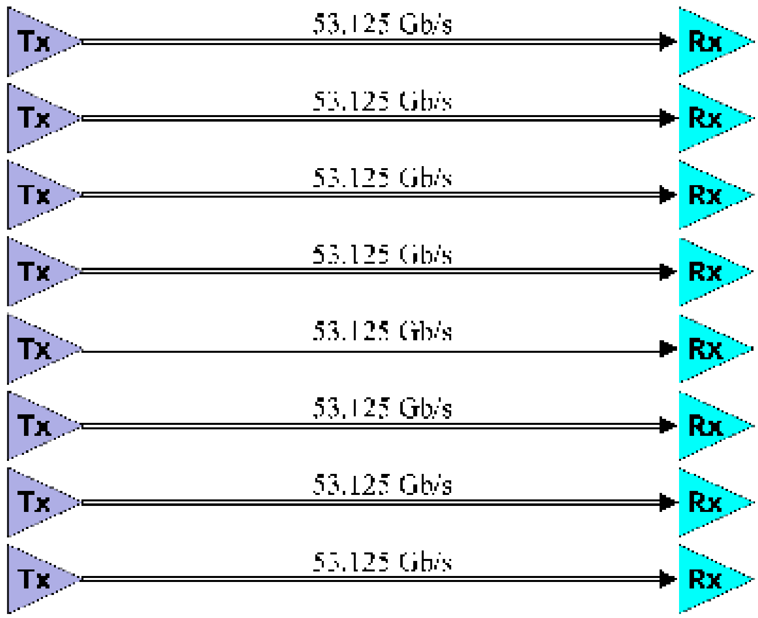

Figure 2: Diagram of the 8×53 Gb/s version of 400 GbE.

The solution is to pack twice as many bits into each transmitted symbol. Where NRZ encodes one bit, either 0 or 1, in each symbol, PAM4 uses four separate levels to encode two bits—00, 01, 10, 11—in each symbol, Figure 3. By encoding two bits per symbol, a 53 Gb/s signal can be transmitted at 26.5 Gbaud with a fundamental frequency of 13.25 GHz.

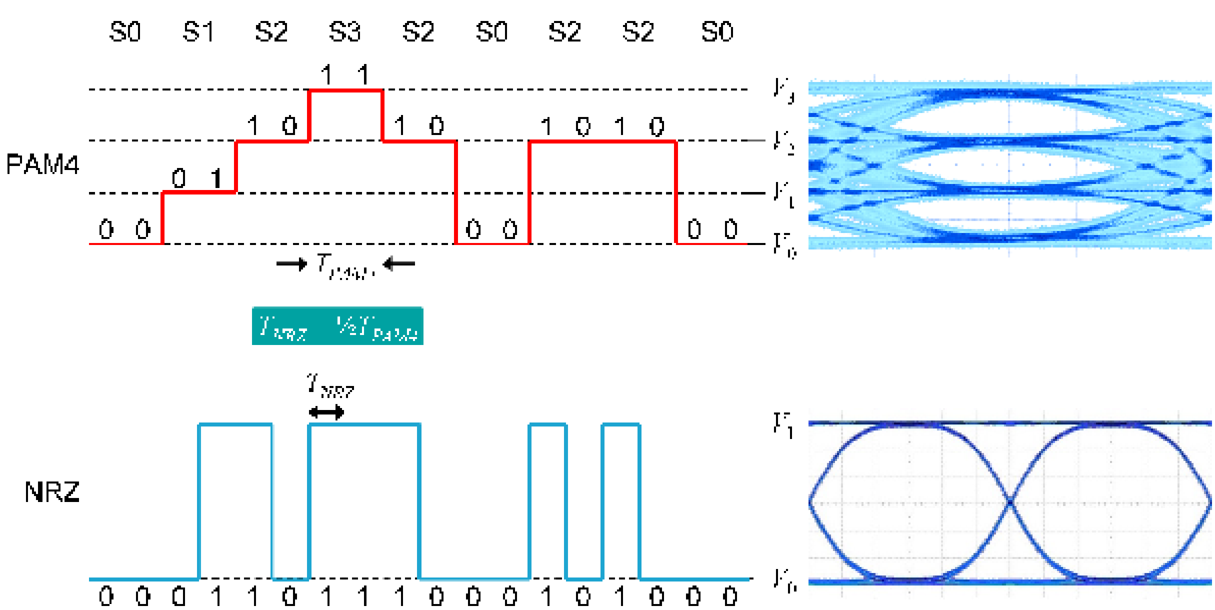

Figure 3: NRZ compared to PAM4.

2.1 PAM4 Complexity

PAM4 gets us back to loss levels that we can handle, but it brings an incremental leap in signal complexity: 12 distinct symbol transitions where NRZ has 3 eye diagrams per bit period where NRZ has 1, and SNR (signal to noise) degraded by at least 9 dB.

PAM4 also introduces potential nonlinearities that we've never faced. Variations in the eye heights of the three eye openings is called "eye compression." Similarly, "timing skew" occurs when the centers of the three eyes are misaligned. Figure 4 shows an eye diagram that suffers both eye compression and timing skew.

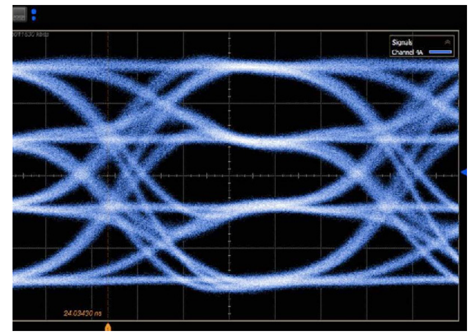

Figure 4: Potential PAM4 nonlinearities: (a) eye compression and (b) timing skew.

And we still have to contend with jitter, noise, and crosstalk; we still need super-clean clocks, robust clock recovery circuits, and sensitive symbol decoders/voltage slicers—now three slicers to decode the four possible symbols; and we need even better equalization schemes at both the transmitter and receiver to deal with ISI (inter-symbol interference) caused by the channel.

2.2 PAM4 ISI and Equalization

The frequency content of both NRZ and PAM4 signals consist of square wave harmonics and sub-harmonics from sequences of identical symbols. The frequency components have precisely coordinated amplitudes and phases. Since PAM4 has 4 levels, its amplitude-phase-frequency relationships are more complicated than those of NRZ.

Take another look at Figure 1. The channel affects the amplitudes and phases of each frequency component differently. By de-correlating its frequency and phase structure, the channel causes ISI.

Equalization removes ISI by inverting the channel response. Like everything else, equalization in PAM4 systems is more complicated than in NRZ systems.

Transmitter equalization, called emphasis or FFE (feed-forward equalization) pre-distorts the signal to compensate for the gross features of the channel response. Two types of equalization operate at the receiver: CTLE (continuous time linear equalization) and DFE (decision feedback equalization). CTLE attenuates low frequency components and amplifies high frequency components to compensate for the channel's low pass nature, Figure 5. The CTLE typically precedes the clock recovery circuit in the receiver as shown in Figure 6.

DFE, a type of IIR (infinite impulse response) filter, feeds symbol decisions back into the decision circuit. Since PAM4 has 4 symbols, DFE becomes more complicated but can also provide greater flexibility and the possibility for innovations that could improve PAM4 DFE compared to its NRZ counterpart.

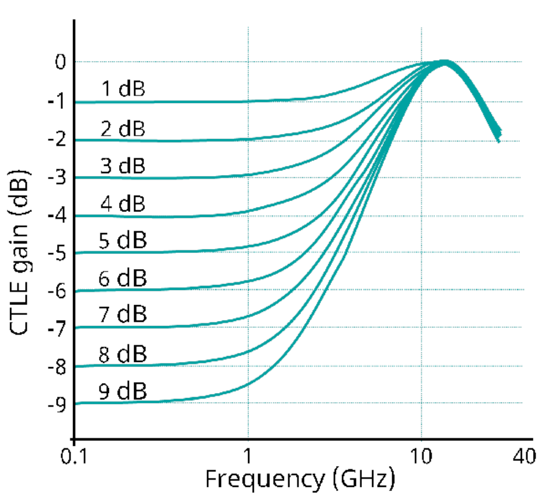

Figure 5: CTLE frequency response for different levels of gain.

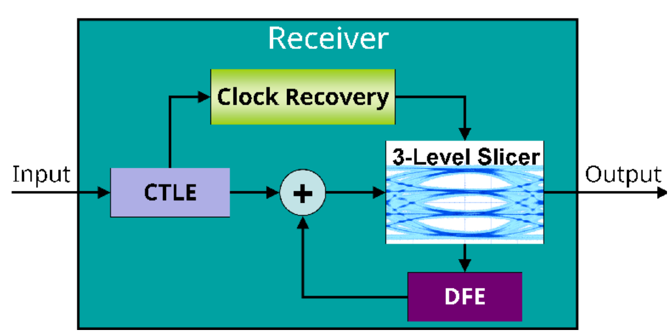

Figure 6: A PAM4 receiver.

3. The importance of accurate S-parameters in development and manufacture of PAM4 components and systems

Our goals include optimizing and verifying designs, reducing time to market, manufacturing high quality components, and assembling systems that perform at least as well as the required BER.

While equalization can compensate for the ISI caused by channel frequency and phase response, it causes trouble in the presence of crosstalk. Crosstalk is predominantly high frequency interference. To compensate for the low pass channel response, transmitter emphasis and CTLE amplify high frequency signal components, but in so doing they also amplify crosstalk noise. Since DFE operates at the symbol level, it doesn't affect crosstalk either way.

The parameters that go into ISI, crosstalk, and equalization present us with a multi-variable optimization problem. We can think of it as a multi-dimensional design space where we search for that point where BER is minimized.

S-parameters provide the tools we need to predict the combined effects of ISI, equalization, and crosstalk.

3.1 ISI and Impulse Response

Lets start with ISI. Figure 7 shows how a sharp impulse is affected by a channel. A perfect, infinite amplitude, infinitesimal width—Dirac delta function—impulse propagates through the channel. Different frequency components experience different loss and phase variation. The output is the "impulse response."

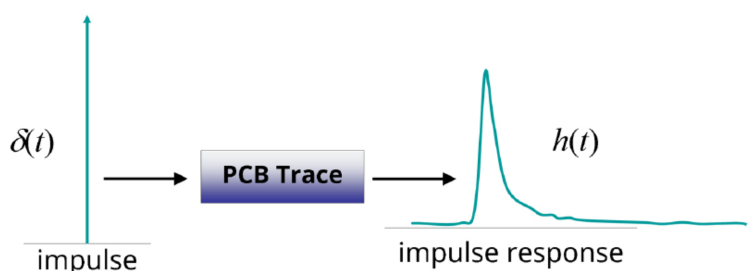

Figure 7: Impulse response.

Every characteristic of the channel is encoded in the impulse response: every impedance variation, all the reflections and losses however tiny. The impulse response is the time-domain equivalent of the frequency response. The output of any channel can be calculated by convolving the impulse response with the input:

$$Rx(t) = \int Tx(t')h(t-t')dt'$$

where *h(t)* is the impulse response, *Tx(t)* is the transmitted signal, and Rx(t) is the received signal. Think of the impulse response as a smearing function: *h(t)* smears every symbol in the transmitted waveform, causing them to overlap, and by that, to interfere.

Of course we can't produce an impulse signal that has infinite bandwidth and infinitesimal width. An alternative is to use a step function rather than an impulse since the derivative of an ideal, zero rise-time step is a Dirac delta function. By transmitting a very small rise-time step and measuring the transmitted and reflected responses, we can approximate the impulse response. This technique is used in TDT/TDR (time domain transmission/time domain reflectometry) measurements. Since TDT/TDR are broadband measurements and the noise of an instrument is proportional to the instrument bandwidth, they tend to have high noise floors and low dynamic range. It can be an effective analysis tool for low rate, low loss NRZ systems but it's too noisy and lacks the dynamic range necessary for high data rate, high loss PAM4 analysis.

There's a better approach. What makes an ideal impulse response is that it has a wide, flat frequency spectrum from DC to infinite frequency, Figure 8.

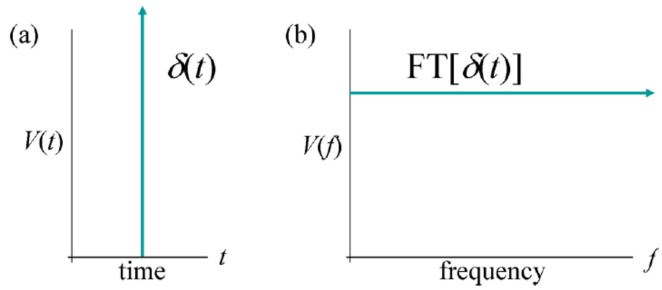

Figure 8: Perfect impulse response, Dirac delta function, in (a) the time domain, and (b) the frequency domain.

It's impossible to create a perfect impulse or step, but we can create a signal with almost the same spectrum: It's fairly easy to generate sinusoidal waveforms from very low to very high frequencies. Instead of measuring the impulse response in the time domain, we measure it one frequency at a time in the frequency domain.

VNAs characterize channels with unprecedented accuracy by transmitting a series of sine waves and measuring the resulting signals and their reflections. VNAs build up the frequency-domain equivalent of the impulse response one frequency at a time. Since their measurements are performed in sequential steps, each in a narrow intermediate frequency bandwidth, (as low as 1 Hz on VectorStar), VNAs have extremely low noise floors, (typically –100 dBm), high bandwidths, (up to 145 GHz in a single coaxial connection), and broad dynamic range, (over 100 dB).

In the frequency domain, a convolution reduces to a simple product:

$$Rx(f) = Tx(f)H(f)$$

where *Tx(f)* and *Rx(f)* are the frequency-domain expressions of the transmitted and received signals respectively and *H(f)* is the transfer function, the frequency domain equivalent of the impulse response.

Once we've measured an accurate transfer function, *H(f)*, we can transform it to the time domain and get the impulse response, *h(t)*.

3.2 Introduction to S-Parameters

S-parameters describe how a signal is "scattered" by an object. For us, the object is a transmission channel. S-parameters are functions of both amplitude and phase over frequency.

A two port network, Figure 9, is characterized by a 2×2 matrix of S parameters. The channel frequency response in Figure 1, is also called the insertion loss as a function of frequency, *IL(f)*, and is given by the magnitude of an S-parameter, |S21|. The return loss as a function of frequency, *RL(f)*, is given by |S11|.

In addition to being straightforward to measure accurately, S-parameters account for all the partially reflected and transmitted signals that bounce around a network, including crosstalk. Since they're matrices, they fit perfectly in models.

2 Port S-parameters

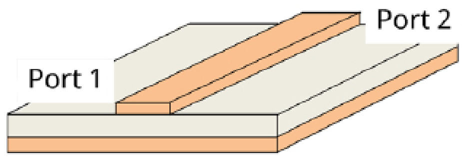

		Stimulus	
Response	Port #	1	2
	1	**S_{11} = Port 1 RL**	**S_{12} = IL**
	2	**S_{21} = IL**	**S_{22} = Port 2 RL**

$$\begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix}\begin{pmatrix} V_1^{in} \\ V_2^{in} \end{pmatrix} = \begin{pmatrix} V_1^{out} \\ V_2^{out} \end{pmatrix}$$

$$S_{11}V_1^{in} + S_{12}V_2^{in} = V_1^{out}$$
$$S_{21}V_1^{in} + S_{22}V_2^{in} = V_2^{out}$$

Figure 9: A two port network and its S-parameters.

3.3 Crosstalk S-Parameters

Figure 10 shows two channels, a “4 port network” in S-parameter jargon: two input ports plus two output ports. Consider S23, the signal emerging from the right side of the top channel, port 2, due to a signal transmitted into port 3, the left side of the bottom channel. Since port 3 is on the far side of the network from port 2, S23 gives the FEXT (far end crosstalk) phase and frequency response. Similarly, S24 parameterizes the NEXT (near end crosstalk) response of the output on port 2 due to a signal transmitted into port 4.

4 Port S-parameters

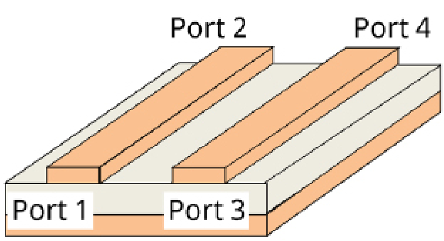

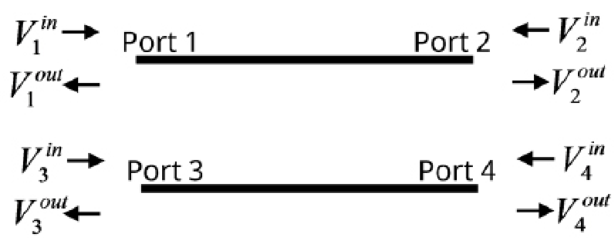

$$\begin{pmatrix} S_{11} & S_{12} & S_{13} & S_{14} \\ S_{21} & S_{22} & S_{23} & S_{24} \\ S_{31} & S_{32} & S_{33} & S_{34} \\ S_{41} & S_{42} & S_{43} & S_{44} \end{pmatrix}\begin{pmatrix} V_1^{in} \\ V_2^{in} \\ V_3^{in} \\ V_4^{in} \end{pmatrix} = \begin{pmatrix} V_1^{out} \\ V_2^{out} \\ V_3^{out} \\ V_4^{out} \end{pmatrix}$$

$$S_{11}V_1^{in} + S_{12}V_2^{in} + S_{13}V_3^{in} + S_{14}V_4^{in} = V_1^{out}$$
$$S_{21}V_1^{in} + S_{22}V_2^{in} + S_{23}V_3^{in} + S_{24}V_4^{in} = V_2^{out}$$
$$S_{31}V_1^{in} + S_{32}V_2^{in} + S_{33}V_3^{in} + S_{34}V_4^{in} = V_3^{out}$$
$$S_{41}V_1^{in} + S_{42}V_2^{in} + S_{43}V_3^{in} + S_{44}V_4^{in} = V_4^{out}$$

Figure 10: A four port network and its S-parameters.

Notice that NEXT and FEXT calculated from S-parameters include all multiple reflections, not just the forward wave contribution to FEXT and the reverse wave contribution to NEXT, but the net signal that emerges from port 2 due to signals transmitted in port 3 for FEXT and port 4 for NEXT.

3.4 *Differential S-Parameters*

Differential channels are made from pairs of conductors with minimal skew. Differential pairs reduce the crosstalk both radiated by and picked up by a channel. We derive the S-parameters for a differential channel from the two channel system by transforming the simple 4-port matrix to a combination of differential and common mode S-parameters. Differential mode signals are those that transmit a signal on one trace and its opposite on the other. Common mode signals have identical signals on both traces. Any signal can be composed of a combination of differential and common mode signals. The transformation is equivalent to a coordinate rotation in geometry: no information is gained or lost, it’s just rotated. The resulting S matrix is shown in Figure 11.

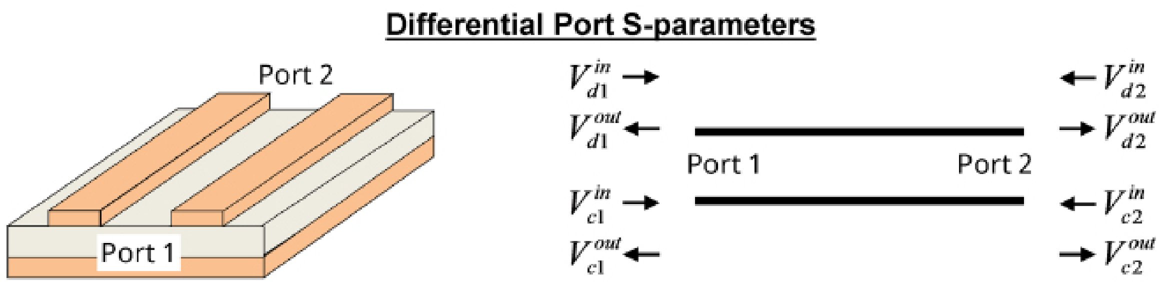

$$\begin{pmatrix} S_{dd11} & S_{dd12} & S_{dc11} & S_{dc12} \\ S_{dd21} & S_{dd22} & S_{dc21} & S_{dc22} \\ S_{cd11} & S_{cd12} & S_{cc11} & S_{cc12} \\ S_{cd21} & S_{cd22} & S_{cc21} & S_{cc22} \end{pmatrix} \begin{pmatrix} V_{d1}^{in} \\ V_{d2}^{in} \\ V_{c1}^{in} \\ V_{c2}^{in} \end{pmatrix} = \begin{pmatrix} V_{d1}^{out} \\ V_{d2}^{out} \\ V_{c1}^{out} \\ V_{c2}^{out} \end{pmatrix}$$

$$\begin{aligned} S_{dd11}V_{d1}^{in} + S_{dd12}V_{d2}^{in} + S_{dc11}V_{c1}^{in} + S_{cd12}V_{c2}^{in} &= V_{d1}^{out} \\ S_{dd21}V_{d1}^{in} + S_{dd22}V_{d2}^{in} + S_{dc21}V_{c1}^{in} + S_{dc22}V_{c2}^{in} &= V_{d2}^{out} \\ S_{cd11}V_{d1}^{in} + S_{cd12}V_{d2}^{in} + S_{cc11}V_{c1}^{in} + S_{cc12}V_{c2}^{in} &= V_{c1}^{out} \\ S_{cd21}V_{d1}^{in} + S_{cd22}V_{d2}^{in} + S_{cc21}V_{c1}^{in} + S_{cc22}V_{c2}^{in} &= V_{c2}^{out} \end{aligned}$$

Figure 11: A differential channel and its S-parameters.

The transformed S-parameter matrix consists of four sectors. The top left 2×2 submatrix consists of purely differential elements and the bottom right 2×2 submatrix consists of purely common mode elements. For example, Sdd21 describes the response of a differential signal that emerges on the right, port-pair 2, from a differential signal transmitted on the left, port-pair 1. Similarly, Scc21 describes the transmission properties of the pair of conductors for common mode signals.

The top right and bottom left 2×2 submatrices consist of mixed terms. For example, Sdc21 describes the differential response of the channel at port-pair 2 due to a common mode signal transmitted into port-pair 1.

Since crosstalk is picked up by both differential lines, it is primarily a common mode signal. A key feature of differential signaling is the ability to cancel common mode signals at the receiver. The CMRR (common mode rejection ratio) is measured in decibels and is given by:

$$\mathrm{CMRR} = 20 \log \frac{|Sdd21|}{|Scc21|}.$$

3.5 S-parameter Requirements of Interconnects

Many standards specify masks for channel S-parameters to assure minimum acceptable performance. Figure 12 shows PAM4 masks for insertion loss, differential return loss, and differential-common mode conversion.

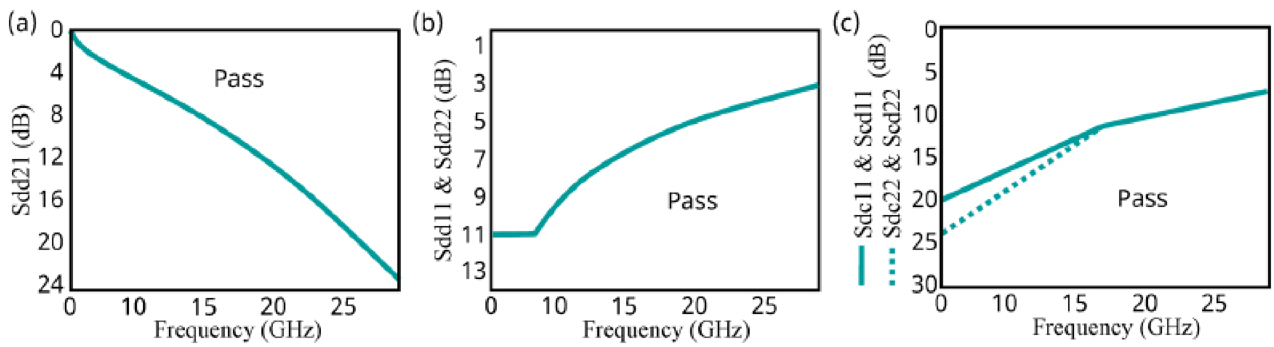

Figure 12: Examples of S-parameter mask requirements: (a) channel insertion loss, (b) return loss, (c) differential to common and common to differential mode conversion.

3.6 S-Parameters for a Many Channel System

In the 400 GbE system that consists of eight separate differential channels operating at 53 Gb/s, Figure 2, the entire system is described by a 32×32 S-parameter matrix. It includes every pair-wise crosstalk response. It's unlikely that you'll ever look at the full matrix, after all it's made of 1024 separate complex functions. But by plotting the magnitudes of specific elements you can learn a great deal about how a system will perform, the balance of crosstalk and insertion loss that it can tolerate, and the degree to which your equalization schemes will be challenged.

Figure 13 shows the insertion loss, return loss, NEXT and FEXT for a differential channel in a two channel differential system.

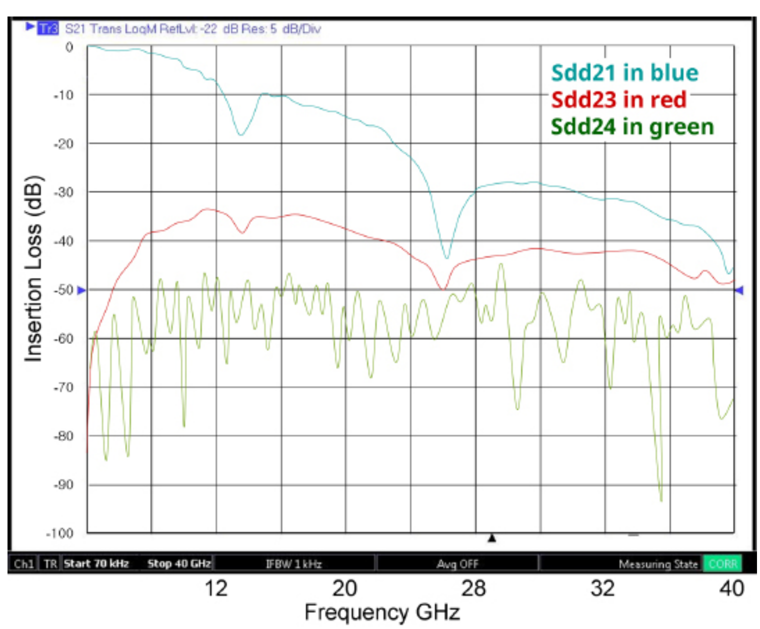

Figure 13: Sdd21, Sdd24 (NEXT), and Sdd23 (FEXT).

4. Modeling and Simulation of SERDES, Interconnects, and Circuits

The use of models to optimize and verify designs radically reduces cost and development time. By modeling a system, we can find the optimum ISI-crosstalk tradeoff in the huge design space that includes the number of transmitter emphasis and DFE taps and CTLE gain, plus differential pair coupling, channel spacing, and all of the other circuit design considerations like reference planes, via locations, coupling capacitors, etc.

To model the eye diagrams eye height and width for each of PAM4's three eye diagrams and the SER (symbol error rate) and BER performance of each design, we need the response of a system to millions of PAM4 symbols.

4.1 Essentials of IBIS-AMI Modeling

IBIS-AMI models integrate analog simulation with algorithm calculations that run fast enough for designers to try many design variations.

IBIS stands for Input/output Buffer Information Specification. IBIS specifies the interface between ASCII files that describe system components and simulation platforms. "Touchstone" files provide the model's network parameters. S-parameters are written to ASCII Touchstone files that have the extension .s*N*p, where *N* indicates the number of ports. For example, the S-parameters of a single differential channel are stored in <filename>.s4p and the S-parameters of the eight channel 400 GbE, 8×26.6 Gbaud PAM4 topology are stored in <filename>.s32p.

AMI stands for Algorithmic Modeling Interface. AMI extends IBIS so that algorithm driven SERDES subsystems like clock recovery, CTLE, DFE, and symbol decoding can be tested.

Other IBIS-AMI input files include the symbol scheme, PAM4 or NRZ, the baud rate, signal characteristics like voltage levels, rise and fall times, and the data pattern. Many IC manufacturers provide IBIS models for their components, but since PAM4 is still emerging few hardware models are available.

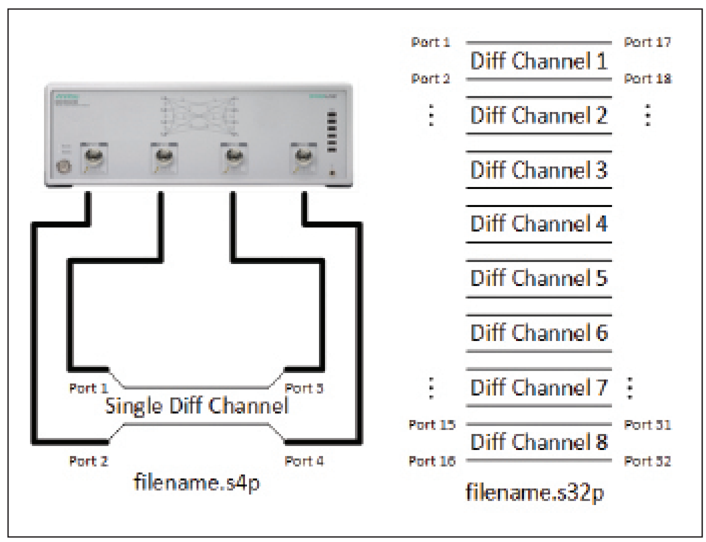

Figure 14: Example mapping of ports to differential channels

IBIS-AMI models allow you to analyze performance in the time and frequency domains at any point from the transmitter to the interconnect and inside the receiver. For example, you can look at eye diagrams before and after CTLE and/or DFE.

Two categories of IBIS-AMI modeling provide different insights. The "statistical flow" processing mode assumes LTI (linear time invariant) behavior. While statistical flow offers quick high statistics results, it's not capable of predicting the nonlinear problems like voltage compression and timing skew that can plague PAM4 systems. The "bit-by-bit" mode is considerably slower, but it can simulate nonlinearities.

Since IBIS-AMI is standardized, Anritsu sNp files can be used seamlessly with any IBIS-based EDA (electronic design automation) simulator. Anritsu's VNAs can be purchased with optional signal integrity applications. VectorStar VNAs have the ability to do live updating of eye diagram simulations from S-parameter data on a sweep by sweep basis (Option 047). ShockLine VNAs have Advanced Time Domain software (Option 022) which has the ability to post process S-parameter files to do eye diagram simulations as well as perform signal integrity analysis from S-parameter files.

4.2 IBIS-AMI Evaluation of Transmitter and Receiver Equalization Schemes

Since eye-closure due to ISI varies widely for different sequences of symbols, we need large statistical samples to evaluate equalization schemes. Fortunately, interconnects are passive LTI systems so we can use statistical flow IBIS-AMI models.

Clock recovery is tricky because the clock circuit derives the data rate clock used to time the logic decision circuit from edges in the PAM4 signal. With four symbol levels, those edges are not nearly as well defined as in NRZ systems and nonlinear timing skew aggravates the problem. A bit-by-bit IBIS-AMI model should be used to check the accuracy of the clock-timing in the equivalent statistical flow model. The DFE performance relies on the accuracy of the four-level symbol decoder. If there's timing skew, the DFE taps derived from the statistical flow model come into question and should be confirmed with a bit-by-bit model. With four different symbols to feedback, tuning DFE is more complicated for PAM4 than NRZ.

Most first generation PAM4 symbol decoders consist of three synchronized voltage slicers, each with its own slice threshold: lower, middle, upper. In IBIS-AMI you can offset the relative timing of each slicer to determine whether or not a more flexible receiver is necessary for a given transmitter-interconnect-receiver combination.

4.3 Accurate Models Require Accurate S-parameters

The accuracy of IBIS-AMI models depends on the accuracy of the sNp input files. A few simple concepts will help you understand the strengths and weaknesses of S-parameter measurements.

4.3.1 Frequency Resolution and Aliasing

VNAs make the most accurate S-parameter measurements possible, but the measurements are made at discrete frequencies. Accurate modeling requires high resolution S-parameters, so the step-size, Δf, between adjacent S-parameter frequency measurements should be small. Another reason to use small spacing has to do with aliasing.

When a discrete function is transformed from the frequency domain to the time domain, the time domain function becomes periodic: It repeats, or aliases, in time intervals called the alias-free range, $T_{alias} = 1/(2\Delta f)$. If the time-delay length of the channel is $T_{channel}$, the alias free range must be at least $T_{alias} > T_{channel}$ which translates to a step size $\Delta f < 1/(2T_{channel})$. To account for impedance mismatches that could cause multi-path effects, the frequency step size should include at least one point per 90 degrees of return loss phase angle change to assure accurate interpolation, the maximum frequency step size should be $\Delta f \leq 1/(8T_{channel})$.

4.3.2 Causality

The biggest potential weakness in transforming frequency-domain measurements to the time domain comes from the limited bandwidth of measurement equipment.

The Fourier transform from frequency to time requires integration over frequency from 0 to infinity, that is, from DC to very high frequency. Since the frequency content of a real PAM4 signal rarely exceeds the fifth harmonic, $5/2\ f_{baud}$, where f_{baud} is the baud rate and ½ fbaud is the fundamental, we can approximate the upper integrand, $\infty \approx 5/2\ f_{baud}$ or larger. The trouble comes at the bottom end. Since we can't generate a DC sine wave, we can't measure the S-parameters all the way down to DC: the DC limit of the S-parameters must be extrapolated. The best that we can do is use equipment that gets as close to DC as possible.

High start frequencies cause causality problems that give the appearance of effects preceding their causes, like an output that occurs prior to its input, and convergence problems in the transformation to the time domain. Anritsu VNAs reduce the causality and convergence risks by reaching down to the lowest frequency in the industry.

Anritsu's VectorStar ME7838x series broadband VNAs offer the widest single sweep coverage from the lowest start frequency of 70 kHz to the highest stop frequency of 145 GHz. This industry-leading widest single sweep coverage ensures that DC extrapolation errors and causality issues are minimized and your simulations match reality.

4.3.3 Reciprocity

The S-parameters of passive devices are reciprocal, for example S12 should be the same as S21 in both phase and magnitude. S-parameters measured on VNAs rarely suffer these “reciprocity” problems. One way to estimate the uncertainty of S-parameter measurements is to compare reciprocal S-parameters. If the uncertainty is high, you may need to recalibrate.

Reciprocity is a much greater concern for TDT/TDR-based S-parameters. In addition to their timebase jitter and noise problems, the fast rise-time step voltage used by TDT/TDR equipment can cause synchronization problems in the forward and reverse directions.

4.3.4 Passivity

Passive devices like interconnects shouldn’t have any gain but gain can appear when the receiver is saturated. It’s easy to eliminate passivity problems when you measure S-parameters on a VNA. VNAs are precise instruments that require careful calibration. Passivity problems can be eliminated simply by exercising good lab techniques. Assure that the VNA receivers are not compressed during calibration and measurement, check that all connectors are in good condition and properly torqued and verify that any cables used in the measurement are in good condition and are properly de-embedded.

4.4 De-embedding and Embedding

S-parameters can also be used to remove the effects of cables, connectors, and traces from measurements. The process is called de-embedding. Say we’re interested in how a signal looks at the compliance point in Figure 15. If we measure the S-parameters of everything between the test equipment and the compliance point, we can “de-embed” the effect of the test fixture and reveal how the signal appears at the compliance point.

Anritsu VectorStar and ShockLine VNAs come with advanced de-embedding network extraction software. Most high bandwidth oscilloscopes can be purchased with de-embedding software compatible with Anritsu sNp files.

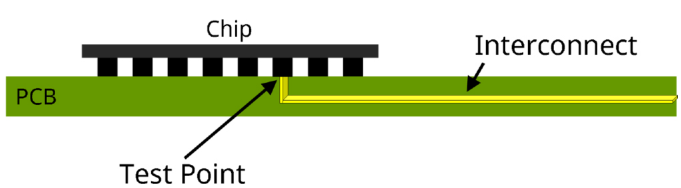

Figure 15: The test fixture must be de-embedded from measurements so that we can evaluate the signal at the compliance point.

De-embedding requires S-parameters measured with dynamic range well in excess of the test fixture’s loss profile; each de-embedded element brings the measurement closer to the noise floor. With 100+ dB of dynamic range, Anritsu VNAs can de-embed quite long backplanes, cables, and traces.

Embedding is the process of adding the frequency response of a network element to a measurement. Embedding is useful for including the effects of internal components like equalizers and clock recovery to see what a signal looks like inside a SERDES.

4.5 Accurate COM Requires Accurate S-parameters

At high data rates, we can't reasonably expect open eye diagrams at the receiver so we need sophisticated tools to analyze signal quality. COM (channel operating margin) is a signal-to-noise-like quantity that combines jitter, noise, crosstalk, and ISI plus the effects of equalization into a single figure of merit. It's like the effective SNR at the receiver input.

Measuring COM is an involved process based on S-parameters. The differential S-parameters of the channel, Sdd21, are used to calculate the ISI signal impairment and the extent to which it can be removed by equalization. The resulting signal amplitude is given by A_s.

FEXT is calculated from the crosstalk elements of the system S-parameters. Along with crosstalk and the remaining effects of ISI after equalization, the transmitter distortion, random jitter and noise are all combined and used to estimate the vertical eye closure defined with respect to a specified SER, A_{ni}. COM is the ratio of the signal amplitude to the vertical eye closure,

$$\mathrm{COM} = 20\log\frac{A_s}{A_{ni}}$$

By specifying COM high speed standards permit design flexibility. Most standards require COM > 3 dB.

Since COM neglects common mode noise and NEXT and approximates the interaction of equalization and crosstalk, it's not a substitute for a proper model.

5. PAM4 Analysis requires VNA Measurements

The introduction of PAM4 signaling complicates design and test of high speed components and systems at every step from design to development to manufacture. The key measurements demand accurate S-parameters that can only be measured on VNAs. Since VNAs are high quality precision instruments that can be major investments, it's important for you to get the equipment you need.

Anritsu Company has a complete line of VNAs from the ultra-high performance 4-port VectorStar VNAs with bandwidths up to 145 GHz to more economical ShockLine VNAs with 2-4 ports that are available in a variety of bandwidth configurations.

For SERDES and interconnect design, you probably need a full featured VectorStar VNA. For system design and development, a 4-port ShockLine Performance VNA, MS4652xB, is probably more than adequate, and for verification and manufacture, a ShockLine Economy VNA might be all you need. Table 1 shows important features for PAM4 applications for two VNA models; the VectorStar and ShockLine families of VNAs offer many other configurations, Table 1 is merely indicative of what's available. The Anritsu field team can provide the guidance you need to make the best choice for your applications and budget.

	VECTORSTAR ME7838A4	**SHOCKLINE MS46524B**
Bandwidth	125 GHz	43.5 GHz
Typical Dynamic Range	110 dB	125 dB
Typical Noise Floor	–100 dBm	–115 dBm
Minimum frequency	70 kHz	50 kHz
Number of measurement points	100,000	20,000

Table 1: Features of a VectorStar and a ShockLine configuration that are important for PAM4 signaling.

5.1 The Advantage of VNA over TDT/TDR

We've already discussed several reasons why VNA measurements are more accurate in both the time and frequency domains than TDT/TDR measurements:

- Noise floor: The noise floor of any measurement is proportional to the bandwidth over which the quantity is measured. By converting a step response to an impulse response, TDT/TDR systems make measurements over 35+ GHz whereas VNA's make many narrowband measurements each over tens of hertz making TDT/TDR measurement noise floors inherently higher.
- Dynamic range: The typical dynamic range of a TDT/TDR test set is around 40 dB. The dynamic range of Anritsu's VectorStar and ShockLine VNAs are typically around 100 dB—a million times that of TDT/TDR.
- Reciprocity: With their short rise-time voltage steps and intrinsic timebase uncertainty, TDT/TDR measurements suffer synchronization problems that can lead to inconsistent S-parameters and make it impossible to model PAM4 problems like timing skew.

While TDT/TDR can be used to measure crosstalk S-parameters, at least in principle, their limited dynamic range and high noise floors make it difficult for them to access the weak coupling between victims and aggressors with any accuracy. Without accurate NEXT and FEXT S-parameters, we can't make accurate IBIS-AMI models or COM measurements.

5.2 The Most Accurate and Economical VNAs You Can Find

Since PAM4 signaling is used at high data rates in high loss environments, you need the dynamic range and accuracy of a VNA. The VectorStar and ShockLine Performance VNAs offer typical dynamic ranges over 100 dB and noise floors below –70 dBm. Figure 16 shows a measurement of the VectorStar ME7838D dynamic range from 70 kHz to 145 GHz.

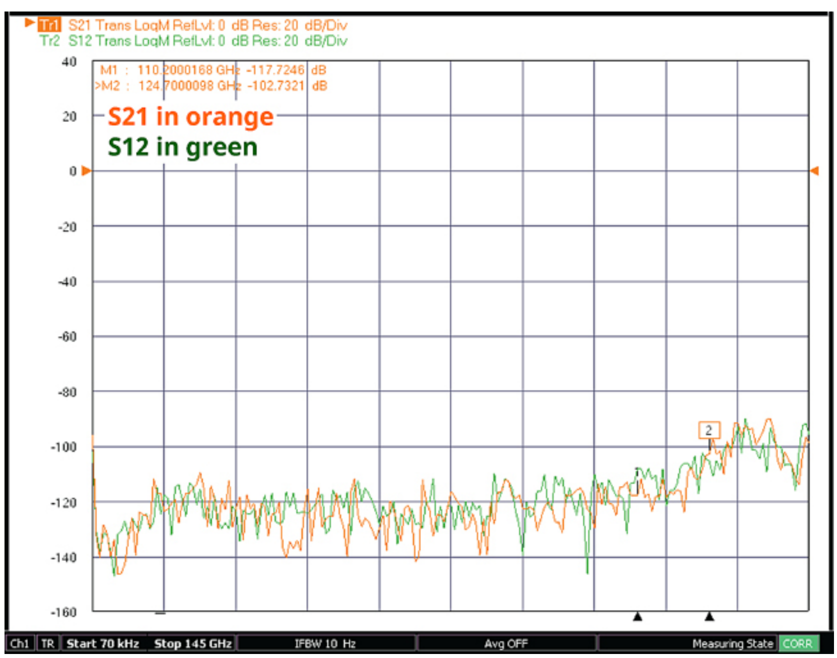

Figure 16: Measurement of the VectorStar ME7838D dynamic range from 70 kHz to 145 GHz.

VectorStar VNAs achieve the best time domain measurement accuracy available. The combination of unique low-frequency coverage and with up to 100,000 measurement points make VectorStar VNAs all but immune to aliasing and causality problems. The more economical ShockLine VNAs also have excellent low frequency coverage and up to 20,000 measurement points.

All of the ShockLine and VectorStar VNAs share common test software that makes it easy to verify a device with a performance VNA model on the bench, and test it with an economy model in manufacturing.

Specifications are subject to change without notice.

• United States
Anritsu Company
1155 East Collins Boulevard, Suite 100,
Richardson, TX, 75081 U.S.A.
Toll Free: 1-800-267-4878
Phone: +1-972-644-1777
Fax: +1-972-671-1877

• Canada
Anritsu Electronics Ltd.
700 Silver Seven Road, Suite 120,
Kanata, Ontario K2V 1C3, Canada
Phone: +1-613-591-2003
Fax: +1-613-591-1006

• Brazil
Anritsu Electrônica Ltda.
Praça Amadeu Amaral, 27 - 1 Andar
01327-010 - Bela Vista - Sao Paulo - SP - Brazil
Phone: +55-11-3283-2511
Fax: +55-11-3288-6940

• Mexico
Anritsu Company, S.A. de C.V.
Av. Ejército Nacional No. 579 Piso 9, Col. Granada
11520 México, D.F., México
Phone: +52-55-1101-2370
Fax: +52-55-5254-3147

• United Kingdom
Anritsu EMEA Ltd.
200 Capability Green, Luton, Bedfordshire LU1 3LU, U.K.
Phone: +44-1582-433280
Fax: +44-1582-731303

• France
Anritsu S.A.
12 avenue du Québec, Batiment Iris 1-Silic 612,
91140 Villebon-sur-Yvette, France
Phone: +33-1-60-92-15-50
Fax: +33-1-64-46-10-65

• Germany
Anritsu GmbH
Nemetschek Haus, Konrad-Zuse-Platz 1
81829 München, Germany
Phone: +49-89-442308-0
Fax: +49-89-442308-55

• Italy
Anritsu S.r.l.
Via Elio Vittorini 129, 00144 Roma Italy
Phone: +39-06-509-9711
Fax: +39-06-502-2425

• Sweden
Anritsu AB
Kistagången 20B, 164 40 KISTA, Sweden
Phone: +46-8-534-707-00
Fax: +46-8-534-707-30

• Finland
Anritsu AB
Teknobulevardi 3-5, FI-01530 VANTAA, Finland
Phone: +358-20-741-8100
Fax: +358-20-741-8111

• Denmark
Anritsu A/S
Kay Fiskers Plads 9, 2300 Copenhagen S, Denmark
Phone: +45-7211-2200
Fax: +45-7211-2210

• Russia
Anritsu EMEA Ltd.
Representation Office in Russia
Tverskaya str. 16/2, bld. 1, 7th floor.
Moscow, 125009, Russia
Phone: +7-495-363-1694
Fax: +7-495-935-8962

• Spain
Anritsu EMEA Ltd.
Representation Office in Spain
Edificio Cuzco IV, Po. de la Castellana, 141, Pta. 5
28046, Madrid, Spain
Phone: +34-915-726-761
Fax: +34-915-726-621

• United Arab Emirates
Anritsu EMEA Ltd.
Dubai Liaison Office
P O Box 500413 - Dubai Internet City
Al Thuraya Building, Tower 1, Suite 701, 7th floor
Dubai, United Arab Emirates
Phone: +971-4-3670352
Fax: +971-4-3688460

• India
Anritsu India Pvt Ltd.
2nd & 3rd Floor, #837/1, Binnamangla 1st Stage,
Indiranagar, 100ft Road, Bangalore - 560038, India
Phone: +91-80-4058-1300
Fax: +91-80-4058-1301

• Singapore
Anritsu Pte. Ltd.
11 Chang Charn Road, #04-01, Shriro House
Singapore 159640
Phone: +65-6282-2400
Fax: +65-6282-2533

• P. R. China (Shanghai)
Anritsu (China) Co., Ltd.
27th Floor, Tower A,
New Caohejing International Business Center
No. 391 Gui Ping Road Shanghai, Xu Hui Di District,
Shanghai 200233, P.R. China
Phone: +86-21-6237-0898
Fax: +86-21-6237-0899

• P. R. China (Hong Kong)
Anritsu Company Ltd.
Unit 1006-7, 10/F., Greenfield Tower, Concordia Plaza,
No. 1 Science Museum Road, Tsim Sha Tsui East,
Kowloon, Hong Kong, P. R. China
Phone: +852-2301-4980
Fax: +852-2301-3545

• Japan
Anritsu Corporation
8-5, Tamura-cho, Atsugi-shi,
Kanagawa, 243-0016 Japan
Phone: +81-46-296-6509
Fax: +81-46-225-8359

• Korea
Anritsu Corporation, Ltd.
5FL, 235 Pangyoyeok-ro, Bundang-gu, Seongnam-si,
Gyeonggi-do, 13494 Korea
Phone: +82-31-696-7750
Fax: +82-31-696-7751

• Australia
Anritsu Pty Ltd.
Unit 20, 21-35 Ricketts Road,
Mount Waverley, Victoria 3149, Australia
Phone: +61-3-9558-8177
Fax: +61-3-9558-8255

• Taiwan
Anritsu Company Inc.
7F, No. 316, Sec. 1, Neihu Rd., Taipei 114, Taiwan
Phone: +886-2-8751-1816
Fax: +886-2-8751-1817

11410-00984, Rev. A Printed in United States 2016-09