

Technical Article:

HFTA-06.0

First Presented January 28, 2003, at

DesignCon 2003

High-Performance System Design

Designing a Simple, Small, Wide-band and Low-Power Equalizer for FR4 Copper Links

(Featuring the MAX3785, 1Gbps to 6.4Gbps Equalizer)



Abstract

In recent years, a great deal of effort has been expended on materials and equalizing circuits to push the upper range of copper links up to 10Gb/s. The glamour of the 2Gb/s to 3Gb/s performance is fading as economic concerns emerge. Legacy compatibility issues, cost of ownership and longevity direct the focus to the range of operation rather than just the top speed.

This discussion examines the design considerations, tradeoffs and a technique used to produce an equalization solution. The design technique suggests an alternative to developing complex models of the environment. Data are recorded from the target environment in the lab, replayed in the simulator, and altered by the circuit. Results from a working implementation are presented. This equalization approach is realized in a tiny chip-scale package that permits operating from 100Mb/s to 6.4Gb/s, uses less than 50mW and is insensitive to coding.

Introduction

In recent years, a great deal of effort has been expended on materials and equalizing circuits to the push the upper range of copper links up to 10Gb/s. The glamour of devices operating from 2Gb/s to 3Gb/s is fading as economic concerns emerge. Legacy compatibility issues, cost of ownership and longevity direct the focus to the range of operation rather than just the top speed. Products that are eye-catching performers don't always score very high in the "backward compatibility" and "extensibility" categories. Designers need signal integrity solutions that can operate over a wide bit-rate range, use little power, consume no real estate and are coding/frequency content independent. This helps designers use the same subassemblies for a variety of applications and forestall the need for that "fork-lift upgrade!"

This discussion examines some of the history of transmission line compensation, some design considerations and tradeoffs and presents a method to produce a passive equalization solution.

A Brief History

Before diving into the latest architecture with bold claims, it is worth looking to the past for clues and perhaps some inspiration. In consulting the fossil record, some astonishing things were discovered. What follows is a brief, albeit incomplete, tour of the past through the eyes of the US Patent Office. For the purposes of this discussion, only analog or linear techniques are examined.

Automatic/Adaptive Techniques

An early all electronic and automatic (i.e. self regulating) equalizing circuit is found in the 1936 US patent no. 2,054,657 "Automatic Selective Padding Control Circuits." Invented by Hans Mayer while at Siemens Germany in 1934, it employs bandwidth selective detection and regulation of bandwidth selective attenuators. See figure 1. Two cascaded attenuators, a low pass and a high pass, have their respective gains automatically adjusted by accompanying detector circuits. This notion was further embellished in the 1952 patent no. 2,719,270 "Transmission Regulation" by Ketchledge of Bell Labs NY. Ketchledge's invention uses an analog computer to control a plurality of cascaded adjustable equalizers. Earlier electro-mechanical versions are presented by Affel in 1924 with patent no. 1,511,013 and by Green in 1930 with patent no. 1,743,132.

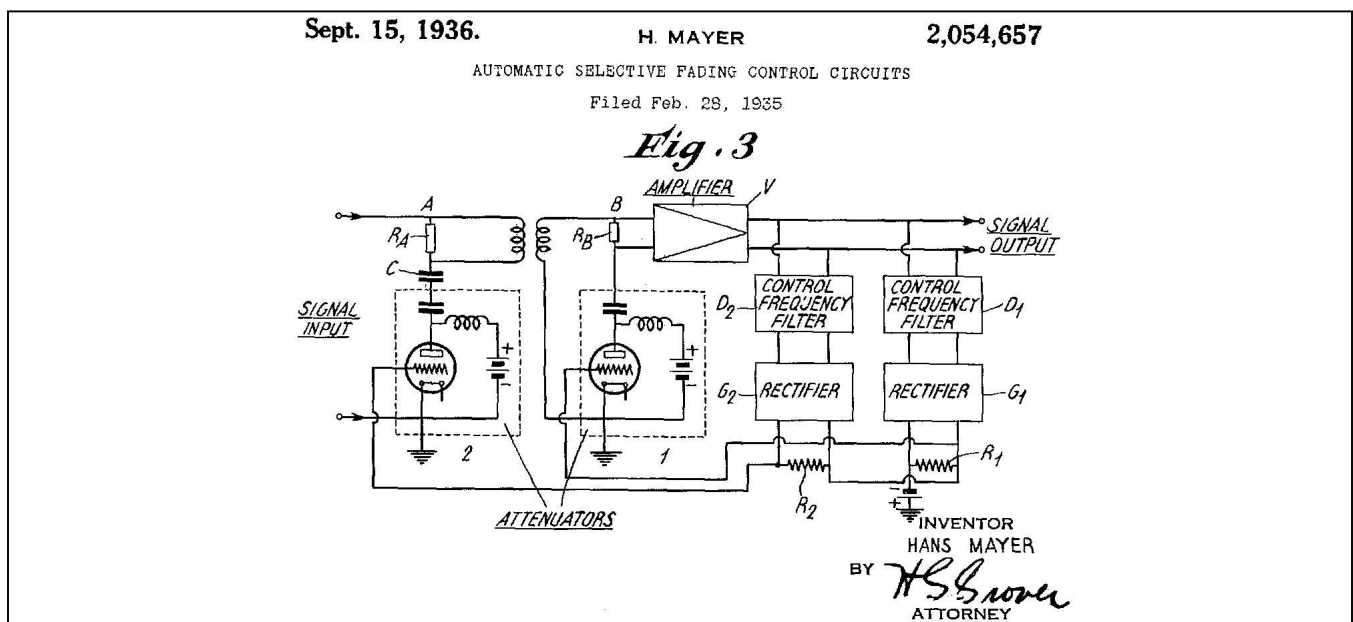


Figure 1. An early automatic, adaptive equalizer invented by Hans Mayer in 1934. The US patent was granted in 1936.

Split-Path or Feed-Forward Amplifier

The split-path or feed-forward amplifier is an innovation that simplifies the analysis and improves the circuit performance. Rather than cascading elements with differing transfer functions, they operate in parallel and do not interact. This allows the transfer functions to add linearly without severe penalties when compensation regions overlap.

In his 1928 US patent no. 1,686,792 "Translation System," Harold Black of Western Electric shows us a crosstalk cancellation circuit in a repeater employing what we today call "feed forward" or "split path." See figure 2(a). This "Translation System" serves to "... repeat electrical waves without distortion... (and) suppress the distortion and modulation produced in an amplifier circuit." In 1979, Pat Quinn of Tektronix patented the "Feed-Forward Amplifier" specifically to correct for distortion in bipolar amplifiers in order to achieve a high degree of linearity and very high bandwidth, dc to >1GHz. This was used to correct and compensate the first high-gain, 1GHz oscilloscope system from Tektronix in the late 1970s. In this real-time analog scope, the 7104, the input signal had to be amplified as much as 1000 times to drive a high-writing rate CRT. This path included circuit boards, interface connectors, twelve meters of delay line cable and a distributed CRT deflector plate assembly. The active equalization technique used to compensate this path is the subject of the 1979 US patent no. 4,132,958 by John Addis and Bruce Hoffer. It is called the "Feedbeside Amplifier." See figure 2(b).

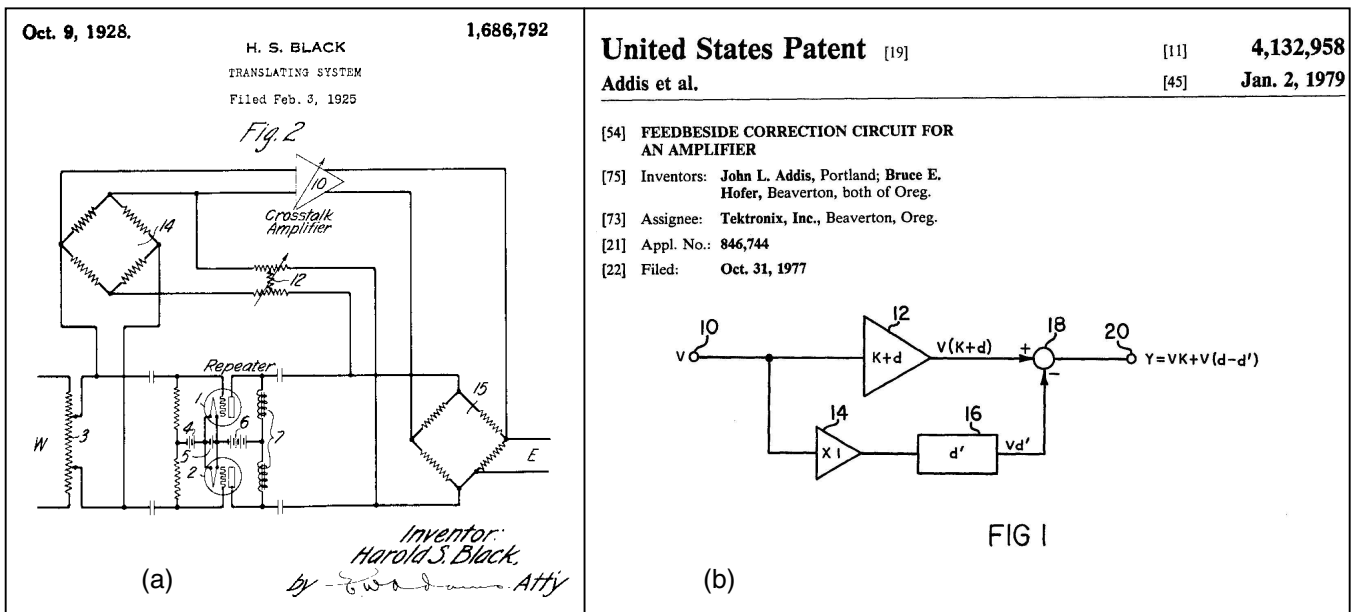


Figure 2. An early example of feed-forward by Harold Black to correct for crosstalk is shown in (a). A similar idea by John Addis to compensate for high frequency loss in a Tektronix oscilloscope is shown in (b).

Modern equalizers continue to employ these early techniques in today's highest performance copper transmissions. Maxim's first 3.2Gb/s cable equalizer, introduced in 2001, uses approaches similar to Mayer, Quinn and Addis. Figure 3 shows the block diagram of the MAX3800 equalizer. Maxim has extended the split-path/feed-forward technique to 12.5Gb/s with the MAX3804. It is a manually adjusted device making it suitable for a variety of applications from 1Gb/s to 12.5Gb/s. Figure 4 demonstrates that the MAX3804 can compensate up to 30inches of FR4 with outstanding performance at 5Gb/s and 10Gb/s.

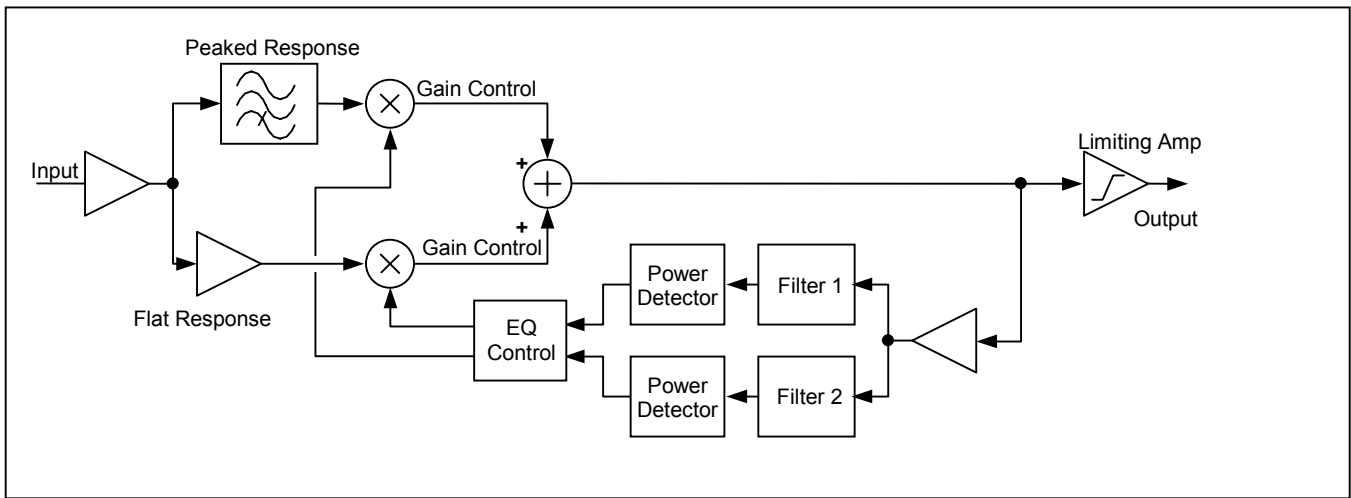


Figure 3. Maxim's first 3.2Gb/s automatic/adaptive cable equalizer, the MAX3800, uses the split-path approach similar to Quinn and Addis (1979), but uses a regulating scheme similar to Mayer (1936).




Figure 4. The manually adjusted MAX3804 equalizer restoring signals at 5Gb/s (left) and 10Gb/s (right) over a 30in, 6mil FR4 line.

Passive Compensation

In his 1936 patent no. 2,096,027 the "Attenuation Equalizer," Hendrik W. Bode shows us a variety of constant-resistance compensation networks including the familiar bridged-T network. This is accompanied by an extensive mathematical foundation for the design of such networks. Bode improves upon and elucidates the work of O. J. Zobel. Zobel's 1926 patent, no. 1,603,305, introduces the "Constant Resistance Network" for use in equalizing transmission lines.

Today the bridged-T network and its differential/balanced version, the bridged-H network, are found in cable assemblies to compensate a particular length of cable. Example applications include Fibre Channel and InfiniBand. An early example of a passive compensation network used to correct for high-speed (>1GHz) losses in copper is found in the Tektronix 7104 Instruction manual published in 1978. This 1GHz analog oscilloscope used a twelve-meter delay line cable in the signal path with a passive compensation network. As mentioned earlier, the "Feedbeside Amplifier" did the rest.

In Bode's 1941 patent no. 2,242,878 the "Design of Broad Band Repeaters," we see the use of passive compensation on both ends of the transmission line (i.e. source and destination) in order to realize a flat response for the span. See figure 5. This is an early form of pre/de-emphasis. Being a linear system, the filter can be placed anywhere in the path and accomplish the same results. If we could observe the output at T2, we would see lots of overshoot characteristic of de-emphasis and pre-emphasis.




Figure 5. Hendrik Bode shows that the line attenuation (center) can be corrected by placing passive filters T1 and T2 on each side of the line.

To illustrate the concept, figure 6(a) shows the output of a 30inch FR4 link with the compensation network located at the end of the link. Figure 6(b) shows the same result but with the network at the beginning of the link. As expected with a linear system, the location of the network makes no difference to the total response. Such a network offers the advantage of being placed wherever it is convenient, the source, destination or at an interface point such as a connector. Figure 6(c) shows the output of the compensation network near the transmission source. Note the low frequency "de-emphasis" and the over compensation of the high frequency content. For this particular network, values were chosen such that there is no significant jitter penalty for over compensating the shorter lengths. Compare the jitter at 10in and 30in. A network similar to this was incorporated into a fixed (non-adjustable) equalizer from Maxim called the MAX3785.

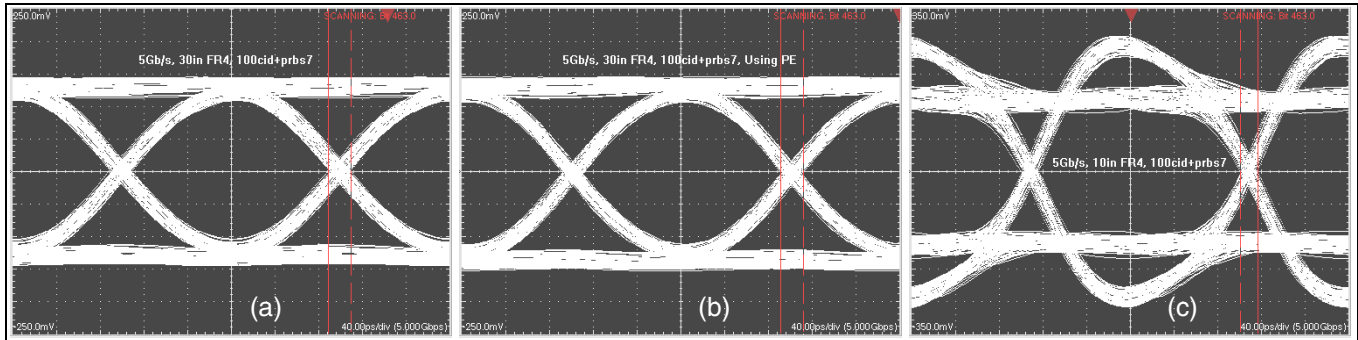


Figure 6. Response of an experimental passive network operating at 5Gb/s. The result of spanning 30in of FR4 with the network placed at the destination is shown in (a). The result with the network placed at the source is shown in (b). Observing the signal near the source at 10in is shown in (c). Note that there is no jitter penalty for over compensating.

Comparison Summary

A summary comparing three categories, Automatic, Manual and Fixed analog compensation methods, is found in table1. For reference, example Maxim products of each type of equalizer are included.

Table 1

Feature/Spec	Relative Score		
	Automatic adjustment	Manual adjustment	Fixed (not adjustable)
Example device	MAX3800	MAX3804	MAX3785
Complexity	High	Medium	Low
Coding independence	Medium/Low	High	High
Bit Rate Range	Medium/Narrow	Wide	Wide
Package size	Small	Smaller	Smallest
Compensation range	High	High	Medium
Power	Medium	Medium	Low
Gain Optimized SNR	High	High	Low
Tolerance to model errors	High	High	Low

High gain, automatic and sophisticated techniques have most of the attention these days. The big advantage with an automatic equalizer is that the automatic/adaptive nature corrects for the customer's environment variations. In many cases, the correction goes beyond the customer's environment. It includes the vendor's environment too! The closed-loop correction found in automatic equalizers can correct for process variations in the manufacture of the chip as well as small but significant errors in modeling the environment. These designs can tolerate the use of models that are less than perfect, yet are reasonable approximations. Additional adjustment range can be added to reduce the risk of modeling errors and correct for manufacturing process variations. The biggest challenge to building an automatic equalizer is the control system. Assumptions about the nature of signals and the circuits ability to detect and control must be carefully evaluated. Pathological conditions that will disrupt the control loop must be anticipated. Dynamic range, spectral characteristics, signal-to-noise ratio and mark/space balance (or lack thereof) are but a few considerations to integrate into a design. For example, incorrectly anticipating key differences in the spectral content of scrambled data and 8b/10b coded data might cause regulation failures. A wide bit rate range and coding tolerance can be accommodated in an automatic equalizer, but with added complexity and some compromises. As legacy issues and multi-standard use take the spotlight, creating a universal and automatic equalizer becomes a big challenge.

An alternative is to implement a subset of the automatic device. If the control loop is eliminated, but the gain control is made available to the user, a manually adjusted device is created. Such a device is much simpler and can be more robust. As long as the equalization response is a good match for the environment, it should be coding independent and support a wide bit-rate range. The compromise here is that the customer must take a little more care to know his environment variation in order to choose the optimum settings.

With the pressure to produce low power devices and smaller devices, passive networks deserve some more attention. Passive compensation has been relegated to applications that correct at one specific length and bit rate and tend not to be applied to a wide range of conditions. Indeed, there are compromises. An adjustable equalizer permits controlling the gain to optimize the Signal-to-Noise ratio, whereas a fixed does not. Irrespective of this, passive equalizers can work extremely well when designed around a well-defined environment. The versatility of a fixed, passive network is demonstrated in the following example.

Building a Passive Equalizer

Consider the following requirements for a transmission link.

Medium: FR4 transmission line, $\tan\delta=0.02$, $\epsilon_r=4.0$, edge-coupled, 6mil wide lines, 0in to 30in.

Bit rate range: 1Gb/s to 5Gb/s

Data format: NRZ, 8b/10b and scrambled

Impedance: maintain the characteristic line impedance within $\pm 5\%$ up to 2.5GHz.

It is obvious that the target environment must be represented as accurately as possible to conduct valid simulations. Loose model approximations will be limiting and the true potential of a solution will not be revealed. Without adaptation to mitigate an inaccurate model, capturing the environment is vital when designing fixed compensation. Accurate models are often very complex. This complexity has the impact of longer execution times during simulation. The variety and complexity of data patterns are often compromised to yield acceptable simulation times. Although a good model offers the promise of a universal library component that can be applied nearly anywhere, it quickly becomes unique to a specific case when a connector and wire type are added to the mix. Furthermore, the model must be verified (and tweaked) by building prototypes of the environment and gathering frequency and time domain data to compare to simulations. Attempts to simplify this process by only considering the magnitude and ignoring phase will lead to poor results and it is not worth the savings. Both phase and magnitude must be analyzed and corrected simultaneously. Crafting, measuring and adjusting models is time consuming. Several bit rates, transmission line lengths, data patterns and data sources of varying quality must be rendered for simulation and model verification. The method suggested below is a way to avoid the complex model entirely.

Methodology

The key to success in this alternate method was the recording of the time-domain signals from the target environment in the lab and replaying them in the simulator as a voltage source. This eliminated the lengthy execution time and development time associated with complex models. Many combinations of bit rates, line lengths and data patterns representing the application range were gathered and replayed. Because the ultimate measure of performance is a time-domain spec (i.e. jitter), the compensation was crafted in the time domain. This simplified the task by compensating for both the magnitude and the phase simultaneously.

For simplicity in this demonstration, a singled-end design is presented. The actual implementation used a duplicate network to form an "H" network for the differential signaling. See figure 7(c).

Start with a basic “symmetric-T” attenuator network that provides attenuation equal to the loss encountered at about half to three-quarters of the bit rate. There are a number of ways to estimate this loss: 3-D solvers, published data, t-line calculators, scope measurements and VNA measurements. For this case, it was determined that the high frequency loss is approximately 16dB at 3GHz over 30 inches of FR4. Figure 7(a) shows the familiar impedance matched, symmetric-T attenuator network. The next step is to add some bypassing, figure 7(b), to get the mid-band and high-band portions of the spectrum to pass with less and no attenuation respectively. This forms a network that resembles the classic “bridged-T”; however, it does not conform to the “bridged” rule that $R3=R4=Z_0$, where Z_0 is the characteristic line impedance. Instead, the symmetric-T attenuator values are chosen as:

$$R3 = R4 = Z_0[1 - 2/(K+1)]$$

$$R5 = 2 Z_0/(K-1/K)$$

where K is the ratio of input voltage to output voltage across the network.

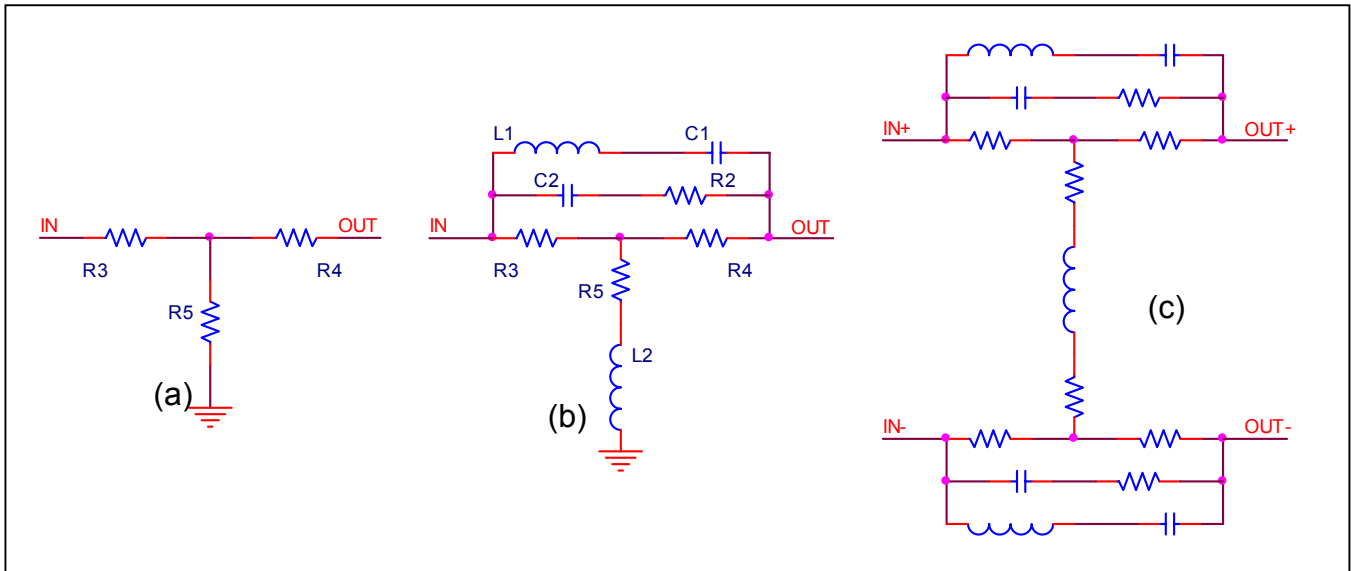


Figure 7. Simple compensation networks (b) and (c). Starting with the familiar impedance matched, symmetric-T attenuator (a), bypassing is added (b) for the needed frequency response. For differential/balanced applications, the network is duplicated (c) for the complimentary signals IN- and OUT-.

Selecting the circuit values starts with restoring the transient response. Figure 8 shows the step response of the network when driven by a voltage source that is a recording of the step response of the 30in long line. For these simulations, a demonstration version of OrCAD™ was used. Figure 8(a) shows the effect of varying the low frequency compensation (R2 and C2). Figure 8(b) shows the mid-band compensation (L2). Figure 8(c) shows the high frequency compensation (L1 and C1). These elements interact; so, converging upon the best step response will take some tinkering. Inductor L2, in figure 8, plays a critical role in restoring the characteristic impedance as well as providing the mid-band compensation. Figure 9 shows the input and output impedance as L2 is varied. Using L2 for impedance matching is critical to minimizing return loss. In figure 9, "Best Z" shows that the impedance can be maintained to within $\pm 5\%$ well past 3GHz.

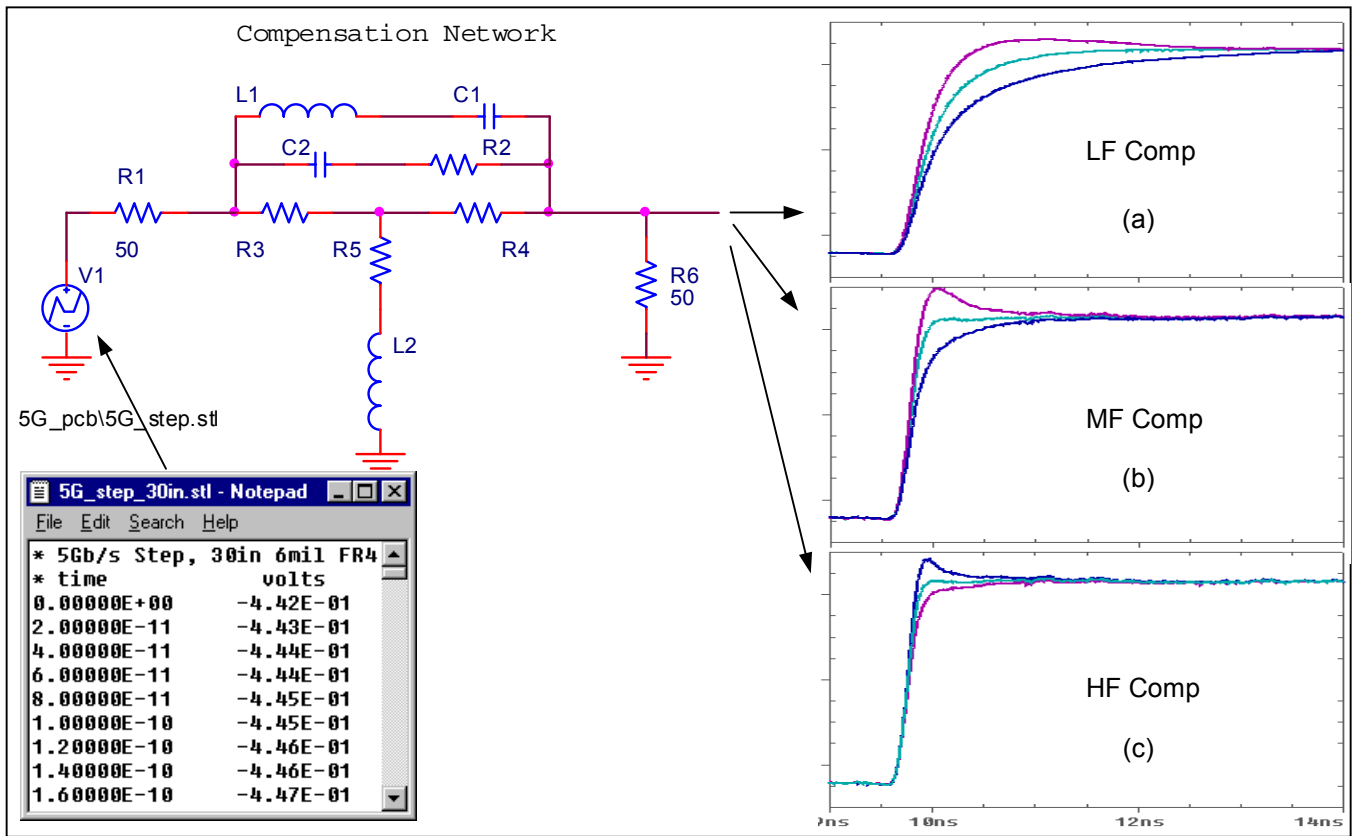


Figure 8. Examples of sweeping the time constants for the network of 7(b). The first goal is to get a flat step response.




Figure 9. Sweep varying L2. L2 also changes the impedance of the network in a way that is opposite to that of the bypassing caused by series R2, C2 and series L1, C1.

After getting the front corner of the step response as sharp and flat as possible, it was a matter of trying other patterns and checking the jitter. If the maximum compensation is not too aggressive (e.g. less than 15dB), there will be little or no jitter penalty for over compensating the shorter line lengths. Figure 10 illustrates how the voltage source, V1, replayed a data file that was acquired at the output of the transmission line. The low frequency content was attenuated such that it balanced with the high frequency information. An ideal limiting amplifier was used in the simulation to restore the signal amplitude and clip the overshoot that results in less lossy environments. The example simulation of figure 10 shows 5Gb/s operation with a 30in FR4 transmission line. The eye diagram at the right shows remarkable fidelity with very little jitter.

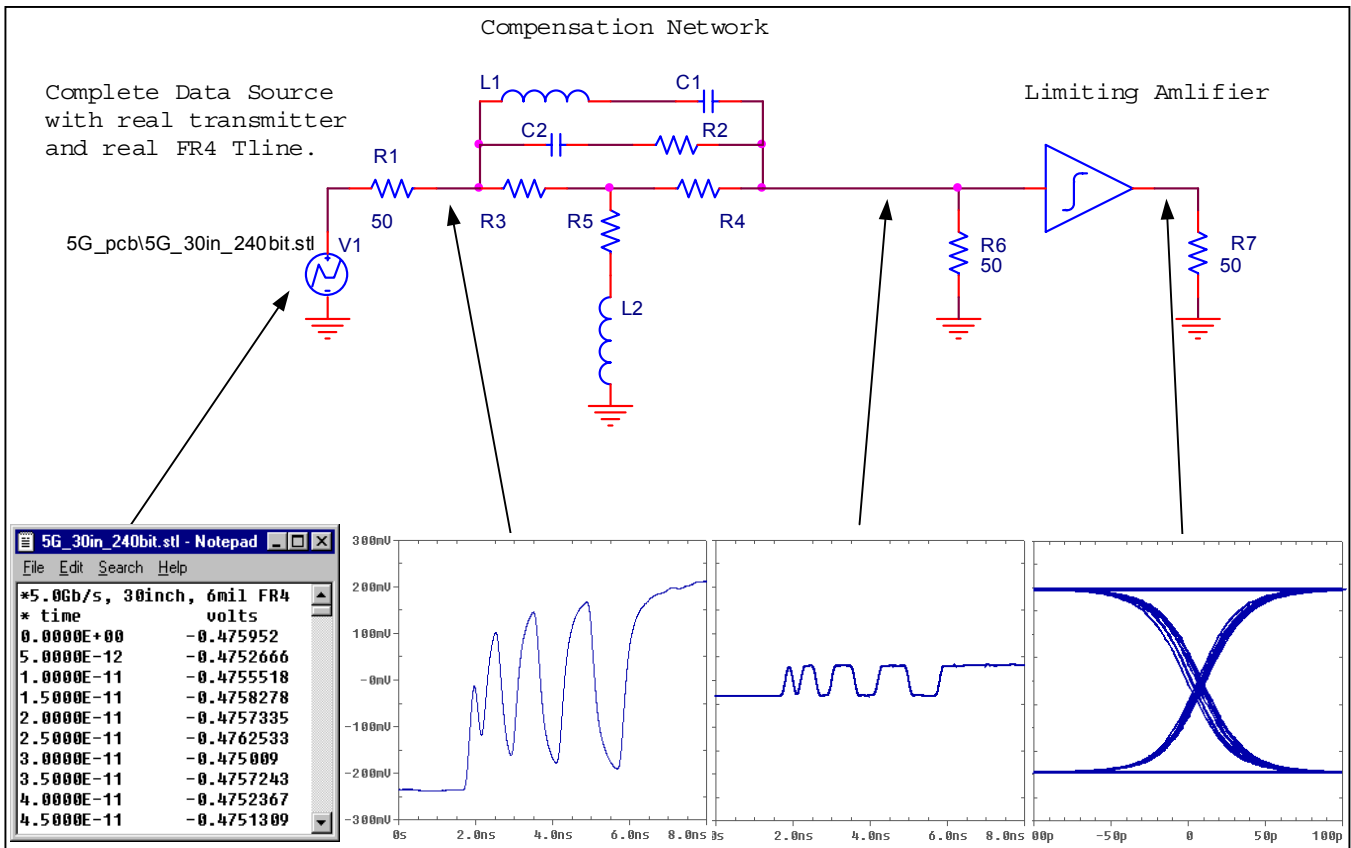


Figure 10. Simulations using a stressful pattern with 100CIDs. A time-domain signal from the target environment is replayed as a voltage source, V1. The eye diagram at the right is the output of the limiting amp.

Figure 11 shows the bottom view of a new equalizer from Maxim that uses one such network to provide the compensation. In the MAX3785, a limiting amplifier follows the passive network. This restores the full swing and clips the overshoot. The device is fabricated using Maxim's proprietary SiGe, bipolar process. With the compensation reduced to passive components, the remaining active portion permits the use of a low 1.8V supply for a total dissipation of approximately 50mW. The size of the package is 1.5mm square, the size of two 0603 resistors! The jitter performance over a wide variety of patterns and bit rates is shown in figure 12. More evidence of the MAX3785 performance is shown in figure 13. These eye patterns were captured using Tektronix' FrameScan™ in the CSA8000. The deterministic jitter of the data source was approximately 10ps. This fixed compensation is capable of restoring the signal over a variety of conditions often thought to be excessively challenging.




Figure 11. Bottom view (pad side) of the MAX3785. This 1.5mm square, chip-scale package is the size of two 0603 resistors.

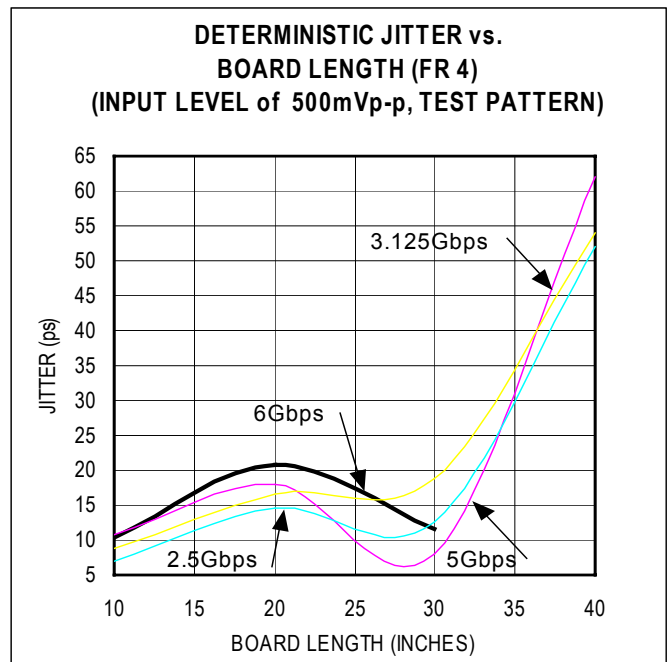
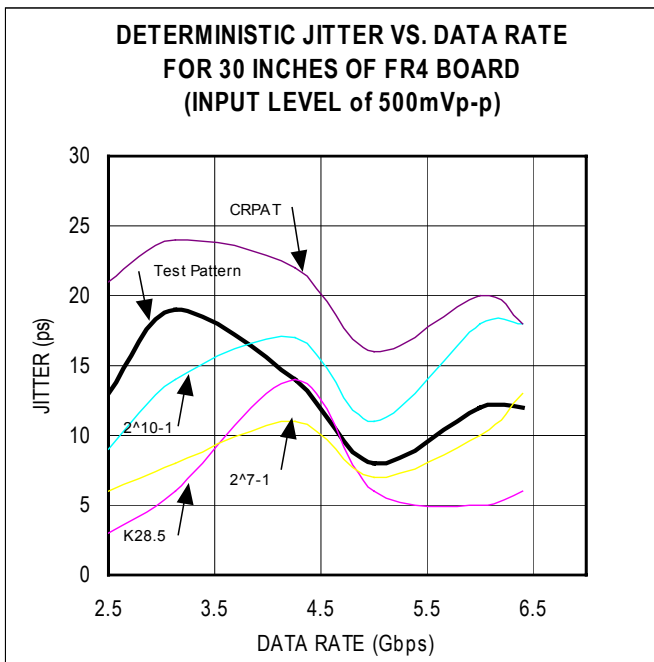


Figure 12. Jitter performance summary of the MAX3785.

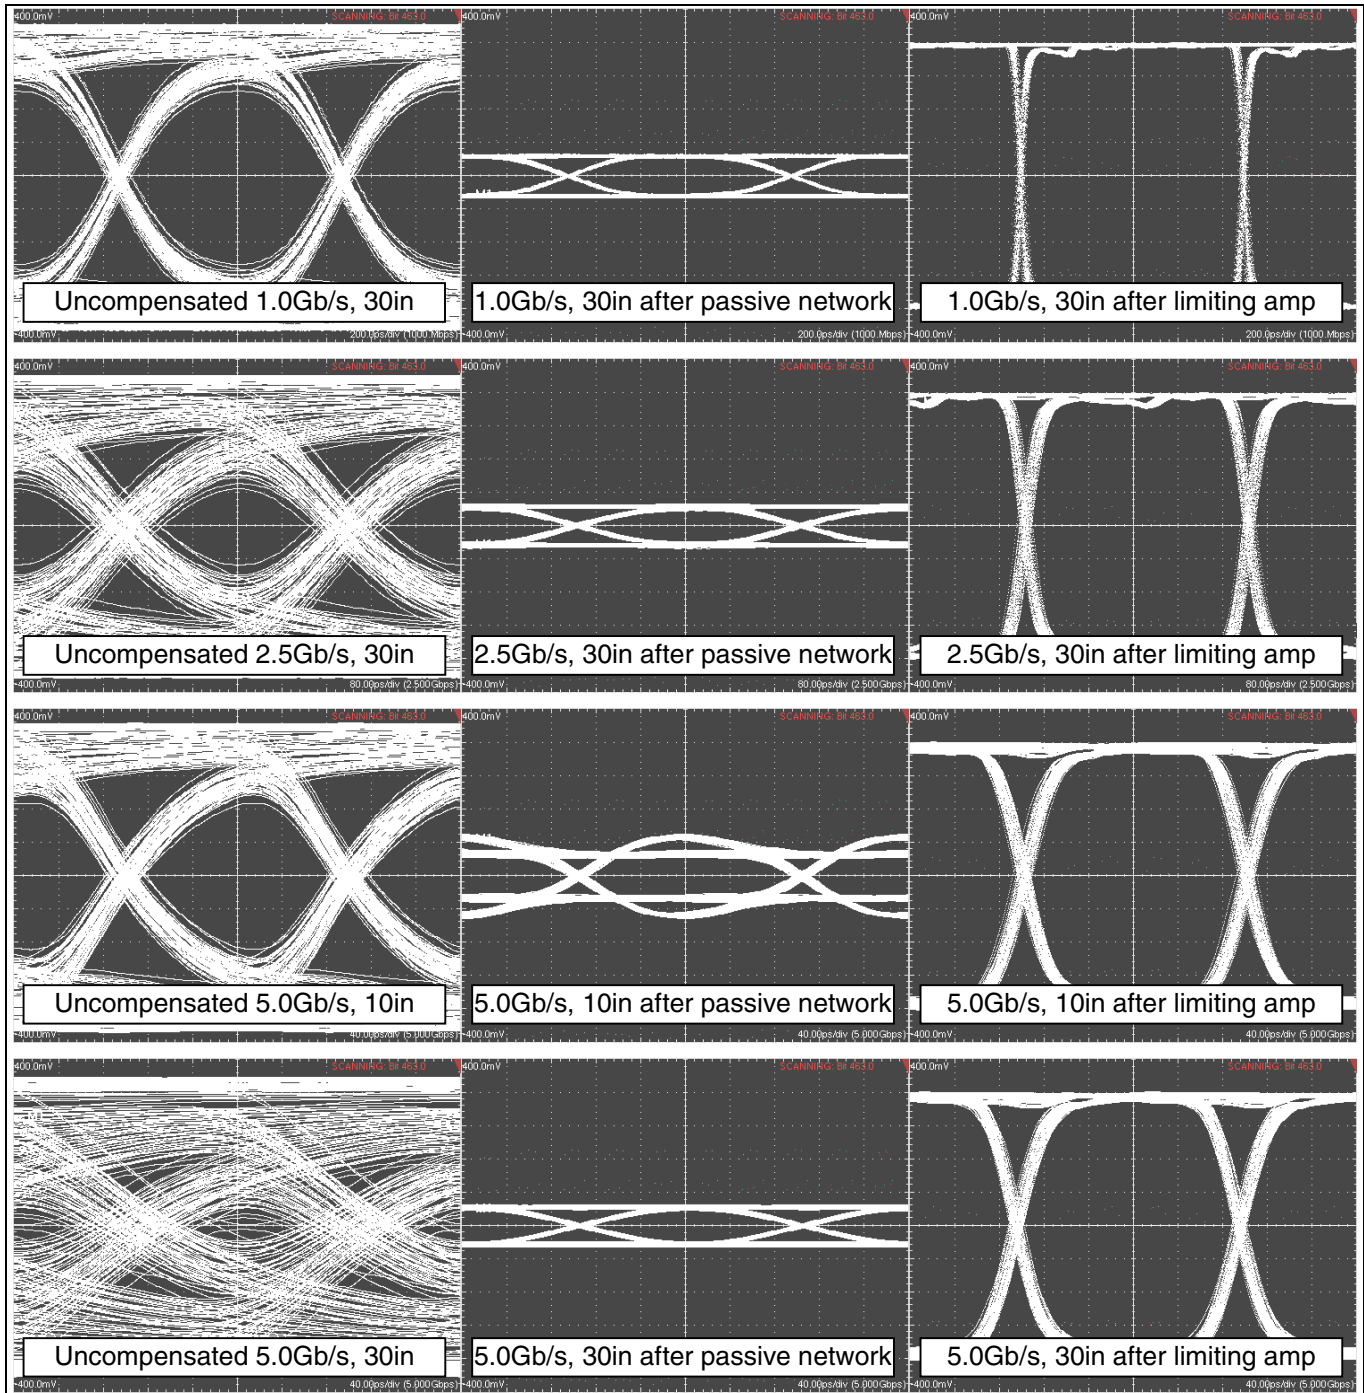


Figure 13. Sample of MAX3785 eye performance at different rates. An extremely stressful pattern of 100-zeros, 1, 0, 1, 0, prbs7, 100-ones, 0, 1, 0, 1, prbs7 was used. The displays to left are the "uncompensated" inputs to the equalizer. The center displays are the outputs of the passive compensation network. The right displays are the outputs of the MAX3785 after the limiting amplifier. The vertical scale is 80mV/div for all waveforms.

Conclusion

This discussion has examined equalizer design considerations and tradeoffs and presented a technique used to produce a passive equalization solution. The passive network was chosen as a vehicle to demonstrate this method because it is imperative to simulate the environment as accurately as possible. There are many subtleties and refinements required to reduce such a network to a complete solution that were not mentioned; however, this design technique provides:

1. efficient and accurate representations of the environment including connectors, noise and non-ideal data sources,
2. fast execution time,
3. simultaneous correction of magnitude and phase by operating in the time domain, and
4. direct evaluation of jitter characteristics.

This equalization approach is realized in a tiny chip-scale package that operates from 100Mb/s to 6.4Gb/s, uses less than 50mW and is not sensitive to coding. This passive approach is but one more tool designers can add to their inventory of tricks. Ultimately, this allows designers to use the same subassemblies for a variety of applications and, in turn, produce products with longer life and affordable upgrade paths.

References

US patents:

- 1,511,013 "Equalization of Carrier Transmissions," 1924, Herman A. Affel
- 1,603,305 "Electrical Network and Method of Transmitting Electrical Currents," 1926, O. J. Zobel
- 1,686,792 "Translation System," 1928, Harold Black
- 1,743,132 "Equalization of Carrier Transmissions," 1930, Estill I. Green
- 2,054,657 "Automatic Selective Padding Control Circuits," 1936, Hans Mayer
- 2,096,027 "Attenuation Equalizer," 1936, Hendrik W. Bode
- 2,242,878 "Design of Broad Band Repeaters," 1941, Hendrik W. Bode
- 2,719,270 "Transmission Regulation," 1952, Ketchledge
- 4,132,958 "Feedbeside Amplifier," 1979, John Addis and Bruce Hoffer
- 4,146,844 "Feed-Forward Amplifier," 1979, Pat Quinn

Tektronix' 7104 Oscilloscope Instruction Manual, Tektronix, Inc., published 1978