Repeaters: Learn to Love‘em

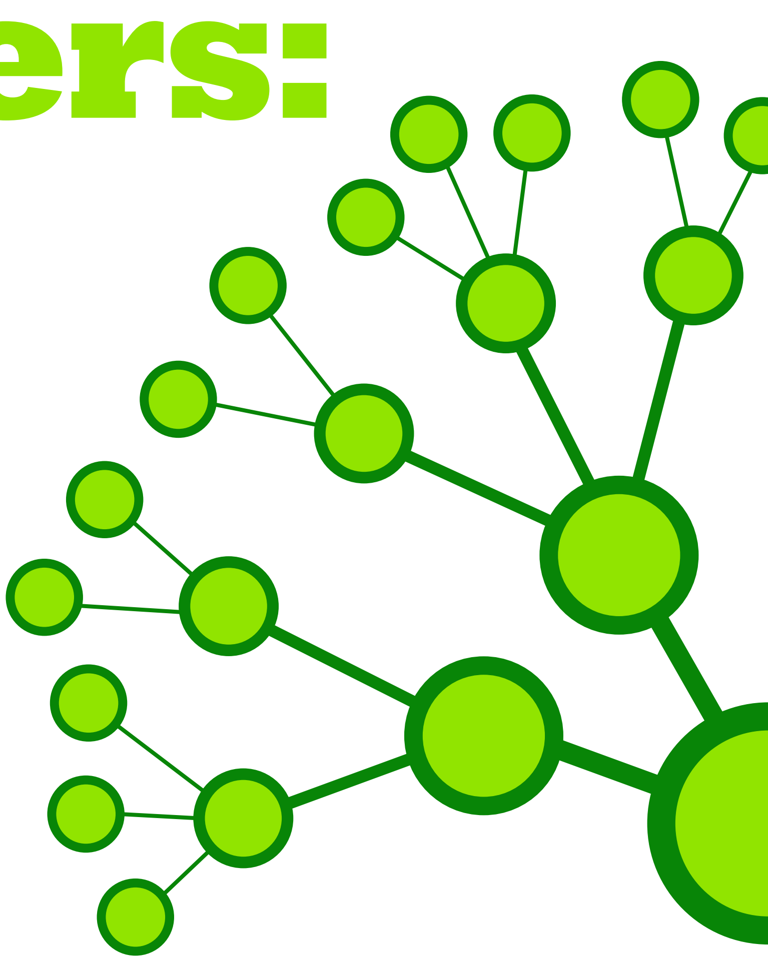

Michael Steinberger
Lead Architect, Serial Channel Products
SiSoft

Eventually, any stationary system you're working on will use electrical repeaters. This article explains why and demonstrates the type of analysis that's required when designing with repeaters.

1.0 The Need for Speed

Many years ago, a director at Bell Labs told me that the cost of a piece of equipment was roughly proportional to its weight. “We sell equipment by the pound,” he said. While that was a slight exaggeration, he had a point. Printed circuit boards, power supplies, connectors, and sheet metal represent a significant percentage of the total cost of any shelf or rack of equipment.

Some years later, I had an opportunity to explore this principle in detail. I was the circuits manager working with the manager of the mechanical design group, helping to choose the technology for a very high-performance system. Instead of trying to reduce the cost of the system at a fixed performance level, we decided to evaluate the cost and performance for each possible set of technology choices. Because the performance of the system was known to be directly proportional to the bandwidth of the interconnect, my group was responsible for estimating the maximum achievable data rate, and therefore the performance. My partner's group was responsible for estimating the cost. We put all our data into a sophisticated spreadsheet that plotted a scatter graph of the performance versus the cost for each one of the ten thousand possible technology choice combinations.

Our study demonstrated that the highest performance-to-cost ratio was consistently achieved by the technology combinations that used electrical repeaters to support a higher data rate in the system interconnect.

I have since lost access to that data, and so can no longer estimate cost accurately. Nonetheless, Figure 1 gives some sense of what those results looked like.

Figure 1 assumes an equipment shelf containing sixteen line

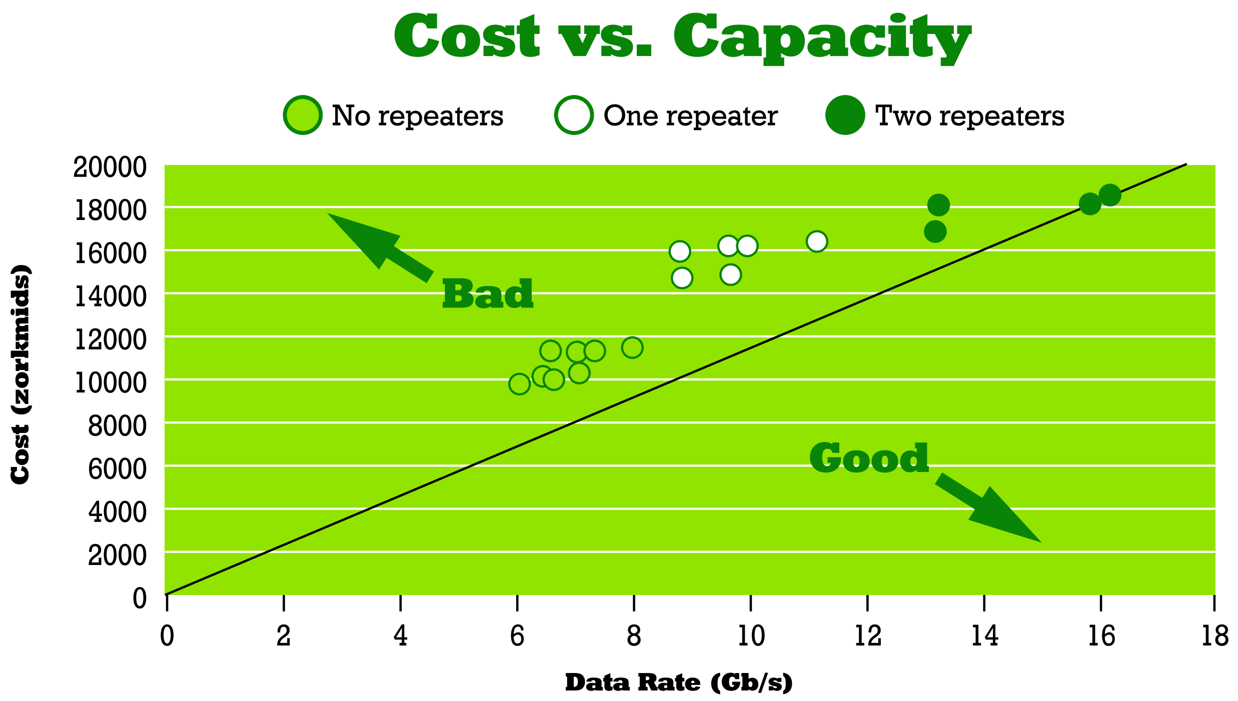

Figure 1: Capacity vs. relative cost for a hypothetical shelf of equipment

cards and two switch cards. The technology choices were as follows:

1. The material for the line card, switch card and backplane could be either FR4 or a lower loss material. The material was chosen independently for each of the three printed circuit boards.
2. There could either be no electrical repeaters in the system, repeaters on the line cards only, or repeaters on both the line cards and switch cards.

The capacity of the processing ICs for the line and switch cards was held constant for all configurations. It was assumed that the capacity of the processing ICs was high enough, and that it was practical to populate the cards with enough ICs to consume the maximum interconnect bandwidth. If this assumption isn't valid, then repeaters aren't required in the first place.

Figure 1 shows that the most cost effective configurations are the ones that use repeaters on both the line and switch cards—in other words, the configurations that can support the highest data rate. Most studies of this type reach a similar conclusion.

In other words, the most cost effective system designs are the ones which achieve a high enough interconnect bandwidth to pack as much processing power as possible onto a printed circuit board. If the highest processing density requires electrical repeaters in the interconnect, then repeaters are going to be an indispensable part of the optimal solution.

Note that the above conclusion applies primarily to high-capacity stationary systems such as core data routers and high-performance computers. The power limitations in portable devices typically reduce the processing power to such an extent that interconnect bandwidth is not a limiting factor.

2.0 Rules of the Road

It's tempting to analyze a channel with electrical repeaters by breaking it into individual segments and analyzing those segments independently. Unfortunately, this approach cannot be relied upon to produce accurate results. Analyzing the entire channel in a single analysis is much more reliable, and there are products out there such as SiSoft's Quantum Channel Designer™ which are able to perform such analyses.

There are several effects that need to be considered:

1. Many repeaters are either linear

1. or quasi-linear. That is, they do not recover the clock and regenerate the data. The effect of such repeaters is cumulative across the entire channel and is not isolated to any individual segment of the channel.

2. If the repeater is nonlinear, then the placement of the repeater in the channel has an additional constraint in that placing the repeater too close to the transmitter could reduce the effectiveness of the repeater.

3. If the repeater has clock recovery and data detection, then the individual segments of the channel are more nearly independent. That may or may not be a good thing depending on the placement of the repeater. Furthermore, the recovered clock from the repeater becomes the clock that the receiver must track. The resulting increase in clock phase noise should be accounted for in the analysis.

The following sections briefly illustrate each of these effects.

2.1 Linear Repeater

Figure 2 is a schematic for a simplified channel with repeater.

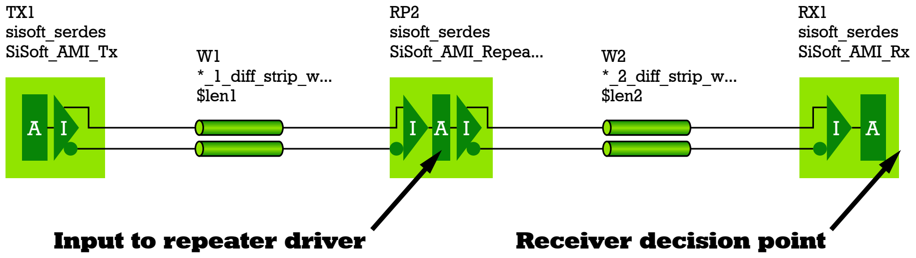

Figure 2: Simplified channel with repeater

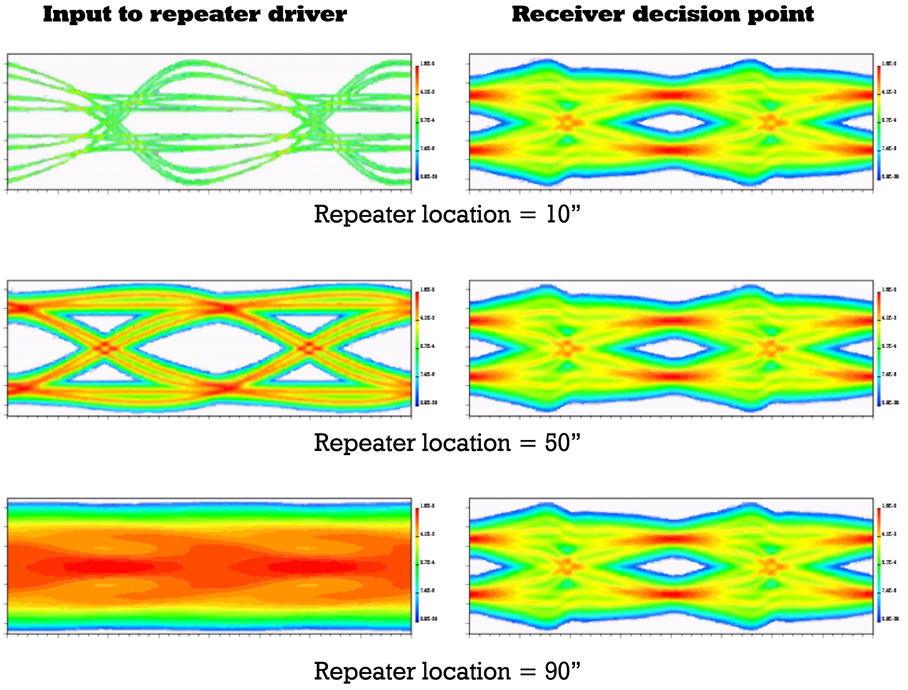

Figure 3: Repeater and receiver eye diagrams for three different repeater locations

In Figure 2, the length of transmission line from the transmitter to the repeater and from the repeater to the receiver are variable. To make observations about repeater placement clearer, the sum of the transmission line lengths was constrained to be a constant 100". Observations were made at the output of the repeater and at the decision point of the receiver. Both the transmitter and the repeater output driver had 3dB of precursor deemphasis, both the repeater input and receiver input had a continuous time linear equalizer (CTLE) and the receiver had five taps of decision feedback equalization (DFE). The data rate was 5 Gb/s.

Figure 3 shows eye diagrams at the input to the repeater driver and at the receiver decision point for three different repeater locations: 10", 50", and 90" from the transmitter.

As shown in Figure 3, even though the eye diagram at the input to the repeater driver varies considerably as a function of repeater location, the eye diagram at the receiver remains almost constant. This is consistent with the assumption that the repeater is linear.

There is a slight variation of the eye diagram at the receiver due to changes in reflected waves between the repeater, the transmitter driver, and the receiver buffer amplifier. However, these variations are relatively small and therefore difficult to see in Figure 3.

Note that, as demonstrated in Figure 3, the eye diagram at the output of the repeater is no indication whatsoever of the performance of the end-to-end channel. The eye diagrams at the output of the repeater are very different, and yet the end-to-end channel performance is almost identical in all three cases.

For later comparison, Figure 4 shows the eye height for a 10-12 BER at the input to the repeater driver and the receiver decision point, each as a function of repeater location.

2.2 Nonlinear Repeater

In general, repeaters are not perfectly linear; and in many applications, the non-linearity of the repeater is significant. One of the primary functions of a repeater is to insert gain into the channel; a number of repeater designs achieve this by providing a high gain saturating amplifier.

The analysis of this type of system is perhaps the most complex. The repeater is not linear, so cascading transfer functions is not valid. Yet the repeater does not recover the clock and regenerate the data, so the channel cannot be broken into independent segments.

The most reliable way to analyze a channel with nonlinear repeaters is a time domain simulation. However, as demonstrated in [1], time domain simulations cannot produce a statistically significant sample of the data waveform. There are approximate methods in statistical analysis and statistical extrapolation that can produce relatively accurate results; however, those methods are beyond the scope of this article.

As a demonstration of the practical considerations associated with the application of nonlinear repeaters, Figure 5 shows the eye height versus the repeater location for a nonlinear repeater in the same way that Figure 4 does for a linear repeater.

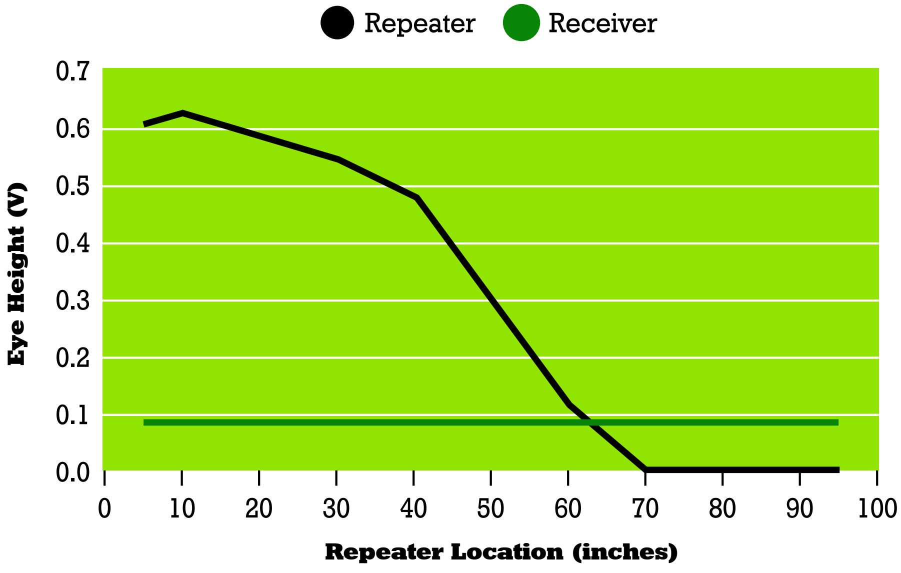

Figure 4: Eye height as a function of repeater location

In particular, Figure 5 demonstrates that if the repeater is placed too close to the transmitter, then the saturation of the repeater will severely degrade performance.

2.3 Retiming Repeater

As mentioned above, there are also repeaters that include clock recovery and data detection. These are often called retimers, or in telecommunications transmission systems they're often called regenerators.

While a channel with retimers can be analyzed as the concatenation of independent segments, there are practical considerations and cumulative effects which should be understood.

Figure 6 is analogous to Figure 4 and Figure 5 in that it shows eye height as a function of repeater location. In this case, however, the eye height at the input to the repeater driver is particularly important in that for a retimer, it is the node where the data is detected.

Note in Figure 6 that for repeater locations beyond 60", the data detection in the repeater fails, causing failure of the entire channel. Even though the receiver has more than enough equalization capacity for the failing repeater locations, that excess capacity cannot be used to improve performance because it occurs after the bit errors have already been made. Figure 6 also shows that with a retimer, the channel length could be extended

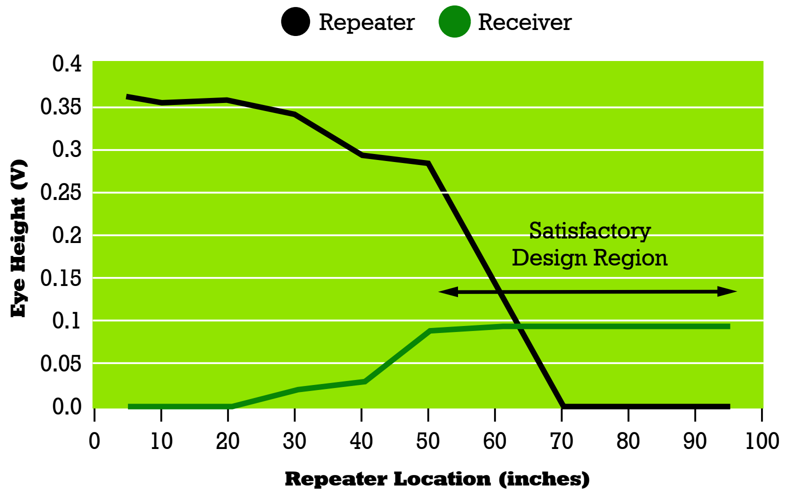

Figure 5: Eye height vs. repeater location for a nonlinear repeater

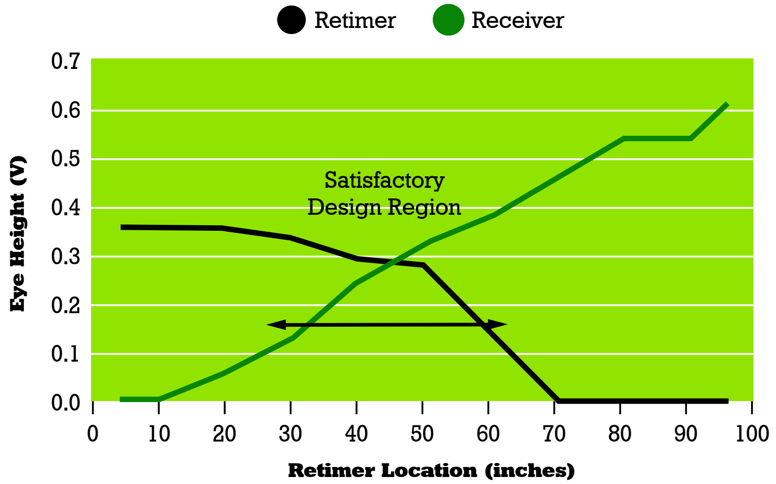

Figure 6: Eye height vs. repeater location for a retimer

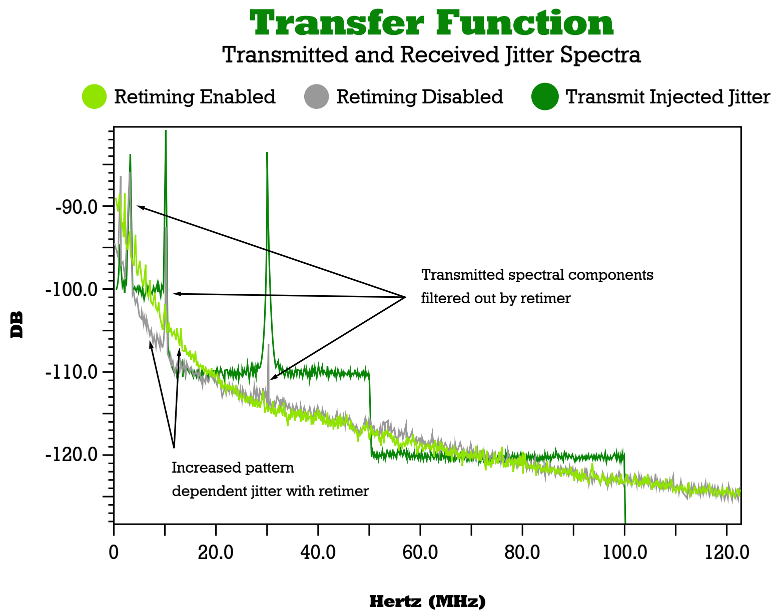

Figure 7: Clock phase noise spectra with and without retimer

to a total of at least 130"—60" from the transmitter to the retimer and 70" from the retimer to the receiver.

When using retimers, clock phase noise also becomes a more complex phenomenon. Rather than recovering the transmitter clock, the receiver recovers the clock from the retimer, which in turn recovers the transmitter clock. Thus, there are two clock recovery loops in series. There is an advantage in that any clock phase noise at the transmitter gets filtered twice, and so its effect at the receiver is considerably reduced. The disadvantage is that each clock recovery loop introduces pattern dependent jitter.

Figure 7 shows the clock phase noise spectral density produced in a simulation in which a specific phase noise spectrum was injected into the transmitter and then the receive phase noise spectrum was measured with and without a retimer in the channel. In Figure 7, the transmit spectrum is shown in green, the spectrum without retimer is shown in grey, and the spectrum with retimer is shown in light green.

The transmit phase noise spectrum in Figure 7 was generated to demonstrate a behavior and does not attempt to represent any real system. In particular, the spectral components in the transmit spectrum are used to make it easy to distinguish between transmit phase noise and pattern dependent jitter. When the retimer is disabled, the transmit spectral components are still readily visible in the receiver clock spectrum. However, when the retimer is enabled, then the transmit spectral components are no longer visible. The trade-off, however, is that there is increased pattern dependent jitter when the retimer is enabled.

3.0 Conclusion

For high performance stationary systems, electrical repeaters are eventually going to become necessary to achieve a competitive performance-to-cost ratio. Designing with repeaters has its own set of considerations which at first may not be intuitively obvious. When incorporating electrical repeaters into the system design, the performance analysis must consider the entire channel, end to end.

4.0 References

[1] Mike Steinberger, "Accuracy of the Computational Experiments called Time Domain Simulations", EEWeb, http://www.eeweb.com/blog/michael_steinberger/accuracy-of-the-computational-experiments-called-time-domain-simulations, July 4, 2011.

About the Author

Michael Steinberger, Ph.D., is responsible for leading SiSoft's ongoing tool development effort for the design and analysis of serial links in the 5-30 Gbps range. Dr. Steinberger has over 30 years experience in the design and analysis of very high speed electronic circuits. Dr. Steinberger began his career at Hughes Aircraft designing microwave circuits. He then moved to Bell Labs, where he designed microwave systems that helped AT&T move from analog to digital long-distance transmission. He was instrumental in the development of high speed digital backplanes used throughout Lucent's transmission product line. Prior to joining SiSoft, Dr. Steinberger led a group of over 20 design engineers at Cray Inc. responsible for SerDes design, high speed channel analysis, PCB design and custom RAM design. ■