SIMPLE INTRODUCTION TO GIGABIT COMMUNICATIONS OVER POF

Document Number, Revision 1.1

TABLE OF CONTENTS

Communicating at a speed of 1 Gbps over a simple and inexpensive medium like standard SI-POF, while using non-sophisticated devices like LEDs and silicon photodiodes, is a difficult engineering

1 OBJECTIVE

The objective of this document is to explain in simple terms how to communicate at Gigabit speed over POF using the VDE-DKE standard V 0885-763:2013-09.

This document will prioritize clarity and simplicity avoiding technical details at the probable cost of some lack of accuracy.

This document focuses on the "how's" and avoids explaining the "why´s"; it explains how the different blocks and subsystems work within the overall communication system, at a high level. It does not, however, explain why this particular set of blocks and techniques were chosen over other alternatives nor how are they configured to operate together.

The intended audience for this document includes electronic, electrical, CS, and telecom engineers with a minimum background in digital communications and communication systems. Many more details on each technique are available in scientific literature, a sample of which is provided in the References appendix.

Please note that this document describes the transmitter part of the communications system for two reasons:

- To avoid an excessive length with redundant information that does not add any significant incremental value
- To adjust our description as much as possible to the standardized system as captured in the VDE-DKE document. Typically communication standards only describe the transmitter side of systems leaving the receptor to the choice of the implementer.

2 THE CHALLENGE

Communicating at a speed of 1 Gbps over a simple and inexpensive medium like standard Step-Index Plastic Optical Fibre (SI-POF), while using non-sophisticated devices like LEDs and silicon photodiodes, is a difficult engineering challenge.

To demonstrate the size of this engineering challenge we will try to show how a "simple" digital communication system would have difficulty in accomplishing this task. By doing this we will reveal the different challenges and barriers Gigabit technology needs to overcome.

First of all, let us define the key components of our "simple" communication system:

Information to transmit: 1 Gigabit of data per second with negligible probability of error and enough power margin for tolerances

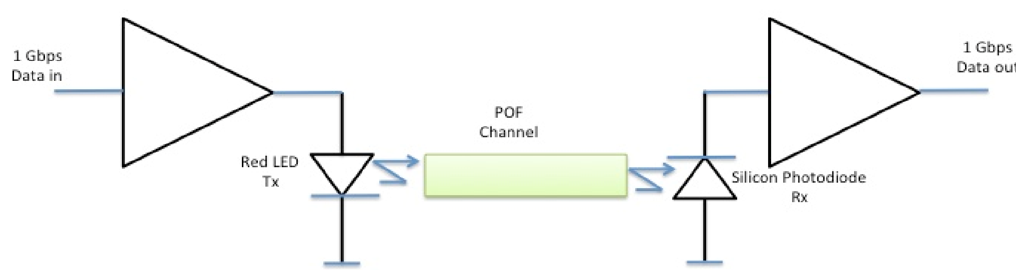

FIGURE 1. SCHEMATIC OF A SIMPLE COMMUNICATIONS SYSTEM

Transmitter: Our light source will be a standard red LED

Channel: Standard SI-POF, 50 meters long

Receptor: Silicon photo detector

Figure 1 provides a schematic of this simple communication system.

2.1 Speed

The first challenge we encounter is the limitation of speed of the source of light (LED) and the POF.

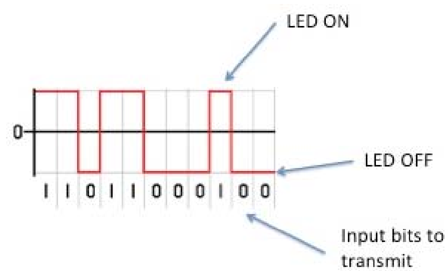

FIGURE 2. SIMPLE MODULATION

To explain this limitation imagine that in this simple communication system we modulate the LED light with data using a simple line coding like NRZ where we switch ON the LED when the bit to transmit is a 1 and switch OFF the LED when the bit to transmit is 0.

With this simple modulation if we try to transmit 1 Gbps we will need to switch on and off the LED at least 10^9 times per second, which translates into a device bandwidth requirement of around 500 MHz. This requirement highlights the first limitation, speed, as a standard red LED used for communications has a typical (electrical to electrical -3 dB) bandwidth of around 80 MHz. You might be able to operate the LED at a faster rate (350 MHz) using workarounds like very high current peaking, for example. However, that solution has side effects including lower reliability and lifetime or smaller available optical power. As you can see, it is clear that it is not possible to find devices that can transmit light at the speeds required by digital communications at 1 Gbps.

If you take into account just the LED, it is not possible to transmit 1 Gbps using a simple modulation scheme like NRZ.

Moreover, the POF itself also imposes a severe bandwidth limitation that will prevent 1 Gbps communication with NRZ modulations even if the light source is able to switch light at 1 Giga symbol per second (GSps). Typical SI-POF has a 3dB bandwidth of 40 MHz at 100 m, making it difficult to transmit a 1 GSps signal even if the length is 50 meters.

In summary, the “speed limitation” means that it is not possible to use simple line codes like NRZ to transmit 1 Gbps over SI-POF using red LEDs.

One workaround would be to put more than just one bit in each light pulse sent over the POF. One way to do this is using more than two levels of light than just the 1/0, light-no light scheme. This method is known as PAM (Pulse Amplitude Modulation).

For instance, if we have four levels of light (fully OFF, fully ON and, for example, an additional two levels of light intensity between fully ON and fully OFF light), we could transmit groups of two bits. In other words, each symbol would carry two bits of information as shown in Figure 3.

Using PAM each time we switched the LED, we would transmit two bits of information. In practical terms, using PAM we would multiply by two the information sent through the fibre using the ON-OFF NRZ scheme.

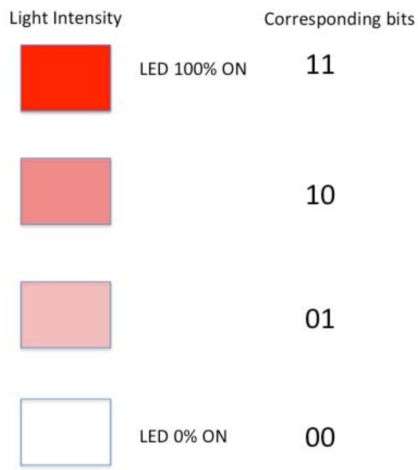

FIGURE 3. PAM SAMPLE

In addition to the speed limitation, other limitations require complex algorithms and techniques to be overcome.

2.2 Channel Distortion

The next challenge we will discuss is channel distortion and the need for an equalizer. Equalizers are always needed in digital communication systems to overcome the Inter Symbol Interference (ISI). ISI occurs in communication channels that are not "fast" enough to react to the signals that travel through it. As a consequence symbols that travel along the channel arrive to the end distorted.

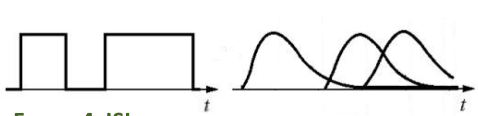

FIGURE 4. ISI EFFECT

An example of this ISI effect is shown in the following Figure 4.

On the left of this illustration you see a digital signal transmitted through a channel (in this case the 1011 sequence) that is not "fast enough" to accurately reproduce the square shape. As a consequence the received signal at the other end, shown at the right, is distorted from square signals into bell shaped curves. The tails of these received signals may overlap and confuse the distinction between a 1 and 0 at reception.

The different techniques used to compensate for the ISI effect are generically called "equalization".

An eye diagram is a way of visualizing ISI. The eye diagram set-up simply over-imposes the received symbols on top of each other so that the final picture appears to be a set of "eyes". The more open the eyes seen in the oscilloscope, the less ISI effect the channel introduces, requiring a less powerful equalizer to compensate.

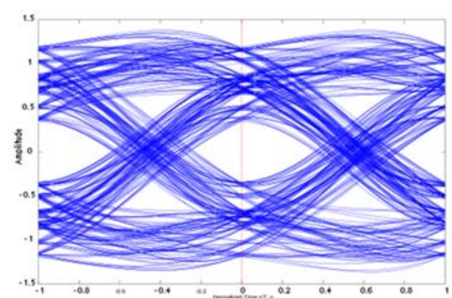

FIGURE 5. SAMPLE EYE DIAGRAM

There are a lot of different ways to equalize a communication channel (such as Zero Forcing, DFE, MMSE, etc.), each with its own set of pros and cons.

DFE, which is a non-linear technique used to equalize linear ISI, is considered one of the "canonical" equalization techniques thanks to its optimality and simple structure. Unfortunately for the purpose of transmitting 1 Gbps over POF with simple LED sources, DFE suffers from several limitations (such as error propagation, difficult to combine with coding, etc.) that make it sub-optimal for our needs.

The solution requires the selection of individual techniques and the fine-tuning and optimization of each one in the context of the overall system.

In our case we need to use more complex techniques where the channel is not only equalized at the receiver, in an effort to compensate for the ISI, but at the transmitter side as well. These techniques are generally called "precoding techniques". The particular precoding technique used in our case, called THP (Tomlinson-Harashima Precoding), will be reviewed in the next section.

2.3 Error Correction System

The last challenge we will discuss is the error correction system. Even in the presence of equalizers, errors happen at reception so we need to have a means of detecting and, if possible, correcting these errors. Correcting for errors translates into an increase in link power budget and helps to compensate for implementation losses and other effects.

The simplest error correcting systems use redundant digits to detect and correct errors. The main drawback of this approach is the obvious increase of the number of bits transmitted. With the limitations of the 1 Gbps POF communication system, we need to have a mechanism to detect and correct errors and, at the same time, avoid losing channel capacity or unnecessarily increasing the complexity of the system.

The solution adopted is based on an approach that combines modulation and coding in what is now known as "coded modulation". Coded modulation is able to improve the performance of the communication channel at the price of a very small bandwidth expansion. Obviously, the trade-off comes from a potentially more complex signal decoder. But, a smart decoder implementation can result in smaller computational cost, power consumption and area.

The KDPOF solution uses MLCC (Multilevel Coset Codes) which are explained in the next section. MLCC splits the signal into three levels and then encodes each level separately. The error correction code chosen for each level, known as "component code", has a correcting power adjusted to the expected level of noise in reception. Levels with more noise will require stronger component codes; levels with less noise will require less powerful component codes or even no coding at all. All in all, the net result is an error correcting system that minimizes the "overhead" of the correction.

In summary, to overcome the speed, attenuation and distortion characteristics of the components of 1 Gbps POF communication system, we need to use elaborated, well-engineered coding, modulation and equalization techniques. Each of these techniques must be selected and adapted with the overall system performance in mind with the goal to optimize the global performance to achieve robust, 1 Gbps, error-free, communication.

3 THE SOLUTION

The solution to the challenge of transmitting 1 Gbps over POF, error-free and with enough robustness and margin, requires the selection of individual techniques and the fine-tuning and optimization of each one in the context of the overall system. All these blocks are integrated into a circuit, generically called ASSP (Application Specific Standard Product), that interfaces with the LED in Tx and the photodiode in Rx.

The MLCC technique is a "coded modulation" technique, meaning that the modulation stage of the digital communication channel is an integral part of the encoding process

The fine-tuning of the settings for each block to approach the Shannon limit as much as possible is a complex multi-variable problem that requires intensive simulation and iterations. The final result is the proposal of a communication system specified by the VDE-DKE standard V 0885-763:2013-09.

- **Pilot Data-Path:** Pilots generated in this data-path consist of training sequences transmitted through the channel. The modulation of these signals depends on the function each performs and, as will later be explained, is either simple NRZ sequences for pilots of type 1 or a signal with a continuum range in the case of pilots of type 2. At this point we will just mention briefly that pilots are needed to perform functions like link start-up, channel estimation, timing recovery and equalization adaptation.

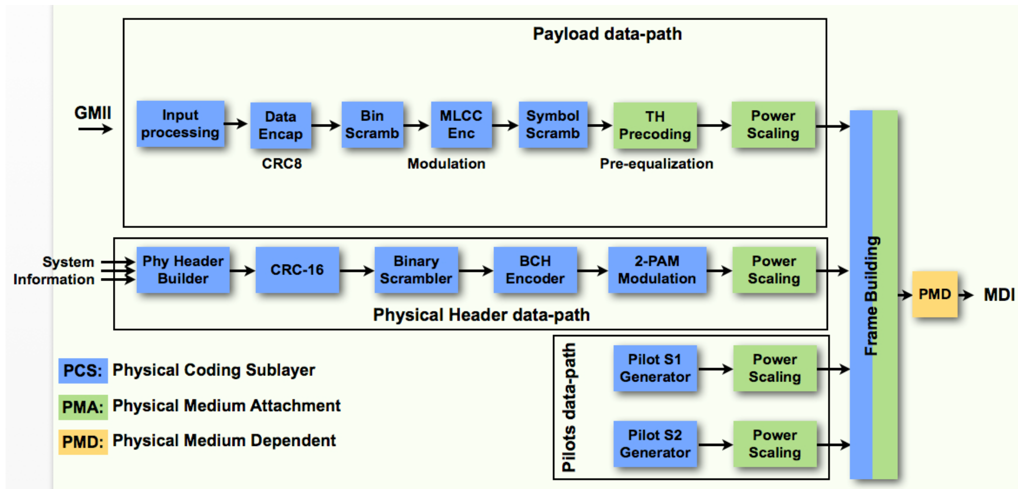

FIGURE 6. TX SIDE OF THE SYSTEM

The overall Tx side of the system can be summarized in the block diagram shown in Figure 6.

From top to bottom the system has three main data-paths:

- **Payload Data-Path**: Path followed by the user data being transmitted through the POF.
- **Physical Header Data-Path:** Header assembled in this path is transmitted through the channel that offers the most robustness against the noise. Its 2D-BPSK modulation scheme and BCH coding maximizes the probability of correct detection regardless of the status of the payload path. The physical header is responsible for precoding and coding negotiation between both sides of the link as well as for data signaling.

At this stage we will focus our attention on the payload data-path as headers and pilots will be described in another section in the context of the frame and data encapsulation. Another reason is that the main blocks that represent the core of the communication system reside in the payload data-path and as such will be briefly covered in the rest of this section leaving a more detailed explanation for later sections of the document.

In the payload data-path, and from left to right, digital data moves from the input of the communication channel to the LED driver or PMD (Physical Medium Dependent), following these steps:

1. Input data is encapsulated in a container that travels through the rest of the system. The VDE standard specifies an Ethernet MAC-PHY interface of the type GMII.

2. The coded modulation stage (MLCC) comes once the data from the GMII interface is encapsulated, and after a scrambling of bits to avoid constant chunks of binary bits. It is at this point that the binary stream of information is transformed into symbols, each of which contains several bits of information. Each of the coded streams is modulated and symbols from each level combine together into a 16 PAM modulation.

3. After another scrambling step, now at symbol level, to ensure that every symbol sees the same noise variance (optoelectronics will eventually introduce non linear distortion) and as a consequence of the same error probability in detection, the THP precodes the signal to perform the channel equalization. It is important to note that the equalization is split between the Tx and the Rx, taking place partly in the Tx side of the link and partly in the Rx side. As the THP is “split” between the Tx and Rx it needs a “back channel” to operate. The THP net effect is to convert the previous 16 PAM signal in a quasi-continuous PAM signal of light levels avoiding the phenomena of error propagation that is characteristic of other solutions based on DFE.

4. The pre-coded symbols are scaled in power to ensure that all the constellations, regardless of their position within the frame, are transmitted with the same Optical Modulation Amplitude and sent to the LED driver to be converted in light pulses injected into the fibre.

4 CODING WITH COSETS

We will start with a deeper view of the communication system with the MLCC/Modulation block immediately after the binary scramble of the input bits.

At this early point of the communication system path we have a stream of bits to transmit. The first thing we must do is to encode these bits in order to protect the information from the noise present in the channel and transform these coded bits into symbols to be transmitted. These two functions are performed jointly in a single block that uses a technique called Multi-Level Coset Coding (MLCC)

The MLCC technique is a “coded modulation” technique, meaning that the modulation stage of the digital communication channel is an integral part of the encoding process and is designed in conjunction with the code to increase the minimum Euclidean distance between pairs of coded signals. The key to this integrated modulation and coding approach is to devise an effective method for mapping the coded bits into signal points such that the minimum Euclidean distance is maximized.

MLCC is best understood with a real-life analogy. MLCC works in a way similar to how an individual person is found through his postal address.

Postal addresses are an easy and economical way to identify people. If we take a look at the way the postal address system works we find that the address information is divided or "partitioned" in blocks or "cosets". First we partition the people in countries, then into cities and postal codes within the cities, into streets, and, finally, into house numbers within the street. The very last partition includes the full name of the addressee which provides us the complete codification for the person we want to address.

The postal address system, shown in Figure 7, simply partitions sequentially the identification of each individual into cosets and labels each position within each coset.

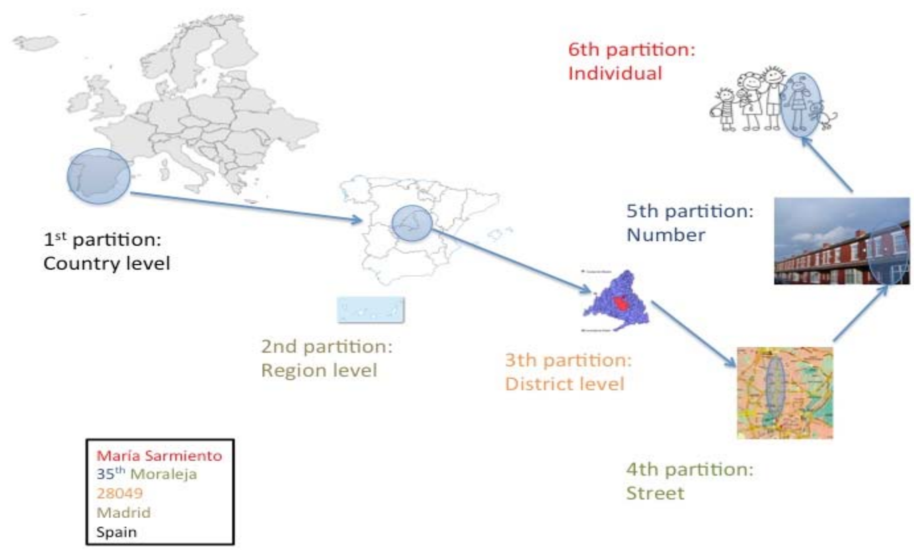

FIGURE 7. EXAMPLE OF PARTITIONING USING A POSTAL ADDRESS

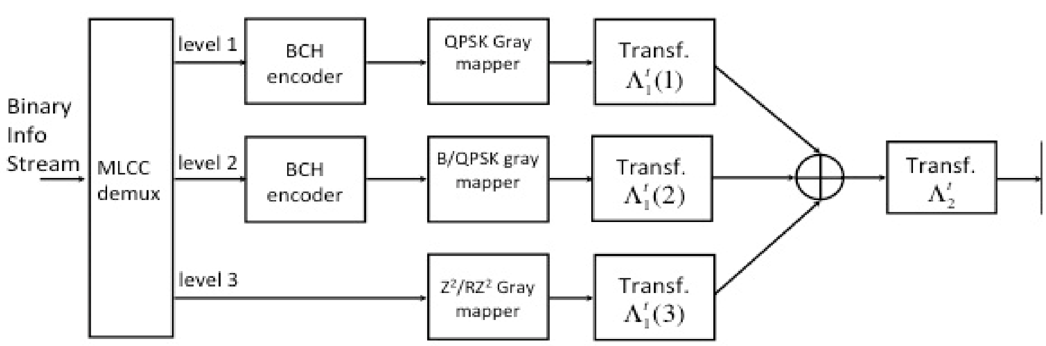

FIGURE 8. SIMPLIFIED BLOCK DIAGRAM OF THE MLCC PROCESS

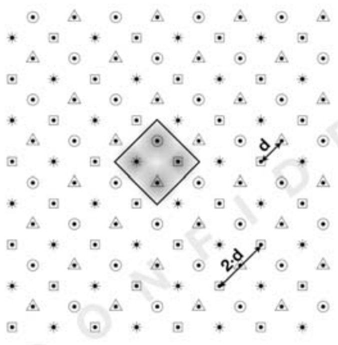

FIGURE 9. FULL POTENTIAL CONSTELLATION OF POINTS TO TRANSMIT

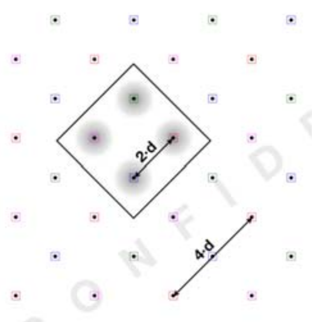

FIGURE 10. CONSTELLATION OF POINTS BELONGING TO SQUARE COSET

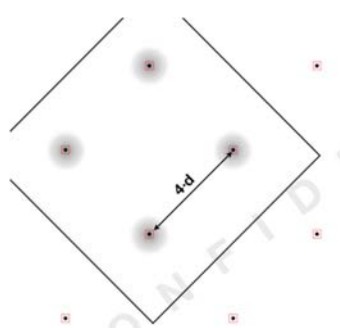

FIGURE 11. CONSTELLATION OF POINTS BELONGING TO THE RED SQUARE COSET

When you write full postal address on an envelope for a letter, you are actually writing each coset on a different line in the front of the envelope.

MLCC works in a similar way. We encode each symbol to be transmitted in a sequential way. We make several partitions into cosets, one after the other, until we identify completely the symbol to transmit or, in reception, we decode the symbol we have received in a sequential way, coset after coset, identifying each level until we have fully identify the symbol.

Figure 8 shows a simplified block diagram of the MLCC process.

The input binary stream is divided into three streams that we will call “levels”. The first level consists of two bits of the binary input stream which are used to “classify” or “map” each of the constellation of points to transmit into one of four cosets. This first step is shown in Figure 9.

This illustration shows the full potential constellation of points to transmit. The first level, with its two bits just assigns the specific point to transmit to one of the four possible cosets (square, triangle, asterisk or circle). As shown in the figure, the spatial separation between different cosets is “d”. If we compare this separation with the noise level that can potentially affect the point during transmission (see the grey cloud around the points) we see that the chance of mixing in reception one coset with another is rather high. So we will need a fairly strong and powerful code component at level 1 to avoid errors in reception.

Now let’s move to the second level.. Let us suppose that, after the first level, our point to transmit was assigned to the coset of “squares”. In this case we can forget all the points of the constellation except those belonging to the “square” coset and move to the second level of encoding. Figure 10 shows the constellation of points belonging to the “square” coset.

In this second level we create another partition into new cosets. Let us, for example, use colors to identify the next coset within the square coset. In our case we can have four square cosets, one in red, the other green, the third blue and the last one magenta. Partitioning at this level means that we must decide if our point to transmit belongs to the red, blue, magenta or green coset. Once decided, we encode this information with the two bits of the second level. At this second level the distance between different cosets is higher than in level 1; now it is “2d”. As the “cloud noise” is the same as in the previous case this means that we can use a less powerful code component than in level 1 to avoid mixing cosets in reception.

Finally let’s go to the last stage, level 3. Let’s suppose that the point to transmit belongs to the coset of red squares. We can forget about the other squares in blue, green and magenta and draw the constellation at this level shown in Figure 11.

At this point, we can see only red squares and we must encode to which specific red square our point to transmit belongs. We do that with the last two bits to transmit. As can be seen in Figure 11, the separation between dots is the greatest “4d” and in comparison with the noise cloud makes any coding unnecessary. Level 3 bits are not coded.

Altogether we have the full stream of bits to encode 10110001111…. We take the first two bits and based on these two bits we decide to which of the "shape cosets" we will eventually code our symbol (circle, square, triangle, etc.) We code these bits with strong coding. We take the next two bits and, based on these two bits, decide to which of the "color cosets" we will eventually code our symbol. Once again we code but with a slightly "softer" code. Finally we take the last three bits and we decide to which specific point of the color coset we will map our string of bits and we are done! We began with a stream of bits as the input and we have output a set of symbols separated into cosets with different levels of protection. The final full constellation with each of the symbols and its associated binary labels is shown in the Figure 12.

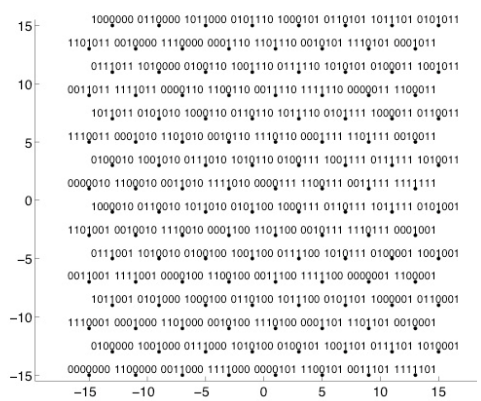

FIGURE 12. FINAL FULL CONSTELLATION

In this figure you can see 128 points in a 2D constellation. We need seven bits to encode each of these symbols: two bits belong to the first level, two bits to the second level and three bits to the third level. Now you understand how we move from bits to symbols in a multi-level coset way.

Let us briefly summarize this process.

From the binary input stream we take the first two bits and at this first level use the bits to identify the first coset. The constellation of symbols at this point is shown in the Figure 13.

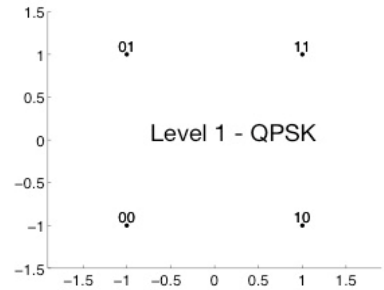

FIGURE 13. LEVEL 1 CONSTELLATION OF SYMBOLS

Next, we take the next two bits and, at this second level, partition the second coset with a similar constellation but with symbols farther apart, as shown in Figure 14.

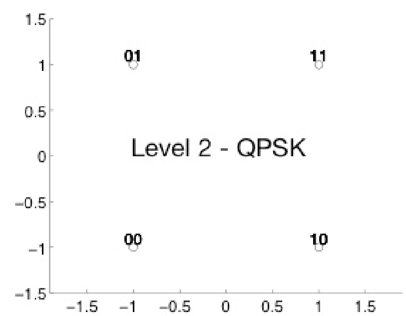

FIGURE 14. LEVEL 2 CONSTELLATION OF SYMBOLS

Finally, at the increasing third and last level we take the last three bits and partition the last coset as shown in Figure 15.

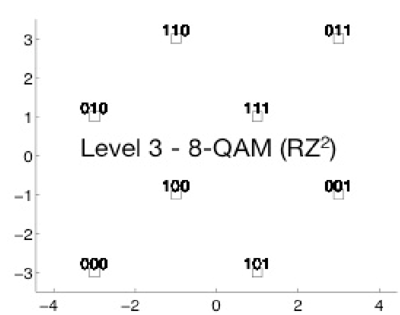

FIGURE 15. LEVEL 3

We are almost done. We add (in a vectorial fashion) all three constellations to get the overall constellation as shown in Figure 9.

Finally we use geometrical operations to get the 128 symbols 2D constellation (shown in Figure 9) prepared for the next step of the transmission system.

When handling coding and modulation this way, in steps or levels partitioning on cosets and with code components for each level, the overall coding performance, if done smartly, is optimum and more efficient than just coding the full stream of bits together and mapping them into symbols. This is the magic of MLCC!

5

MULTI-LEVEL MODULATION

For the next step we want to transform the 128 points 2D constellation from the previous MLCC step into something more suited for an optical transmission based on an LED. Remember that the LED is going to be modulated in intensity so we have just one dimension to work with. We know from earlier that this 1D LED modulation will be achieved by converting the symbols into (initially) 16 levels of intensity. This modulation is called 16-PAM. In other words we must move from a 2D constellation to a 1D constellation.

The way we do this is interleaving, or transmitting in an alternating way, the "x" and "y" coordinates of each point of the 2D constellation. An example is shown in Figure 16.

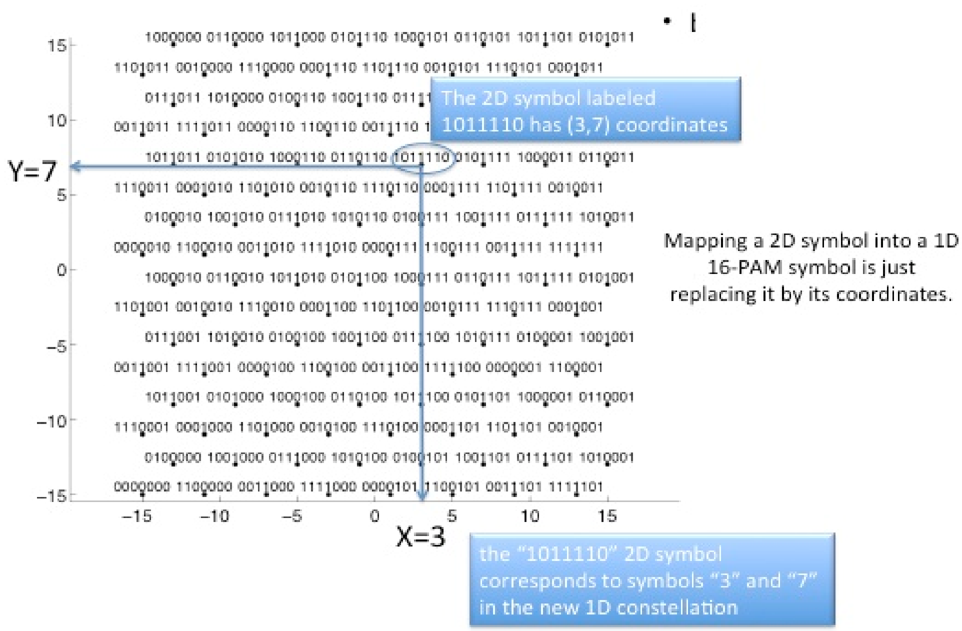

FIGURE 16. 2D CONSTELLATION

The “x” and “y” coordinates of each point in the 2D constellation can take 16 different values (only odd values from -15 to +15). By interleaving or transmitting alternately, at double rate, the “x” and “y” of each point, we will have a 1D 16-PAM modulation suitable for our purposes.

After all these transformations, we can recap a bit and do the math on symbols and bits to further clarify the concepts.

Remember that at the 2D constellation we had seven bits/symbol. Now we have interleaved each of the two dimensions so with the 16-PAM 1D constellation we end up with 7/2 = 3.5 bits/symbol. This figure might be strange for a 16-PAM modulation that requires 4 bits/symbol, but remember the way we crafted the 16-PAM constellation with an interleaving of dimensions on a 2D constellation. That is why we have ended up with a non-integer number of bits per symbol.

Furthermore, remember that we have used error correction in each of the MLCC levels, so some of the bits are used for this purpose. If we take this into account, the final figure we are transmitting for the 1 Gbps system is 3.3145 bits (of information) per each symbol sent through the channel. This figure is usually called “spectral efficiency”.

In the case of the 1 Gbps over POF system 16 levels of light are used.

This choice of 16 levels is the “sweet point” is shown in the Figure 17. Here we show a map of Optical Power budget (dBo) as a function of symbol rate (MHz).

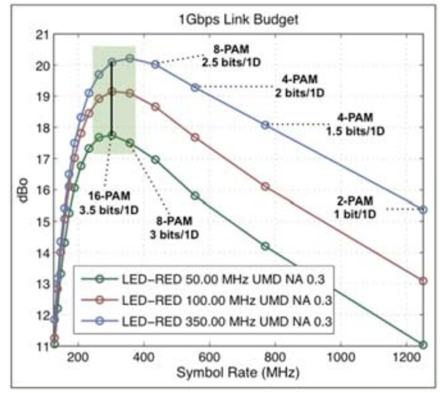

FIGURE 17. THE “SWEET POINT”

The three curves shown correspond to three different LEDs with three different electrical-to-electrical 3 dB *bandwidths. The optimum perfor*mance for the three devices, shown in the green shaded area, is attained with a 16-level modulation. At this optimum point, and taking into account the 3.3145 bits/symbol transmitted as well as the header and encapsulation overheads, the transmission of 1 Gbps of binary data translates into a symbol rate of 312.5 Mega Symbols per second (MBaud).

6 EQUALIZING THE LINEAR PART

Our channel has a response with two very different parts: one is linear and the other is non-linear. We need to compensate both effects to realize an efficient and capacity-achieving communication system.

In this section we will describe how we compensate (equalize) the linear part of the channel response. The non-linear part will be briefly mentioned in the Q&A section.

As already mentioned we use a technique code “precoding”; specifically the technique known as Tomlinson-Harashima Precoding (THP).

As mentioned before, precoding techniques equalize the channel at the transmitter side using information sent through the back channel in a non linear (modular arithmetics) way. The key point here is that THP avoids the disadvantages of DFE as no error propagation occurs and it is much easier to combine equalization with coded modulation.

Precoding is applicable only if the channel transfer function is known at the transmitter. To determine the channel response at the transmitter stage, we use physical layer headers to dynamically adapt the channel impulse response. The precoder at the transmitter side uses the estimation of the channel response made at the receiver side. This channel estimation is sent from Rx to Tx using the return channel. Once this channel response is known at the Tx side, the feedback part of the precoder filter gives the desired impulse response to compensate for the channel ISI.

Figure 18 shows a simplified diagram of the THP (Tx side).

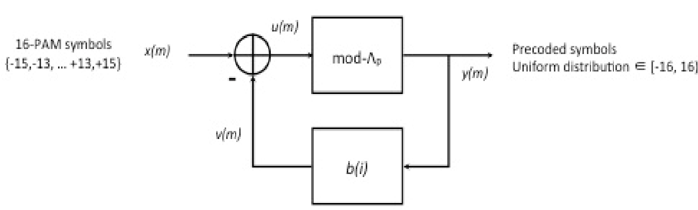

FIGURE 18. SIMPLIFIED DIAGRAM OF THP (TX SIDE)

The THP block receives the 16 PAM symbols from the previous MLCC step. The precoder filter consists of two parts. The first part is a module 2M reductor (mod-Λ_p) and the second is a feedback filter (b(i)).

As already mentioned, the coefficients of the precoder filter are dynamically adapted using the physical layer headers sent from the Rx to the Tx using the return channel. It effectively subtracts the "tail" or causal part of the ISI, (v(m)), due to previously transmitted signals from the input signal (x(m)).

The module 2M block reduces the pre-equalized symbols (u(m)) (with a discrete distribution between -15 and +15) to the continuum region (-16,16) with and almost i.i.d, uniform distribution so the power spectrum of the transmitted signal is approximately flat.

From an intuitive standpoint, THP eliminates the causal part of the ISI before transmission, in a different and better way than DFE. DFE would try to eliminate the ISI causal tail at the receiver based on hard decisions over the received symbols. DFE feedback eliminates the ISI under the assumption that previous detected symbols do not contain errors. If this is not the case, these decision errors propagate through the feedback filter with catastrophic consequences.

On the contrary, THP operates with "a priori" known symbols (not affected by any impairment) so the feedback is implemented without errors and avoids the typical DFE error propagation behavior.

7 ENCAPSULATION AND FRAME STRUCTURE

So far we have described how the digital information at the input is transformed into symbols to be transmitted in the form of light pulses over the fiber.

We have seen the need to characterize the channel periodically to allow the precoding block operation. We do this in part thanks to the back channel transmitting information to help the THP adapt itself to the channel.

All these tasks are done using "overhead" bits as opposed to "payload" or "information" bits.

- Clock frequency tolerance up to +/- 500 ppm

With the aid of Figure 19, the frame structure is described as follows:

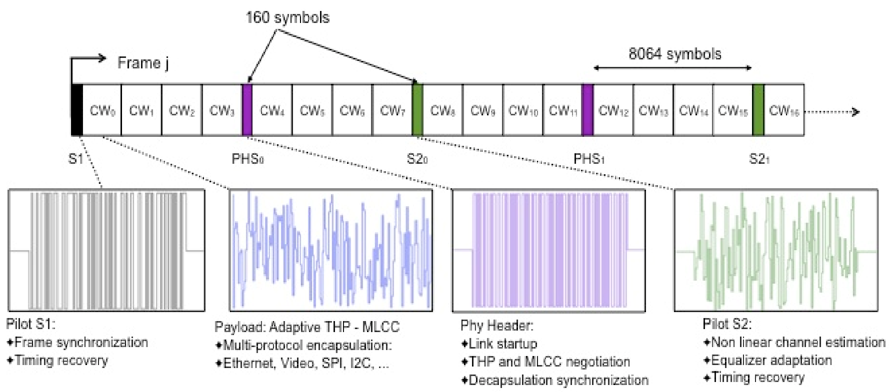

FIGURE 19. FRAME STRUCTURE

The task of the overhead bits is to help to achieve the desired communication system performance. Moreover, overhead bits also support other basic functions required in any communication system, like frame synchronization, timing recovery and link start-up negotiation.

All this overhead information along with the payload has to be encapsulated in a frame structure that is well suited for the system performance, that is, to achieve 1 Gbps with no errors and a minimum latency under any supported condition.

In summary, the encapsulation and frame scheme chosen allows:

- Continuous timing and channel estimation and equalization tracking
- Robust negotiation among link partners
- Fast link start-up
- Optimum performance for large core fibers

The payload in each frame consists in 28 chunks of Code Words (CW). Each chunk contains four CW. Each CW is 2016 symbols long. So in total we have 225792 information symbols in each frame.

The area in blue shows a typical trace of the payload symbols (random sequence from [-16,16) at the symbol frequency).

Besides the payload, every frame contains a pilot S1 training sequence at the beginning for synchronization and timing recovery. Shown in black, this sequence consists of a 2-PAM sequence of symbols.

Other pilots labeled as $S2_0$, $S2_1$, … $S2_{12}$ are sent at different places within the frame to help the equalizer adaptation and timing recovery. An example of these pilots is shown in green. Pilots consist of sequences of 160 symbols in length exercising the full [-16,16) range of the channel. The beginning and end of each pilot sequence consists on a group of 16 zeroes. This is used to accommodate the channel impulse response and to operate the THP.

Finally, each frame includes physical headers, as well that are distributed along the frame. These are labeled as PHS_0, PHS_1, … PHS_{13} in diagram. Headers support the link start-up and THP and MLCC negotiation, and synchronize the decapsulation process. A typical header chunk, shown in purple, consists, like in the PHS case, in a 160-symbol wide sequence with zeros at the beginning and end.

It is worth mentioning at this point that the physical header uses the most robust channel of the overall system as it uses a BPSK 2-PAM modulation with the corresponding BCH specially tailored for it. Headers are not precoded with THP so they are always detected in the Rx independently of the Tx state. Due to THP, headers are 6 dB in variance over the payload.

As it can be seen the frame structure chosen allows the optimal operation of the communication system giving support to the modulation and equalization schemes adopted at the same time that maximizes the payload proportion within the overall frame.

8 Q&A

Q: Why can´t I see an eye diagram when I inspect the symbols transmitted through the channel?

A: Remember that in our case we use THP precoding and as a result symbols consist of a random sequence of levels of light uniformly distributed along the [-16,16) range. When trying to visualize this signal in an eye-diagram scope set-up, you are over-imposing random continuum signals one on top of the other so there will be not any eye at all. On other communication systems based on M-PAM modulations without precoding, symbols belong to single discrete levels within the constellation set; over-imposing these symbols on the scope screen allows the visualization of eyes. To diagnose this system you will need to use tools other than eye diagrams.

Q: Why are there two data-paths in addition to the payload path?

A: Data-paths transmit pilots and the header of the physical layer. Pilots are training sequences and as such don´t need any type of encoding. On the contrary, the header needs a very robust coding and modulation scheme to ensure a correct detection. Because they have specific needs other than the payload, the pilots and header deserve their own path.

Q: How does the overall complexity of the 1 Gbps POF system compare to 1 Gbps Ethernet copper?

A: None of the techniques used in the current 1 Gbps POF system has been newly developed. MLCC, THP or PAM are all well known and have been frequently employed in wireless, mobile or copper 1 G and 10 G communications for several years. The novelty of the 1 Gbps for POF approach resides on applying a well-tailored mix of these techniques to optical communications for the first time.

Using the silicon area of its implementation, with the same production technology, as a fair measure of complexity, the 1 Gbps POF ASSP from KDPOF occupies 50% less area than a typical 1000BASE-T PHY ASSP for UTP cabling.

Q: Where, within the hierarchy of communication channels, is the 1 Gbps system positioned?

A: The system described in this document is a Layer 1 system within the OSI system interconnection model. Generically called a Physical Layer, this layer performs the following tasks:

- Defines the electrical and physical specifications of the data connection.
- Defines the relationship between the ASSP and the physical transmission medium (in our case a pair of POF). This includes optical parameters like the transmitted Optical Modulation Amplitude and extinction ratio or the receiver sensitivity, POF specifications, signal timing and more.
- Defines the protocol to establish and terminate a connection between two directly connected nodes over the POF.
- Defines the protocol for flow control.
- Defines a protocol for the provision of a reliable connection between two directly connected nodes, and the modulation or conversion between the representation of digital data in user equipment and the corresponding signals transmitted over the physical communications channel.

At this level of the OSI hierarchy the ASSP receives and sends information from the level above, Layer 2 or Data-Link Layer, which in turn provides a reliable link between two directly connected nodes by controlling the medium access and switching and prioritizing packets.

Q: Is this 1 Gbps physical layer compatible with Ethernet?

A: Ethernet is the general term to describe the suite of IEEE standards under the 802.3x denominations. In regards to the OSI hierarchy, Ethernet specifies only Layers 1 and 2 of the stack. The 1 Gbps POF system, as mentioned earlier, is a Layer 1 physical layer and has been designed and implemented to be seamlessly connected with any Ethernet Layer 2 system. The interface used to connect both layers is the set of xGMII specifications, which are defined by the industry for Ethernet connectivity. In that sense the 1 Gbps system is fully compatible with any Ethernet Layer 2 system.

Q: Is this a full duplex system?

A: A full duplex communication system is a point-to-point system composed of two connected parties or devices that can communicate with one another in both directions, simultaneously. An example of a duplex device is a telephone.

Another different but somehow related concept is the number of wires used to establish the link in a full duplex fashion. Full duplex systems can use one single cable, one pair or multiple pairs to operate. In the case of the 1 Gbps POF system the implementation uses two wires of POF to send information in both directions simultaneously. So this is a full duplex over one pair (of POF) communications system.

Q: How fast is the system able to adapt itself to channel changes?

A: The frame structure is equipped with pilots and a physical header distributed along the frame that allow the fine tuning and real-time adaptation to varying channel conditions motivated by external factors like temperature swings, fiber bending, punching, clipping or pressure, components aging, etc. With this automatic continuous channel tracking, the system is able to adapt itself in less than 50 milliseconds, which is more than good enough for any external event that might occur and affect the channel response.

Q: How fast can the system establish the link?

A: The time it takes the 1 Gbps POF system to establish the link (negotiation with the other side of the POF, channel equalization set up, etc.) is less than 50 milliseconds. In practical systems, for example on media converters, the overall system linking time also (and mainly) comprises the response of the rest of the components. For example, the copper interface and application layers typically increase the overall start up to three seconds or more depending on the components and higher level protocols used.

Q: What is the power budget of a typical 1 Gbps over 50 m of standard SI-POF?

A: For the system shown in the figure we will calculate the overall power link budget. The system comprises the PHY Tx/Rx ASSP connected to the Optical transceiver (PMD) and POF connector (MDI) at both sides of the 50 m duplex SI-POF link. We transmit 1 Gbps of payload using a symbol frequency of 312.5 MBaud.

In these conditions, and taking into account component, assembly and temperature variations, the link power budget (at an ambient temperature of 25°C) calculation is summarized in the following steps:

- On the Tx side, the minimum output power injected into the POF is -3.15 dBm.
- 50 meters of POF, depending on the specific fibre used, will introduce an attenuation that lies between 11.5 and 8.5 dBo.
- On the Rx side, right before the photodetector, we have a lens, which introduces an additional loss due to optical coupling of around 2 dBo.
- Finally, the last figure to take into account is the sensitivity at the photodetector, which lies in the range of -21 to -19 dBm.
- Adding all these figures together gives us an overall **link margin between 2.2 and 5.3 dBo**. With this margin we get a bit error rate (BER) less than 10^{-10} which is more than enough for the 50 m of a link with standard SI-POF for home networking applications.

Repeating the same calculations but at a higher temperature (70°C is the typical maximum operating temperature for home applications), the link margin reduces and lies between 2.2 and 3.2 dBo.

Q: Is the system capable of operating at speeds other than 1 Gbps?

A: The VDE DKE standard defines several annexes. The 1000BASE-P normative annex, with a fixed operation at 1 Gbps, is described in this document. However, the standard defines optional annexes as well. One of these, xGBASE-P, specifies a mode of operation called ABR (adaptive bit rate) that allows the system to operate at speeds lower than 1 Gbps. With ABR enabled, the system will detect on the fly if the measured link power budget is enough to maintain error-free communication. If the system detects that the link power margin is insufficient it will gradually reduce the link speed to accommodate the specific inferior power conditions to maintain the error free performance. All this is automatically done without any user intervention. ABR is capable of adapting the transmission speed to as low as 250 Mbps if required by the link conditions. Typically, for common commercial fibres, ABR will operate at 1 Gbps at 50 m of POF at 850 Mbps in a 75 m POF link, and will reduce the speed to 550 Mbps when the link extends to 100 meters.

In addition to 1000BASE-P and xGBASE-P the standard specifies the 100BASE-P optional annex, which is a full duplex mode but at 100 Mbps fixed bit-rate. 100BASE-P uses a lower symbol frequency (62.5 MSps) resulting in a lower spectral efficiency (1.8145 bits/Symbol). The modulation now is a 4-PAM (right before THP).

The advantage of this "long-reach" operational mode is that we trade off speed for reach. In this mode the link can be extended up to 100 m of SI-POF in a worst-case scenario and 150 m of SI-POF in a typical scenario. In the equivalent ABR version for the 100BASE-P annex, called xFBASE-P, the bit-rate can be reduced automatically from 100 Mbps to 50 Mbps.

Q: Is the system compatible with current 100 Mbps POF systems?

A: Yes. Current POF systems on the market operate at 100 Mbps using much simpler communication systems. The most common protocol used is 100 BASE-FX that uses a 4B5B and NRZI line coding.

When a 1 Gbps system is connected in a link face to face with a 100 BASE-FX system at the other side, a detection process occurs. If the 1000BASE-P detects that its partner on the other side of the link is a 100 BASE-FX system it automatically switches to this mode of operation. The result to the user will be a perfectly functional 100 BASE-FX end-to-end system in perfect operation without any interaction needed from his side.

Q: I am confused about the "user data", "net rate" or "payload speed" speeds that we are actually transmitting: 1 Gbps, 312.5 MSps, 1.25 Gbps, …?

A: Let us clarify the different reported speeds at different stages of the communication system.

When the physical layer is connected to a Layer 2 device through, for example, the RGMII Ethernet interface, it will exchange data at a binary regime of 1 Gbps. At the other side of the link the physical layer will deliver exactly the same 1 Gbps data transported over the POF. This is the "user data" or "net rate" speed; the speed the user will care about for the application and the warranted speed by the standard.

But internally, as described in this document, it's a bit more complex and we move things internally faster than this 1 Gbps.

In reality, the communication system transmits symbols and not bits over the POF (remember the 16 PAM levels of light that are latter transform in a uniform distribution of levels of lights thanks to the THP). As each symbol contains several bits of information (the 3.3145 bits/symbol figure calculated before), the rate of symbols crossing the channel is less than the rate of information bits (312.5 Million symbols per second or MBaud). If your eye was fast enough to see the light intensity changes, you would see blinking at this frequency.

If we were to speak in bits instead of symbols we could say that the system, internally, is able to move 3.3145 bits/symbol x 312.5 MBaud = 1035 Mbps. Since this is more than the 1 Gbps for which the customer has paid, why do we need this extra speed? As described earlier, it is necessary to transmit through the channel not only payload but overhead as well. All this “extra weight” uses up some of the 1035 Mbps available speed and decreases it to 1015 Mbps.

But still it seems as if we were delivering more than that paid by the customer. Not really, as this remaining 15 Mbps is used for data encapsulation which finally puts the rate at 1000 Mbps which is the net rate the user will enjoy from end-to-end of the communication system.

Please forget about 1.25 Gbps, which is a well-known and popular figure of other one Gigabit communication systems. These systems are based on interfaces likes SGMII or SerDes MAC and use a simple 8b10b line coding enabling them to operate at 1.25 GSps. In our case, as it has been described, 1.25 Gbps does not correspond to anything.

Q: What do you mean by an “error-free” communication system? Is it really a system that never will fail?

A: A physical layer is supposed to ensure reliable transmission of bits end-to-end over the communication channel. Typical 1 Gbps Ethernet systems consider a system “error-free” when the bit error rate (BER) is less than 10^{-10}. In consequence, our system is designed to deliver 1 Gbps over 50 m of POF with a BER of less than 10^{-10}. For any practical purpose this is another way to say that the system is “error-free”.

Q: Is the system open to other silicon vendors that want to develop products?

A: Yes. Any silicon vendor that wants to develop its own 1 Gbps over POF ASSP can do that by following the publicly available specification published by the German VDE DKE V 0885-763:2013-09.

Although KDPOF owns IP on key “must have” techniques, these are offered under RAND terms to any interested party. KDPOF firmly believes in fair competition as a mean to develop and grow the market to the user’s benefit. Single supply, monopolistic solutions are much less attractive to customers than well-served competitive markets, of course, if all the devices offered are compatible regardless of the vendor.

Q: Do you support LPI modes?

A: Yes, the 1 Gbps over POF system includes LPI (Low Power Idle) operation modes. In these modes the system adapts itself to the link load and reduces the power consumption accordingly. The final effect of the LPI mode is shown in Figure 20.

This figure shows the relative power consumption of the overall system (Phy ASSP plus Optical Transceiver) versus the link load (0 load with no packets travelling the POF and one load when the maximum rate is transmitted). The three curves show how

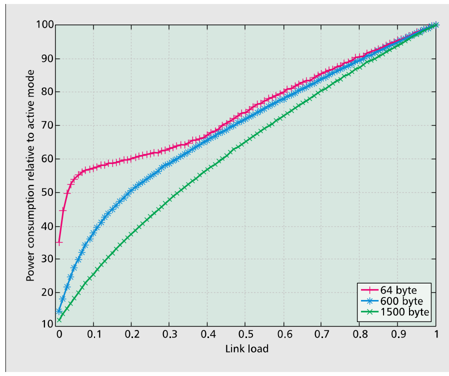

FIGURE 20. FINAL EFFECT OF THE LPI MODE

the system behaves for different packet sizes. As shown, the power consumption behaves in a fairly lineal way against the link load. There is a "fixed" term of power consumption even at 0 load, but from there on the system adjusts itself pretty well to the load of the link. Longer packets achieve better performance than shorter ones as will be explained.

As the pilots and physical layer headers are needed to keep the receiver aligned to the channel conditions, they should be transmitted even when in low-power mode. Therefore, the mechanism implemented allows stopping transmission for groups of four CWs. This is illustrated in the Figure 21.

The transitions to low-power mode (sleep), which occur when there is no user data to transmit, take place immediately after a pilot or physical layer header. Conversely, wake transitions occur just before a pilot or physical layer header. Since the receiver knows the timing of the frame, there is no need to signal the transitions other than the presence or absence of light.

The point here is that the system is awake only when it is needed to transmit payload. If payload is not present the link and all its associated power consumption go to sleep.

One last comment on LPI. In the case when the system is not transmitting for a large chunk of time and is in sleep mode, we wouldn´t be sending pilots or headers over the line with the effect of losing track of the channel (equalization, timing, etc.). This wouldn´t be a problem unless we require that the next time there is available information to send the system didn´t take a long time to again be up and running. But, if the system has lost track of the channel and the timing, for example, we need to synch again with all the time involved in this operation. To avoid this delay at start-up after sleep, the system, from time to time, wakes up itself, sends some pilots to maintain the channel and clock and goes to sleep again (this is responsible for the 10% fixed term shown in Figure 21 and for the higher efficiency of longer packets versus shorter ones). Figure 21 shows this as well.

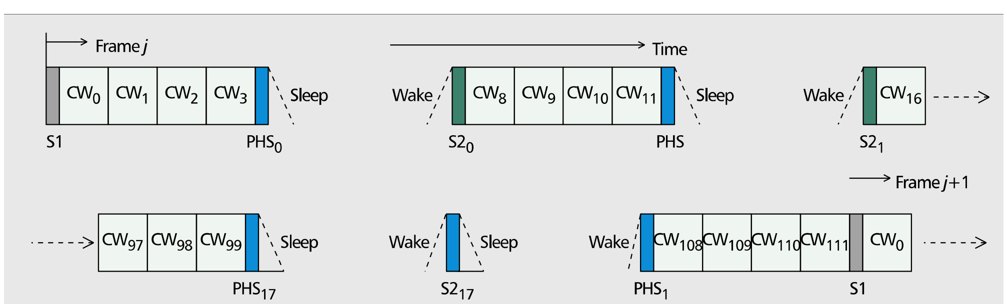

FIGURE 21. STOPPING TRANSMISSION FOR GROUPS OF FOUR CWS

FIGURE 22. SAMPLE COMMERCIAL OPTICAL TRANSCEIVER

On every S1, S2 and PHS the system wakes up transmission to keep track of the clock and equalizers. However, the price to pay for LPI is an increase in latency and jitter. During the sleep time (four CW long), the transmitter has to store any user data that might arrive from the upper MAC layer until it is allowed to transmit during the next pilot or PHS chunk. The latency, and therefore jitter, maximum increase thus corresponds to the maximum time required to transmit four CW which is around 26 microseconds when operating at 1 Gbps.

Q: Do you need to consider the LED non-linearity?

A: Yes, the LED is an inexpensive, robust and easy to operate device; however, the drawback is that it is a bandwidth-limited and non-linear device when considering its current to light intensity characteristic. Non-linearity on a component of the system at this point of the channel has a direct effect on the system performance as it translates itself directly into symbol distortion. One way to reduce distortion would be to reduce the LED operating OMA (optical modulation amplitude) but this has the side effect of reducing the power link budget. The approach to solving this issue is tricky and relies on a deep understanding characterization and modeling of the non-linearity effect to be able to compensate the non-linear channel response in an optimum way and under any device and condition.

Q: What type of optical transceiver do you need?

A: The optical transceiver is the device that interfaces the physical layer ASSP with the POF. It consists of a source of light in the Tx direction (an LED with its driver) and a light detector on its Rx side (a silicon photodiode with its trans-impedance amplifier). Figure 23 is an example of a commercial optical transceiver.

The ASSP sends the symbols in the form of electrical signals to the driver and the LED transforms the symbols into light levels. On the Rx side, at the other side of the POF channel, light pulses are transformed back into electrical signals by the silicon photodiode and amplified before being sent back to the ASSP for demodulation and decoding.

So far there doesn´t seem to be anything special or different from other current optical transceivers already available in the market for POF-based communication systems. But there is a very important point to take into account and it is that the 1 Gbps POF system transmits a continuum of light levels, or in other words, the electrical signal delivered to the LED driver spans a continuum range. The same comment applies to the receptor side and the photodiode with its amplifier; these receive and amplify continuum signals. This means that the driver and the amplifier need to be linear devices. In other words, its input-output characteristic must follow a linear curve as much as possible. As mentioned earlier, current 100 Mbps POF systems are based on NRZI line codes; that is, light is either on or off over the POF and the LED will only need to send light or shut down.

The same at reception; the photodiode will either get a light pulse or darkness. This means that neither the driver nor the amplifier need to be linear devices, but basically two-stage, non-linear devices operating from cutoff to saturation. This is very different from the linear devices required by 1 Gbps POF.

In summary, the 1 Gbps POF system requires a tailored optical transceiver. However, the photonic devices (LED and photodiode), which are the critical parts in terms of cost, reliability and qualification, are carried over from current POF systems.

Where can I procure compatible optical transceivers?

Optoelectronic vendors produce optical transceivers for POF. There are several vendors offering transceivers for 100 Mbps POF systems (non-linear devices). Due to the novelty of the 1 Gbps POF technology, only a few vendors offer commercial products for 1 Gbps POF systems (linear devices). Contact KDPOF for inquiries on the list of qualified transceiver vendors that operate with 1 Gbps ASSPs.

Please note that typically optoelectronic vendors are not ASSP vendors; if you want to build a 1 Gbps POF system you will need to procure the ASSP from a silicon supplier (like KDPOF) and the compatible optical transceiver from a different vendor. Always make sure both meet the VDE-DKE standard specification.

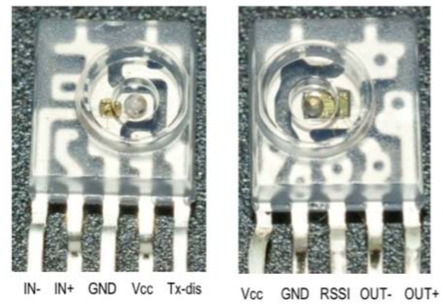

FIGURE 23. SAMPLE OPTICAL TRANSMITTER AND RECEIVER

APPENDIX A: REFERENCES

The complete standard specification for all the operating modes over POF can be found (in German) at the German VDE website:

VDE 0885-763, "Physical Layer Parameters and Specification for High Speed Operation over Plastic Optical Fibres Type HS-BASE-P," Oct. 2013.

Forney set up the mathematical formulation of MLCC in a famous paper:

G. D. Forney et al., "Sphere-Bound-Achieving Coset Codes and Multilevel Coset Codes," IEEE Trans. Info. Theory, vol. 46, no. 3, May 2000, pp. 820–50.

Tomlinson and Harashima discovered THP in parallel in the early 70´s:

M. Tomlinson, "New Automatic Equalizer Employing modulo Arithmetic," IET Electron. Letters, vol. 7, no. 5, Mar. 1971, pp. 138–39.

H. Harashima and H. Miyakawa, "Matched-Transmission Technique for Channels with Intersymbol Interference," IEEE Trans. Comm., vol. 20, no. 4, Aug. 1972, pp. 774–80.

A great review article justifying the precoding technique and describing THP was written by Forney:

G. D. Forney er al., "Combined equalization and coding using precoding", IEEE Comm. Magazine, Dec 1991, pp. 25-34

Companies need to protect their developments and findings regardless of standardization. Pérez de Aranda holds the following patents which cover key aspects of the communication system:

R. Pérez-Aranda et al., "Adaptive error correcting code for data communications over a plastic optical fibre", European Patent Application, EP2498432A1

R. Pérez-Aranda et al., "Frame structure for adaptive data communications over a plastic optical fibre", European Patent Application, EP12171346

APPENDIX B: ACRONYMS

ABR: Adaptive Bit Rate

ASSP: Application Specific Standard Product

AWGN: Additive White Gaussian Noise

BER: Bit Error Rate

BCH: Bose, Hocquenghem and Chaudhuri codes

BPSK: Binary Phase Shift Keying

DFE: Decision Feedback Equalization

DKE: Deutsche Komision Elecktrotechnik, Elektronik Informations Technik

ETSI: European Telecommunication Standards Institute

IEEE: Institute of Electrical and Electronic Engineers

I.I.D: Independent and Identically Distributed

ISI: Inter Symbol Interference

IP: Intellectual Property

KDPOF: Knowledge Development for POF

LED: Light Emitting Diode

LPI: Low Power Idle

MAC: Media Access Control

MDI: Media Dependent Interface

MLCC: Multi-Level Coset Coding

MMSE: Minimum Mean Square Error

NRZ: Non-Return to Zero

NRZI: Non-Return to Zero Inverted

OMA: Optical Modulation Amplitude

OSI: Open Systems Interconnection

PAM: Pulse Amplitude Modulation

PMD: Physical Media Dependent

POF: Plastic Optical Fibre

PSK: Phase Shift Keying

RAND: Reasonable And Non Discriminatory

RGMII: Reduced Gigabit Media Independent Interface

Rx:	Reception
RZ:	Return to Zero
THP:	Tomlinson-Harashima Precoding
Tx:	Transmission
UTP:	Unshielded Twisted Pair
VDE:	Verband Der Elektrotechnik Elektronik Informations Technik