Optical Wireless Communication Conference

Reducing Signalling Overhead and Power Consumption in OWC DCO-OFDM with Adaptive Bit Loading

J. P. M. G. Linnartz[1,2] and T. E. Bitencourt Cunha[1],
[1]Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands
[2]Signify (Philips Lighting) Research, 5656 AE Eindhoven, The Netherlands
{J.P.linnartz , t.e.bitencourt.cunha}@tue.nl, j.p.linnartz@signify.com

Post-conference version

Light Fidelity (LiFi), using light-emitting diodes (LEDs) to communicate wirelessly via light, predominantly uses orthogonal frequency division multiplexing (OFDM) with adaptive bit loading. In contrast to Radio (RF) communications, on the optical wireless communications (OWC) the gains are substantial while the signalling overhead can be limited. An example is the International Telecommunication Union (ITU) G.9991 standard. This paper shows that the bit loading scheme can be further optimized for LED communication. The LED channel is characterized by a transfer function that is predominantly declining with frequency and does not exhibit severe multipath nulls. In this paper, we describe an algorithm that allows the exchange of a signal-to-noise ratio (SNR) margin on some subcarriers, to improve the SNR on subcarriers that experiences poorer channel responses. It combines a number of subcarriers to get a more equally distributed SNRs. Subcarriers with an excess margin in SNR are used to help subcarriers with insufficient SNR.

Optical wireless communication (OWC) employs the visible light or infrared spectrum for data transmission. The light spectrum can be used license free. OWC attracted a lot of attention as it promises to become a key solution to mitigate the congestion on the radio spectrum for internet-of-things (IoT) applications. However, the light-emitting diode (LED) channel, particularly in visible light communication (VLC), exhibits a low-pass frequency response with a bandwidth that is limited by LED properties, such as the junction capacitance. Orthogonal Frequency-Division Multiplexing (OFDM), yet with a DC-bias or in a unipolar variant, is popular to accommodate the low-pass nature of the channel, as it allows the selection of an optimized power and bit density on every sub-carrier of its intensity-modulation frequency spectrum. In this way, heavily attenuated portions of the LED channel bandwidth can still be used, though with lower bit rates per Hz.

This low-pass nature of LEDs deteriorates major portions of the spectrum while these can still contribute to the throughput. OFDM allows the use of different parts of the spectrum with a specific power and signal constellation. In more heavily attenuated parts

of the spectrum, larger or lower powers can be used (i.e., power loading), and signals at higher-attenuated frequencies, which contain a relative larger amount of noise, more robust signals can be used. That is, smaller constellations, that carry fewer bits can be used at higher frequencies (i.e., bit loading).
This paper proposes a more effective use of the link-budget, thus allowing a lower power consumption or enabling a higher bit rate, it includes

- Adaptive choice of bandwidth, using a bandwidth optimized for the given link budget. We use insights from [1].
- Exchange of signal-to-noise ratio (SNR): a SNR-trading approach that combines a number of subcarriers to get more equally distributed SNRs. Subcarriers with a margin in SNR are used to help subcarriers with insufficient SNR.
- Efficient data format to identify the chosen modulation per subcarrier.

1. Power Loading Strategies

Pre-emphasized power loading attempts to invert the channel attenuation by boosting signals at high frequencies. In contrast to this, we learned from waterfilling theory that this principle is not optimum. Waterfilling is an optimum strategy for loading power in every frequency bin of the OFDM signal [2]. The number of bit/s/Hz is chosen in accordance with the SNR of the various frequency bins. Typically, waterfilling solutions invest less power at highly attenuated, thereby under-performing higher frequencies. This is thus in sharp contrast to pre-emphasis, that invests most power on underperforming subcarriers. While waterfilling is based on a theoretical Lagrange optimization, the iterative Hughes-Hartogs (HH) loading algorithm approaches this theoretical result in practical implementation. However, it is regarded as a relatively compute-intense solution. Uniform loading transmits the same power on all frequencies but only adapts the constellation. Our earlier results show that, for the LED channel, the choice of power loading does not matter so much and that a uniform loading up to a well-chosen cut-off frequency performs equally well as waterfilling [1].

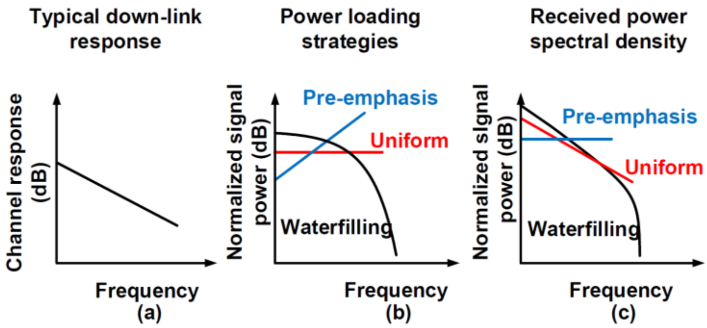

Fig. 1: Graphical comparison of three power loading strategies [1].

However, both uniform power loading and waterfilling lose 1.5 to 3 dB in performance if one cannot select a constellation of an arbitrary size. In practice, quadrature amplitude modulation (QAM) is used with an integer number of bits, (M-QAM, with $M = 2, 4, 8, 16, \dots$ bits per complex symbol) or even an integer number of bits per dimension, (square M-QAM constellations, with $M = 4, 16, 64, 256, \dots$). Round-off the number of bits per subcarrier makes 3- or 6-dB steps in effectively using the SNR response. Thus, on average, uniform power loading wastes 1.5 to 3 dB response. However, in this document we show how this loss can be repaired.

2. ITU G.9991 vs IEEE 802.11bb

These two practical strategies (uniform power and pre-emphasis) also appear to be relevant to the current debates in standardization, in particular, in ITU G.9991 [3] and in IEEE 802.11bb. In fact, the recently released ITU G.9991 (or G.vlc) applies a uniform power loading, but it allows an adaptation of the constellation size per subcarrier. In this way, the ITU standard can reuse chipsets that already exist for communication over a power line, where the main constraint is the power spectral density (PSD), typically requiring the same power on all subcarriers to satisfy electromagnetic interference (EMI) regulations [4]. In fact, power line communication (PLC) uses the mains wiring, which typically leads to unwanted radiation of the signals. As these regulations prescribe different tolerable emission spectra on different bands, a specific spectral mask is prescribed that adapts this uniform PSD depending on the frequency range. Thus, deviating from a uniform PSD in some frequency ranges where more stringent requirements apply. The algorithm then dynamically assigns more or fewer bits to each subcarrier. The system is robust against a low pass filtering by the LED and / or optical indoor multipath by assigning fewer bits at high frequencies.

Alternatively, IEEE 802.11bb tends to opt for reusing its legacy physical layer (PHY), which was designed for RF systems that do not require adaptive power loading. In fact, deep RF fades are rare and narrow in radio multipath reception. This is particularly the case with a sufficiently long delay spread. So, the approach is to overcome fades by appropriate error correction coding, combined with interleaving, with a fixed bit constellation on all subcarrier frequencies. Yet, when used over an LED channel, this approach is not so effective [1, 5], while adaptive bit-loading can more effectively guarantee adequate throughput. In fact, a pre-emphasis is needed to avoid that highly attenuated high frequencies force the system to use only small constellation sizes. Adaptive bit loading is more effective if wide portions of the channel are heavily attenuated.

3. Overhead in signaling bit loading profile

In RF communication, adaptive bit loading failed to be adopted in many standards. It would be very inefficient, for every subcarrier to describe separately which constellation is used. As systems typically use 512, 1024 or even 2048 or more subcarriers, it may require excessively many parameters to be exchanged. This may have been one of the reasons that historically prohibited the use of adaptive bit loading in radio frequency (RF) systems. In RF systems, adaptive bit loading is less a necessity, because typical fades experience in RF communication are not very wide, thus only affect a few subcarriers. Error correction, particularly in combination with bit interleaving adequately solves the problem of fades in most RF systems.

However, for OWC the typical transfer functions have wide bandwidths of deep attenuation. These can be caused by a decline of responsivity of the LED at high modulation frequencies. Another cause of wide fades can be multipath reception with a short delay spread, as for instance seen in indoor communication in which the direct path and the reflected path only have a minor length difference. Thirdly, if multiple light sources carry the same data, but are connected via cables of slightly different length, say a few meters differences, a wide null may occur at modulation frequencies of a few tens of MHz.

These can cause a problem in LiFi using LEDs, as wide fades cannot easily be resolved by error correction coding, because too many subcarrier signals may be lost to recover the signal. Adaptive bit loading, in which these poorly performing subcarriers, i.e., those in a fade, are skipped or modulated by only a small, thus robust constellation.

3.1. Data Format for how subcarriers are loaded

It is inefficient in terms of communication signaling overhead, describe for every subcarrier which constellation is used. Unless a smarter optimization is used, the transmitter must inform the receiver about the constellation size and code rate used on every subcarrier. This requires a few bits per subcarrier for such overhead, while every subcarrier itself can only carry a few (1 to, say 12) bits. Bit strings to describe the bit loading are preferably included in the header of a sequence of OFDM blocks.

We exploit the fact that the LiFi channel is a monotonously declining function, and that many neighboring subcarriers have the same constellation. In order to simplify the overhead in signaling what constellations are used, ITU G.9991 already uses a grouping of neighboring subcarriers, that all have the same constellation. Then, a single identifier can be used to communicate the constellation for all subcarriers that are part of the group. Yet, further identifiers are needed to describe which subcarriers belong to the group. Here, one can exploit the monotonously declining frequency response of the LED.

ITU G.9991 already describes "grouping" of neighboring subcarriers that use the same constellation. Yet, ITU G.9991 does not exploit the property that on an LED channel, constellation sizes mostly monotonously decrease as illustrated in Fig. 2.

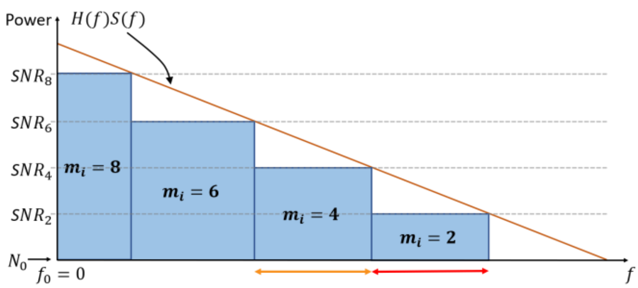

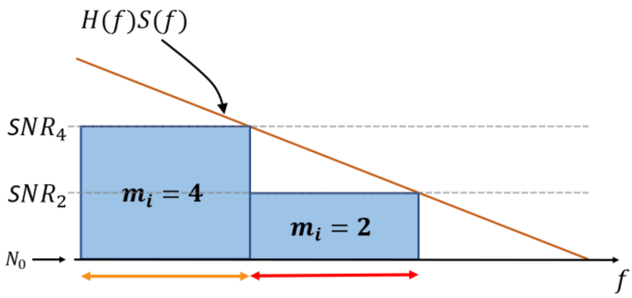

Fig. 2: Regardless of the SNR, for a typical LED profile, the number of subcarriers that can carry 2 bits is a system constant (red arrow bandwidth). Similarly, a constant number of subcarriers can carry 4 bits (orange arrow bandwidth).

In particular, for an LED response that declines exponentially with frequency, which approximates well many measured LED channels. If we also consider gentle multipath and PD capacitances at the detector, further monotonically declining effects are added.

Using the commonly reported LED channel, we have noticed a particular property: the number of adjacent subcarriers that uses a particular constellation is constant, irrespective of the link budget. In other words, one common pre-agreed profile can describe the bit loading used, while, as the signal-to-noise level changes, this profile is shifted up or down in frequency. This leads to a very effective, highly compressible data format, and simple (one-bit) updates if during operation of the link SNR changes gradually.

4. Exchange of SNR scheme

In a system with uniform power loading, the power per subcarrier does to exactly match the SNR required for a constellation with an integer number of bits. Some subcarriers have an excess SNR, that ranges from 0 to (3 or 6) dB, depending on the set of constellations used in the system. This headroom in SNR is not used to effectively achieve the maximum throughput as no suitable constellation exists or can be used that fully utilizes the available SNR.

As we will show next, this inefficiency gap can be repaired. We show that the excesses of SNR on some subcarriers can be exchanged with other subcarriers, improving performance. This allows subcarriers which have just-too-little SNR to be boosted to reach an adequate SNR. This can be done without changing the power spectral density of the subcarriers, thus without deviating from certain power loading.

It appears possible to exchange an excess in SNR on one frequency to repair a shortage in SNR at another frequency. A special form of multiplexing can be used. In a basic form, one can transmit the sum two data symbols on one frequency and the difference of the data symbols at another frequency. That is, a simple pre-processing before the transmission on the i −th fast Fourier transform (FFT) is proposed. If the receiver does the inverse, the original symbols are recovered, but the noise level is averaged over the two data symbols.

These sums and differences can be generalized as a unitary matrix operation over N incoming data symbols, into N outgoing subcarriers. Preferably the matrix just contains +1 (addition) and -1 (subtraction) weight factors and more preferably a Walsh Hadamard (WH) Matrix is used.

5. Formulation of a model

For an OFDM system with channel $H(f)$, the system can reliably decode this constellation on the i-th subcarrier at frequency f_i if the channel attenuation $H(f_i)$ ensures that the SNR suffices, where i ranges from $0, 1,$ to $N-1$ with N being the size of the FFT block. In a typical monotonically deteriorating LED channel, that also means that at any lower frequency $f < f_i$ reliable detection is possible.

The transmitter is aware about how $H(f)$ and the SNR deteriorates with frequency. To this end, a feedback mechanism is used. The receiver estimates the channel response $H(f_i)$ as $\widehat{H}(f_i)$. As we may assume that, for proper operation of the receiver, in good accuracy $\widehat{H}(f_i) = H(f_i)$. In an OFDM receiver, before slicing, i.e., before quantizing the noisy received signal, the receiver equalizer normalizes the signal by multiplying every subcarrier signal by $1/\widehat{H}(f_i)$.

We use the following notation: at the i-th subcarrier, for a QAM constellation to carry m_i bits, thus for $M = 2^{m_i}$ possible symbol values (2^{m_i}-QAM), a SNR equal to SNR_{m_i} is needed. The coefficient d_i corresponds to the QAM symbol that is going to be transmitted on the i-th subcarrier. In our proposed system, we will mix signals of a set multiple subcarriers to create a new set of symbols, with the same average power.

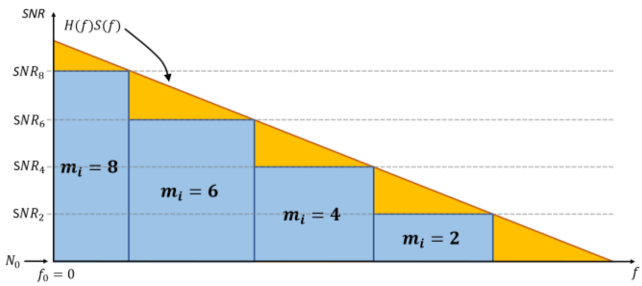

Fig. 3: Typical gradually decreasing received signal power, due to declining $H(f)$. Conventional bit loading is illustrated. Orange triangles: Unused signal power.

The original SNR profile is shown in the upper drawing as an orange declining line. The blue blocks indicate the bit loaded achieved by ITU G.9991. An improvement here, can be to "trade and exchange SNR" between subcarriers. The underlying idea can be illustrated in an abstract fashion as follows

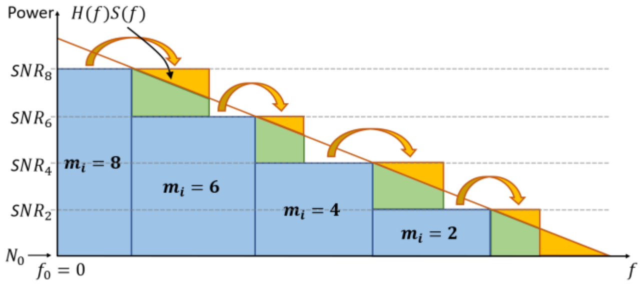

Fig. 4: Illustration of the trade and exchange SNR scheme.

The green area plus the white triangles indicates unused power. Our algorithm intends to hand over power associated with the white triangles into the location of the orange triangles. This allows a better bit loading, including blue, green, and orange areas. This is redrawn and shown in Fig. 5.

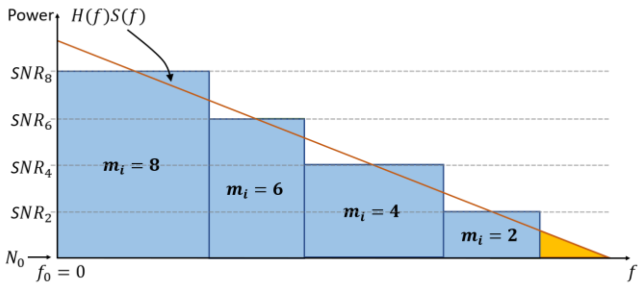

Fig. 5: Illustration of the achieved bitloading distribution after trade and exchange SNR scheme.

As can be observed in Fig. 5, the trade and exchange SNR scheme allows for better bit-loading distribution over the LED channel. Unused excess of SNR on higher subcarriers can be combined with the unused SNR of lower frequency subcarriers. in order to boost bit allocation and to increase data rate.

5.1. A 2-by-2 SNR trade and exchange scheme

In its simplest form, the system takes two (complex valued) subcarrier signals d_1 and d_2 and transmits the sum of the two data symbols $(d_1 + d_2)/\sqrt{2}$ at subcarrier 1, with an amplitude renormalization of $\sqrt{2}$ to retain the same power, and it transmits the difference of the two data symbols $(d_1 - d_2)/\sqrt{2}$ at subcarrier 2. The receiver first applies the per-subcarrier channel equalization that is common practice for OFDM. Then, the two subcarrier signals are added up to recover the first data symbol d_1. The two subcarrier signals are subtracted to recover the second data symbol d_2. The signal-to-noise levels then become the average of the signal-to-noise levels at the two frequencies. Thus, an excessively large SNR on a lower subcarrier can be used to boost the SNR at a higher frequency. Both symbols can be taken out of the same size of a constellation. A block diagram of the procedure is shown in Fig. 6.

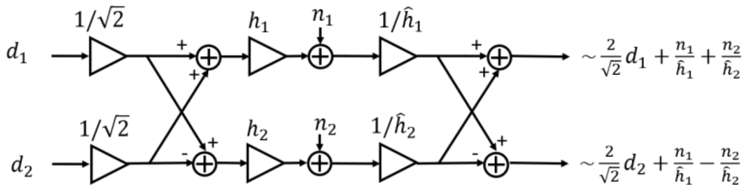

Fig. 6: Block diagram of the 2 by 2 exchange scheme.

The "SNR exchange" idea can be extended across a block of subcarriers. A basic approach is to do a 2-by-2 exchange, but also larger blocks can be used. As depicted in Fig. 7, an example of the concept is to share the subcarriers two by two, round a central frequency where the SNR is just adequate.

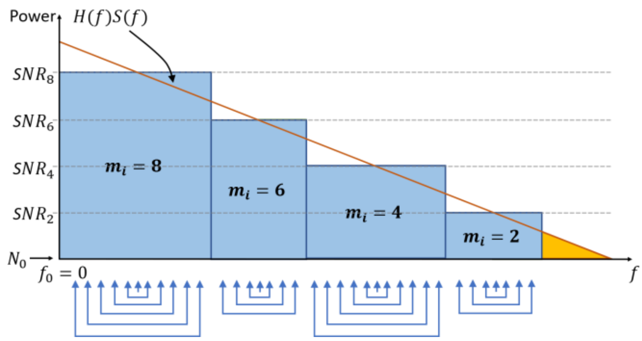

Fig. 7: Illustration of the 2-by-2 SNR trade and exchange scheme . Regions where $m_i = m_2, m_4, m_6$ and m_8, are at frequencies at which the SNR is just adequate for 2, 4, 6 and 8 bits.

In this case, the signal on the next lower subcarrier $f_{c,m_i} - k\Delta_B$and the next higher subcarrier $f_{c,m_i} + k\Delta_B$ are combined, where f_{c,m_i} is the central frequency of the set of subcarriers whose are able to carry m_i bits, $k \in [1, \dots, (N_{m_i} - 1)/2]$ and Δ_B corresponds to the subcarrier bandwidth. The coefficient N_{m_i} represents the number of subcarriers in the set. Similarly, farther sets of signals are combined, two by two.

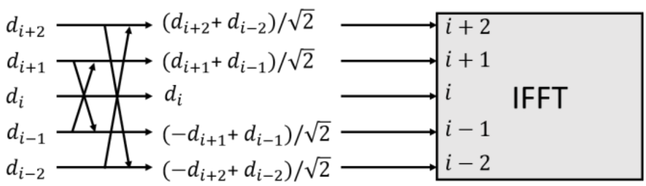

Fig. 8: Two-by two subcarrier multiplexing (adding and subtracting), before the OFDM inverse FFT (IFFT) block.

The 2-by-2 multiplexing serves the purpose, but it has disadvantages: firstly, the normalization by multiplication by $1/\sqrt{2}$ requires a full multiplication. Secondly, we see some loss in SNR. The two combined subcarriers do not exactly give the same SNR as the SNR at the central subcarrier. Both disadvantages can be mitigated by using 4-by-4 combining.

5.2. A 4-by-4 SNR trade and exchange scheme

The combining can be done with 4 data symbols, spread over 4 subcarriers.

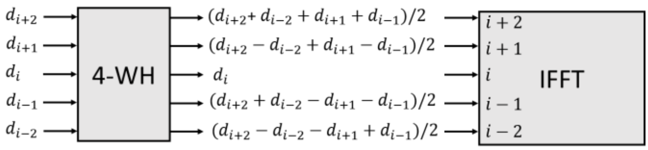

Fig. 9: 4-by-4 subcarrier multiplexing.

In the 4-by-4 trade and exchange scheme we use a Walsh-Hadamard Matrix of size 4, as given by

$$\mathrm{WH}=\frac{1}{2}\begin{bmatrix}1 & 1 & 1 & 1\\ 1 & -1 & 1 & -1\\ 1 & 1 & -1 & -1\\ 1 & -1 & -1 & 1\end{bmatrix} \tag{1}$$

where each column represents the weight factor for the fourth data symbols. Every subcarrier has its own specific column to generate its payload. That is, one subcarrier carries the sum of a data symbols, another one the sum of all signal with alternating signal, a third subcarrier carries the sum of the first two symbols minus the sum of the last two symbols, and a fourth subcarrier carries the four symbols weighted by 1, -1, -1, 1.

6. Performance Analysis

We denote the channel response at the *i-th* subcarrier by $H(f_i)$ and the response as estimated by the receiver by $\widehat{H}(f_i)$. As we may assume that for proper operation of the receiver, in good accuracy $\widehat{H}(f_i)\sim H(f_i)$. Yet, for intuitive understanding of the performance, we use separate symbols. For the 2-by-2 trade and exchange scheme, it is possible to compute the SNR over a set of subcarriers able to carry m_i bits as

$$SNR_{m_i,2x2}=\frac{2P}{\left(\frac{1}{\widehat{H}^2\left(f_{c,m_i}-\Delta_B\right)}+\frac{1}{\widehat{H}^2\left(f_{c,m_i}+\Delta_B\right)}\right)\sigma_n^2} \tag{2}$$

The frequency transfer function of an LED channel decreases with frequency. In this work we considered an exponential model, such as [1] :

$$|H(f)|^2 = H_0^2\mathrm{e}^{-\ln(2)\frac{f}{f_0}} \tag{3}$$

Where the coefficients H_0 and f_0 are, respectively, the DC channel gain and the LED 3-dB cut-off frequency. By inserting (3) in (2), we obtain

$$SNR_{m_i,2x2} = \frac{2\,H_0^2 P}{\left(e^{\ln 2 \frac{f_{c,m_i}+\Delta_B}{f_0}} + e^{\ln 2 \frac{f_{c,m_i}-\Delta_B}{f_0}}\right)\sigma_n^2} \quad (4)$$

After a math simplification, it possible to relate (4) with the SNR at the central frequency f_{c,m_i} such as,

$$SNR_{m_i,2x2} = \frac{SNR(f_{c,m_i})}{cosh\left(\frac{\ln(2)\,\Delta_B}{f_0}\right)} \quad (5)$$

This is somewhat less than $SNR(f_{c,m_i})$. For better understanding of (5), we go further. The exponentially declining $H(f)$ allows us to write $H^2(f-x) = \alpha H^2(f)$ and $H^2(f+x) = H^2(f)/\alpha$ with α being some value close to unity, that depends on x. To quantify the relation between the frequency separation between neighboring subcarriers and α, we have

$$\boldsymbol{\alpha} = \frac{\boldsymbol{H^2(f-\Delta_B)}}{\boldsymbol{H^2(f)}} = \mathbf{e}^{\frac{\mathbf{ln(2)}\boldsymbol{\Delta_B}}{\boldsymbol{f_0}}} \quad (6)$$

Or equivalently, $\Delta_B = \ln(\alpha) f_0/ln(2)$. Then, it is possible to rewrite (2) as

$$SNR_{m_i,2x2} = \frac{2\alpha}{\alpha^2+1} SNR(f_{c,m_i}) \quad (7)$$

Filling in some number shows for this is a significant improvement: It gives, after the 2-by-2 SNR exchange, an SNR deterioration of only 5.5% for write $\alpha = 1/\sqrt{2}$, while the SNR without exchange varies from 70% to 141% (1.5 dB variation). This is typical for a constellation step of 3 dB.

In a 4-by-4 SNR-exchanging system, the new SNR becomes

$$SNR_{m_i,4x4} = \frac{4}{\alpha^2+\alpha+\frac{1}{\alpha}+\frac{1}{\alpha^2}} SNR\left(f_{c,m_i}\right) \quad (8)$$

For a fair comparison with the 2-by-2 strategy, we consider the same total span for both 2-by-2 and 4-by-4 strategies. In this case, we get $\alpha^2 = 1/\sqrt{2}$ for the 4-by-4 SNR-exchanging system, and this gives an SNR deterioration of 3.6%. Thus, the 4-by-4 SNR-exchanging system is more effective than the 2-by-2 exchange.

In an example, as we take constellations with an integer number of bits per subcarriers (including odd integers) and if a 3 dB step per required SNR is made, then the SNR-exchange algorithm operates up to 1.5 dB difference from the central frequency. This leads to an SNR reduction of 0.11 dB and 0.07 dB for a 2-on-2 and 4-on-4 exchange, respectively. It corresponds to frequency shift of about 2%. In other words, the new SNR, after processing is in very good approximation identical to the SNR of the subcarrier in the center of this range.

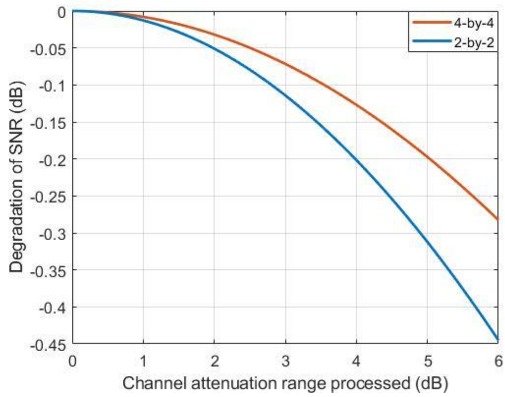

Figure 10: Degradation of SNR (y axis) in dB, versus the attenuation at which the highest subcarrier is taken for a size 2 and size 4 WH matrix. Typically, constellations can be taken in 3 dB steps, then the SNR for the highest subcarrier in a group is 1.5 dB worse than in the center of the group.

6.1. Exchange over many subcarriers

The transmitter spreads a single data symbol over N_k subcarriers, where N_k is substantially smaller than the FFT size (N_k < N). To keep the transmit power unchanged, every participating subcarrier get the symbol attenuated by $1/\sqrt{2}$. At the receiver, the N_k subcarriers are recombined coherently, which amplifies the received signal by a factor N_k. To recover the same received signal strength, we multiply the receiver again by $1/\sqrt{N_k}$. Thus, the received SNR becomes

$$SNR = \frac{\left(\frac{1}{\sqrt{N_k}\sqrt{N_k}}\right)^2 \left(\sum_{i=1}^{N_k} \frac{H(f_i)}{\widehat{H}(f_i)}\right)^2 P}{\left(\frac{1}{\sqrt{N_k}}\right)^2 \left(\sum_{j=1}^{N_k} \frac{1}{\widehat{H}^2(f_{i-1})}\right) \sigma_n^2} = \frac{N_k}{\sum_{j=1}^{N_k} \frac{1}{H^2(f_{i-1})}} \frac{P}{\sigma_n^2} \tag{9}$$

In fact, these expressions resemble the SNR averaging in Multi-Carrier CDMA [6, 7]. Here, we use these WH principle over a well-chosen subset of the subcarriers and with the purpose of overcoming the penalty cause by the use of integer constellations.

7. Conclusion

We identified improvements for G.9991 that significantly enhance the power efficiency, or that can be used to increase bit rate. In contrast to tracking the rapid multipath fading in RF, in OWC the overhead on adaptive bit loading can be contained to a reasonable amount. Thereby one can get the benefits, which are high on OWC LED channels. An exchange of SNR can be used for a LiFi System with a low pass frequency response and specifically exploit excessive SNR on some subcarriers that would otherwise be wasted.

The LED channel gives quite severe attenuation at higher frequencies. That differs from a typical radio (Rician) multipath channel. We show that for the LED channel, uniform power loading can improve its link budget by 1.5 to 3 dB by using the SNR trade and exchange approach.

References

[1] S. Mardanikorani, X. Deng and J. P. M. G. Linnartz, "Sub-Carrier Loading Strategies for DCO-OFDM LED Communication," *IEEE Transactions on Communications,* vol. 68, no. 2, pp. 1101-1117, 2020.

[2] R. G. Gallager, Information Theory and Reliable Communication, New: Wiley, 1968.

[3] International Telecommunication Union Recommendation G.9991-202104 Amd.2, "High-speed indoor visible light communication transceiver – System architecture, physical layer and data link layer specification," 2021.

[4] Maxlinear, "88LX5152, 88LX5153 Wave-2 G.hn Digital Baseband (DBB) Processor," 14 July 2020. [Online]. Available: https://assets.maxlinear.com/web/documents/88lx5152_88lx5153.pdf. [Accessed 14 12 2021].

[5] S. Mardani, X. Deng and J. P. Linnartz, "Efficiency of Power Loading Strategies for Visible Light Communication," in *2018 IEEE Globecom Workshops (GC Wkshps)*, 2018.

[6] N. Yee, G. Fettweis and J. P. M. G. Linnartz, "Multi-Carrier CDMA in indoor wireless Radio Networks," in *IEEE Personal Indoor and Mobile Radio Communications (PIMRC) Int. Conference*, Yokohama, Japan, 1993.

[7] J. P. Linnartz, Wireless communication: the interactive multimedia CD-ROM, Kluwer Academic Publishers, 2001.