When Does Unequal Power-loading Make Sense?

Sarah Kate Wilson
Department of Electrical Engineering
Santa Clara University
Santa Clara, CA 95053
Email: skwilson@scu.edu

Milica Stojanovic
Department of Electrical and Computer Engineering
Northeastern University
Boston, MA
Email: millitsa@ece.neu.edu

***Abstract*—In Orthogonal Frequency Division Multiplexing (OFDM) unequal power-loading, i.e. water-filling, is used to selectively assign power to each subcarrier, thus maximizing the capacity for a fixed power-budget. Unequal power-loading is used successfully in applications such as digital subscriber lines (DSL), where the channel effectively acts as a time-invariant low-pass filter. In such conditions, the receiver can send the transmitter timely and lasting information about the channel. In contrast, wireless channels suffer from time-variability, which implies outdated feedback and inevitable channel estimation errors. This paper reexamines the power-loading methods proposed for time-varying wireless channels, and determines regions of Doppler and delay spreads where this approach has benefits.**

I. Introduction

Unequal power-loading is a technique for optimizing the rate of an Orthogonal Frequency Division Multiplexing (OFDM) signal in frequency-selective channels. It is based on the water-filling solution for capacity [1]. If the transmitter knows the receiver's channel, the water-filling solution dictates how it can allocate power to each subcarrier based on its strength. Unequal power-loading has been used successfully in digital subscriber line (DSL) transmission to increase the bit rate. There have been several attempts to apply power-loading to wireless channels, e.g. [2], [3], but the major practical difficulties stem from the fundamental difference between DSL channels and wireless channels. Namely, DSL channels are relatively static and inherently low-pass, with a large spectral variability. Wireless channels rarely exhibit such properties. Generally they are not static due to a moving transceiver or a changing environment and their frequency response changes with the multipath composition. However, the frequency selectivity that results from multipath does not often cause the deep and wide fades found in DSL channels.

This paper looks at the nature of wireless channels and the implications it has on power-loading. We first present some background material on unequal power-loading and discuss water-filling and related implementation issues. We then discuss the likelihood that a wireless channel will experience a deep and wide fade within the allocated spectrum. Finally we present rate results for frequency selective channels with channel estimation errors; examining the effect of channel estimation on rate and its impact on the effectiveness of power-loading at the transmitter. We show that in general, the improvement due to unequal power-loading is minimal.

II. Power-loading and the nature of wireless channels

With an OFDM system, power-loading according to the water-filling principle is specified by [1]

$$P_k = \max\left(\lambda - \frac{\sigma^2}{|H_k|^2}, 0\right) \tag{1}$$

where P_k is the power on the k^{th} subcarrier, H_k is the frequency response of the channel on the k^{th} subcarrier, and σ^2 is the variance of the additive noise. In addition, λ is chosen so that $\sum_{l=1}^{N} P_l = P$, the power budgeted to the system.

As has been noted by others, including [2], [3], as the SNR increases, the power allocated to each subcarrier approaches a constant value. An increase in the total power – with its consequent increase in the SNR – leads to an increase in the water-level, λ. In turn, as λ increases, the effect of the term $\frac{\sigma^2}{|H_k|^2}$ is minimized. The bit rate on each subcarrier is then upper-bounded by

$$\begin{aligned} b_k &= \log_2(1 + P_k \frac{|H_k|^2}{\sigma^2}) && (2) \\ &= \begin{cases} \log_2(\lambda \frac{|H_k|^2}{\sigma^2}) & \text{if } P_k > 0 \\ 0 & \text{otherwise} \end{cases} && (3) \end{aligned}$$

A. Toy example

In the power-loading equation (1), when $(\lambda - \frac{\sigma^2}{|H_k|^2}) < 0$, the subcarrier has zero power and is not used. When a subcarrier is not used, its power is redistributed to the other subcarriers. For example, if there are 4 subcarriers, with the SNR of $1, 2, 1$ and 0 and the power budget is 4 Watts, then λ is determined from

$$3\lambda - 1/1 - 1/2 - 1/1 = 4$$

which yields $\lambda = 13/6$ and power allocations of $7/6, 10/6, 7/6$ and 0. (Note that regardless of the value of the power budget, the fourth subcarrier, in this example, will always be zeroed out.) One quarter of the power has thus been transferred to the other subcarriers. Given this redistribution of power, how much has been gained? If we had uniform power distribution on all subcarriers, the effective bit rate would be on the order of

$$\begin{aligned} b_{unif} &= 2\log_2(1+1) + \log_2(1+2) \\ &= 3.585 \text{ bits} \end{aligned}$$

With water-filling it would be

$$\begin{aligned} b_{power} &= 2\log_2(1+7/6)+\log_2(1+2\times 10/6) \\ &= 4.35 \text{ bits} \end{aligned}$$

which is a rather large increase. In this toy example we used the standard water-filling solution which redistributed the power of the unused subcarrier based on the remaining subcarriers' SNR. Had we used the simpler on-off water-filling method [4], which redistributes the power equally among the remaining subcarriers, the gain in bit rate would have been almost as much. It is the act of power redistribution – not so much *how* it redistributed – that provides the bulk of the benefit from power-loading. This can be seen by noting that

$$\begin{aligned} b_{redisr} &= 2\log_2(1+4/3)+\log_2(1+2\times 4/3) \\ &= 4.32 \text{ bits} \end{aligned}$$

This provides an example of how on-off water-filling is very nearly as effective as standard water-filling.

Our toy example's effective gain in bit rate has, however, been exaggerated relative to what can be expected in a wireless radio channel. We redistributed the power from one quarter of the subcarriers and increased the power of the remaining subcarriers by a factor of $1/3$, a more significant increase than can be expected in wireless radio channels. In most wireless channels the fraction of subcarriers that will have zero power is small and the amount of power available to redistribute will be too little to significantly increase the bit rate.

One can estimate this effect by noting that for higher SNR, $\log_2(1+\beta P_l SNR_l) \approx \log_2(\beta P_l SNR_l) = \log_2(\beta)+\log_2(P_l SNR)$ where β is the increase in power. If only 5% of the subcarriers are zeroed out and power is reallocated, then there is an increase per subcarrier of less than a tenth of a bit.

B. *Practical limitations of power-loading*

Water-filling is a fluid method of assigning power that does not take into account the fact that power allocation must adapt to a fixed number of bits per symbol. In addition, there is a back-off factor that compensates for such factors as mis-estimation [1]. In our toy example, $7/6$ Watts were allocated to two subcarriers, while $10/6$ Watts were given to the third subcarrier. The respective SNRs on the two subcarriers is 0.7 dB, while the SNR on the third subcarrier is 5.2 dB. In practice these values would be adjusted to accommodate a specific constellation size. To accommodate a round number of bits on each subcarrier, power must be redistributed from one subcarrier to another. Because the SNR on two of the subcarriers is so small, all the power could be added to the subcarrier with the largest SNR. For this case the resulting bit rate would be $\log_2(1+4\times 2) = 3.2$ bits, which is less than the original equal power-loading rate. Hence small variations in power across subcarriers dictated by water-filling may be evened out when practical constellation assignment is taken into consideration. This further reinforces the notion that on-off water-filling [4] is more practical than fluid power allocation.

C. *Nature of the wireless channel*

In studies such as [4], [2], [3], power-loading benefits are shown for channels with an SNR of less than 0 dB, rather than larger SNRs. This is due to the fact that the variation in the channel determines the extent of the effect of power-loading. Given a multi-tap channel with impulse response

$$h(t) = \sum_{l=1}^{L} \alpha_l \delta(t-\tau_l)$$

where α_l are independent, zero-mean complex Gaussian random variables with variance $E[|\alpha_l|^2]$, the mininum value of the channel strength, $|H_k|^2$ is upper-bounded by $(\sum_{k=1}^{L}|\alpha_k|)^2$ and lower-bounded by some combination of the square of the sum of $\pm|\alpha_k|$s. This lower bound depends on the value of the channel taps, their relative phases and positions.

For a two-tap channel, the square of the frequency response $|H_k|^2 = |\alpha_1|^2+|\alpha_2|^2+2|\alpha_1||\alpha_2|\cos(2\pi(\frac{(\tau_2-\tau_1)k}{NT}-\theta))$ where θ is the angle between the complex tap values. A two-tap channel can exhibit a deep fade when taps have comparable magnitudes. Unequal power-loading will then be effective. However, this will occur only when $|\alpha_0| \approx |\alpha_1|$ and $\tau_2-\tau_1$ is such that $\frac{(\tau_2-\tau_1)k}{NT}$ is an odd integer. How often will this happen in a multi-tap channel?

To answer this question in general, we consider the likelihood that a given number of subcarriers will have no power allocated to them. It is convenient to find this in terms of the average gain of the channel, $M_h = \sum_{k=0}^{N-1}|H_k|^2$. For a given λ, subcarriers will be assigned zero power if $|H_k|^2 < \frac{\sigma^2}{\lambda}$. By setting $\lambda = \frac{\sigma^2}{AM_h}$, where $A<1$ is a constant, then all subcarriers where $|H_k|^2 < AM_h$ will be assigned zero power. For example, if $A = 1/2$, then all subcarriers where $|H_k|^2$ is 3 dB below the mean gain will have zero power.

Given this value for λ, the SNR on the subcarriers with non-zero power will be

$$\begin{aligned} SNR_k &= (\lambda - \frac{\sigma^2}{|H_k|^2})\frac{|H_k|^2}{\sigma^2} \qquad (4) \\ &= \frac{|H_k|^2}{AM_h} - 1 \qquad (5) \end{aligned}$$

This means that after water-filling, the average SNR on the remaining subcarriers will be

$$\text{SNR}_{\text{ avg of rest}} \geq \frac{1}{A} - 1 \qquad (6)$$

The question is: what percent of subcarriers are zeroed out in wireless channels? The probability that a given fraction of the subcarriers are less than a specific ratio of the mean of the gain of the frequency response can be calculated by simulation. For this simulation, 10^4 channel realizations were used to calculate each probability that a specific number of subcarriers was less than $A \times \text{SNR}_{avg}$. There were 128 subcarriers; the channel has 10 independent taps with an RMS delay spread of 5% of the payload of the OFDM symbol, i.e. 128 samples.

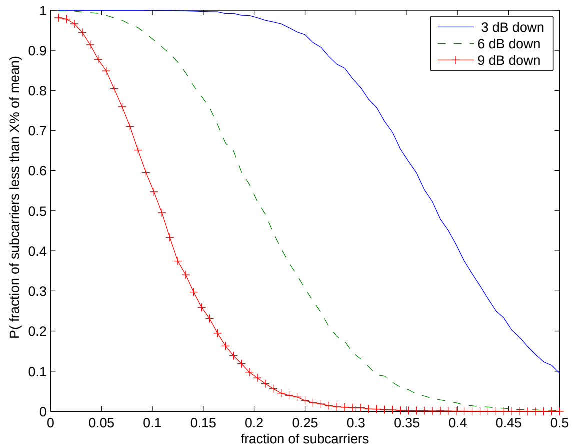

Fig. 1. This plot shows the probability that a fraction of the subcarriers are X dB or more below the mean gain of the frequency response.

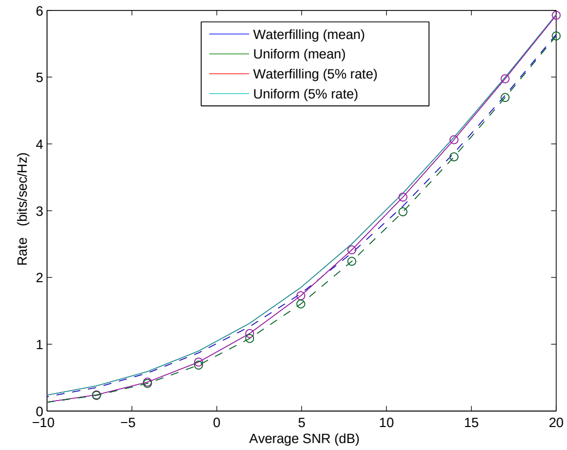

Fig. 2. The mean and 5^{th} percentile of the rate of a 128 subcarrier OFDM signal with water-filling and with uniform power distribution. The channel has an exponential RMS delay spread that is 5 % of the length of OFDM symbol with 16 sample-spaced taps. The mean and percentile value are taken over 10^4 channel instances. The mean capacity for water-filling ranges from .4 bits at -7 dB, to 1 bit at 0 dB. At 20 dB, the mean capacity for water-filling is approximately 6 bits .

Figure 1 shows the probability that a fraction of the subcarriers are X dB or more below the mean gain of the frequency response. From the plots one can see that it is likely that a large percentage of subcarriers are less than 3 dB below the mean of the SNR, but it is much less likely that a significant fraction of the subcarriers will 9 dB or more below the mean. This finding is significant as it shows the likelihood that significantly different power allocations will occur is small in a high SNR case. For example, if 10 % of the channels are 3 dB or more below the mean of the SNR, then as shown in (6) the average SNR of the remaining subcarriers has a lower bound of 1 (0 dB). Only when the fraction A is smaller, e.g. 9 dB as shown in Figure 1, do we get a lower bound on the average SNR of the remaining subcarriers of 7, (8.54) dB. The probability of this occurring in a diversity rich multipath channel is quite small.

III. Rate, SNR and Channel Estimation

Before examining how channel estimation and Doppler affect power-loading, it is important to have a baseline view of the system with perfect knowledge of the channel. To determine how much gain one can actually achieve in a wireless multipath channel, we simulated the system using uniform power-loading and water-filling over 10^4 channels with a RMS delay spread that was 5% of the OFDM symbol payload. The difference between the rate achievable using uniform power-loading and water-filling is shown in Figures 2 and 3. Figure 2 shows the mean and the 5^{th} percentile rates for water-filling and uniform power-loading. Figure 3 shows the mean difference between water-filling and uniform power distribution for a range of average SNRs. This result is in accordance with e.g. [4], [2], [3], where water-filling in wireless channels is shown to be effective only at low SNRs.

From the viewpoint of designing a system, the increase of water-filling over uniform power-loading at -7 dB SNR of 0.1 bits seems small. However, it is somewhat significant when one considers that the capacity at that point is 0.4 bits/sec/Hz for water-filling and 0.25 bits/sec/Hz for the uniform power distribution, an increase of over 50%. However, implementation issues such as synchronization and channel estimation in the presence of time-variability can limit this performance. The effect of synchronization on the rate, though important, is beyond the scope of this paper. Instead we focus on Doppler fading and its effect on channel estimation.

A. Channel estimation in a static channel

To isolate the effect of channel estimation on the rate, we first focus on the case of a static channel, i.e. no variability due to Doppler fading. In this case, channel estimation has two effects on the rate: the water-filling will have some mismatch due to channel estimation error and the bit error will be increased due the channel estimation error.

Given a minimum mean-squared error (MMSE) channel estimate $\hat{H}_k = H_k + \tilde{H}_k$, where the channel error $\tilde{H}_k$ is orthogonal to $\hat{H}_k$, a lower bound on the mutual information is [5], [6]

$$I_{lower} \quad = \quad \frac{1}{N}\sum_{k=0}^{N-1}\log_2(1+\frac{P_k|\hat{H}_k|^2}{\sigma^2+P_kE[|\tilde{H}_k|^2]}) \qquad (7)$$

where the power, P_k assigned to k^{th} subcarrier is based on the estimate of the channel, and $E[|\tilde{H}_k|^2]$ is the variance of the channel estimation error. The above expression indicates two sources of rate loss: increase in the noise variance and the possible mismatch of the power, P_k. Because the additional channel estimation error is a function both of the channel and the allocated power, an ill-chosen large power value can decrease the rate.

Channel estimation error depends on the number of pilots in the OFDM symbol and the number of taps in the channel.

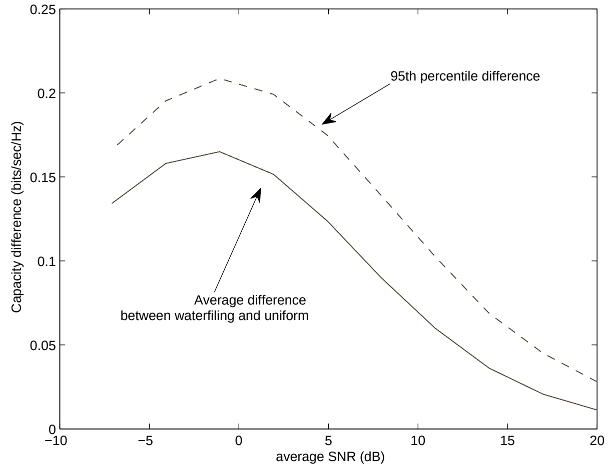

Fig. 3. The mean and 95^{th} percentile of the differences between capacity of a 128 subcarrier OFDM signal with power-loading and with uniform power distribution. The channel has an exponential RMS delay spread that is 5 % of the length of OFDM symbol with 16 sample-spaced taps. The mean and percentile value are taken over 10^4 channel instances. The mean capacity for water-filling ranges from .4 bits at -7 dB, to 1 bit at 0 dB. At 20 dB, the mean capacity for water-filling is approximately 6 bits .

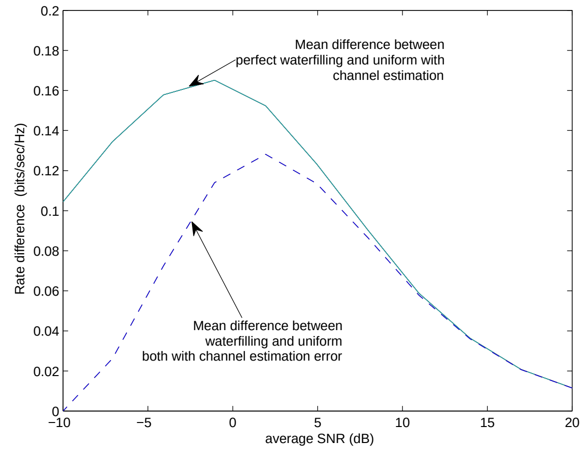

Fig. 4. The mean difference between water-filling and uniform power-loading of a 128 subcarrier OFDM signal with channel estimation. The number of pilots $N_p = N$; this is a static channel case. The channel has an exponential RMS delay spread that is 5 % of the length of OFDM symbol with 16 sample-spaced taps. The mean capacity for water-filling ranges from .4 bits at -7 dB, to 1 bit at 0 dB. At 20 dB, the mean capacity for water-filling is approximately 6 bits .

The more pilots, the better the estimate; the more taps in the channel, the wider the variability in the frequency domain and the harder it is to smooth it. This notion was captured in [7] where it was noted that the MMSE error is proportional to $\frac{\sigma^2 L}{N_p}$ where L is the number of taps and N_p is the number of pilots.

For the static case, one can assume a single OFDM pilot symbol with $N = N_p$ pilots. If the channel is not changing, there is no need to resend the pilot signal. Hence, the overhead for the pilots will not be considered in this case. Fig. 4 shows the difference in achievable rates with uniform loading and water-filling. In terms of channel estimation error, the static case is the best case. However, as shown in Figure 4, as the SNR decreases, the channel estimation will increase, which will in turn lead to mismatch in the power-loading.

B. Fading channel

Doppler fading has two consequences: outdated channel estimation and the need for more pilots, both of which reduce the achievable rate. We are interested in assessing the rate that is achievable in the presence of outdated feedback, i.e. in the presence of channel estimation errors. To this end, i.e. for assessing the *fundamental* limitations imposed by imperfect channel knowledge, we are focusing on power-loading (water-filling) as a capacity-maximizing technique. In practice, however, rate maximization can be accomplished by methods that take into account the error in the channel information available at the transmitter, e.g.[3]. However, if the estimate of the channel is so outdated as to be meaningless, then it is not clear that an intermediate method, though elegant, will be worthwhile. In addition, the gain between water-filling and uniform power-loading is so small at such low SNRs that it is not clear whether it is worth the complexity to implement intermediate methods.

In a fading channel, the effort needed for channel estimation depends not only on the required rate and the amount of diversity in the channel, but on the fading rate. In a static channel, one can neglect the overhead due to channel estimation. In contrast, with a fading channel one must assign enough pilots to minimize the channel estimation error, but not too many that the information rate is significantly reduced. As in [6], including the pilots in the rate estimation clarifies the push and pull of minimizing the impact of channel estimation errors while maximizing the rate.

The estimated rate for a set number of pilots N_p out of K OFDM symbols is

$$I_{lower} = \frac{1}{N}\sum_{k=0}^{N-1}\log_2\left(1+\frac{P_k|\hat{H}_k|^2}{\sigma^2+P_kE[|\tilde{H}_k|^2]}\right)\frac{KN-N_p}{N}$$

For example, if every OFDM symbol had $N_p = N/2$ pilots, the overall rate would be decreased by a factor of 2. As the number of pilots decreases, the rate can increase, but only as long as the channel estimation error does not significantly impair the performance.

Fading in the wireless channel can be modeled by noting that the correlation between two symbols T seconds apart with Doppler frequency f_d is $\rho = J_0(2\pi f_d T)$ [8]. We can consider the effect of fading on the achievable rate by breaking up the noisy signal as follows:

$$H_k x_k + n_k = (\rho(\hat{H}_k + \tilde{H}_k) + \sqrt{1-\rho^2}Z_k)x_k + n_k$$

where x_k is the data symbol value assigned to the k^{th} subcarrier, Z_k is a complex, zero-mean Gaussian random value with the same variance as H_k, and n_k is the Gaussian noise. Because Z_k is uncorrelated with H_k, we can model the extra term $\sqrt{1-\rho^2}Z_k$ as uncorrelated noise as well. In this case, the achievable rate is given by

$$I_{lower,\rho} = \frac{1}{N}\sum_{k=0}^{N-1}\log_2(1+\frac{P_k\rho^2|\hat{H}_k|^2}{\sigma^2+E_n}) \tag{8}$$

where $E_n = P_k(E[|\tilde{H}_k|^2) + \sqrt{1-\rho^2}P_kE[|Z_k|^2]$. P_k, the power on the k^{th} subcarrier is the average power of the symbol value x_k.

We simulated the effect of channel estimation on 10^4 channels with 5% RMS delay spread as earlier. Each OFDM symbol was approximately 72 μs long, on par with some cases of 3G LTE [9]. We assumed a 200 Hz Doppler frequency, f_d, (also compatible with a cellular system and a user traveling at around 60 mph) which corresponds to a relative Doppler from one OFDM symbol to the next of about 1%. A relative Doppler of 1% is acceptable from the viewpoint of ignoring intercarrier interference. We tried several pilot configurations and found the following to be the best for this scenario. One OFDM symbol where $(1/8)^{th}$ of the subcarriers were pilots was repeated every third symbol. This arrangement provided adequate channel estimation without severely reducing the information rate. The rate was averaged over three symbols: one with the pilots and two without. For the OFDM symbol with pilots, $\rho = 1$; for the two adjacent symbols, $\rho = J_0(2\pi 200\times 72\times 10^{-6}) = 0.998$.

Figure 5 shows the difference in performance between uniform power allocation and water-filling. The dashed line shows the improvement of water-filling over uniform power distribution without considering channel estimation errors at the receiver, but including the effect of outdated channel information at the transmitter due to Doppler fading. The solid line shows the effect of using imperfect channel knowledge using (8) Note that we are not considering the processing and propagation time required to send the information back to the transmitter. So these curves can be considered upper bounds on performance.

When channel estimation is considered, the available improvement due to water-filling is quite small. This is because with perfect channel estimation, a significant improvement only occurs at small SNRs. When the SNR is less than 0 dB, it is impractical to operate the system, let alone apply water-filling.

The loss in performance is due to outdated and imperfect channel estimation. The dashed and dotted lines in Figure 5 are in a sense upper and lower bounds on how well one can do. Though there are methods that can mitigate imperfection in the channel estimation, this indicates that there is very little wiggle-room for these methods.

IV. Conclusions

Water-filling is an attractive power-loading method when there are large variations in the frequency response of the channel and the channel is relatively static. However, for many wireless applications, the benefits of water-filling are minimal. If a channel has specific characteristics, e.g. a known low-pass filtering effect within the band, then water-filling makes sense. However, in channels where these variations are fleeting and not so large in magnitude, uniform power-loading, though unexciting, is the preferred method.

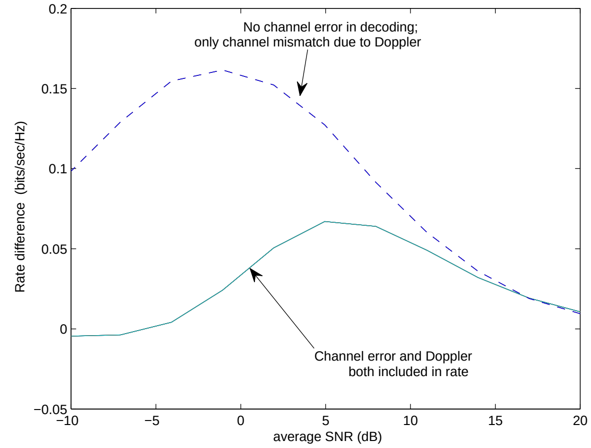

Fig. 5. Rate difference between uniform loading and water-filling with channel estimation errors in a fading channel. There is one OFDM symbol where one-eighth of the subcarriers are pilots surrounded by two symbols with no pilots. The dashed curve shows the difference in rate between uniform loading and water-filling where we only consider the mismatch due to Doppler for power-loading, but not the additional noise due to channel estimation error. The solid line in the plot shows the difference between perfect water-filling, but with the rate loss due to pilots and channel estimation error. The bottom solid curve shows the effect of imperfect water-filling.

References

[1] J. Cioffi, "A multicarrier primer,(online)," http://www.stanford.edu/group/cioffi/documents/multicarrier.pdf.

[2] D. Love and R. Heath, "OFDM power loading using limited feedback," *IEEE Transactions on Vehicular Technology*, pp. 1773–1780, September 2005.

[3] Y. Yao and G. B. Giannakis, "Rate-maximizing power allocation in OFDM based on partial channel knowledge," *IEEE Transactions on Wireless Communications*, pp. 1073–1082, May 2005.

[4] W. Yu and J. Cioffi, "Constant-power waterfilling: Performance bound and low-complexity implementation," *IEEE Transactions on Communications*, vol. 54, no. 1, pp. 23–28, January 2006.

[5] T. Yoo and A. Goldmsith, "Capacity and power allocation for fading mimo channels with channel estimation error," *IEEE Transactions on Information Theory*, pp. 2201–2213, May 2006.

[6] B. Hassibi and B. Hochwald, "How much training is needed in multiple-antenna wireless links?" *IEEE Transactions on Information Theory*, pp. 951–963, April 2003.

[7] M. Morelli and U. Mengali, "A comparison of pilot-aided channel estimation methods for OFDM systems," *IEEE Transactions on Signal Processing*, pp. 3065–3073, December 2001.

[8] W. Jakes, *Microwave Mobile Communications*. Wiley, 1974.

[9] E. Dahman, S. Parkvall, J. Sköld, and P. Beming, *3G Evolution: HSPA and LTE for Mobile Broadband*. Academic Press: Elsevier, 2008.